i-SAIRAS 2024

Information Gathering Path Planning with Minimal Environmental Disruption

[1,2*]Strooper, K., [2]Lawrance, N., [2]Bandyopadhyay, T., [2] Talbot, F., [2]Fiamingo, R. & [1,2]Chung, J. J.

*lead presenter

[1]k.strooper@ uq.net.au, The University of Queensland, Australia

[2] CSIRO Robotics, Australia

Abstract

Ground robots are often tasked to autonomously explore and characterise unknown and sensitive environments. While performing these searching tasks, for applications such as interplanetary exploration, it is important that these robots also minimise how much they disturb the area that they are exploring to decrease contamination of the original environment. This work develops a Minimising Disturbance Informative Path Planner (MDIPP) which models an unknown environment using discrete observations whilst prioritising finding high concentrations of samples of interest. Disturbance to the environment is also incorporated into the path planner to minimise how much of the environment is being disturbed. MDIPP uses Gaussian processes (GPs) to model the unknown environment and Bayesian optimisation (BO) to determine the next best position to move to. MDIPP reduces the disturbance to the environment by at least 48.6% compared to a baseline informative path planner (BIPP), without sacrificing accuracy in the GP prediction generated from its online observations. The results have been validated in hardware on a tracked ground vehicle at CSIRO's lunar testbed in Queensland, Australia to mimic planetary exploration. The project website is available here.

I. INTRODUCTION

Consider an interplanetary rover equipped with a probe that must autonomously find and locate mineral samples. The robot must predict the locations of substances of scientific interest (e.g. searching for high concentrations of hydrogen) based on measurements taken online–an important and common objective in space exploration [1], [2], [3], [4]. While performing this primary searching task, these robots must also minimise how much their motions disturb the surrounding environment to preserve the original geological formations on the planetary body's surface and avoid potential contamination [1]. This can be formulated as a multi-objective informative path planning (IPP) problem where the planner must balance the potential information gain from moving to and revealing new parts of the scene to the sensor, with the cost incurred from disturbing the environment.

Indeed, few works explicitly consider the effect of the robot's movement on its environment. Those that do are typically applied to needle-steering robots [5] or crowd-navigating robots [6]. The former assumes shared control via live teleoperation, while the latter requires a goal pose and computes velocity commands to match the overall dynamics of the crowd flow. In contrast, environment disturbance in our problem is defined as the swept footprint of the robot, presenting the unique challenge of trading off the need to visit new locations to gather high-quality information, while simultaneously attempting to reduce the total swept area.

We propose the Minimising Disturbance Informative Path Planner (MDIPP) which uses Gaussian processes (GPs) to model the unknown environment and Bayesian optimisation (BO) for active perception to determine the next best position to move to. The GP provides a probabilistic regression model of mineral density, where density is considered as a continuous function from which the robot makes noisy observations. Disturbance to the environment is considered binary, i.e. once an area has been disturbed, it does not matter if that same area is disturbed again, thus the robot would ideally plan routes through previously disturbed regions when moving between viewpoints. Our proposed BO acquisition function models such disturbances as a function of the **new** area that would be disturbed during movement, and weights this against the exploitation and exploration elements from the GP mean and variance.

Results comparing MDIPP against a Baseline Informative Path Planner (BIPP) show that MDIPP reduces environmental disturbance by at least 48.6% without sacrificing accuracy in the GP predictions generated from its online observations. Adjusting the exploration and disturbance hyperparameters within MDIPP results in a trade-off between model accuracy and the amount of disturbance, which can be tuned to suit the application. Further testing in a high-fidelity simulator allowed us to quicky evaluate different tuning parameters and provided a stepping stone to testing on hardware. Finally, we validated MDIPP on a ground robot in CSIRO's lunar testbed in Queensland, Australia to mimic a planetary exploration mission.

