Pilot Surveys for Adaptive Informative Sampling

Stephanie Kemna[1], Oliver Kroemer[1,2] and Gaurav S. Sukhatme[1]

Abstract— **Adaptive sampling has been shown to be an effective method for modeling environmental fields, such as algae concentrations in the ocean. In adaptive sampling, a robot adapts its sampling trajectory based on data that it is collecting. This data is often aggregated into models, using techniques such as Gaussian Process (GP) regression. The (hyper-)parameters for these models need to be manually set or, ideally, estimated from data. For GP regression, hyperparameters are typically estimated using prior data. This paper addresses the case where initial hyperparameters need to be estimated, but no prior data is available. Without prior data or accurately pre-defined hyperparameters, adaptive sampling techniques may fail, because there is no good model to base path planning decisions on. One method of gathering data is to perform a pilot survey. This survey needs to select informative samples for initiating the model, but without having a model to determine where best to sample. In this work, we evaluate four pilot surveys, which use a softmax function on the distance between waypoints and previously sampled data for waypoint selection. Simulation results show that pilot surveys that maximize waypoint spread over randomization lead to more stable estimation of GP hyperparameters, and create accurate models more quickly.**

I. INTRODUCTION

Robotic approaches to environmental sampling and monitoring can speed up sampling times, provide more data, and enable frequent and persistent monitoring. One application of robotic use for environmental sampling is in spatial modeling, where the robot has to create a spatial model of a field, such as a temperature field or Chlorophyll distribution. One standard method of sampling a field is to run an off-line coverage approach over the area, such as a lawnmower survey, to then reconstruct the field after the survey. Off-line approaches are planned before execution, and the survey does not adapt to collected data. When the vehicle plans its own paths during execution, incorporating newly sampled data to select future waypoints, then the approach is known as an on-line approach, i.e. adaptive sampling. When paths are planned using information-theoretic metrics for optimization, we speak of informative sampling.

In adaptive informative sampling, the vehicle typically learns a model during execution. This model is used by the vehicle to decide where to sample next. Model hyperparameters, such as a kernel length scale, can be estimated from prior data, but this data may not be available. We are interested in the problem of adaptive informative sampling where no prior data is available. If the hyperparameters are estimated poorly, then the adaptive sampling approach will optimize for a wrong model and fail to collect good samples. To overcome the model initialization problem, one can run a pilot survey. A pilot survey is a short survey that is executed before running any actual survey and adaptive sampling approaches. However, it is not clear which pilot survey is the most suitable for collecting representative samples that lead to reasonable estimates of the hyperparameters.

In this work, we evaluate four pilot surveys for adaptive sampling. We call these 'integrated pilots' because the pilot has been integrated into the adaptive sampling mission, subtracted from the overall survey time. This is done to provide a fair comparison to non-adaptive sampling methods that do not require hyperparameter optimization before running the mission. The pilots we evaluate use a softmax distribution over the distance between the chosen waypoints and previously sampled locations, for waypoint selection. The four pilot surveys roughly correspond to a cross trajectory, random waypoints, and two intermediary solutions. As a baseline, we compare the performance of adaptive sampling to running a standard lawnmower survey, and find that in general the adaptive sampling surveys provide a good model more quickly. Simulation results further show that the pilot surveys that focus more on spreading out the initial waypoints, rather than randomizing waypoint locations, on average lead to better estimation of the hyperparameters.

II. RELATED WORK

Informative sampling was pioneered by Krause, Guestrin and Singh [1]–[3]. Low et al. [4] and Singh et al. [5] extended these informative sampling methods into adaptive informative sampling approaches. In all these papers, and many thereafter, Gaussian Process (GP) regression is used for modeling the spatial fields. The GP is fully specified by its prior mean and covariance function [6]. A common choice for the mean function is zero mean, and a common choice of covariance function is the isotropic squared exponential (SE) function, i.e. the Gaussian kernel. The kernel specifies the smoothness assumption between data points. While the GP model has no direct parameters, it does have hyperparameters: the kernel's parameters, as further explained in Section III. These hyperparameters can be estimated from data using, for example, maximum likelihood estimation [6].

For off-line methods, e.g. [2], [3], [7]–[9], the hyperparameters can be estimated after all of the data has been collected. For example, the works in [2], [3], [9] use a subset of all collected data for hyperparameter optimization. For on-line estimation of the model, i.e. in active learning and for

*This work was funded in part by the Office of Naval Research, ONR N000141410536

[1]Stephanie Kemna and Gaurav S. Sukhatme are with the Computer Science department, University of Southern California, SAL104, 941 Bloom Walk, Los Angeles, CA, USA {kemna,gaurav}@usc.edu
[2]Oliver Kroemer is with the Robotics Institute, Carnegie Mellon Institute, okroemer@andrew.cmu.edu

adaptive sampling, the hyperparameters should be estimated before or during execution.

Some previous works estimate the hyperparameters before running their adaptive sampling path planning: Hitz et al. [10] estimated hyperparameters based on prior data for the sampling region. Binney et al. [8] estimated hyperparameters by using data from an "initial run", what we call a pilot survey, which was executed prior to running any other sampling routines. However, their paper does not specify the shape or length of the pilot survey [8]. Other works estimate the hyperparameters during the sampling: Thompson et al. [11] estimated the hyperparameters initially by starting every adaptive mission with a $10-20\,s$ straight line drive, and then periodically re-estimated the hyperparameters during adaptive sampling. Their approach assumes that the data collected within this straight line drive is representative for the whole field and leads to reasonable hyperparameters. For some of the scenarios we are considering, this assumption does not hold, see for example Figure 3. Garg and Ayanian [12] estimated the hyperparameters during execution by keeping a belief distribution over the hyperparameters, initialized randomly, and using particle filtering for determining the hyperparameters at any time. This approach also allowed them to account for spatio-temporally varying fields. However, the random initialization could still lead to problems with model learning. To the best of our knowledge, there are no other works that explicitly use pilot surveys and/or investigate how best to design a pilot survey for model initialization.

