Active Sensing for Target Tracking in Dynamic Environments using Autonomous Surface Vehicles

Sanjeev Ramkumar Sudha, Marija Popović, and Erlend M.Coates

***Abstract*— Mobile robots are becoming increasingly common for monitoring and inspection tasks. The cleanup of pollutants with autonomous surface vehicles in rivers and the oceans is one such important application. However, monitoring in marine environments is challenging due to the dynamic nature of the environment. Planning informative paths in such environments requires maintaining an accurate map by predicting the spatial and temporal variations. Therefore, we present our active sensing approach for tracking moving targets in dynamic environments, in real time. We introduce a spatiotemporal prediction network to predict the uncertain future target distributions. Our adaptive planning approach leverages predictions from this network with a new planning utility for target tracking. It is also validated through field deployments with an autonomous surface vehicle.**

I. Introduction

Autonomous robot platforms offer an efficient alternative for information gathering [1] and environmental monitoring as compared to traditional approaches such as manual and teleoperated surveys [2], or sensors placed at fixed locations [3]. Some applications in environmental monitoring include pollutant cleanups [4], search and rescue [5], and wildlife monitoring [6]. Timely mapping and subsequent collection of floating pollutants such as plastic litter is crucial, as they can break down into toxic microplastics and prove harmful to humans and marine animals if allowed to persist. However, monitoring in dynamic environments is a challenging problem, as the robot has to predict and react to changes in the map to collect newly acquired information in an efficient way.

This study focuses on active sensing for tracking freely floating targets of interest, such as plastic litter, with an autonomous surface vehicle (ASV). Such targets generally drift due to external disturbances such as wind and currents. We aim to keep an updated global map of target positions through a dynamic occupancy mapping approach. We also use an adaptive planning strategy to plan informative trajectories for the robot, using predictions from our proposed spatiotemporal prediction network.

In active sensing, the goal is to find robot actions that maximise the information gathered, subject to constraints such as maximum allowed mission time or energy [7], [8]. The informativeness or utility is generally quantified by entropy-based measures in the map, which is updated as the robot acquires new sensor observations. For static environments, reducing entropy achieves efficient spatial coverage [9], [10]. However, in spatiotemporally varying environments such as target tracking applications, predicting the changes in the map over a certain time horizon is crucial.

Sanjeev R.S. and Erlend M. Coates are with the Department of ICT and Natural Sciences, Norwegian University of Science and Technology (NTNU), Ålesund, Norway. {sanjeev.k.r.sudha, erlend.coates}@ntnu.no. Marija Popović is with the MAVLab, Faculty of Aerospace Engineering, TU Delft, The Netherlands. m.popovic@tudelft.nl

Fig. 1: Our autonomous surface vehicle (ASV) during field deployments for active mapping of dynamic floating targets of interest, such as the buoy in the picture. We use four buoys during the tests, which are freely drifting under the influence of winds and currents.

Traditional target tracking techniques [11] involve identifying and estimating the states of individual targets through a series of observations. Occupancy grids provide a more efficient alternative for target tracking by maintaining a single global map with the target positions represented as occupancy probabilities [12], [13]. Previous works on occupancy grid mapping in dynamic environments [12], [14] rely on observations to learn the state transition probabilities. In contrast, we leverage the fact that the drift of the targets can be approximated [15], [16] given the measured wind velocity.

Prior studies in informative planning for monitoring in dynamic environments [4], [17], [18] assume that the dynamics can be reliably predicted in advance or assume perfect localisation of the targets of interest. In reality, the environmental forces are stochastic and hard to predict in advance. To address this issue, we propose using a physics-informed prediction network that is independent of the number of targets to predict their positions represented as uncertain spatial distributions. This also enables probabilistic reasoning for effective decision-making by considering these uncertain predictions over longer horizons.

To overcome these limitations, we have proposed a framework for active mapping of freely drifting targets, such as plastic litter, with an ASV. A key component of our approach is a spatiotemporal prediction network that predicts uncertain future target position distributions. We then combine these predictions with our adaptive planning utility. We validate our approach through field deployments with a custom ASV. This paper follows our recent work [19] by providing additional details about the experimental setup, training procedure of the network, and an extended discussion about the limitations of this study and possible future directions for this line of research.