II. RELATED WORK

Previous methods of information gathering for unknown environments typically favour previously planned offline full coverage solutions [7], [8], [9]. In contrast to these offline methods, online methods actively integrate live observations to improve performance using next-best-view (NBV) paths [10], [11], [12]. However, while these planners integrate online observations to improve performance, their methods assume that the information about the environment is uniformly distributed, and that each observation is equally weighted. In contrast, our solution aims to provide an online information gathering approach to optimise observations to maximise finding high concentrations of samples of interest (e.g. minerals beneath extraterrestrial surfaces). Similar approaches have been presented in [13], [14], [15], which feed observation information back into the planner to efficiently identify the maxima of the distribution, rather than attempting to accurately reconstruct the full distribution.

To model such spatial correlations, GPs [16] are a popular regression method and have been successfully applied to predicting distributions of an information field (e.g. temperature distribution) [15], [17]. However, these methods typically have unconstrained coverage and plan on a time/distance or computational budget, which means they do not consider the impact of the robot's motion through the environment which is particularly important for planetary exploration scenarios. To combine our multiple objectives of exploration (finding high concentrations of samples of interest) and minimising disturbance to the environment our solution uses BO to select informative samples from GPs [18].

BO is a process for finding the extrema of an acquisition function that is expensive to observe or evaluate [19]. It is applicable in situations where you don't have a closed-form expression for the objective function, but where you can obtain values at sampled locations with optimisation methods such as genetic algorithms [14], [20], [21], [22]. Much of the efficiency comes from the ability of BO to incorporate prior belief about the problem to help direct the sampling, and to trade off exploration and exploitation of the search space [23].

III. MINIMISING DISTURBANCE INFORMATIVE PATH PLANNER

Our proposed Minimising Disturbance Informative Path Planner (MDIPP) uses Gaussian process (GP) regression to model the predicted field from local, noisy observations. To determine 'informative' sample locations, we use acquisition functions within a Bayesian optimisation (BO) framework to balance the objectives of information gathering and disturbance minimisation when determining the next best position to move to and take a new observation. MDIPP maintains a binary 'disturbance array' over the environment. Each cell within the grid is approximately the size of the robot's footprint. For the purpose of path planning, the disturbed cells have a lower traversal cost of 0.1 compared to the undisturbed cells which have a cost of 1. MDIPP creates an 8-connected grid cost map from these values and uses A* over the resulting graph to plan and quantify the disturbance induced by travelling through the environment to a target observation location. This cost is then incorporated into the BO acquisition function for evaluating potential viewpoints.

Figure 1 gives an overview of MDIPP which also explains each section of the decision logic in more detail. Disturbance minimisation is considered in two parts of our algorithm where it has two distinct objectives. Including disturbance information in the acquisition function influences the high-level decision of selecting the next most valuable observation location. In the A* search path planner, the disturbance information sets the traversal costs, affecting low-level decisions of how the robot will move to a viewpoint.