We are interested in developing approaches that assume no prior data is available. This means that we cannot estimate the hyperparameters off-line, prior to running our sampling routines. If possible, the hyperparameters should be set to reasonable values based on expert knowledge, or knowledge about the area size or phenomena. An initial guess can decrease chances of estimating the hyperparameters incorrectly, e.g. by avoiding local maxima. We recommend the following steps for estimating the hyperparameters during an adaptive sampling survey:

1) Start adaptive sampling with an integrated pilot survey, to estimate the hyperparameters and initialize the model.
2) Re-estimate the hyperparameters every X minutes, to update the model based on new data.

We make the case for using an integrated pilot, where the pilot is an integral part of the adaptive sampling routine. This is recommended because adaptive sampling methods will not work well without a good estimation of the hyperparameters. The pilot should be integrated into the mission and subtracted from the overall mission time, rather than being a separate mission, when comparing to approaches that do not need a pilot survey, e.g. lawnmower surveys. The data from the pilot is kept in the model and used for the subsequent adaptive selection of waypoints for further sampling.

In this paper we evaluate four integrated pilots created using the softmax function, with a 'temperature' parameter τ, which is set to $\tau = \{1, 6, 30, 100\}$. This roughly corresponds to a cross trajectory over the sampling area, two intermediary solutions, and one pilot of randomly drawn waypoints, respectively. Simulation results show that the pilots with lower values of τ, which spread out the waypoints more, are on average more successful in obtaining a good estimate on the hyperparameters of the model.

III. THEORY

In this section, we discuss the theory behind our adaptive informative sampling approach, and the choice of pilots. This theory follows explanations in prior works, e.g. [13], [14]: The robot internally constructs a model of environmental phenomena, e.g. Chlorophyll abundance, using Gaussian Process regression. This model is used to pick waypoints with maximum entropy for further sampling. The contributions in this paper lie in the development of the pilot surveys and methods of handling hyperparameter estimation.

A. Gaussian Process Regression for Model Creation

Gaussian Process (GP) regression is a standard method for spatial field modeling [6]. The GP is specified by its prior mean and covariance function. We use a zero mean prior and for the covariance function a combination of an isotropic squared exponential (SE) kernel, and white noise covariance function. The SE covariance function is given by [6]:

$$k(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp\{-\frac{1}{2l^2}|\mathbf{x} - \mathbf{x}'|^2\} \tag{1}$$

where $\mathbf{x}$ and $\mathbf{x}'$ are two training sample locations, $\mathbf{x} \in \mathcal{X}$, $\mathcal{X} \subset \mathbb{R}^2$, σ_f^2 is the signal variance (or amplitude), and l is the kernel's length scale. σ_f^2 and l are hyperparameters. We combine the SE kernel with a white noise kernel, to better model the expected noise in the data. This kernel has one hyperparameter σ_n^2, noise variance. Following Low [13], we use a log Gaussian Process (ℓGP) to model the field: the vehicle takes the log of the measurements before incorporating them into the GP. This approach considers that biological data from fields with 'hotspots' tend to follow a log-normal distribution, due to large areas with low values and small areas with high values [15].

We follow previously introduced notation for the ℓGP model [13], [14]: Let Y_x denote an ℓGP, which is used to model the sensor value $y_{\mathbf{x}}$ at location $\mathbf{x} \in \mathcal{X}$. Let $Z_x = \ln Y_x$, denote a GP. Then we can create the ℓGP using GP regression by utilizing the fact that $z_{\mathbf{x}} = \ln y_{\mathbf{x}}$. The GP's posterior mean and variance, $\mu_{Z_x|d_i}$ and $\sigma^2_{Z_x|d_i}$ for sampled data d_i, are used to calculate the posterior mean and variance of the ℓGP [13] as:

$$\mu_{Y_x|d_i} = \exp\{\mu_{Z_x|d_i} + \sigma^2_{Z_x|d_i}/2\} \tag{2}$$

$$\sigma^2_{Y_x|d_i} = \mu^2_{Y_x|d_i}(\exp\{\sigma^2_{Z_x|d_i}\} - 1) \tag{3}$$

The vehicle thus creates an ℓGP model on board, while it is sampling the environment.

The hyperparameters of the ℓGP are estimated every $500\,s$. To reduce computational complexity for hyperparameter optimization, we downsample the data by a factor of four, keeping every fourth sample, for the optimization only. The model created on the vehicle contains all measurements. We test four pilot surveys that gather data for initial hyperparameter estimation, as further explained in section III-C. We use the *libgp* library [16] and the conjugate gradient method for hyperparameter optimization, with 100 iterations.

B. Path Planning

The vehicle uses the ℓGP model to decide where to sample next. We take a state-independent approach to path planning: based on the measurements made so far, we find the most informative sampling location across the whole space, which is then chosen as the next waypoint. The vehicle makes straight line movements between these waypoints via a standard waypoint behavior. We measure informativeness using the posterior entropy on the model, as derived in [13]:

$$\mathbb{H}\left[\mathbf{Y}_{x_{i+1}}|d_i\right] = \log\sqrt{2\pi e\sigma^2_{Z_{x_{i+1}|d_i}}} + \mu_{Z_{x_{i+1}|d_i}} \tag{4}$$

The next waypoint is thus an (unvisited) location with the highest posterior map entropy, anywhere in the sampling space, regardless of the distance from the current location. For the log-GP this sampling approach means that we maximize both for locations with high posterior variance, as well as locations with high expected sensor values. For example, if we are sampling for algae abundance, then high sensor values are measured in areas where algae blooms exist. These areas are likely to be most interesting to the data customer, e.g. biologists or oceanographers.

