*Information-Driven Path Planning*



**As Published:** https://doi.org/10.1007/s43154-021-00045-6

**Publisher:** Springer International Publishing

**Persistent URL:** https://hdl.handle.net/1721.1/136804

**Version:** Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

# Information-Driven Path Planning

**Cite this article as:** Shi Bai, Tixiao Shan, Fanfei Chen, Lantao Liu and Brendan Englot, Information-Driven Path Planning, *Current Robotics Reports* doi: 10.1007/s43154-021-00045-6

# Information-Driven Path Planning

Shi Bai*[1], Tixiao Shan†[2], Fanfei Chen‡[3], Lantao Liu§[4], and Brendan Englot¶[3]

[1]Wing, Alphabet Inc. Mountain View, CA 94039
[2]Senseable City Lab, Massachusetts Institute of Technology, Cambridge, MA 02139
[3]Department of Mechanical Engineering, Stevens Institute of Technology, Hoboken, NJ 07030
[4]Luddy School of Informatics, Computing, and Engineering, Indiana University-Bloomington, IN 47408



**Abstract**

**Purpose of Review** The era of robotics-based environmental monitoring has given rise to many interesting areas of research. A key challenge is that robotic platforms and their operations are typically constrained in ways that limit their energy, time, or travel distance, which in turn limits the number of measurements that can be collected. Therefore, paths need to be planned to maximize the information gathered about an unknown environment while satisfying the given budget constraint, which is known as the informative planning problem. This review discusses the literature dedicated to information-driven path planning, introducing the key algorithmic building blocks as well as the outstanding challenges.
**Recent Findings** Machine learning approaches have been introduced to solve the information-driven path planning problem, improving both efficiency and robustness.
**Summary** This review started with the fundamental building blocks of informative planning for environment modeling and monitoring, followed by integration with machine learning, emphasizing how machine learning can be used to improve the robustness and efficiency of informative path planning in robotics.
**Keywords** informative path planning, environmental sensing and modeling, mapping and exploration

## 1 Introduction

Path planning is one of the most critical capabilities for autonomous robots operating in complex unstructured environments [57]. Particularly, in many outdoor missions the robots (shown in Figure 1) need to leverage information perceived from environments to look ahead and plan actions for the subsequent steps. Such information-driven path planning has been extensively

*E-mail: {baishi}@wing.com. Corresponding author.
†E-mail: {shant}@mit.edu
‡E-mail: {fchen7}@stevens.edu
§E-mail: {lantao}@iu.edu
¶E-mail: {benglot}@stevens.edu

(a) (b) (c)

Figure 1: (a) Autonomous surface vehicle for water quality sensing and modeling. (b) Unmanned aerial vehicle for air quality monitoring. (c) Autonomous ground vehicle for 3D mapping.



investigated by integrating with information theoretic and decision theoretic paradigms as well as learning based frameworks. The goal of this paper is to provide a survey on the most recent research of the information-driven path planning, and also to provide insights on future challenges of related directions. Since the information-driven planning involves interactions with environments, we will base our discussions on two parts: i) environmental sensing and modeling where informative knowledge is continuously retrieved and utilized for planning during the long-term autonomy; ii) exploration in unknown environment where a map is built autonomously for robot accurate navigation.