The rest of the paper is organised as follows. First, we formally define the problem in Sec. II. Our approach from [19] is summarized in Sec. III, while Sec. IV provides additional details about the experimental setup and discusses results. We discuss some challenges and potential future research directions in Sec. V. Finally, Sec. VI concludes this study.

II. Problem Statement

We consider the active sensing problem for resource-constrained robots, where the aim is to find an optimal robot sequence of actions $\mathcal{T}^*$ that maximises the following objective:

$$\mathcal{T}^* = \underset{\mathcal{T}\in\Psi}{\arg\max}\, I(\mathcal{T}) \text{ s.t. } C(\mathcal{T}) \leq \ B\,, \tag{1}$$

where $\mathcal{T}$ is a trajectory from the set of all feasible action sequences Ψ, with associated cost $C(\mathcal{T})$, which does not exceed the budget B. An information-theoretic measure, $I(\mathcal{T})$ is used to quantify the informativeness of the measurements by performing these actions, based on the map state.

We aim to build and maintain an accurate occupancy map of dynamic target positions with a limited time budget. We consider the application of mapping floating targets of interest, such as plastic debris, using an ASV. Such targets on the water surface are subject to environmental forces like wind, current, and waves. Mapping in such scenarios requires the ASV to trade off between two potentially conflicting objectives: explore still unvisited regions in the environment, or redetect already detected targets to improve confidence in their positions. Therefore, we have defined a novel informative planning utility for dynamic target tracking scenarios that caters for both objectives to achieve efficient decision-making.

III. Our Approach

An overview of our active sensing approach is shown in Fig. 2. We use a stereo camera onboard the ASV to detect targets of interest and use a dynamic mapping approach, as described in Sec. III-B. We introduce a spatiotemporal network for predicting the uncertain target positions over a given time horizon. The training procedure of the network is also discussed in Sec. III-C. Our proposed informative planning utility, as described in Sec. III-D, leverages predictions from this network. We open-source our code at github.com/sanjeevrs2000/ipp_dyntrack.

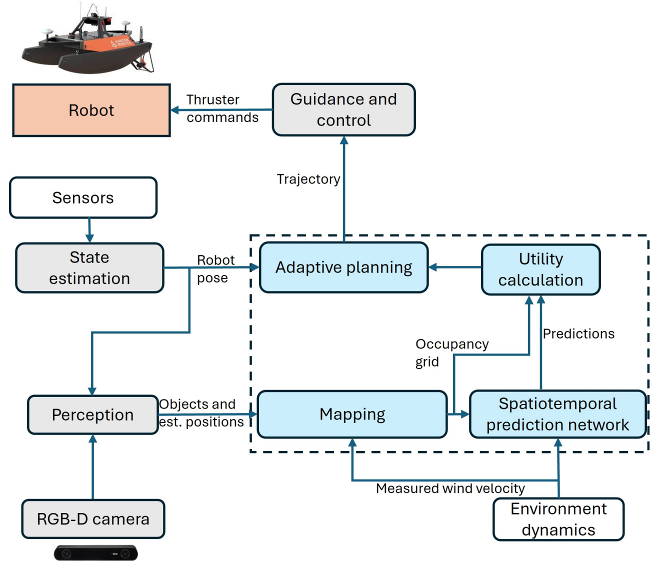

Fig. 2: An overview of our framework. A stereo camera is used to detect and map targets of interest onto a global occupancy map. We employ a dynamic occupancy grid mapping approach to maintain an accurate map of targets that are drifting due to environmental disturbances. We then introduce a neural network to predict the spatiotemporal target position distributions. The predictions from this network are used to guide our adaptive planning strategy.

A. Perception

Targets of interest are detected through object detection with a stereo camera onboard the ASV. We employ YOLOv8 [20] for object detection in this work. It is trained on a custom dataset that is collected from our experiments. The positions of the targets in the occupancy map are computed by first estimating the depth of the object by averaging the depth of pixels inside the bounding boxes of the detections. It is then transformed from the camera frame to the global frame. We then define distance-based inverse sensor models for mapping the targets onto the global occupancy map [19].