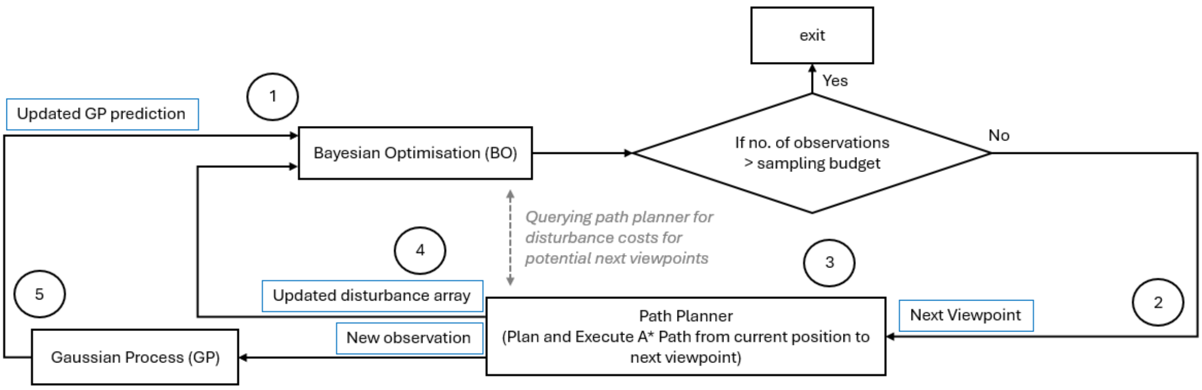

Figure 1: Logic overview of MDIPP. 1) The current GP prediction of the field and disturbance array are passed into the BO, which determines the next best viewpoint to move to based on the acquisition function (Section IIIA). 2) If the sampling budget is not exceeded, this next viewpoint, along with the robot's current position is passed to the path planner. 3) The path planner for MDIPP uses A search where the edge traversal costs are drawn from the disturbance array over the environment. 4) After the robot reaches the next target observation location, it takes a new observation with its sensor and the disturbance array is updated to reflect newly disturbed cells. 5) The GP prediction is updated with the new observation.*

A. Acquisition Function

While a complete treatment of GP regression is beyond the scope of this work (interested readers are referred to [24]), we briefly outline the process for consistent notation. GP regression provides a tool for estimating an underlying function $f(x): R^n \to R$ based on discrete, noisy observations $y = f(x) + \varepsilon$, where $\varepsilon \sim \mathcal{N}(0, \sigma_n^2)$. The GP provides a posterior distribution over functions, conditioned on the set of observations y, by relating the input points of those observations X via a covariance function $K(\cdot)$. The resulting model can be queried to provide the predictive likelihood $f_*(x_*)$ at a location x_* as a normal distribution,

$$p(f_*(\boldsymbol{x}_*)|X, \boldsymbol{y}, \boldsymbol{x}_*) = \mathcal{N}\left(\hat{f}_*(\boldsymbol{x}_*), \sigma_*^2\big(f_*(\boldsymbol{x}_*)\big)\right), where \quad (1)$$

$$\hat{f}_*(\boldsymbol{x}_*) = K(\boldsymbol{x}_*, X)[K(X, X) + \sigma_n^2 I]^{-1\boldsymbol{y}} \quad (2)$$

$$\sigma_*^2\big(f_*(\boldsymbol{x}_*)\big) = K(\boldsymbol{x}_*, \boldsymbol{x}_*) - K(\boldsymbol{x}_*, X)[K(X, X) + \sigma_n^2 I]^{-1}K(X, \boldsymbol{x}_*). \quad (3)$$

We use the common squared exponential covariance function, with fixed length scale and noise hyperparameters learned by maximising marginal likelihood over a representative sample dataset.

BO takes advantage of the probabilistic GP predictions via an acquisition function that balances finding high values in the field (high estimated GP mean $\hat{f}_*(x)$), with uncertain places in the field (high $\mathrm{cov}\big(f_*(x_*)\big)$). Equation (4) shows the acquisition function used by MDIPP to determine the next best viewpoint $x^* = \mathrm{argmax}[f_{MDIPP}(x)]$, where,

$$f_{MDIPP}(\boldsymbol{x}) = \hat{f}_*(\boldsymbol{x}) + K_{EXP} \times \sigma_*\big(f_*(\boldsymbol{x})\big) - K_{DIST} \times D(\boldsymbol{x}). \quad (4)$$

K_{EXP} and K_{DIST} are the 'exploration' and 'disturbance' scaling factors and $D(x)$ is the disturbance cost for traversing to location x from the robot's current position according to the planned A* path over the disturbance cost map. These parameters strongly affect how MDIPP behaves and it is important to identify which values of K_{EXP} and K_{DIST} provide appropriate scaling between the three terms for different applications.

B. Baseline Informative Path Planner (BIPP)

To provide an understanding of the effectiveness of MDIPP, we compare the performance against a baseline planner, BIPP, which does not incorporate disturbance minimisation in its planning. BIPP is very similar to MDIPP, except for two main differences. First, BIPP assumes uniform edge traversal costs during the A* search (hence returning the shortest path on the graph). Second, the BIPP acquisition function does not consider the number of disturbed cells, effectively equivalent to MDIPP$K_{DIST} = 0$:

$$f_{BIPP}(\boldsymbol{x}) = \hat{f}_*(\boldsymbol{x}) + K_{EXP} \times \sigma_*\big(f_*(\boldsymbol{x})\big). \quad (5)$$

Comparing the performance of BIPP and MDIPP provides an ablation study on the impact of including the disturbance minimisation components in our planner. It also reflects the current state in IPP solutions, which do not consider the impact of the robot's motion on the environment when moving to take samples.