C. Integrated Pilot Surveys

To estimate the initial hyperparameters of the ℓGP model, we use an integrated pilot survey. We want to collect data that is representative of the field, i.e. including the full spectrum of data values, but we have to choose sampling locations without having prior data from the field. To optimize for coverage and spread over the area, one can use the area's corners as waypoints. However, this may leave large gaps between sampling locations, which could mean missing out on small hotspot areas. Alternatively, one can visit random locations in the area. However, the risk then is that only part of the field may be covered, and features in other parts of the field will be missed. Therefore, we want to balance between random sampling and maximizing the spread of waypoints.

To obtain a pilot with points spread across the area, we evaluate a utility function D. This function D is the minimum distance between waypoints and previously visited paths, which we want to maximize:

$$D(\mathbf{x}_i) = \min(d(\mathbf{x}_i, s_j)) \ , \ \forall s_j \in S \tag{5}$$

where $\mathbf{x}_i$ is the location of a waypoint candidate, d is the distance function for the Euclidean distance between possible waypoints and previous line segments, and s_j is a line segment from the set S of line segments between previously chosen waypoints.

To choose waypoints, we evaluate the probability of choosing a possible waypoint location using a softmax equation [17] on the utility function:

$$p(\mathbf{x}_i) = \frac{e^{D(\mathbf{x}_i)/\tau}}{\sum_j e^{D(\mathbf{x}_j)/\tau}} \tag{6}$$

where τ is the 'temperature' factor [17]. If τ is high, all actions become nearly equally probable and we have a uniform random sampling approach. For $\tau = 1$, it maximizes the minimum distance between potential waypoints and sampled paths, choosing waypoints at corners of the area. Paths are generated using a maximum length equal to the path length for $\tau = 1$. The number of waypoints is unconstrained.

Choosing Parameter τ: In order to use equation 6, we need to choose a value for the 'temperature', τ. Therefore we evaluated different metrics for the value function for the field. Remember that the value function considers the minimum distance between possible waypoints and previously sampled locations. Because the softmax sampling is probabilistic, we ran 500 trials for every value of τ. For every run, we calculated the minimum distance of all grid points to the chosen pilot path. Based on this, we calculated the expected minimum distance for each grid point location (expected min distances) over all runs, and we calculated the maximum minimum distance for each field (maximum of min distances). The metrics are, for each τ, the maximum (max) and average (avg) over the expected minimum distances, and the average of all maxima, as shown in Figure 1. Figure 1 also shows the average number of waypoints, which is not used for deciding τ. However, it shows a general trend of increasing number of waypoints with increasing values of τ. This means that waypoints often end up being closer together, and more waypoints can be chosen within the same path length budget.

We want to choose a value for τ that both reduces the chance of leaving large areas systematically unvisited for $\tau = 1$, or not covering the whole area for $\tau \rightarrow \inf$. Based on Figure 1 we chose $\tau = 6$, which corresponds to a minimum in the graphs for average over expected minimum distances (black dash-dot) and average over maxima (magenta dashed). We also chose $\tau = 30$, which corresponds to one of the

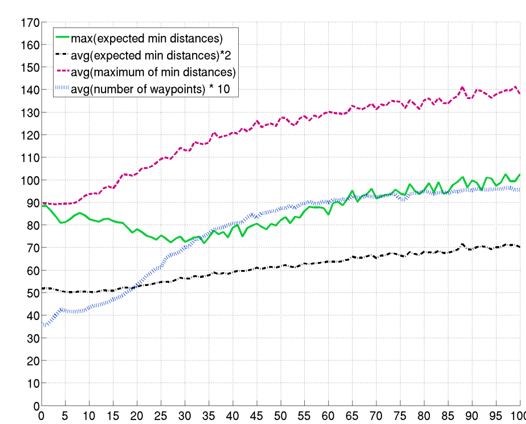

Fig. 1: Determining τ: The maximum (green solid) and average (black dash-dot) over the expected minimum distances, the average over all maxima (magenta dashed), and the average number of waypoints (blue dotted), for 500 runs per choice of τ. Y axis as per labels: (scaled) distance or (scaled) number waypoints. Best viewed in color.

minima in the graph for maximum over expected minimum distances (green solid). By choosing these values, we aim to balance between spreading out waypoint locations and randomizing sampling locations. A value of $\tau = 6$ is a conservative minima, and corresponds to pilots that are still very similar to the cross pilot. A value of $\tau = 30$ is closer to uniform random sampling, but would still try and spread out the waypoints to some extent. Figure 2 shows an example of waypoints and paths chosen by the vehicle for the different pilot survey cases; $\tau = 1$ (cross), $6, 30$, and 100 (random).

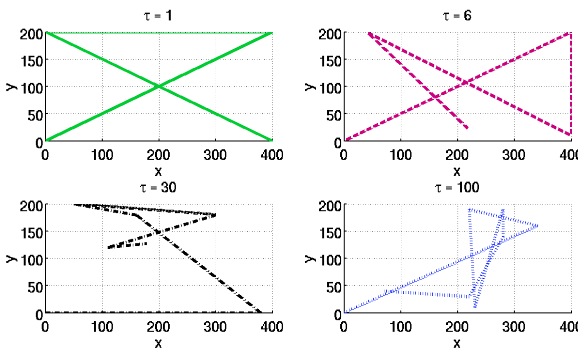

Fig. 2: Example trajectories for $\tau = 1$ (top left), $\tau = 6$ (top right), $\tau = 30$ (bottom left) and $\tau = 100$ (bottom right).

IV. EXPERIMENTAL SET-UP

For our evaluation of the integrated pilots, we ran simulation studies. We briefly explain the implementation and set-up details, for each type of experiment, in this section.

We simulate algae abundance for six scenarios, as shown in Figure 3. The first two scenarios were used in prior work [14], [18]. The other four scenarios were added to test performance in case there would be less pronounced blooms, or non-Gaussian shapes. Work is underway to create scenarios from data obtained in the field. All scenarios assume a rectangular sampling region, i.e. a 2-D grid space, of $400\text{x}200\,m$. Data are simulated at $10\,m$ grid spacing with output values ranging from $0 - 40 \mu g/L$, as a proxy for high Chlorophyll values. Random Gaussian noise is added, with a noise amplitude up to $10 - 20\%$ of the data range. The simulated vehicle takes samples with added Gaussian noise (signal variance $= 1.5$). The vehicle is not given information about the sensor noise.