B. Dynamic occupancy grid mapping

We use an occupancy grid for mapping the targets of interest. We update the map in a two-stage fashion at every time step. First, the map is updated with the probabilistic occupancy grid mapping algorithm [21]. In the prediction step of the mapping, the map is updated to compensate for the drift of the targets due to external disturbances such as wind and currents. For this, we use the instantaneous wind speed, assuming that measurements are available at each time step. Following [15], [16], we assume that the drift is approximated to be linearly proportional to the wind velocity, neglecting wave effects. Therefore, we assume each target to drift by $(dx, dy) = \Big(\gamma\, v_w \cos(\psi_w)\, \Delta t\,, \gamma\, v_w \sin(\psi_w)\, \Delta t\Big)$, given the measured wind speed v_w and direction ψ_w. For more details of this mapping approach, we refer the reader to [19].

C. Spatiotemporal network

Planning informative paths for the ASV in dynamic environments requires predicting changes in the map over a time horizon. To make efficient predictions of the spatiotemporal variations of target occupancies in the map, we introduce our spatiotemporal network, which uses a modified UNet structure [22]. The input to this network is a binary occupancy grid denoting target positions, and a vector of the measured instantaneous wind speed, and the prediction time step t_k are added as additional inputs into the latent space of the UNet. The output is a grid of the same size with the predicted spatial distributions of target occupancies at time $t + t_k$ in the future.

A binary mask is applied to the current state of the global occupancy grid. This masked binary grid is an input to the network. Each occupied cell in this masked grid is assumed to translate at the current wind speed over the planning horizon. The uncertainty region for each target in the map is represented as a 2D Gaussian kernel with standard deviations proportional to the time horizon and the wind speed. The principal axis of each kernel is assumed to be along the direction of the wind measured at that instant.

To train the network, we generate a synthetic dataset simulating random target positions and compute corresponding output distributions for various combinations of wind conditions and prediction time intervals. We consider $v_w = \{0, 3, \ldots, 12\}$, $\psi_w = \{0, \frac{\pi}{4}, \ldots, \frac{7\pi}{4}\}$, and prediction intervals $t = \{0, 5, \ldots, 25\}$. We generate 100 binary grids of sizes 100×100 with varying numbers of targets and positions. scenarios. To further generalise the training set, in each of these scenarios, the number of targets is chosen between $0 - 20$ with their positions being drawn from a uniform distribution inside the map. This leads to a training set of size 24000 images and associated vector inputs, for diverse combinations of wind conditions and prediction intervals. The network is trained in minibatches with mean absolute error (MAE) as loss, for 100 epochs with an Adam optimiser at a learning rate 0.001. Training is performed on a computer equipped with an Intel Core i7-12700H processor, 32 GB DDR5 RAM, and an NVIDIA RTX 3050 Ti GPU.

D. Adaptive planning

For target tracking in *a priori* unknown environments, the ASV is required to explore the map and attain better estimates of target positions by redetecting them at frequent intervals. Therefore, the ASV also needs to replan at frequent intervals to account for the dynamic nature of the map. We propose a new utility, as shown in Eqs. (2)-(3) for adaptive informative planning that combines the two objectives, i.e. exploration and target tracking.

$$I(\mathcal{T}) = \Big[H(\mathcal{M} \mid z_{\mathcal{T}}) - H(\mathcal{M})\Big] + \frac{w}{|\mathcal{T}|} \sum_{t} J(\boldsymbol{x}^t)\,, \quad (2)$$

$$J\left(\boldsymbol{x}^t\right) = \sum_{i,j} \frac{e^{-2 \cdot p_{ij}(t)}}{|\texttt{fov}(\boldsymbol{x}^t)|},\ \forall (i,j) \in \texttt{fov}(\boldsymbol{x}^t)\,. \quad (3)$$

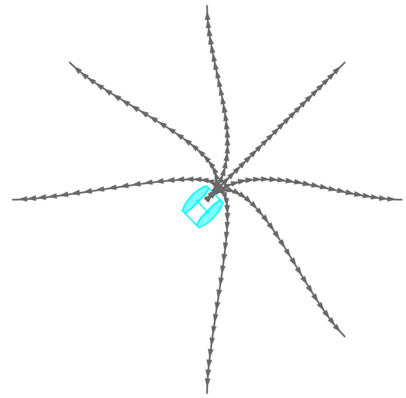

Fig. 3: An illustration of the candidate trajectories with sampling points at regular intervals assuming a constant speed for the ASV. Each trajectory corresponds to a change in heading relative to the current heading of the ASV. The trajectories are then parameterised with Bézier curves to generate smooth and continuous trajectory for the vessel. The utility as described in Eq. (2) is evaluated for each of these trajectories in the replanning stage.