IV. RESULTS

A. Trade-Off Between Exploration and Disturbance

To provide statistical results on each planner's performance, we first evaluated the two methods on a set of simple simulated environments where the underlying ground truth field is generated by sampling from a GP prior. The robot takes point observations to simulate taking measurements with a ground penetrating radar (GPR) or a probe. Simulated robot observations of the ground truth include additive Gaussian noise to emulate real-world sensor measurements, and the robot motion is deterministic and assumed to travel in straight lines.

We performed a parameter sweep testing K_{EXP} values in a range of 0 to 20 with increments of 0.2, and K_{DIST} values of $\{0, 0.01, 0.1, 0.5, 1\}$, since $D(x)$ can be quite large compared to the other terms motivating a smaller scaling factor. 20 trials (20 different environments) were completed for each combination of values of K_{EXP} and K_{DIST}. The results after 15 observations are shown in Figure 2a. The top plot displays the percentage of the environment that was disturbed at the end of each run. The middle plot shows the root mean square error (RMSE) of the final GP prediction, and the bottom plot shows the RMSE weighted by the absolute value of the field, to help analyse the ability of the planner to accurately predict regions of high concentration.

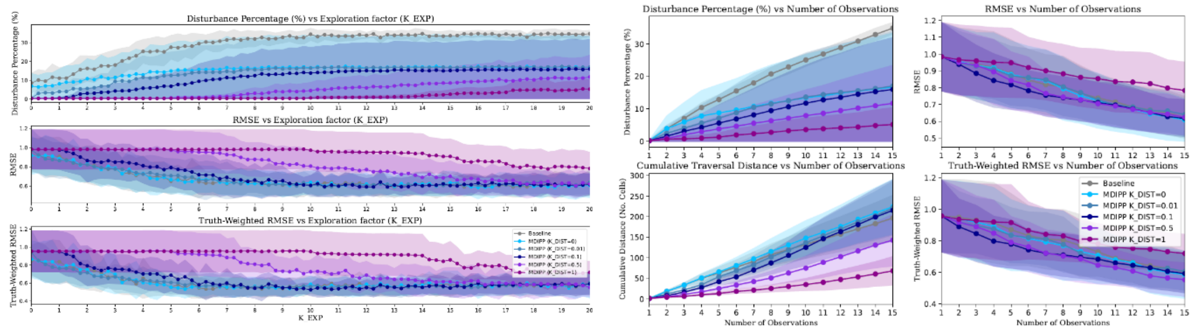

(a) (b)

Figure 2: Results from parameter experimentation for BIPP and MDIPP over 20 trials for the (a) full parameter sweep after 15 observations, and (b) intra-episode results over the 15 observations for $K_{EXP} = 20$. The solid lines represent the mean values and the shaded area represents one standard deviation above/below the mean.

The results show that MDIPP disturbs less of the environment than BIPP and continues to disturb less as K_{DIST} increases. Even without accounting for the disturbance in the acquisition function ($K_{DIST} = 0$), MDIPP already reduces the disturbance to the environment by 48.6%. This shows that our method of path planning over the disturbance cost map already meets our objective of reducing environmental disturbance during exploration. As K_{DIST} increases, there remains very little difference in the RMSE between the path planners ($\leq 7\%$). Therefore, the robot's prediction of the environment does not suffer when incorporating this disturbance minimisation objective. Furthermore, the results show that various combinations of $\{K_{DIST}, K_{EXP}\}$ lead to the same RMSE after 15 observations, with $\{0.5, 20\}$ causing the least disturbance.