We ran simulations for five different surveys:

- lawnmower survey (non-adaptive)
- adaptive survey, $\tau = 1$, 'cross pilot'
- adaptive survey, $\tau = 6$
- adaptive survey, $\tau = 30$
- adaptive survey, $\tau = 100$, 'random waypoints pilot'

The lawnmower survey is a composite of both a vertical and horizontal lawnmower over the survey area, with $20\,m$ track spacing, implemented via a waypoint behavior. The other pilot surveys use equations 5 and 6 for waypoint selection.

For every simulation run, the vehicle starts in the bottom left corner of the field. Variability in scenarios approximately covers for potential different starting locations, though future studies could investigate the effects of vehicle starting positions. For a real world scenario, the starting location may also be restricted by possible deployment locations.

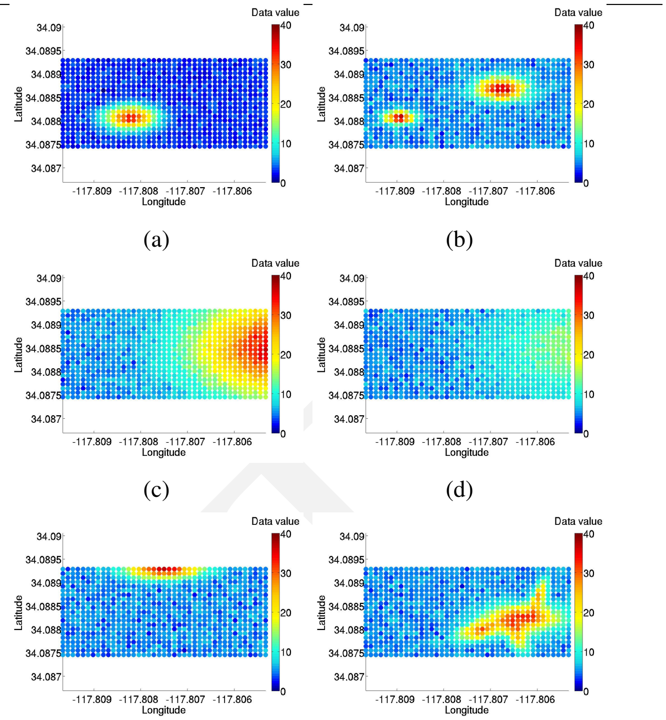

(a) (b) (c) (d) (e) (f)

Fig. 3: Six scenarios with generated data. Theoretically the data can represent anything, in this paper the 'data value' is Chlorophyll, $\mu g/L$.

We ran 30 simulations for each survey type for each scenario. Each simulation was time-limited to the duration of the lawnmower survey. Our focus is on improving initial sampling performance to provide an anytime prediction capability. During every simulation, we recorded the latest model created on the vehicle every 10 minutes ($600\,s$).

For our implementation we used the MOOS-IvP middelware [19], which incorporates a behavior-based autonomy. We use standard behaviors such as; waypoint, loiter, and constant depth. For the simulations, we use a simple vehicle dynamics model with PID control to simulate an autonomous underwater vehicle (AUV), and the biological data (Chlorophyll) simulator described at the start of this section.

V. RESULTS

We evaluate performance in terms of Root Mean Squared Error (RMSE) between the model and the ground truth, and in terms of the estimated hyperparameters versus their ground truth values.

A. Root Mean Squared Error

We ran 30 simulations for each survey type for each scenario. Figure 4 shows the simulation results in terms of the RMSE between the vehicle's model and the generated data, i.e. the ground truth (Figure 3). Each subfigure shows the results for a single scenario, for all five surveys. RMSE is

