

Comparison of Path Planning Approaches for Harmful Algal Bloom Monitoring

Chris Denniston*, Aravind Kumaraguru* and Gaurav S. Sukhatme

Abstract—Harmful algal blooms are an important phenomenon to monitor because of their potential for damaging ecosystems that depend on the affected body of water. A useful instrument to monitor these are Autonomous Underwater Vehicles with attached sensors. In this work, we explore methods to plan vehicle routes that maximize the coverage of sensor readings in algal blooms, as opposed to the traditionally used lawnmower coverage surveys. This is done in order to save time and cover blooms with large spatial extents. To this end, we focus our search on areas of high concentration, ignoring the monitored quality of low concentration areas. We also look to monitor areas of high rates of change, such as the continuous boundary of an algal bloom, to inform marine ecologists of the extent and boundary conditions of a bloom. We use Gaussian process regression to build a map of the concentration of chlorophyll and greedily plan paths using a supplied metric. Our simulation results show that a simple metric can be used to monitor the concentration of chlorophyll as well as the rate of change of the concentration. The results are validated on real and synthetic datasets.

I. INTRODUCTION

Harmful algal blooms (HABs) have become an important area of research. There is growing scientific consensus that algal blooms are worsening and require more advanced methodologies and technologies to study the composition and the surrounding water quality [1]. Autonomous Underwater Vehicles (AUVs) have been shown to be useful in studying harmful algal blooms as they enable large spatial sensing of blooms because of their ability to autonomously sense characteristics such as chlorophyll, salinity and temperature [2]. A key point in good monitoring of algal blooms is adequate coverage of spatial features. In taking these measurements, it is important to model areas of high concentration more accurately than areas of low concentration, even though areas of low concentration may be more abundant.

Traditionally, a lawnmower path or some full coverage path is planned because of its guarantee to cover the entire sensed area [3]. In this work, we aim to examine harmful algal bloom sensing using informative path planning (IPP), instead of the traditional lawnmower pattern or transects. The lawnmower coverage method is static and does not adapt to in-water sensor readings, potentially wasting time in areas of low interest. In IPP, a model of the environment is created and information theoretic measures are used to determine where the next sample should be taken. This method leads to better




Fig. 1: From left to right: (Scale is in meters) (1) Original path taken by the Ecomapper AUV (2) Gaussian process predictions built from original samples (Samples below mean removed) (3) Gaussian process variance built from original samples (Samples below mean removed)




Fig. 2: From left to right: (Scale is in meters) (1) Generated Field (2) Gaussian process predictions built from original samples (Samples below mean removed) (3) Gaussian process variance built from original samples (Samples below mean removed)

results in situations where full coverage is not needed or is impossible.

In this work, we evaluate different metrics for IPP for sampling of underwater HABs using a non-holonomic AUV.

We compare metrics for 3D IPP on a real data set taken from the chlorophyll sensor of a YSI Ecomapper AUV. One model can be seen in Fig. 1 which shows the original data set taken by the Ecomapper AUV as well as the mean and variance of the model that was fit to the dataset. We also use generated models that are mixtures of Gaussians, which can be seen in Fig. 2. In both figures we show evenly spaced samples of the mean and variance but remove values that are below the mean so the structure can more easily be seen. In both, color intensity represents chlorophyll concentration.

We show the effectiveness of these methods by comparing

*Equal Contribution

This work was supported by USDA/NIFA award 2017-67007-26154, ONR. Chris Denniston, Aravind Kumaraguru, and Gaurav S. Sukhatme are with University of Southern California, Los Angeles, USA {cdennist, akumarag, gaurav}@usc.edu

how well the samples can be used to create a model of the chlorophyll concentrations. In this work, we do not segment blooms into harmful and non-harmful components, but instead aim to build a model of chlorophyll concentration to enable further observations of the blooms to be conducted by specialists. We also aim to show that extensions to known metrics which take into account the distance to travel to a candidate point can outperform metrics that do not take this into account.

In the future, we plan to extend these methods to field trials to study their effectiveness in bodies of water that may have blooms that are too large for an AUV to accurately sample using non-adaptive methods.