Environmental sensing and modeling is a process that given a certain environmental attribute of interest (e.g., elevation of an uneven terrain, concentration of some pollutant), corresponding measurement samples are collected from different locations so that a continuous map describing the levels and variations of environmental attribute can be constructed [42]. In other words, the map is not simply the averaged value of collected samples, but a continuous map that describes or predicts the variance across an entire spatiotemporal field. The environmental sensing and modeling allow scientists to assess the processes of a particular environment, and have been used in a broad range of applications. For instance, stationary sensor networks [1] have been used to perform fixed-location sampling to detect forest fires [78] and volcanic activity [125]. However, such sensor networks lack flexibility for capturing high resolutions in critical regions that might change spatially and temporally, and this limitation can be addressed by using mobile robots, which can also significantly reduce the costs of sensor deployment and maintenance. To date, unmanned aerial vehicles (UAVs) have been used to estimate yields of crops or fruits [94, 131] and to study spatial ecology and its spatiotemporal dynamics [3]; with a capacity for performing long-range and long-duration tasks, marine robots can collect large-area ocean data [113, 29] and trace chemical plumes [46, 59]; autonomous boats have been used to monitor fish schools [121]. Our application examples are also shown in Fig. 1. However, a drawback that cannot be neglected when using robots is that the collection of environmental data requires a series of sequential motion actions (and measurement operations), and the whole course takes time because the measurements typically need to spread across many different spatial locations [36]. This requires that a robot's sensing must be adaptive to environmental attributes (e.g., spatiotemporal variations or distributions). Also, because sensor data is continuously collected, learning methodologies therefore also play an important role in equipping robots with intelligent sensing behaviors to aid the measurement (or sampling) process.

 

# 2 Informative Planning

Navigating robots to collect environmental data samples offering the largest amount of new information is called informative planning [13, 87, 112]. The goal is to maximize information gain (or *informativeness*), which may be derived from a robot's evolving map, its estimation uncertainty, and/or prediction models of environmental phenomena being sampled. Compared with lawnmower-based sampling of the environment, which focuses on spatial resolution, informative planning methods tend to achieve spatial coverage quickly, while managing estimation uncertainty [112]. Due to these reasons, informative planning has been widely used for spatiotemporal environmental sensing, modeling, and monitoring.

## 2.1 Environment Modeling

To model spatial phenomena in continuous domains, a widely adopted method is the Gaussian Process (GP), which is a generic supervised learning method capable of solving regression and probabilistic classification problems [102, 112, 95]. Built on an environmental model learned with the aid of GPs the path planning and motion control can be developed. The regression capability of GPs has been proven a powerful tool for predicting environment states based on the subset of the environment that can be monitored. In the geostatistics or spatial statistics literature, the GP regression technique is often called *kriging* [60], and is mostly used to analyze spatial properties. Oftentimes kriging relies on the knowledge of certain spatial structure, which is modeled via second-order properties, e.g., the variogram or covariance, of the underlying random function [74]. There are many applications where GPs have been utilized as a framework for modeling the environment. For instance, GPs have been used to design placement patterns of static sensors in a sensor network so that the environment model can be predicted with a solution that is near-optimal [56]. GP-aided optimization of static sensor placement has been applied to modeling indoor 3D environments [45] and outdoor urban environments [77] using appropriate kernels (covariance functions).

For robotics information-driven planning, the GP modeling is usually combined with Bayesian learning property to assimilate data collected online. For example, GPs have been used on a mobile robot to build a spatial model describing gas distribution [116], to provide a measure of uncertainty to guide sensor-centric robot localization [17], and by modeling mutual information with the aid of Bayesian optimization, GPs have been used to guide robots to explore unknown static environments [5]. In dynamic settings, variants of GPs have also been employed to learn uncertainty models of environmental processes, to aid the operation of autonomous underwater vehicles (AUVs) in the ocean [62, 61]. By integrating with vehicle routing and communication constraints, methods have been developed for informative ocean sampling and monitoring in complex ocean environments with multi-robot systems [84, 65, 51].

To improve the prediction accuracy of GPs, the choice among different prior covariance functions and the update of its hyperparameters are crucial, especially in scenarios involving spatiotemporal dynamics. This problem is typically referred to as *model selection and adaptation of hyperparameters* [102]. Specifically, the adaptation of hyperparameters can be achieved using a data-driven approach. The most common approach is to maximize the marginal likelihood, or minimize the generalization error, using a cross-validation approach. For the case of GP classification, other optimization criteria such as *alignment* [34] can also be adopted.