The utility consists of two terms: the predicted change in entropy H and the target tracking component, reflecting the two planning objectives mentioned in Sec. II. The target tracking information gain $J\left(\boldsymbol{x}^t\right)$ at a single time step is computed with predictions from the spatiotemporal network $p(t)$ for poses on the trajectory $\boldsymbol{x}^t$. The factor w is a weighting factor to trade-off between the exploration and target tracking objectives, that is tuned manually.

During replanning, we consider a finite set of candidate trajectories by keeping the speed of the ASV constant, as shown in Fig. 3. Each trajectory corresponds to a heading change from the current heading of the vehicle. They are then parameterised with Bézier curves to generate feasible trajectories for the ASV. We evaluate the utility Eq. (2) for each of these trajectories during the replanning stage and perform planning in a finite-horizon way.

IV. Experiments

A. Experimental setup

The Virtual RobotX simulator [23] is used for simulation experiments. The simulator also models the physics of hydrodynamic effects due to waves and wind, making it a suitable choice for our study. We use the WAM-V, a 4.8 m long ASV, as available in the simulator with a differential thrust configuration in the simulation experiments. Buoy markers are generated in the simulation environment and used as targets of interest. The wind and wave effects on the objects are modelled according to Fossen [24].

We compare our adaptive planning strategy against the following baselines: (i) Sampling-based planner as discussed in Sec. III-D with exploration utility, i.e. only Shannon entropy as utility, to evaluate our proposed utility that balances both exploration and target tracking objectives; (ii) A greedy planner that selects waypoints greedily as opposed to

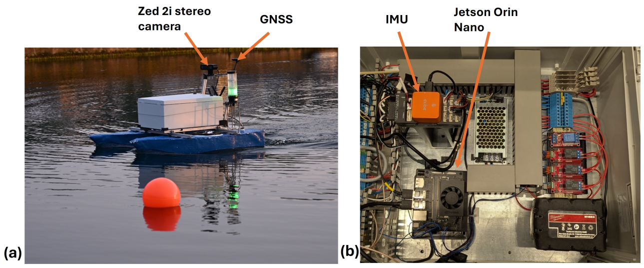

Fig. 4: Sensors onboard our ASV. We use a CubePilot Cube Orange autopilot with integrated IMUs, and GNSS for navigation. All our software is implemented on a Jetson Orin Nano with ROS2 as middleware.

evaluating the utility value over the entire path. We evaluate our approach that uses predictions over longer horizons against a greedy planner that computes the utility with a single time step in the future. For our planner, we adaptively increase the tradeoff factor w (as described in Sec. III-D) linearly from 0 to 5 during the progression of the mission, as this is found to work best from experiments. During replanning, consider seven candidate paths with changes in heading $\Delta\psi$ at intervals $\psi/4$, and a planning horizon of $25\,\mathrm{s}$ with the utility being computed as $1\,\mathrm{s}$ intervals. To compute the utility as described in Eq. (2), we get predictions from the network as a batch as simulate the occupancy grid updates at the same frequency. We note that the replanning for the above-mentioned discretisations and planning horizon takes $\approx 3\,\mathrm{s}$ on our hardware.

We also perform field experiments with an ASV (as shown in Fig. 1), for actively mapping floating targets of interest that are freely drifting. It is a catamaran hull that is 1.2 m long, 0.9 m wide, and weighs 37 kg with payload. The ASV is overactuated with four Blue Robotics T200 thrusters in the X-configuration. It also has an emergency stop button to kill all thrusters in the event of a communication failure. It is equipped with a Jetson Orin Nano developer kit, and our framework is implemented with ROS2 [25] as middleware. An integrated IMU and GNSS are used with the autopilot for navigation of the vehicle. We use line of sight guidance [24] to generate reference trajectories and send the velocity commands to the PID speed controller of the autopilot. Additionally, we use a ZED 2i stereo camera with 72° horizontal field of view. Fig. 4 shows the sensors onboard the ASV.