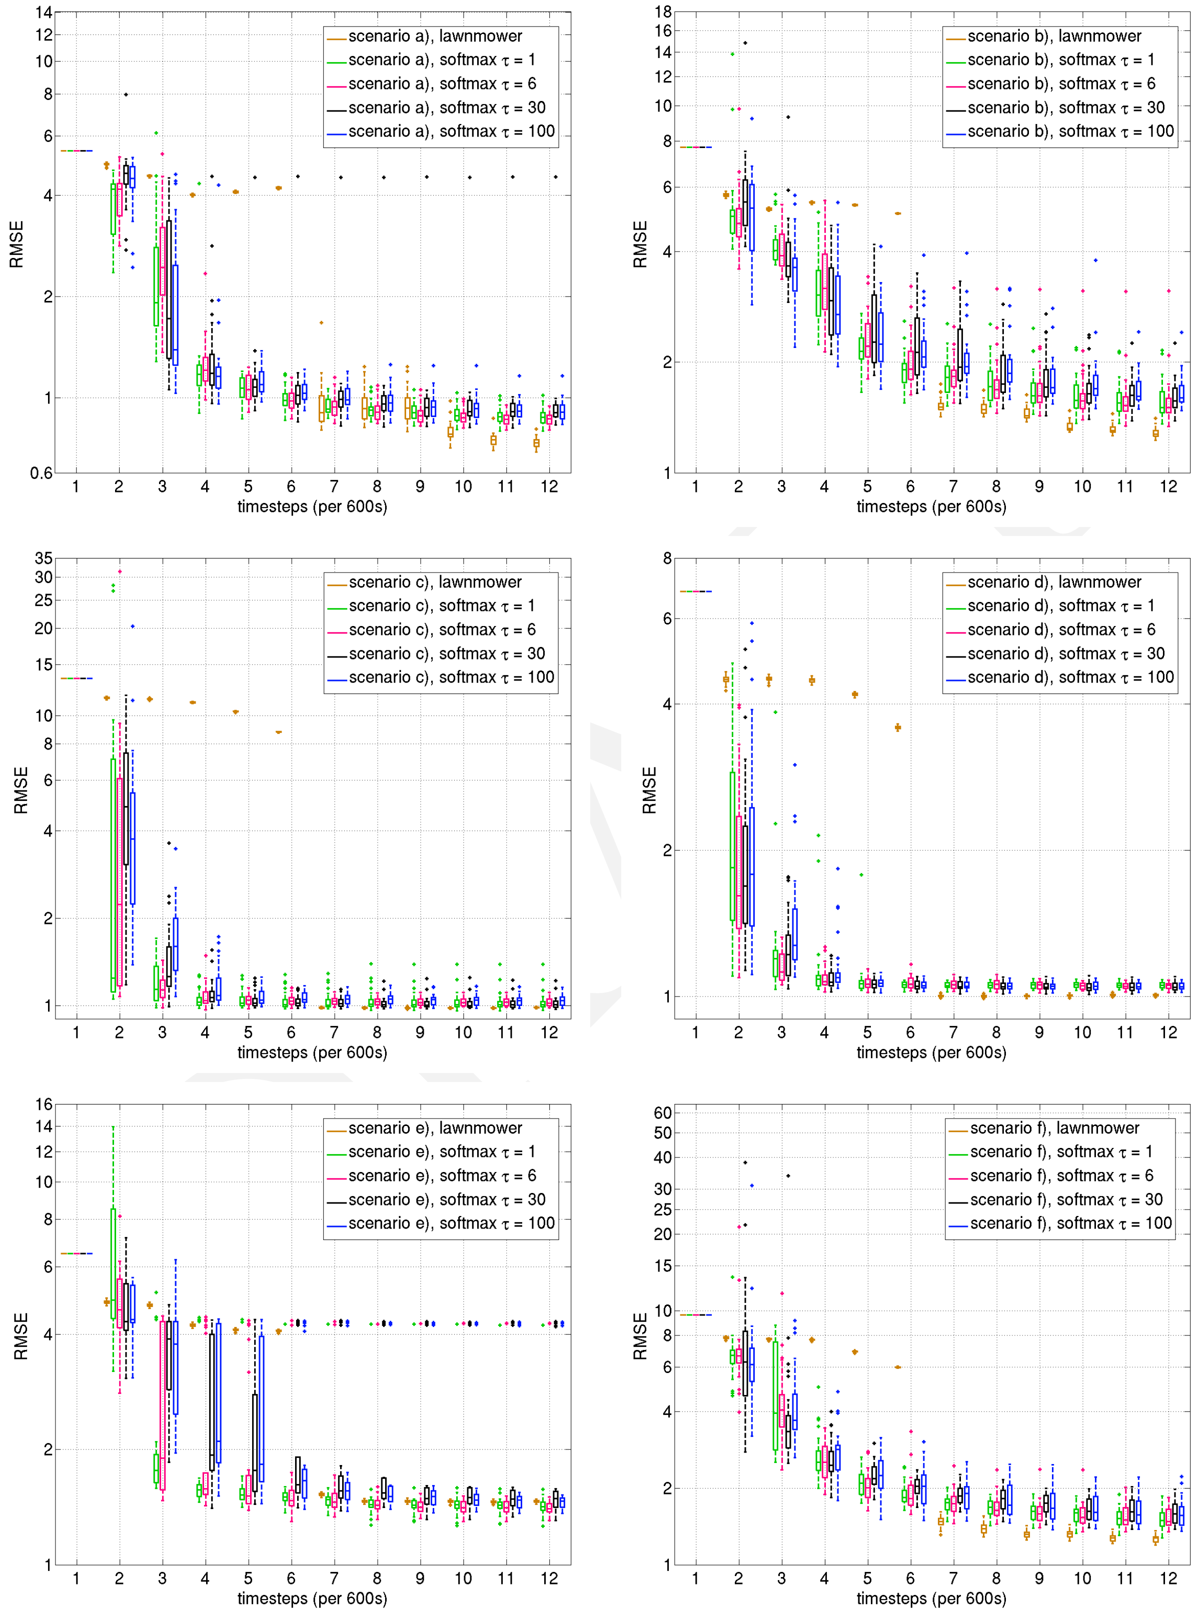

Fig. 4: Boxplots (median, 25th and 75th percentiles) on RMSE for scenario a) - f), 30 simulations per survey. Crosses are outliers. Evaluations are done based on models that were saved every 10 minutes (600 s).