## 2.2 Myopic vs. Non-myopic Planning

A variety of methodologies have been proposed to tackle the informative planning problem, among which the most investigated approaches belong to the *nonmyopic* framework. Formally, the term *myopic* means that the path waypoints are computed individually and greedily, without considering the cost and consequences of making observations over a long horizon into the future. Instead, a nonmyopic strategy performs optimization and computes a series of waypoints by considering the effect of later time-steps [87, 134]. Representative nonmyopic informative planning approaches include, for example, a *recursive-greedy* based algorithm [112] where the informativeness is generalized as a *submodular set function*, upon which a sequential-allocation mechanism is designed in order to obtain subsequent waypoints. This recursive-greedy framework has been extended by taking into account the avoidance of shipping lanes [12] and diminishing returns [13] in the marine planning environments. Differing from the above mechanisms where the path waypoints are built by separate searching techniques with the informativeness as a utility function, Low [79] proposed a differential entropy-based planning method in which a batch of waypoints can be obtained through solving a dynamic program. Such a framework has been extended to approaches incorporating mutual information [22] and Markovian [80] optimization criteria. An informative planning method based on dynamic programming was recently proposed to compute informative waypoints across an arbitrary continuous space [83]. This nonmyopic method has also been combined with Markov Decision Processes to cope with a robot's action uncertainty caused by external disturbances. In addition, there are also many methods that optimize over complex planning and information constraints (e.g., [114]).

## 2.3 Online Planning

A critical problem for persistent tasks (long-term, or even life-long autonomy) is the large-scale data accumulated. Although "big" data might predict the most accurate model, in practice large amounts of data can exceed a robot's onboard computational capacity. Methods for reducing the computing burdens of GPs have been previously investigated. For example, GP regression can be performed in a real-time fashion where the problem can be estimated locally, with local data [92]. Another potentially suitable framework is a sparse representation of the GP model [35, 85, 29] which is based on a combination of a Bayesian online algorithm together with a sequential construction of the most relevant subset of the data [82]. This method allows the model to be refined in a recursive way as the data streams in. This framework has been further extended to many application domains, such as visual tracking [101] and spatial modeling [116].

To reduce both the length of a path and the probability of collisions, pareto optimization has been used in designing path planners [33]. Recently, a sampling-based method has also been proposed to generate Pareto-optimal trajectories for multi-objective motion planning [72]. In addition, multi-robot coordination also benefits from multi-objective optimization. The goals of different robots can be simultaneously optimized [54, 70]. To balance the operation cost and the travel discomfort experienced by users, multi-objective fleet routing algorithms compute Pareto-optimal fleet operation plans [23]. Related work also includes multi-objective reinforcement learning [103], including multi-objective Monte Carlo tree search (MO-MCTS). However, MO-MCTS is computationally prohibitive and cannot be used for online planning. Vast computational resources are needed in order to maintain a global Pareto optimal set with all the best solutions obtained so far. Recently, a framework was developed which maintains a local approximate Pareto optimal set in each node which can be processed in a much faster way [31]. This approach is also flexible and adaptive with regard to capturing environmental variability [30].

## 2.4 Dealing with Motion Uncertainty

If the environment is complex, unstructured, and even dynamic, then the robot motion/action outcomes can become highly uncertain. When motion uncertainty is considered, stochastic methods, e.g., decision theoretic planning based on the Markov Decision Process [16, 98], stochastic optimal control [11, 47], stochastic model predictive control [108, 100], have been broadly used to cope with the stochasticity.