B. Simulation experiments

The mean entropy in the map H, and the mean detections $\overline{N}$ are chosen as metrics for comparison. The mean detections is calculated as the average number of detections over all time steps up to the current step during a mission. We perform ten Monte-Carlo simulations for a mission time of $250\,\mathrm{s}$ in varying wind conditions and target position distributions. Our complete planning framework with the dynamic map and the spatiotemporal prediction network is evaluated against the two chosen baselines. The resulting

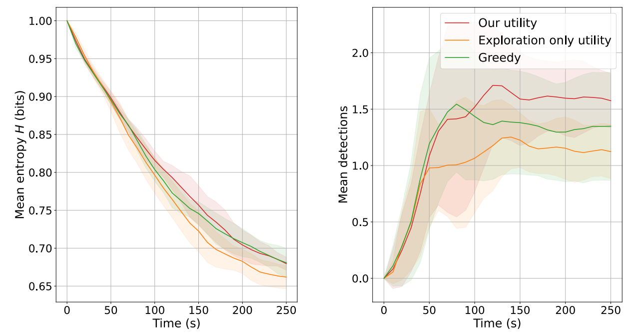

Fig. 5: The progression of the mean Shannon entropy H and the mean target detections $\overline{N}$ plotted over mission time with standard deviation margin for the simulation experiments. The adaptive planner with the purely exploratory reward is most effective at reducing entropy in the map, whereas our utility is most effective at target tracking. This further highlights the trade-off between the two objectives, as outlined in Sec. I.

metrics from the runs are plotted over the mission time in 5. The adaptive planning strategy with the proposed utility is most effective at tracking targets, as it detects up to 17% more targets as compared to the greedy planner on average. The adaptive planner with only entropy as utility is most effective at reducing entropy in the map, but does not track targets as well as the sampling-based planner or the greedy planner, both using our proposed utility. This validates the performance of our proposed planning utility, which uses the spatiotemporal distributions of the targets' positions from our network, thus promoting better target tracking by more frequent redetections. Furthermore, our planning strategy is also better at target tracking as compared to the greedy strategy that computes the utility with a single future time step for planning, rather than considering possible measurements over the entire path.

C. Field tests

For the tests, we use spherical buoy markers, as pictured in Fig. 1, to represent plastic debris in line with our motivating application. Trials are performed on a $25\,\mathrm{m} \times 25\,\mathrm{m}$ area, centered at $(62.469^\circ N,\ 6.237^\circ E)$, for a mission time of $150\,\mathrm{s}$. Our complete framework is run real-time on the embedded compute platform onboard the ASV. The map is updated at 5 Hz, the planning horizon is set to be $25\,\mathrm{s}$, and the speed of the ASV is $0.5\,\mathrm{m/s}$. Replanning is performed once the vehicle reaches the end of the current path in a finite horizon manner. We show results from a field trial in Fig. 6. The targets are initially located on the left edge of the map and slowly drift toward the opposite edge of the map as moderate winds were observed during the test towards the East. We see that the vehicle redetects the targets at several instants during the mission. The video of this field deployment is also released [1].