Our contributions are as follows:

- Present a metric that jointly maximizes the modeling accuracy of high concentration areas as well as high change areas in algal blooms
- Show the importance of making greedy algorithms using this metric aware of traveled distance
- Validate these metrics on real world and simulated datasets

II. RELATED WORK

Marchant and Ramos [4] present an algorithm to maximize a weighted metric using a framework of Bayesian optimization. In that work, a continuous path is planned by maximizing the integral of the metric along the path through the predicted field. Marchant and Ramos also present an alternative algorithm to maximize this metric which takes into account spatio-temporal links using Monte-Carlo tree search [5]. Both of these works differ from ours in that they are not greedy with a single point, but can be costly to evaluate.

Gaussian process Bayesian optimization, used in this work and the works by Marchant and Ramos, has been extensively studied in objective maximization. Srinivas, Krause, Kakade and Seeger [6] show that the GP-UCB metric used in Marchant and Ramos can be bound by cumulative regret in terms of maximal information gain. They also validate their findings on a sensor placement problem.

Guestrin and Krause [7] describe a solution to the problem of optimal sensor placement. Their work discusses the failures of entropy as a metric and explains why mutual information is a better criterion. They also show that greedily maximizing mutual information in a Gaussian process is a good approximation of the optimal placement. The problem of optimal sensor placement is related to that of IPP, but there is no cost for moving between positions or samples between positions.

Binney, Krause and Sukhatmes [8] investigate the problem of planning paths for submodular objective function. Their work focuses on creating a good approximation of an objective evaluated over a graph of points while sampling along the edge. This work differs from ours because it only focuses on maximizing the quality of observation using metrics like mutual information.

Hollinger and Sukhatme [9] investigate path planning for submodular objective function, such as information gain. In that work, the authors propose an extension to RRT* [10]

which allows for multi-step path planning to objective maximization.

Although multi-step path planning is outside the scope of this work, it is important to note that extensions to this work would include more optimal path planning than a single point greedy approach.

Tracking fronts with AUVs is largely related to the problem of tracking the rate of change for chlorophyll since both methods use the gradient of the sensed quality. Ocean front tracking generally uses the temperature gradient and specific algorithms have been designed to only track the front [11].

III. GAUSSIAN PROCESS REGRESSION FOR ENVIRONMENT MODELING

Gaussian processes are widely used modeling tools for aquatic robotics because of their non-parametric and continuous representation of the sensed quantity. Gaussian processes approximate some unknown function f from its known outputs by approximating any points \mathbf{x} in the input space by Gaussians, based on measurements y taken at that input location and nearby measurements [12]. How much other measurements affect the estimate at a certain input location is determined by the kernel $k(\cdot, \cdot)$, that is used for the regression.

Function values y_* at any input location \mathbf{x}_* are approximated by a Gaussian distribution:

$$y_*|y \sim \mathcal{N}(K_*K^{-1}y, K_{**} - K_*K^{-1}K_*^T) \quad (1)$$

where y is training output, y_* is test output, K is $k(\mathbf{x}, \mathbf{x})$, K_* is $k(\mathbf{x}, \mathbf{x}_*)$, K_{**} is $k(\mathbf{x}_*, \mathbf{x}_*)$, \mathbf{x} is training input, and \mathbf{x}_* is test input.

The kernel computes how much the measured value influences predictions on neighboring and far away points. In this work, the radial bases function (RBF, also called squared exponential) kernel is used, which is a stationary, isotropic kernel.

In this work, $\mu(x)$ is used as the mean and $\sigma(x)$ as the standard deviation of the Gaussian process at the point x . For our experiments, a Gaussian process is fit to the underlying data, representing a model of the environment. In Figs. 1 and 2 the mean and standard deviation of the ground truth Gaussian process is shown.

IV. MOTION MODEL

In this work, we use the fixed wing Dubins airplane [13] as the motion model, which expands on the original Dubins curve [14] by adding control over altitude. Dubins airplane calculates time optimal paths for two points with desired planar angles and assumes that the system has independent control over planar turning angles and altitude angles. These constraints loosely match the motion of an AUV because of the limited turn rate and climb rate experienced by an AUV. The fixed wing Dubins airplane calculates paths using an airplane motion model. This work uses Dubins airplane paths to calculate both the route taken between two points and the distance that would need to be traveled to get to that point.

V. METRICS FOR INFORMATIVE PATH PLANNING

In IPP, a measure is used to determine where to sample next. Generally, a position is chosen which maximizes some measure and a path is planned to that point. In this work, different metrics are explored based on the predictive mean and variance as well as the first order methods which rely on the gradient of the predictive mean.