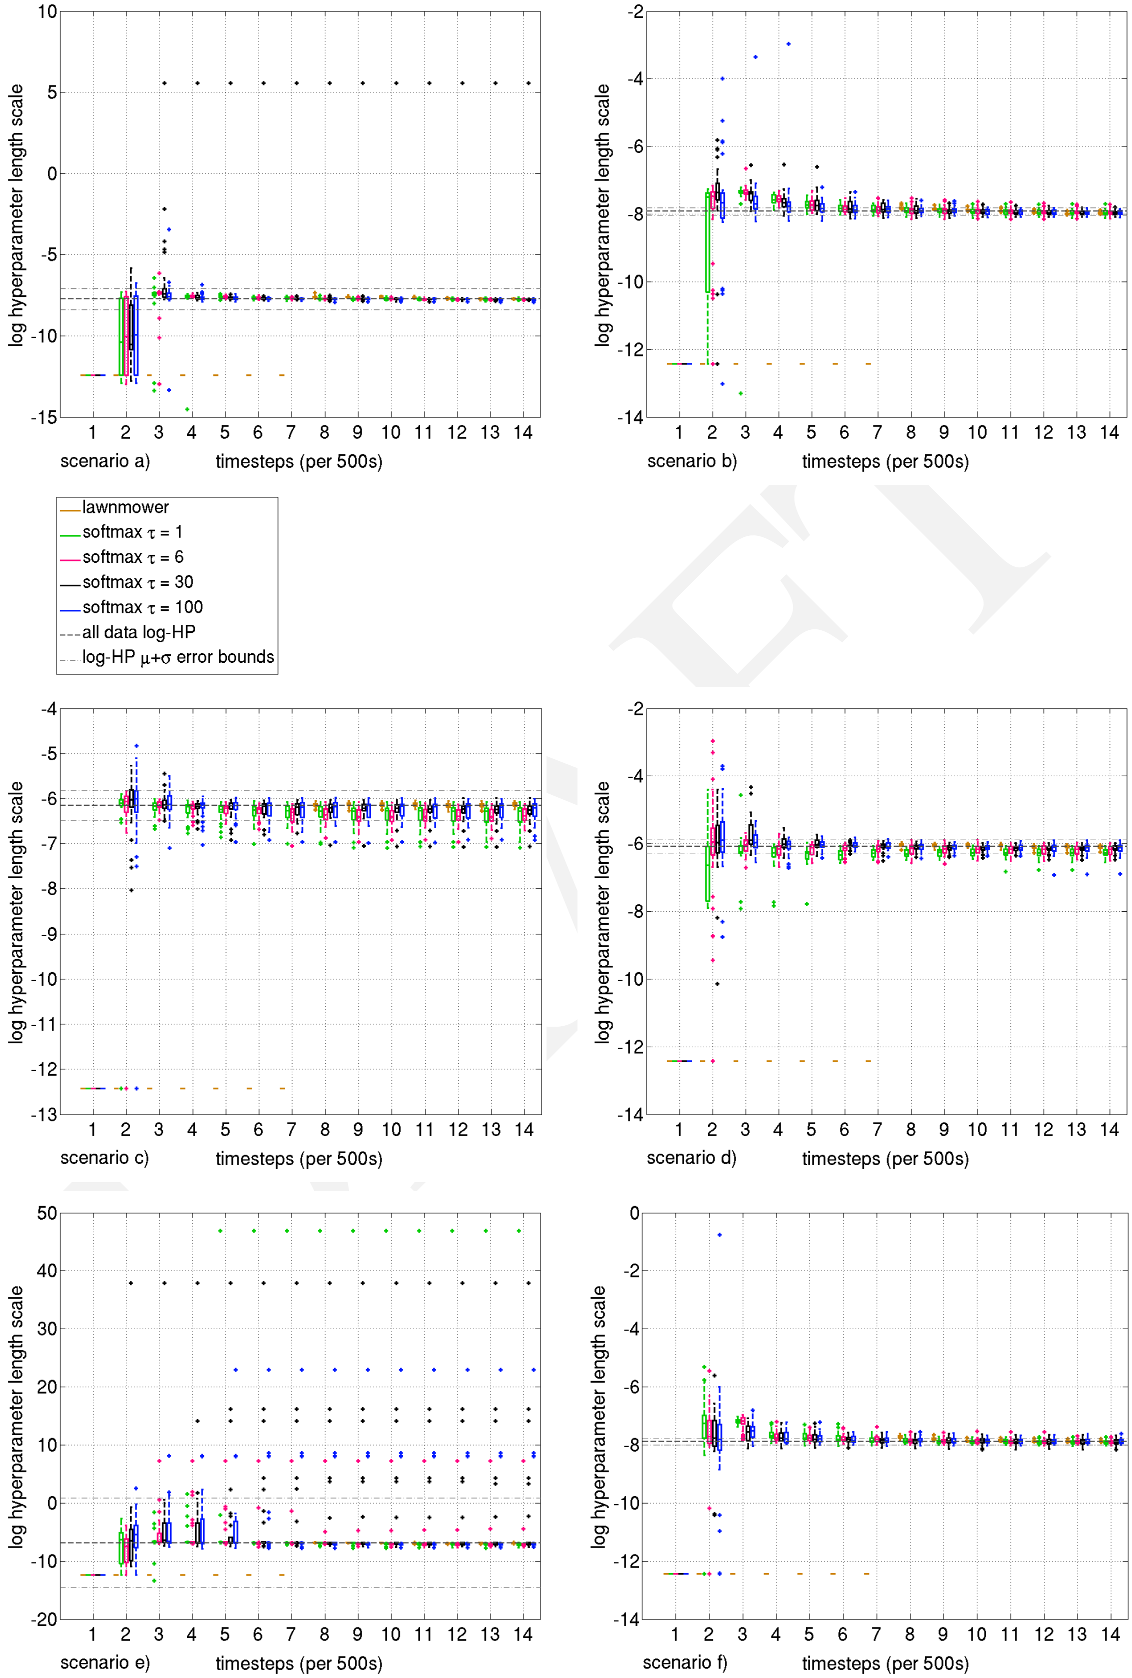

Fig. 5: Boxplots (median, 25th and 75th percentiles) on estimated length scale hyperparameter (HP), for scenarios a-f, 30 simulations per approach. Crosses are outliers. Evaluations are done based on models that were saved every 500 seconds.

Survey type	$\mathbf{t = 2}$	$\mathbf{t = 3}$	$\mathbf{t = 4}$	$\mathbf{t = 12}$
lawnmower	0.67	1.17	0.33	0.33
$\tau = \mathbf{1}$	1.67	2	2.67	2
$\tau = \mathbf{6}$	2	0.67	1.83	0.67
$\tau = \mathbf{30}$	2	2.17	1.67	1.83
$\tau = \mathbf{100}$	2.83	2.33	2.5	1.17

TABLE I: Average number of RMSE outliers per time step, averaged over all scenarios, for every type of survey.

plotted against time steps of $600\,s$, where time step $t = 1$ corresponds to the start of the survey. We show boxplots, rather than averages with standard deviations, to better visualize outliers and general trends. In general, outliers correspond to the runs when no good model was created, and for scenario e) the runs that did not create a good model created a bimodal RMSE performance graph. Our main focus in the evaluation is on the early stages of model creation, to evaluate the pilot surveys, and not the other path planning method.