[1] https://youtu.be/KaOhI2sXhrc

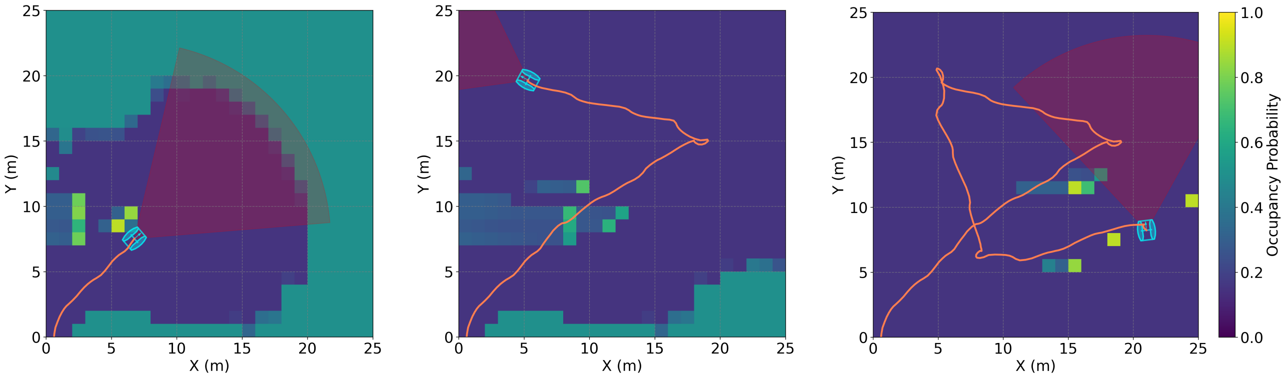

(a) t = 20 s (b) t = 60 s (c) t = 100 s

Fig. 6: Results from a field trial mission with the ASV showing the occupancy grid at various time instants. The ASV first detects the targets at t = 20 s. At t = 60 s, the target position estimates are less confident as the targets go undetected for a while, and later at t = 100 s, the ASV redetects the targets and acquires better target position estimates. A video of the field trial is available at https://youtu.be/KaOhI2sXhrc.

The input grid to the spatiotemporal network is of size 100×100, and we use occupancy grids of size 25×25 in the field tests. We pad the masked occupancy grid with zeros to input a grid of size 100×100, to avoid retraining the network. While the spatiotemporal network is trained entirely on synthetic data, we show that it works in both simulation and field trials without any retraining.

V. Discussion

In this paper, we present an adaptive planning framework as a key step towards active sensing for monitoring in dynamic environments with ASVs, but several open research questions remain. We discuss some possible future directions to address these challenges in this section.

A. *Learning-based techniques for planning*

We perform adaptive planning in a finite-horizon manner. We note that the dynamic nature of the environment necessitates that planning time also be taken into account. The computational burden of our replanning approach complicates decision-making at a higher frequency or using a receding-horizon approach for planning on the fly. In this study, we consider paths with constant speeds for the ASV. To efficiently optimise Eq. (1) for environments that are highly dynamic, considering trajectories with varying speeds is also necessary, further increasing the computational complexity of planning with larger action spaces. Learning-based techniques [26] can help reduce runtime during the replanning stage. Therefore, a direction for future work is to use learning to approximate the information gains [27] to reduce computation times, as evaluating the utility for various trajectories is typically the bottleneck. While several studies have explored the use of reinforcement learning for adaptive planning in active sensing scenarios [4], [18], [28], not many have validated them through field deployments.

B. *Mapping in dynamic environments*

We use a probabilistic occupancy grid with a prediction step that accounts for the drift of the targets. Although we find that this works in practice, it might not be accurate enough or lead to missed target detections for other applications in highly dynamic environments. Using a continuous or feature-based map representation is an alternative approach [5], [18]. Another direction for further research is towards more efficient mapping techniques for dynamic environments to predict the spatiotemporal uncertainties in target positions. Some prior works have explored the use of deep learning based approaches for mapping in dynamic environments [29], [30]. For cases where the dynamics of the environment can be approximated with the known physics, such as in our application, exploring the use of physics-informed learning methods to combine the strengths of both methods could be another research direction.

C. *Multirobot cooperation for active sensing*

Similarly, mapping in larger environments could also lead to inaccurate mapping as the time between successive redetections of a target might be large. Multirobot monitoring presents another promising direction for future research [31], [32]. Furthermore, some studies have explored multirobot planning with a heterogeneous fleet of vehicles where each robot takes different roles [33], [34]. For instance, in a fully integrated system, a team of robots with different hardware capabilities can take diverse roles of scout and cleaner, as proposed in [4].

D. *Foundation models for robot monitoring*

Recently, vision language models (VLMs) are being increasingly used in robot exploration [34], [35]. Thus, it is another avenue of future research towards making autonomous inspection and monitoring more user friendly by acting on language inputs from a human user. They could also help make monitoring tasks generalise well to unseen environments and specific scenarios by using vision and language to get a better understanding of the scene.

VI. Conclusions

In this paper, we present an active sensing approach for mapping moving targets in dynamic marine environments with an ASV. We discuss our experimental setup for both simulation and real-world experiments. Simulation results show that our proposed planning utility improves target tracking as compared to using only entropy reduction as the utility. We show results from field deployments using our ASV with an onboard stereo camera, validating our active sensing approach for tracking an arbitrary number of targets with unknown initial positions. To conclude, we also outline some existing challenges and future research directions.