Conditional Entropy: The conditional entropy is defined as

$$\frac{1}{2} \log(2\pi e \sigma^2(x)) \quad (2)$$

This has been used as the objective for sensor placement [7] and AUV path planning [15].

Conditional entropy has the noted problem of tending to choose locations around the perimeter of the sensed location and not in the places which provide the best information about the environment as a whole [4].

When used as a metric, conditional entropy does not prioritize areas of high concentration and only focuses on the quality of the observation.

Mutual Information: Mutual information builds on the idea of conditional entropy by considering the prediction quality of the sensed locations. Guestrin, Krause and Singh [7] describe a greedy algorithm for sensor placement using mutual information and Gaussian Processes. The algorithm creates an approximation to the optimal sensor placement problem which is related to the adaptive sensing problem but without the consideration for traveled distance.

Mutual information for a Gaussian processes is defined by:

$$\frac{\sigma^2(x) - \Sigma_{xA} \Sigma_{AA}^{-1} \Sigma_{Ax}}{\sigma^2(x) - \Sigma_{xA} \Sigma_{\bar{A}\bar{A}}^{-1} \Sigma_{\bar{A}x}} \quad (3)$$

A is the points already sampled, \bar{A} is the points not currently sampled.

Like conditional entropy, mutual information attempts to maximize the quality of the observation regardless of their concentration values.

Gaussian Process Upper Confidence Bound: Gaussian Process Upper Confidence Bound treats the sampling problem as a multi-armed bandit optimization problem [6]. This formulation has also been used in mobile robotics sampling by Marchant and Ramos [4] who have demonstrated its effectiveness on continuous fields such as light luminosity.

$$\mu(x) + \beta^{1/2} \sigma(x) \quad (4)$$

$\beta^{1/2}$ provides a parameter to tune between exploration and exploitation and makes the tradeoff explicit.

Marchant and Ramos propose a solution that finds an optimal curve to a position by considering all points along the curve, however the solution presented in this paper only finds an optimal position and then plans the quickest Dubins airplane curve.

We extend the formulation by adding a discounting distance term to account for the fact that the AUV must physically move there.

$$\frac{\mu(x) + \beta^{1/2} \sigma(x)}{\gamma \text{distance}(x^*, x)} \quad (5)$$

x^* is the current position and distance is the total length of the Dubins airplane curve [13]. γ provides a parameter to control how much the distance discount should affect the metric. This discount is useful with a greedy planner because it penalizes traveling long distances and, instead, prioritizes modeling the area the AUV is currently in.

First Order UCB: The previous UCB metrics can be extended to take into account the areas with high change in concentration. This can be useful when not only areas with high concentration need to be modeled, but also areas on the boundary of what is being monitored. Because chlorophyll is a continuous value which generally goes from high in the middle of the bloom to low on the outside of the bloom, properly modeling the gradient can give insight into the extent of the bloom and the boundary conditions.

For this work we consider the L2 norm of the gradient which gives the magnitude of the gradient vector, regardless of the direction, since the direction may not be along a coordinate axis.

In this work we use the notation

$$\dot{\mu}(x^*) = \|\nabla_{xyz} \mu(x^*)\|_2 \quad (6)$$

to represent the magnitude of the gradient vector in the x, y and z directions at x^* with respect to the predicted mean.

This leads to an extension of the UCB formula which is First Order UCB:

$$\dot{\mu}(x) + \beta^{1/2} \sigma(x) \quad (7)$$

UCB And First Order UCB: First Order UCB can be combined with UCB by adding an explicit parameter α which controls how much the first order objective contributes to the overall objective. In order for these two objectives to be in the same units, the predictive means are normalized relative to the mean of all of the predictive means and the variance of all the predictive means resulting in the predictive means being mean 0 and variance 1, a common machine learning normalization. This re-scaling also makes the β parameter isolated from the scale of the quantity being sensed.