Figure 4 shows that for most scenarios, the RMSE quickly drops for adaptive sampling. Scenarios b) and f) see a more gradual decline for all methods. For most scenarios, differences between the pilot surveys' performances are not statistically significant, given that all boxplots overlap. We briefly note: For scenarios a) - d) and f), all adaptive sampling surveys get a significantly better model more quickly than the lawnmower survey. The lawnmower performance in most cases drastically improves after 5 time steps at $t = 6$, which approximately corresponds to when the vehicle is executing the second lawnmower pattern (horizontal) and revisits areas. In terms of pilot survey performance, we see: For scenario b), $\tau = 1$ and $\tau = 6$ do better than the softmax surveys with high values of τ. For scenario c), there is a wide spread on the initial estimate. Values of $\tau = \{1, 6\}$ do a little better initially. For scenario e), softmax survey $\tau = 1$ clearly outperforms high τ values from timestep $t = 3 - 5$.

RMSE outlier analysis: From Figure 4 it is clear that: For scenario a), the softmax surveys with high values of τ have more outliers. For scenario d), $\tau = \{1, 100\}$ initially have more outliers. For scenario e), high τ surveys have more outliers. Table I shows the average number of outliers, for each survey type, averaged over all scenarios, for four time steps: $2, 3, 4$ and the final, 12. Time step $t = 2$ corresponds to the predictions that are saved after the first hyperparameter estimation, and time step $t = 12$ corresponds to predictions saved after the final hyperparameter estimation. The number of outliers is initially highest for $\tau = 100$ and lowest for $\tau = 1$. From $t = 2$ to $t = 3$ the number of outliers drops the most for $\tau = 6$, which also ends with the fewest outliers. This indicates more consistent performance across simulations and across scenarios.

B. Hyperparameters

We further investigate performance in terms of the estimation of the hyperparameters, which mainly determine the quality of the model, in particular the kernel's length scale. As previously mentioned, when the hyperparameters are misestimated, the adaptive sampling approach will not work as well. Therefore, we want to make sure that we choose a good initial sampling approach to start the adaptive sampling survey off with. Data that are collected initially should be representative for the field, and lead to good estimation of hyperparameters.

Figure 5 shows the estimated log length scale over time, stored after every hyperparameter optimization (every 500 seconds). As previously mentioned, $t = 1$ corresponds to the start of the mission. Note that all values are distances in longitude-latitude degrees and therefore quite small. The ground truth values of the hyperparameters are calculated from all data in the simulated data files, using the GPML Matlab libraries [20]. These are indicated as dashed grey lines. The dash-dot light-grey lines are error bounds. These error bounds are determined by averaging the final error across all simulations and all scenarios, and taking the average error plus one standard deviation. If the log length scale value of a single simulation is more (or less) than the true length scale plus (or minus) the error bound, then we consider it a misestimation of the hyperparameters.

The results show that the poor performance for scenario e) corresponds to badly estimated length scales. For scenario a), where $\tau = 30$ has one bad run, with a corresponding outlier in the log length scale plot. For the lawnmower surveys, the initial poor estimations are also due to incorrect estimation of the length scale, which does not change from its initial value until $t = 7$. The performance between the different pilot surveys is quite similar. The hyperparameters start to settle down after the second hyperparameter estimation, at timestep $t = 3$. For scenario b) and e), we see that with $\tau = \{30, 100\}$, it takes longer for the log length scale to get close to the correct value. For scenario c), we see that for all surveys the log length scale is slightly underestimated, which corresponds to a length scale that is approximately 20-40 meters shorter than the actual value.

Length scale outlier analysis: From Figure 5 it is clear that: For scenario a), there is one outlier for $\tau = 30$. For scenario e), there are many outliers. While there are some outliers for $\tau = \{1, 6\}$, these are fewer than for $\tau = \{30, 100\}$, meaning that overall they were more successful. For scenario d), $\tau = \{1, 100\}$ have more outliers at the end. For scenario f), $\tau = 30$ initially has more outliers. Table II shows the average number of outliers for the length scale log-hyperparameter (lHP) boxplots, for each survey type, averaged over all scenarios, for four time steps: $2, 3, 4$ and the final, 14. Note that, in comparison to the RMSE plots, because hyperparameter optimization happens every 500 seconds rather than every 600, there are 2 more time steps. Time step $t = 2$ corresponds to the first hyperparameter estimation, and time step $t = 14$ corresponds to the final hyperparameter estimation. Note that while there are no outliers for the lawnmower performance during $t = 2$ to $t = 4$, the length scale is consistently misestimated during this time. We see that after the first hyperparameter estimation, $t = 2$, the number of outliers reduces on average for $\tau = \{6, 30\}$. This downwards trend continues for $\tau = 6$.

Survey type	$\mathbf{t=2}$	$\mathbf{t=3}$	$\mathbf{t=4}$	$\mathbf{t=14}$
lawnmower	0	0	0	1.17
$\tau = \mathbf{1}$	1.33	3	2.33	0.83
$\tau = \mathbf{6}$	3.17	2.67	1.83	0.67
$\tau = \mathbf{30}$	2.83	2.17	1.33	2.33
$\tau = \mathbf{100}$	4	1.5	2.33	2

TABLE II: Average number of lHP outliers per time step, averaged over all scenarios, for every type of survey.

VI. DISCUSSION & FUTURE WORK

The simulation results confirmed that adaptive sampling can improve modeling performance during the early stages of model creation. We investigated four pilot surveys for obtaining representative data for model initialization. Results confirmed that an initial poor estimation of the hyperparameters can be detrimental to the overall modeling performance. We found that for lower values of τ, we obtained better estimated hyperparameters on average and lower model error early on, than for higher values. This suggests that spreading out waypoints is to be preferred over randomly picking waypoints. We recommend using the softmax-based waypoint selection method for pilots, using $\tau = 6$, which on average provided the best and most stable performance.