References

[1] M. Dunbabin and L. Marques, "Robots for environmental monitoring: Significant advancements and applications," *IEEE Robot. & Automat. Mag.*, vol. 19, no. 1, pp. 24–39, 2012.

[2] R. R. Murphy, *Disaster robotics*. MIT press, 2017.

[3] A. Lanzolla and M. Spadavecchia, "Wireless sensor networks for environmental monitoring," *Sensors*, vol. 21, no. 4, 2021.

[4] A. M. Barrionuevo, S. Y. Luis, D. G. Reina, and S. L. T. Marín, "Optimizing plastic waste collection in water bodies using heterogeneous autonomous surface vehicles with deep reinforcement learning," *IEEE Robot. and Automat. Lett.*, 2025.

[5] A. A. Meera, M. Popović, A. Millane, and R. Siegwart, "Obstacle-aware adaptive informative path planning for uav-based target search," in *Proc. IEEE Intl. Conf. on Robot. & Automat. (ICRA)*, 2019.

[6] P. Tokekar, E. Branson, J. Vander Hook, and V. Isler, "Tracking aquatic invaders: Autonomous robots for monitoring invasive fish," *IEEE Robot. & Automat. Mag.*, vol. 20, no. 3, pp. 33–41, 2013.

[7] S. Bai, T. Shan, F. Chen, L. Liu, and B. Englot, "Information-driven path planning," *Current Robotics Reports*, vol. 2, no. 2, pp. 177–188, 2021.

[8] G. A. Hollinger and G. S. Sukhatme, "Sampling-based robotic information gathering algorithms," *Intl. Journal of Robot. Research (IJRR)*, vol. 33, no. 9, pp. 1271–1287, 2014.

[9] M. Popović, T. Vidal-Calleja, G. Hitz, J. J. Chung, I. Sa, R. Siegwart, and J. Nieto, "An informative path planning framework for uav-based terrain monitoring," *Auton. Robots*, vol. 44, no. 6, pp. 889–911, 2020.

[10] B. Moon, S. Chatterjee, and S. Scherer, "Tigris: An informed sampling-based algorithm for informative path planning," in *Proc. IEEE/RSJ Intl. Conf. on Intell. Robots and Sys. (IROS)*, 2022.

[11] C. Robin and S. Lacroix, "Multi-robot target detection and tracking: taxonomy and survey," *Auton. Robots*, vol. 40, pp. 729–760, 2016.

[12] C. Coué, C. Pradalier, C. Laugier, T. Fraichard, and P. Bessière, "Bayesian occupancy filtering for multitarget tracking: an automotive application," *Intl. Journal of Robot. Research (IJRR)*, vol. 25, no. 1, pp. 19–30, 2006.

[13] M. Tay, K. Mekhnacha, M. Yguel, C. Coue, C. Pradalier, C. Laugier, T. Fraichard, and P. Bessiere, "The bayesian occupation filter," in *Probabilistic Reasoning and Decision Making in Sensory-Motor Systems*. Springer, 2008, pp. 77–98.

[14] D. Meyer-Delius, M. Beinhofer, and W. Burgard, "Occupancy grid models for robot mapping in changing environments," in *Proc. Conf. on Advancements of Artificial Intelligence (AAAI)*, 2012.

[15] T. J. Wagner, I. Eisenman, A. M. Ceroli, and N. C. Constantinou, "How winds and ocean currents influence the drift of floating objects," *Journal of Physical Oceanography*, vol. 52, no. 5, pp. 907–916, 2022.

[16] J. Röhrs and K. H. Christensen, "Drift in the uppermost part of the ocean," *Geophysical Research Lett.*, vol. 42, no. 23, pp. 10–349, 2015.

[17] W. Chen and L. Liu, "Multi-objective and model-predictive tree search for spatiotemporal informative planning," in *Proc. Conf. on Decision Making and Control (CDC)*, 2019.

[18] Y. Wang, Y. Wang, Y. Cao, and G. Sartoretti, "Spatio-temporal attention network for persistent monitoring of multiple mobile targets," in *Proc. IEEE/RSJ Intl. Conf. on Intell. Robots and Sys. (IROS)*, 2023.