The resulting equation is:

$$(1 - \alpha) \left(\frac{\mu(x) - \mathbb{E}[\mu]}{\text{var}(\mu)} \right) + \alpha \left(\frac{\dot{\mu}(x) - \mathbb{E}[\dot{\mu}]}{\text{var}(\dot{\mu})} \right) + \beta^{1/2} \sigma(x) \quad (8)$$

This can similarly be extended to add a distance term like Eq. (5) to create the full equation GP UCB Mean and First Order with distance:

VII. EXPERIMENT DESIGN

$$\frac{(1 - \alpha)(\frac{\mu(x) - \mathbb{E}[\mu]}{\text{var}(\mu)}) + \alpha(\frac{\dot{\mu}(x) - \mathbb{E}[\dot{\mu}]}{\text{var}(\dot{\mu})}) + \beta^{1/2}\sigma(x)}{\gamma \text{distance}(x^*, x)} \quad (9)$$

All of the previously discussed UCB-based algorithms can be viewed as Eq. (9) with parameters set to specific values, assuming $\gamma = 0$ is interpreted to mean the distance discount is ignored.

VI. PATH PLANNING ALGORITHMS

We compare two types of IPP algorithms with different information theoretical measures.

Before each planner begins to run, a pre-planned path to initialize the Gaussian process is performed. In this work, a pattern is used which visits the four corners of the work space and crosses in the middle [16].

Both algorithms have a parameter k which determines how many evenly spaced points in each dimension to calculate the metric on which results in k^3 candidate points.

Then, for UCB and Entropy metrics Algorithm 1 is used which greedily takes the current best point according to a metric

```

GP ← fitGP(initial survey)
while AUV.energy > 0 do
  candidatePoints ←  $k^3$  evenly spaced points
  bestPoint ←  $\text{argmax}_{x \in \text{candidatePoints}} \text{metric}(x)$ 
  path ← dubinsPath(bestPoint)
  AUV.follow(path)
  GP ← add(GP, path, samples)
end

```

Algorithm 1: IPP Algorithm 1

This algorithm is used for all metrics besides Eq. (3).

```

GP ← fitGP(initial survey)
path ← emptyList
pathLength ← 0
while pathLength < AUV.energy do
  candidatePoints ←  $k^3$  evenly spaced points
  bestPoint ←  $\text{argmin}_{x \in \text{candidatePoints}} \text{metric}(x)$ 
  path.append(bestPoint)
  pathLength ← TSPSolver(path)
end

```

Algorithm 2: IPP Algorithm 2

The other algorithm, which is more appropriate for Eq. (3), plans a sampling path all at once by greedily selecting points and solving a traveling salesman problem (TSP) between all the points until the total path length is equal to or greater than the distance the AUV can travel. In our experiments, a genetic algorithm based TSP solver is used which only creates an approximation, although this is necessary in some cases because TSP is in general NP-Hard [17].

For all distance considerations, the Dubins airplane curve is used. Dubins airplane curves are also used for motion planning between two points.

Previous runs from the Ecomapper AUV are used to build the ground truth Gaussian Process. The data was collected in the Puddingstone Reservoir in Los Angeles County, which is a 250 acre artificial lake.

These missions are taken across different times of the year and different years. While none of these datasets were taken during a notable algal bloom, they still represent the general concentration of algae that would be expected in a bloom.

Most of these missions are yo-yo transect missions where the AUV moves in a lawnmower pattern in the x and y plane and undulates up and down in the z axis.

One of the missions is a transect which goes back and forth in the x and y plane while slowly getting deeper.

Both of these mission types give an indication of what the whole concentration of chlorophyll would look like if sampled at every point.

From the previous dataset, a Gaussian process is fit which can be seen in Fig. 1.

We also use generated models from a mixture of Gaussians with random parameters where the x and y components of the mean are drawn from a uniform distribution and the depth component is drawn from an exponential distribution. The concentration at any point is the combined probability at that point scaled to an empirically chosen range (0 to 50 ug/L). An example of this can be seen in Fig. 2.

The AUV is simulated for a fixed energy budget which is the total number of meters the AUV can travel. In these experiments, the AUV expends the same amount of energy to travel in any available direction and only the total path length is considered. After the energy is depleted in simulation, a Gaussian process is fit to the sampled points and compared to the ground truth Gaussian process at 1000 regularly spaced points.

In the Dubins path calculations a fixed velocity of 1.2m/s and a max bank of $\pi/4$ are used.

An example planned path with 500 energy can be seen in Fig. 3 using Eq. (9) as the metric and Algorithm 1.

To calculate the gradient for the first order metrics the central difference method is used to numerically calculate Eq. (6).