Comparing with deterministic and many probabilistic motion planners, an advantage of the decision theoretic framework is that it exploits the stochastic structure of the world model to formulate the motion uncertainty using a stochastic state transition function and to enforce rewards/penalties for outcomes of future actions. Building upon the MDP architecture, a partially observable MDP (POMDP) is a generalization of MDPs to situations where an agent cannot reliably identify the underlying environment state [64]. Such an extension dramatically increases the computational complexity, making exact solutions virtually intractable. Planning in dynamic environments in the presence of other non-stationary objects can be modeled with the MDP or POMDP frameworks. For example, existing methods [52, 86, 37] formulate this problem as POMDPs where the behavior of other dynamic objects are not observable but assumed to be selected from a fixed number of closed-loop policies. The deterministic rollouts are then used to determine the best policy to execute. A similar work [8] models this problem as a mixed observability MDP, which is a variant of POMDP [68]. A more general framework is proposed in [40] where the authors combine motion prediction and receding horizon planning to reduce the uncertainty during planning. An important performance metric for decision-theoretic methods is the computational efficiency especially in robot online planning scenarios. Recently, reachability based methods have been exploited [39, 128] to mitigate the computational challenges where the key idea is to identify a small subset of states that contribute most to the reward optimization and then prioritize exploitation of this subset of states.

# 3 Autonomous Mapping and Exploration

The informative planning focuses on the spatiotemporal information of the environment and can work particularly well for non-continuous, sparse and limited sensing data. Apart from the previously discussed environmental modeling and informative planning, there is another type of information gathering framework which focuses on more accurate representation (usually 3D) of the environmental structure. This is based on environment mapping and exploration which are also crucial for precisely controlling robot motion in GPS-denied environments.

Building an accurate map requires the robot to a) localize itself while mapping; b) generate a path consist of informative view points. In this section we survey robot state estimation methods and include different exploration approaches for efficient mapping. We highlight the recent development on machine learning related approaches and extend to outstanding challenges in the next section.

## 3.1 Localization without Prior Map

Mobile robots relies on robust state estimation to perform accurate motion control. The SLAM [20] has been a standard framework to real-time estimate robot state and map environments of all kinds. As a ranging sensor, LiDAR, which stands for light detection and ranging, is invariant to environmental lighting and can capture the environment's fine details in a long range. Therefore, range-based state estimation methods that utilize LiDAR have achieved great success in accuracy. Among them, the LiDAR odometry and mapping method proposed (LOAM) in [135] is the most widely used. LOAM divides the state estimation problem into two sub-problems: a scan-to-scan matching is performed at high frequency but with low accuracy; a scan-to-map matching is performed at low frequency but with high accuracy. A lightweight LiDAR odometry



 

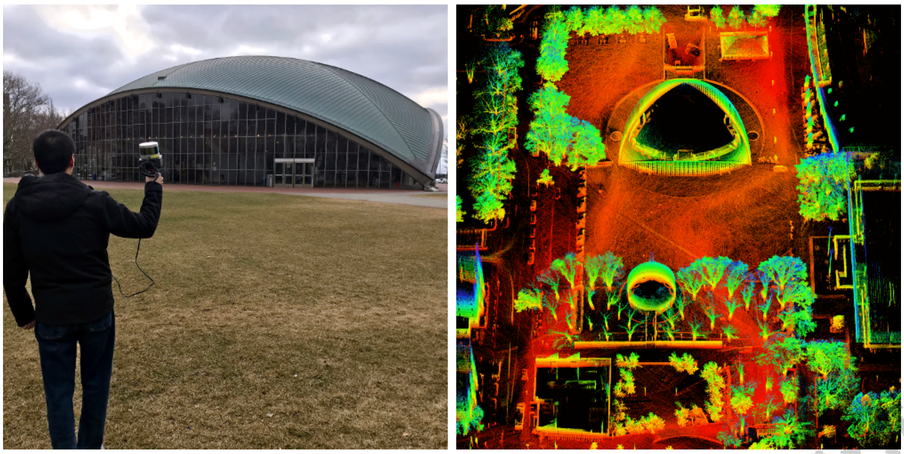

(a) (b)

Figure 2: Representative results of LIO-SAM [111]. (a) An operator carries the sensor suite and performs a mapping task. (b) A point cloud map built using only LiDAR and IMU data. Changes in color indicate elevation change.