B. Informative Path Planning (IPP) Performance within an Episode

We analysed the results after each observation for different K_{EXP} values, focusing on $K_{EXP} = 20$ due to similar trends across all cases. As shown in Figure 2b, MDIPP consistently caused less environmental disturbance than BIPP across all K_{DIST} values, with an accuracy trade-off only at the highest tested K_{DIST} value of 1.

Additionally, we analysed the robot's cumulative traversal distance (number of cells that the robot travels between observation points – regardless of whether it has previously been disturbed) for each observation as a proxy for time. The traversal distance for MDIPP was similar to BIPP, with small K_{DIST} values resulting in almost identical cumulative traversal distance, showing that MDIPP does not significantly increase travel distance despite minimising disturbance. These results suggest that similar outcomes could be expected in larger environments with adjusted K_{EXP} and K_{DIST} values.

C. High-Fidelity Simulation Validation

As an initial step toward hardware validation we tested MDIPP in a high-fidelity lunar simulation with a complete robot navigation stack [25]. The navigation stack has its own integrated path planner for robot navigation including global planning and local obstacle avoidance. To account for the differences in the planner architecture relative to the simple simulation, we modified the main logic for BIPP and MDIPP for the high-fidelity simulation environment. First, the true robot position is used to update the disturbance array and sample observations to capture the localisation and motion planning differences that arise in real-world deployment. Secondly, to minimise disturbance, MDIPP generally does not return the straight-line path to the next sample point, instead it will prefer to reuse pathways. Thus, to integrate with the high-fidelity planner, the computed MDIPP path is first decomposed into straight-line segments and the resulting sequence of sub-waypoints are sent to the planner.

Preliminary trials were conducted for all pareto-front combination values, but the results confirmed that $K_{EXP} = 20$ and $K_{DIST} = 0$ yielded the best performance for our objectives. The results from Figure 3a show that from the start there is a clear distinction between the disturbance to the environment from BIPP and MDIPP. Furthermore, the accuracy remains similar between both methods with consistent overlap between the truth-weighted RMSE in the bottom-right plot. Finally, the cumulative time in the bottom-left plot also shows that the results overlap for both methods even as the mean of the MDIPP cumulative time increases with the number of observations.

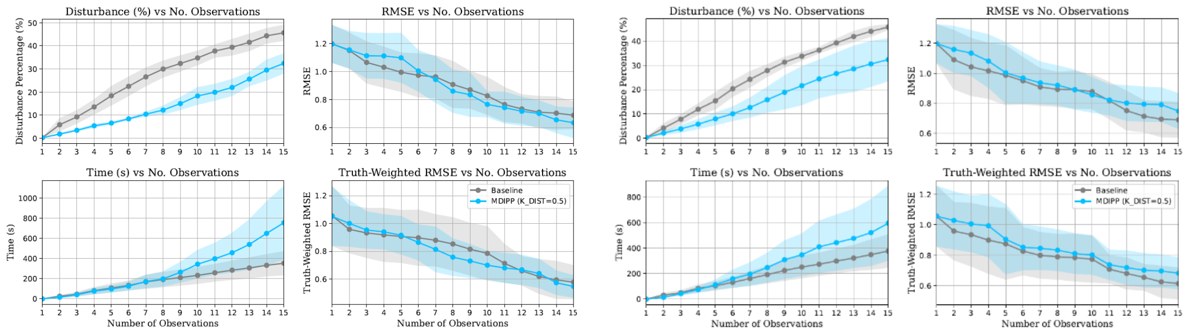

(a) (b)

Figure 3: Results over 15 observations for $K_{EXP} = 20$ and $K_{DIST} = 0.5$ in (a) CSIRO's High-Fidelity Lunar Simulation over 5 trials, and (b) on hardware at CSIRO's Lunar Environment over 5 trials.