We compare using four metrics: Root Mean Square Error (RMSE), Mean Absolute Error (MAE), Weighted Root Mean Square Error (WRMSE), and Weighted Mean Absolute Error (WMAE). WRMSE and WMAE are described by Marchant and Ramos [4] as metrics for evaluating these types of problems.

RMSE and MAE are commonly used metrics for determining the accuracy of models. RMSE is the square root for the average square of the difference between the ground truth and prediction. MAE is defined as the average of the absolute value of the difference between the ground truth and prediction.

WMAE is the absolute error weighted by the value of the reading, while WRMSE is the square root of the mean square error weighted by the value of the reading. These metrics




Fig. 3: An example path planned with Eq. (9) as the metric on the scenario in Fig. 1. Color is chlorophyll concentration intensity. Scale is in meters.

effectively weight points with higher concentration higher than points of lower concentration.

WRMSE is defined as:

$$\sqrt{\frac{1}{\sum_{i=1}^n Y_i} \sum_{i=1}^n (Y_i - \hat{Y}_i)^2 Y_i} \quad (10)$$

WMAE is similar to WRMSE, using the absolute error instead of the square error. WMAE is defined as:

$$\frac{1}{\sum_{i=1}^n Y_i} \sum_{i=1}^n \text{abs}(Y_i - \hat{Y}_i) Y_i \quad (11)$$

For Eqs. (10) and (11) Y_i is the ground truth value at point i and \hat{Y}_i is the predicted value at point i .

Similarly, these metrics can be used for the error in the prediction of the gradient by weighing the error in the predicted gradient by the ground truth gradient.

VIII. EXPERIMENTS

A. Grid search of parameters

Because Eq. (9) has three parameters, a grid search was run across parameters on one real and one generated dataset on three seeds each. For the α parameter the values $[0, 0.2, 0.5, 0.7, 1]$ were tried, for the β parameter the values $[1, 5, 10, 50]$ were tried and for γ the values $[1, 5, 10]$.

It can be seen in Fig. 4 that a good choice of hyper-parameters would be $\alpha = 0.5$, $\beta = 10$, and $\gamma = 10$ based on WRMSE.

B. Comparison to Entropy

In this section we compare Algorithm 1 using two metrics: Eq. (9) with the optimal parameters chosen from the grid search and Eq. (2). The dataset used can be seen in Fig. 1.

Figs. 5 and 6 show a comparison between Eq. (2) and Eq. (9) on a real dataset. In Fig. 5, it can be seen that






Fig. 4: α β and γ results. For each plot the other two parameters were kept at their optimal choices $\alpha = .2, \beta = 10, \gamma = 10$

Eq. (2) tends to outperform Eq. (9) in unweighted metrics, but performs worse in weighted metrics. This fits with intuition because Eq. (2) does not prioritize areas of high concentration but may cover the area as a whole better.

Fig. 6 shows that Eq. (9) slightly outperforms Eq. (2) on

gradient prediction at higher energy budgets.

For extended results, we present results on more real datasets as well as synthetic datasets in Section XI-A.




Fig. 5: Comparison at predicting mean between Eq. (9) and Eq. (2) using Algorithm 1 at different energy budgets, mean of three seeds.




Fig. 6: Comparison at predicting gradient between Eq. (9) and Eq. (2) using Algorithm 1 at different energy budgets, mean of three seeds.

C. Comparison to Mutual Information

We also compare Algorithm 1 with Eq. (9) with the optimal parameters with Algorithm 2 with Eq. (3).

Figs. 7 and 8 show a comparison between Eq. (3) with Algorithm 2 and Eq. (9) with Algorithm 1 on a real dataset, shown in Fig. 1.

Fig. 7 shows that Eq. (9) and Algorithm 1 performs better at unweighted metrics as well as weighed metrics. Eq. (9) and Algorithm 1 only perform slightly better at the weighted metrics which is surprising to us since Eq. (3) does not take into account high concentration areas.

Fig. 8 shows that shows that Eq. (9) and Algorithm 1 performs worse at monitoring the gradient than Eq. (3) and

Algorithm 2. This may be due to it having a tendency to get stuck in local maxima and not monitoring different blooms well.

We only show one experiment because of the computational cost of Eq. (3) and Algorithm 2.




Fig. 7: Comparison at predicting mean between Eq. (9) using Algorithm 1 and Eq. (3) using Algorithm 2 at different energy budgets, mean of three seeds.