and mapping (LeGO-LOAM) approach is proposed in [110] that further improves LOAM's efficiency and accuracy by introducing point cloud segmentation. Both LOAM and LeGO-LOAM use IMU for point cloud de-skew and obtaining motion prior. However, the measurements from IMU are not involved in the optimization stage of the algorithms. Thus these two methods can be categorized as loosely-coupled LiDAR odometry methods. Lately, tightly-coupled LiDAR odometry methods [71, 99, 133] draw more attention as they provide better accuracy and robustness in LiDAR degraded environments, such as a long corridor. A tightly-coupled LiDAR inertial odometry and mapping framework, LIO-mapping, is introduced in [133]. LIO-mapping jointly optimizes measurements from IMU and LiDAR and achieves high robustness when compared with LOAM. However, LIO-mapping's performance suffers in a feature-rich environment as its global voxel map becomes dense for scan-matching. LIO-SAM [111] overcomes the issue by introducing a sliding window-based scan-matching approach. Performing scan-matching at a local scale instead of a global scale significantly improves the system efficiency without sacrificing accuracy. Some representative results of LIO-SAM are shown in Figure 2.

Vision-based SLAM approaches usually incorporate monocular or stereo cameras, inertial measurement units (IMU) and other modes of odometry sensing (such as wheel encoders) when available. A system involves only visual information is called visual odometry [106], which can be solved either by extracting sparse features (indirect method [90]) or use the intensity of each pixel directly (direct method [48]). Indirect methods use feature descriptors such as SIFT [75] or FREAK [2] to find correspondences between different frames, which pays the cost of feature extraction and assumes the motion of the camera can be recovered by the selected feature points. Direct methods however take pixel intensity directly and tend to use more pixels (sometimes all of the pixels) for robustness and accuracy, which can be very computationally expensive. Recently improvements on direct method [49] significantly improved efficiency while preserving the accuracy. After the correspondences of visual appearances are determined in a series of representative frames (key frames), a joint optimization can be performed to further improve on accuracy, this process is called bundle adjustment [117]. Both the 3D location of the visual features and the pose of each key frame are solved in bundle adjustment by minimizing the re-projection error of the features. While visual odometry can recover the motion as well as the 3D location of the feature points, it does not sense the scale [50, 73], which is usually estimated with IMU (and other odometry sensing if available).

Also, the localization uncertainty of the robot may grow unbounded if the SLAM system solely depends on dead reckoning (e.g, through wheel odometry, LiDAR odometry, or visual odometry). Fully developed SLAM frameworks [110, 111] could utilize either global position systems (GPS), or loop closure, to eliminate the drift incurred during motion. Therefore, many planning methods have been proposed in the last two decades to reduce the uncertainty of localization by choosing constrained paths or re-visiting landmarks that was previously seen. Among those planning under uncertainty algorithms, sampling-based methods have drawn great attention due to their capability in high-dimensional settings. The belief roadmap (BRM) [97] finds paths with lower uncertainty by utilizing extended Kalman filter (EKF) covariance factorization when searching for a valid path. By propagating an EKF over rapidly-exploring random graphs (RRG), rapidly-exploring random belief trees (RRBT) [19] yield shortest paths subject to constraints on collision probability. Chance-constrained RRT* (CC-RRT*) [81] incorporates collision risk as an objective during planning and reduces a robot's collision probability. The robust belief roadmap [14] introduces a novel uncertainty metric, which is an upper bound on the maximum eigenvalue of the EKF covariance matrix. This eigenvalue bound metric provides optimal substructure property, which is not available in EKF-based methods. Utilizing this new metric, a min-max rapidly exploring random tree (MMRRT*) [44] proposes a framework with lexicographic optimization for planning under uncertainty. A belief roadmap search method is discussed in [109] that improves the computation efficiency by utilizing a best-first search scheme.