D. Hardware Validation

Our final results validate MDIPP on hardware at CSIRO's lunar testbed in Queensland, Australia in a mock planetary exploration mission. The purpose of the hardware testing was to validate the binary representation of the disturbance compared to the physical disturbance of the lunar testbed, and to demonstrate the system running in real time on-board a robot in a representative environment. We used top-down views of the lunar testbed from an aerial drone and compared it to the disturbed area from the computer-generated binary disturbance array results. The lunar testbed was between trials so that track marks from the rover could clearly be distinguished in the sand.

Five trials were conducted with each planner, using the same ground truth maps for both planners. The quantitative results shown in Figure 3b are very similar to the high-fidelity simulation results. Importantly, MDIPP still disturbs distinctly less of the environment than BIPP while returning similarly accurate predictions.

To validate whether our binary representation accurately captures the true disturbance we collected aerial infrared footage of the trials using a Mavic 2 drone. The videos were processed using frame alignment and differencing to aggregate the robot movements and generate the final images shown in Figure 4 and Figure 5 for BIPP and MDIPP, respectively.

There is significant alignment between the planner's predicted disturbance and the actual disturbance in the lunar testbed. Some regions that do not perfectly overlap, e.g. bottom-left quadrant of Figure 4c, may be attributed to slight differences in the orientation of the computed disturbance and the true disturbance, possibly due to variations in the drone's positioning. Nevertheless, the overall disturbance values were well-aligned in both cases.

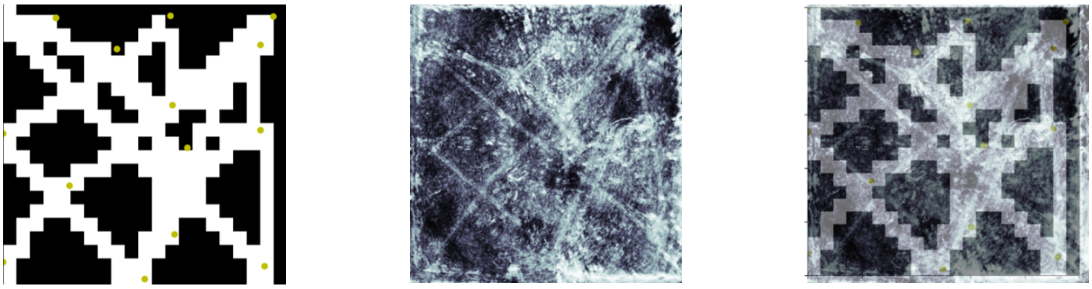

(a) (b) (c)

Figure 4: BIPP disturbance comparison from hardware testing. (a) Top-down view of the computed binary disturbance. (b) Top-down view of the drone footage of the disturbance. (c) Drone footage superimposed over the computed disturbance.

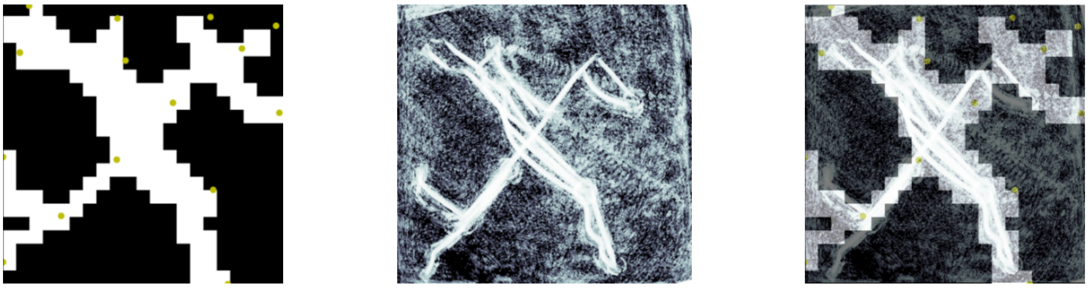

(a) (b) (c)

Figure 5: MDIPP disturbance comparison from hardware testing. (a) Top-down view of the computed binary disturbance. (b) Top-down view of the drone footage of the disturbance. (c) Drone footage superimposed over the computed disturbance.