Fig. 8: Comparison at predicting gradient between Eq. (9) using Algorithm 1 and Eq. (3) using Algorithm 2 at different energy budgets, mean of three seeds.

IX. CONCLUSION

We believe that Eq. (9) is a good metric for monitoring chlorophyll concentrations and spatial changes in concentration. A greedy planning approach gives an indication of how well this metric performs but a more robust and long term planner is needed to evaluate the realistic performance of this metric.

In some experiments the modeling accuracy decreases with longer energy budgets. This may be due to fitting errors in the Gaussian process or local maxima causing the planner to get stuck.

The greedy planner is sensitive to starting conditions as well as local maxima. In this work $\gamma = 10$ is used which provides a good trade off between getting stuck in local maxima and expending energy early to go between algal blooms.

X. FUTURE WORK

In future work, we plan to study how these metrics can be combined with different path planners such as rapidly exploring information gathering trees [9] or generalized recursive greedy [8] planners. It would need to be shown that Eq. (8) is submodular for these types of algorithms to work. Eq. (9) with Algorithm 1 is sensitive to initial starting points and easily gets caught in local maxima but a longer horizon planner may not have this problem. Eq. (9) may also not need a distance discount term if a longer horizon planner is implemented.

We would also like to explore jointly maximizing the modeling accuracy of multiple on-board sensors, such as salinity and temperature, which can provide strong signals of algal bloom conditions. In this work, we have only explored using the chlorophyll sensor which in most cases cannot separate algal blooms into harmful and non-harmful components by itself and requires further study by biologists.

REFERENCES

- [1] J. Heisler, P. Glibert, J. Burkholder, D. Anderson, W. Cochlan, W. Denison, C. Gobler, Q. Dortch, C. Heil, E. Humphries, A. Lewitus, R. Magnien, H. Marshall, K. Sellner, D. Stockwell, D. Stoecker, and M. Suddleson, "Eutrophication and Harmful Algal Blooms: A Scientific Consensus," *Harmful algae*, vol. 8, no. 1, pp. 3–13, dec 2008.
- [2] I. Robbins, G. Kirkpatrick, S. Blackwell, J. Hillier, C. Knight, and M. Moline, "Improved monitoring of habs using autonomous underwater vehicles (auv)," *Harmful Algae*, vol. 5, no. 6, pp. 749 – 761, 2006.
- [3] H. Choset, "Coverage for robotics – a survey of recent results," *Annals of Mathematics and Artificial Intelligence*, vol. 31, no. 1, pp. 113–126, Oct 2001.
- [4] R. Marchant and F. Ramos, "Bayesian optimisation for informative continuous path planning," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, May 2014, pp. 6136–6143.
- [5] R. Marchant, F. Ramos, and S. Sanner, "Sequential bayesian optimisation for spatial-temporal monitoring," *Uncertainty in Artificial Intelligence - Proceedings of the 30th Conference, UAI 2014*, pp. 553–562, 01 2014.
- [6] 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*, ser. ICML'10. USA: Omnipress, 2010, pp. 1015–1022.
- [7] C. Guestrin, A. Krause, and A. P. Singh, "Near-optimal sensor placements in gaussian processes," in *Proceedings of the 22Nd International Conference on Machine Learning*, ser. ICML '05. New York, NY, USA: ACM, 2005, pp. 265–272.
- [8] J. Binney, A. Krause, and G. S. Sukhatme, "Informative path planning for an autonomous underwater vehicle," in *2010 IEEE International Conference on Robotics and Automation*, May 2010, pp. 4791–4796.
- [9] G. A. Hollinger and G. S. Sukhatme, "Sampling-based robotic information gathering algorithms," *The International Journal of Robotics Research*, vol. 33, no. 9, pp. 1271–1287, 2014.
- [10] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *The International Journal of Robotics Research*, vol. 30, no. 7, pp. 846–894, 2011.
- [11] Y. Zhang, J. G. Bellingham, J. P. Ryan, B. Kieft, and M. J. Stanway, "Two-dimensional mapping and tracking of a coastal upwelling front by an autonomous underwater vehicle," in *2013 OCEANS - San Diego*, Sep. 2013, pp. 1–4.
- [12] C. Rasmussen and C. Williams, *Gaussian Processes for Machine Learning*. MIT press, 2006.
- [13] M. Owen, R. W. Beard, and T. W. McLain, *Implementing Dubins Airplane Paths on Fixed-Wing UAVs**. Dordrecht: Springer Netherlands, 2015, pp. 1677–1701.
- [14] L. E. Dubins, "On curves of minimal length with a constraint on average curvature, and with prescribed initial and terminal positions and tangents," *American Journal of Mathematics*, vol. 79, no. 3, pp. 497–516, 1957.
- [15] K. H. Low, "Multi-robot adaptive exploration and mapping for environmental sensing applications," Ph.D. dissertation, Carnegie Mellon University, Pittsburgh, PA, USA, 2009, aAI3374752.
- [16] S. Kemna, O. Kroemer, and G. S. Sukhatme, "Pilot surveys for adaptive informative sampling," in *International Conference on Robotics and Automation (ICRA)*, January 2018.
- [17] M. Jingyan and Z. Kehong, "Research on tsp solution based on genetic algorithm of logistic equation," in *Proceedings of 2012 2nd International Conference on Computer Science and Network Technology*, Dec 2012, pp. 738–742.