Other approaches focusing on reducing both uncertainty of a robot's pose and entropy of the map were developed. The initial attempt [115] uses the particle filter and captured a trajectory's uncertainty using the particle weight. Recent work [122] considered the correlation between localization and information gain to reduce uncertainties. More recently a Expectation-Maximization (EM) exploration algorithm [123] introduces *virtual landmarks* to represent the uncertainty of unknown regions, with a novel utility function it solves for optimal map accuracy given all the possible sensing actions.

## 3.2 Data-Driven Exploration

To quickly build a map of the environment, although frontier-based exploration approaches [130] achieve reasonable efficiency in 2D mapping scenarios (assuming perfect localization), information theoretic approaches have shown great advantages with the aid of map entropy, where the mutual information (MI) between a robot's sensor observations and the cells of its map is a key ingredient used to guide a map-building robot to uncover unknown regions. However, as a robot incrementally reduces the entropy of its map by collecting sensor observations (e.g., from a depth camera or LiDAR), the updates to the map will also update the mutual information, even when multiple observations are gathered from the same viewpoint.

### 3.2.1 Conventional Paradigm

Among the earliest information theoretic exploration strategies are those proposed by Elfes [43] and Whaite and Ferrie [126]. The former focused on maximizing the MI between sensor observations and an occupancy grid map, and the later proposed to explore an *a priori* unknown environment with minimizing the map entropy. More recent works [15], [115] brought up the trade-off between maximizing MI and managing the localization uncertainty in a robot's SLAM process, in addition to selecting trajectories in favor of map accuracy [67]. In order to reduce computational cost of MI, efforts made to consider small, predetermined candidate trajectories, using a skeletonization of the known occupancy map [66] and the evaluation of information gain over a finite number of motion primitives [132], [24], 3D viewpoints [127], or compression on occupancy grid map [91]. There are also learning-based exploration methods which try to make inferences of the information gain associated with candidate view points, e.g. with Gaussian Process Regression as the model [7], or through Bayesian optimization to improve on sampling view points [6].

### 3.2.2 Deep Neural Networks

Deep neural networks have been successfully applied to challenging problems such as image recognition [104], robot manipulation [96] and control [38], [136]. Recent work on obstacle avoidence [118] successfully trained a deep neural network which took RGB images as input and generated steering angles as output, while a robot was moving forward at constant speed. A novel approach for visual navigation proposed in [137] also took RGB images as input, and the learned model was able to recognize the cues for navigation to a target for which only the apperance of the target was exposed to the network. Another work in [21] used a deep neural network to detect exit locations from building blueprints.

More recently, in [4] and [25], high-quality exploration solutions can be trained via deep neural networks with supervised and reinforcement learning using local occupancy maps as input data. Moreover, in [93], the best next view frontier is estimated via a learned DRL policy using global occupancy maps as input data.

While convolution neural networks became more and more popular, graphs based networks can offer generalized topological representations of robot navigation, allowing a much smaller state space than metric maps. Graph neural networks (GNNs) [107] incorporate graphs into neural network models to do learning tasks for many fields. Graph Nets [9] is adopted for solving control problems for dynamical systems [105]. Chen et al. [27], [26] proposed to use GNNs with supervised learning [27] and reinforcement learning [26] to perform robot exploration under localization uncertainty. In [28], a novel neural network model is proposed to solve the robot navigation problem using a camera image and a topological map. Wang et al. [124] proposed a deep graph neural network with reinforcement learning to learn a scalable control policy.

### 3.2.3 Reinforcement Learning

Compared to supervised learning, model-based reinforcement learning methods often accelerate the training process [55] since the agent can obtain training information from a model in addition to rewards from the environment. A pre-trained policy can also improve the performance of learning and increase a robot's learning efficiency [41]. Jaderberg proposed an auxiliary approach to explore the potentially reward-rich areas of an environment, avoiding low-reward areas [63]. Such methods require *a priori* knowledge of either the environments or the policy model.