[19] S. Ramkumar Sudha, M. Popović, and E. M. Coates, "An informative planning framework for target tracking and active mapping in dynamic environments with asvs," *arXiv preprint arXiv:2508.14636*, 2025.

[20] G. Jocher, A. Chaurasia, and J. Qiu, "Ultralytics yolov8," 2023. [Online]. Available: https://github.com/ultralytics/ultralytics

[21] A. Elfes, "Using occupancy grids for mobile robot perception and navigation," *Computer*, vol. 22, no. 6, pp. 46–57, 1989.

[22] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image degmentation," in *Proc. of Image Computing and Computer-Assisted Intervention (MICCAI)*, 2015.

[23] B. Bingham, C. Agüero, M. McCarrin, J. Klamo, J. Malia, K. Allen, T. Lum, M. Rawson, and R. Waqar, "Toward maritime robotic simulation in gazebo," in *Proc. MTS/IEEE OCEANS SEATTLE*, 2019.

[24] T. I. Fossen, *Handbook of marine craft hydrodynamics and motion control*. John Wiley & Sons, 2011.

[25] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, A. Y. Ng, *et al.*, "Ros: an open-source robot operating system," in *ICRA workshop on open source software*, 2009.

[26] M. Popović, J. Ott, J. Rückin, and M. J. Kochenderfer, "Learning-based methods for adaptive informative path planning," *Robot. and Auton. Sys.*, vol. 179, p. 104727, 2024.

[27] L. Schmid, C. Ni, Y. Zhong, R. Siegwart, and O. Andersson, "Fast and compute-efficient sampling-based local exploration planning via distribution learning," *IEEE Robot. and Automat. Lett.*, vol. 7, no. 3, pp. 7810–7817, 2022.

[28] J. Rückin, L. Jin, and M. Popović, "Adaptive informative path planning using deep reinforcement learning for uav-based active sensing," in *Proc. IEEE Intl. Conf. on Robot. & Automat. (ICRA)*, 2022.

[29] K. S. Mann, A. Tomy, A. Paigwar, A. Renzaglia, and C. Laugier, "Predicting future occupancy grids in dynamic environment with spatio-temporal learning," in *Proc. IEEE Intell. Veh. Symp. (IV)*, 2022.

[30] M. Toyungyernsub, M. Itkina, R. Senanayake, and M. J. Kochenderfer, "Double-prong convlstm for spatiotemporal occupancy prediction in dynamic environments," in *Proc. IEEE Intl. Conf. on Robot. & Automat. (ICRA)*, 2021.

[31] S. Manjanna, A. Q. Li, R. N. Smith, I. Rekleitis, and G. Dudek, "Heterogeneous multi-robot system for exploration and strategic water sampling," in *Proc. IEEE Intl. Conf. on Robot. & Automat. (ICRA)*, 2018.

[32] P. De Petris, S. Khattak, M. Dharmadhikari, G. Waibel, H. Nguyen, M. Montenegro, N. Khedekar, K. Alexis, and M. Hutter, "Marsupial walking-and-flying robotic deployment for collaborative exploration of unknown environments," in *Proc. IEEE Intl. Symp. on Safety, Security, and Rescue Robot. (SSRR)*, 2022.

[33] F. Shkurti, A. Xu, M. Meghjani, J. C. G. Higuera, Y. Girdhar, P. Giguere, B. B. Dey, J. Li, A. Kalmbach, C. Prahacs, *et al.*, "Multi-domain monitoring of marine environments using a heterogeneous robot team," in *Proc. IEEE/RSJ Intl. Conf. on Intell. Robots and Sys. (IROS)*, 2012.

[34] Z. Ravichandran, V. Murali, M. Tzes, G. J. Pappas, and V. Kumar, "Spine: Online semantic planning for missions with incomplete natural language specifications in unstructured environments," in *Proc. IEEE Intl. Conf. on Robot. & Automat. (ICRA)*, 2025.

[35] N. Yokoyama, S. Ha, D. Batra, J. Wang, and B. Bucher, "Vlfm: Vision-language frontier maps for zero-shot semantic navigation," in *Proc. IEEE Intl. Conf. on Robot. & Automat. (ICRA)*, 2024.