V. FUTURE WORK AND CONCLUSIONS

This project proposed the Minimising Disturbance Informative Path Planner (MDIPP), which selects and plans paths to observation points that minimise the robot's disturbance to the environment while accurately modelling the underlying field. By treating disturbance as a cost map and planning paths using A* search, environmental disturbance is significantly decreased.

Since this is the first minimising disturbance informative path planner application for ground robots, several key areas for further development have been identified. First, hardware testing showed that the robot generates more disturbance when turning, which should be considered in the planner. Second, adapting MDIPP for continuous observations would allow the robot to dynamically adjust its path based on real-time data, although this would require updates to the acquisition function logic and consideration of time delays to prevent excessive redirection. Third, testing MDIPP with real-world sensor data, such as LIDAR or camera inputs, could help assess its functionality across different sensing modalities, potentially requiring preprocessing of sensor outputs and updating the underlying GP model for more complex sensor models.

ACKNOWLEDGEMENTS

This work was supported in part by CSIRO's Space Future Science Platform. The authors would also like to thank Kartik Yellapantula for his assistance with the trajectory planner.

REFERENCES

[1] M. Allan *et al.*, "Planetary Rover Simulation for Lunar Exploration Missions," in *2019 IEEE Aerospace Conference*, Mar. 2019, pp. 1–19. doi: 10.1109/AERO.2019.8741780.

[2] K. Schilling and C. Jungius, "Mobile robots for planetary exploration," *Control Eng. Pract.*, vol. 4, no. 4, pp. 513–524, Apr. 1996, doi: 10.1016/0967-0661(96)00034-2.

[3] T. Kubota, R. Ejiri, Y. Kunii, and I. Nakatani, "Autonomous behavior planning scheme for exploration rover," in *2nd IEEE International Conference on Space Mission Challenges for Information Technology (SMC-IT'06)*, Jul. 2006, p. 7 pp.-. doi: 10.1109/SMC-IT.2006.22.

[4] A. Arora, P. M. Furlong, R. Fitch, S. Sukkarieh, and T. Fong, "Multi-modal active perception for information gathering in science missions," *Auton. Robots*, vol. 43, no. 7, pp. 1827–1853, Oct. 2019, doi: 10.1007/s10514-019-09836-5.

[5] K. B. Reed *et al.*, "Robot-Assisted Needle Steering," *IEEE Robot. Autom. Mag.*, vol. 18, no. 4, pp. 35–46, Dec. 2011, doi: 10.1109/MRA.2011.942997.

[6] T. Fraichard and V. Levesy, "From Crowd Simulation to Robot Navigation in Crowds," *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 729–735, Apr. 2020, doi: 10.1109/LRA.2020.2965032.

[7] E. Galceran and M. Carreras, "A survey on coverage path planning for robotics," *Robot. Auton. Syst.*, vol. 61, no. 12, pp. 1258–1276, Dec. 2013, doi: 10.1016/j.robot.2013.09.004.

[8] R. Bähnemann, N. Lawrance, J. Chung, M. Pantic, R. Siegwart, and J. Nieto, "Revisiting Boustrophedon Coverage Path Planning as a Generalized Traveling Salesman Problem," 2021, pp. 277–290. doi: 10.1007/978-981-15-9460-1_20.

[9] G. Mier, J. Valente, and S. de Bruin, "Fields2Cover: An Open-Source Coverage Path Planning Library for Unmanned Agricultural Vehicles," *IEEE Robot. Autom. Lett.*, vol. 8, no. 4, pp. 2166–2172, Apr. 2023, doi: 10.1109/LRA.2023.3248439.

[10] C. Connolly, "The determination of next best views," in *1985 IEEE International Conference on Robotics and Automation Proceedings*, Mar. 1985, pp. 432–435. doi: 10.1109/ROBOT.1985.1087372.

[11] A. Bircher, M. Kamel, K. Alexis, H. Oleynikova, and R. Siegwart, "Receding Horizon 'Next-Best-View' Planner for 3D Exploration," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, May 2016, pp. 1462–1468. doi: 10.1109/ICRA.2016.7487281.