For the lawnmower surveys, we saw the RMSE greatly reduce after 6 time steps. This approximately corresponded to the second pass over the area, leading to more samples through the blooms. From Figure 5 we concluded that the initial bad performance was due to bad estimates of the kernel length scale. This highlights the need for quickly collecting representative samples when the actual hyperparameters are not known. Furthermore, if anytime prediction capability is desired for an off-line survey approach, it may be necessary to also run a pilot survey before the main survey.

For future work, there are many possible avenues for further investigations. For one, we did not find theoretical proofs yet to guarantee an initial good model creation, when no prior information is available. For the softmax approach, it is impossible to set the temperature based on the scenario, if the scenario is not known. However, if one allows, for example, for some expert knowledge to guide sampling, this may become feasible. Therefore we recommend investigating approaches for best incorporating expert knowledge. Furthermore, we have chosen a utility function that maximizes minimum distance between sampled points, optimizing for coverage. Again, this is based on the fact that no prior data is available. It would be interesting to explore other utility functions. Finally, in our evaluation of pilots we have chosen to use a set length for the pilot, equal to the distance of a cross trajectory over the area. It would be interesting to evaluate whether representative data could be obtained in a shorter amount of time, i.e. a shorter path.

VII. CONCLUSIONS

In this paper we evaluated four pilot surveys for adaptive informative sampling. These pilots are essential at the start of any adaptive mission where no prior data is available, to collect representative data for hyperparameter estimation. An initial bad estimate on hyperparameters can lead a vehicle to construct a bad model and thus collect non-informative samples. One method of running a pilot is to travel between all corner locations of a sampling area. We explored whether adding a degree of randomness to this approach improves its performance. We evaluated a softmax function for waypoint selection, which balances between spreading out waypoints and uniform random sampling. We showed that, for the given value function, low values of τ, e.g. $\tau = 6$ with spread out waypoints, lead to the most stable performance in terms of initial hyperparameter estimation.

REFERENCES

[1] A. Krause and C. Guestrin, “Nonmyopic active learning of Gaussian processes: an exploration-exploitation approach,” in *International Conference on Machine Learning (ICML)*, 2007, pp. 449–456.

[2] A. Singh, A. Krause, C. Guestrin, W. Kaiser, and M. Batalin, “Efficient Planning of Informative Paths for Multiple Robots,” in *International Joint Conference on Artificial Intelligence (IJCAI)*, 2007.

[3] A. Krause, A. Singh, and C. Guestrin, “Near-Optimal Sensor Placements in Gaussian Processes-Theory, Efficient Algorithms and Empirical Studies,” *The Journal of Machine Learning Research*, vol. 9, no. May, pp. 235–284, 2008.

[4] K. H. Low, J. M. Dolan, and P. Khosla, “Adaptive Multi-Robot Wide-Area Exploration and Mapping,” in *Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2008, pp. 23–30.

[5] A. Singh, A. Krause, and W. J. Kaiser, “Nonmyopic adaptive informative path planning for multiple robots,” in *IJCAI*, 2009, pp. 1843–1850.

[6] C. Rasmussen and C. Williams, *Gaussian Processes for Machine Learning*. MIT press, 2006.

[7] T. Patten, R. Fitch, and S. Sukkarieh, “Large-Scale Near-Optimal Decentralised Information Gathering with Multiple Mobile Robots,” in *Proceedings of Australasian Conference on Robotics and Automation*, December 2013.

[8] J. Binney, A. Krause, and G. S. Sukhatme, “Optimizing waypoints for monitoring spatiotemporal phenomena,” *The International Journal of Robotics Research (IJRR)*, vol. 32, no. 8, pp. 873–888, 2013.

[9] G. A. Hollinger and G. S. Sukhatme, “Sampling-based robotic information gathering algorithms,” *IJRR*, vol. 33, pp. 1271–1287, 2014.

[10] G. Hitz, A. Gotovos, F. Pomerleau, M.-E. Garneau, C. Pradalier, A. Krause, and R. Y. Siegwart, “Fully Autonomous Focused Exploration for Robotic Environmental Monitoring,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2014.

[11] D. R. Thompson, D. S. Wettergreen, and F. J. C. Peralta, “Autonomous Science during Large-Scale Robotic Survey,” *Journal of Field Robotics*, vol. 28, no. 4, pp. 542–564, 2011.

[12] S. Garg and N. Ayanian, “Persistent Monitoring of Stochastic Spatio-temporal Phenomena with a Small Team of Robots,” in *Robotics: Science and Systems*, 2014.

[13] K. H. Low, J. M. Dolan, J. Schneider, and A. Elfes, “Multi-Robot Adaptive Exploration and Mapping for Environmental Sensing Applications,” PhD thesis, CMU, 2009.

[14] S. Kemna, J. G. Rogers III, 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 *ICRA*, May 2017.

[15] E. L. Crow and K. Shimizu, *Lognormal Distributions: theory and applications*. Marcel Dekker, Inc., 1988.

[16] M. Blum, “libgp: Gaussian process library for machine learning,” https://github.com/mblum/libgp, 2016, [Online; accessed Nov 2016].

[17] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. The MIT press, Cambridge, MA, USA, 1998.

[18] S. Kemna, D. A. Caron, and G. S. Sukhatme, “Adaptive Informative Sampling with Autonomous Underwater Vehicles: Acoustic versus Surface Communications,” in *Oceans*, September 2016.

[19] M. R. Benjamin, P. Newman, J. J. Leonard, and H. Schmidt, “An Overview of MOOS-IvP and a Users Guide to the IvP Helm Autonomy Software,” MIT, Tech. Rep. MIT-CSAIL-TR-2010-041, 2010.

[20] C. Rasmussen, “Documentation for GPML Matlab Code version 4.0,” http://www.gaussianprocess.org/gpml/code/matlab/doc/, 2017, [Online; accessed Aug 2017].