XI. APPENDIX

A. Extended Entropy Experiment

We present extended results for Section VIII-B.




Fig. 9: Real world dataset 2 comparison between Eq. (9) and Eq. (2) using Algorithm 1. Metrics prefixed with D are metrics on the magnitude of the gradient, described by Eq. (6), mean of three seeds.




Fig. 10: Real world dataset comparison between Eq. (9) and Eq. (2) using Algorithm 1. Metrics prefixed with D are metrics on the magnitude of the gradient, described by Eq. (6), mean of three seeds.




Fig. 11: Synthetic dataset 1 (shown in Fig. 2) comparison between Eq. (9) and Eq. (2) using Algorithm 1. Metrics prefixed with D are metrics on the magnitude of the gradient, described by Eq. (6), mean of three seeds.

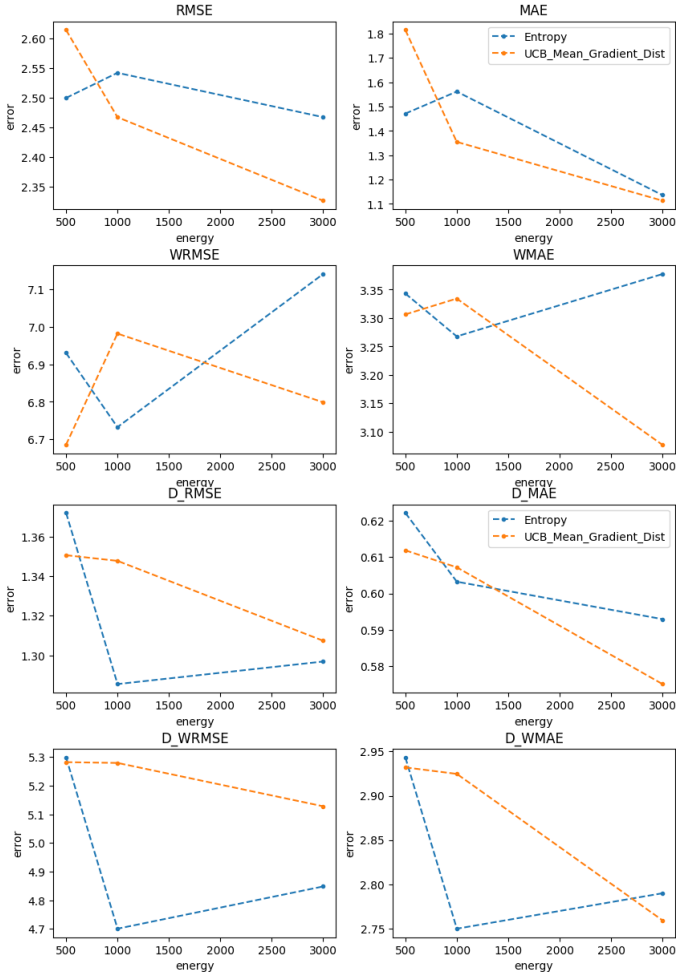


Fig. 12: Synthetic dataset 2 comparison between Eq. (9) and Eq. (2) using Algorithm 1. Metrics prefixed with D are metrics on the magnitude of the gradient, described by Eq. (6), mean of three seeds.