[12] M. Naazare, F. G. Rosas, and D. Schulz, "Online Next-Best-View Planner for 3D-Exploration and Inspection With a Mobile Manipulator Robot," *IEEE Robot. Autom. Lett.*, vol. 7, no. 2, pp. 3779–3786, Apr. 2022, doi: 10.1109/LRA.2022.3146558.

[13] G. Flaspohler, V. Preston, A. P. M. Michel, Y. Girdhar, and N. Roy, "Information-Guided Robotic Maximum Seek-and-Sample in Partially Observable Continuous Environments," *IEEE Robot. Autom. Lett.*, vol. 4, no. 4, pp. 3782–3789, Oct. 2019, doi: 10.1109/LRA.2019.2929997.

[14] M. Popović *et al.*, "An informative path planning framework for UAV-based terrain monitoring," *Auton. Robots*, vol. 44, no. 6, pp. 889–911, Jul. 2020, doi: 10.1007/s10514-020-09903-2.

[15] H. Zhu, J. J. Chung, N. R. J. Lawrance, R. Siegwart, and J. Alonso-Mora, "Online Informative Path Planning for Active Information Gathering of a 3D Surface," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, May 2021, pp. 1488–1494. doi: 10.1109/ICRA48506.2021.9561963.

[16] C. E. Rasmussen and C. K. I. Williams, "Gaussian Processes for Machine Learning: Book webpage." Accessed: Apr. 25, 2024. [Online]. Available: https://gaussianprocess.org/gpml/

[17] K. Jakkala and S. Akella, "Multi-Robot Informative Path Planning from Regression with Sparse Gaussian Processes," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, May 2024, pp. 12382–12388. doi: 10.1109/ICRA57147.2024.10610484.

[18] A. Hebbal, M. Balesdent, L. Brevault, N. Melab, and E.-G. Talbi, "Deep Gaussian process for multi-objective Bayesian optimization," *Optim. Eng.*, vol. 24, no. 3, pp. 1809–1848, Sep. 2023, doi: 10.1007/s11081-022-09753-0.

[19] E. Brochu, V. M. Cora, and N. de Freitas, "A Tutorial on Bayesian Optimization of Expensive Cost Functions, with Application to Active User Modeling and Hierarchical Reinforcement Learning," arXiv.org. Accessed: Oct. 26, 2023. [Online]. Available: https://arxiv.org/abs/1012.2599v1

[20] K. C. T. Vivaldini, V. Guizilini, M. D. C. Oliveira, T. H. Martinelli, D. F. Wolf, and F. Ramos, "Route planning for active classification with UAVs," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, May 2016, pp. 2563–2568. doi: 10.1109/ICRA.2016.7487412.

[21] G. Hitz, E. Galceran, M.-È. Garneau, F. Pomerleau, and R. Siegwart, "Adaptive continuous-space informative path planning for online environmental monitoring," *J. Field Robot.*, vol. 34, no. 8, pp. 1427–1449, 2017, doi: 10.1002/rob.21722.

[22] G. Francis, L. Ott, R. Marchant, and F. Ramos, "Occupancy Map Building through Bayesian Exploration," arXiv.org. Accessed: Sep. 09, 2024. [Online]. Available: https://arxiv.org/abs/1703.00227v1

[23] J. Snoek, H. Larochelle, and R. P. Adams, "Practical Bayesian Optimization of Machine Learning Algorithms," Aug. 29, 2012, *arXiv*: arXiv:1206.2944. doi: 10.48550/arXiv.1206.2944.

[24] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. Accessed: Sep. 25, 2024. [Online]. Available: https://direct.mit.edu/books/monograph/2320/Gaussian-Processes-for-Machine-Learning

[25] N. Hudson *et al.*, "Heterogeneous Ground and Air Platforms, Homogeneous Sensing: Team CSIRO Data61's Approach to the DARPA Subterranean Challenge," *Field Robot.*, vol. 2, pp. 595–636, Mar. 2022, doi: 10.55417/fr.2022021.