Tai introduced a deep reinforcement learning [119] based obstacle avoidance policy that is trained and tested using sensor data from the same environment. However, it is ineffective to use learned policy across heterogeneous environments. Combining an RNN and a DQN, [69] proposed Deep Recurrent Q-Networks (DRQN) to play First-Person-Shooting (FPS) games in a 3D environment. The DRQN can generate appropriate outputs that depend on the temporally consecutive inputs. Meanwhile, in [137], the robot learned a target-driven visual navigation policy in a simulated environment and implemented its policy successfully in a physical environment. However, it is difficult to predict the information gain of many potential future sensing actions using a single camera view and the requirement for the simulation environment is significantly high. ExpLOre algorithm [32] is proposed to gather information using imitation learning. the non-myopic solutions are provided by this algorithm but the training and testing environments are the same, with different view nodes. Besides, for robot navigation problems, learning the topology of an environment [58] and motion planning [120] can be solved by reinforcement learning.

# 4 Challenges and Opportunities

While information driven planning and exploration algorithms have been deployed on real robot platforms [83], it is far from a solved problem. Since the environment is unknown and could be dynamic, it is hard to estimate the environment model, and thus it is challenging to avoid myopic decision-making. In addition, the complex and dynamic nature of the environment brings external disturbances, which makes a robot's motion stochastic. These external disturbances need to be considered when planning, otherwise the actual execution result may be far from the expected result.

An interesting problem for environmental sensing and modeling lies in environmental dynamics which require time-varying Markov transition models [76, 129]. The Markov transition dynamics can be time-varying, and a literature review informs us that time-varying Markov models have been investigated for pattern analysis of the economic growth [89, 10], which aims at understanding the dynamics of growth based on a collection of different states corresponding to different countries. However, these existing models have been constructed on hidden Markov models, and assume that there is no action to control state transitions. One drawback of state-of-the-art decision-making methodologies lies in the fact that, the basic model relies on a fixed and exact form of uncertainty probabilistic distribution. Such frameworks may be upgradable to some sort of online planning methods through catching up to the latest dynamics and a series of repetitive replanning processes. However, replanning strategies still cannot properly take into account future time-varying uncertainty dynamics which if properly considered can be beneficial. In essence, we believe that the planning and decision-making community still lacks a general methodology to compute solutions that consider not only a fully known (current and past) stochastic description, but also a (possibly uncertain) prediction of future dynamics.

To enable a robot to effectively learn to act in complex and especially unknown environments, transfer learning provides a method for avoiding the expensive training cost in real-world environments by doing the training process in simulation environments. An off-line interactive replay recorded from a real-world environment for one-shot reinforcement learning is proposed in [18]. To fill the gap between robot training in real and simulated environments, domain randomization is utilized for transfer learned policy [88]. Zero-shot reinforcement learning for autonomous vehicle driving problem is proposed in [53] by extracting important features from an input image with the attention mechanism. While there are more and more datasets for training a neural network model to perform information gathering tasks, generalization capability remains a key challenge and requires further extensive investigations.



# 5 Conclusion

The purpose of this paper is to provide a survey on the most recent research of the information-driven path planning for mobile robots, and also to provide insights on future challenges of related directions. In many outdoor missions the robots need to leverage sensed information of environments to look ahead and plan actions. Using the scenarios of environmental sensing, modeling and monitoring, we first discuss key information-driven planning modules including environmental modeling with Gaussian process, myopic vs non-myopic planning, online planning and motion with uncertainty. We then discuss the autonomous mapping and exploration, followed by integration with data-driven methodologies, emphasizing how machine learning can be used to improve the robustness and efficiency of existing SLAM frameworks. The enhancement of these these critical components will lead to more robust and adaptive frameworks with increased functionalities that shall allow mobile robots to perform missions in highly cluttered and unstructured environments.

## Compliance with Ethics section

**Conflict of Interest**: The authors declare that they have no conflict of interest.
**Human and Animal Rights and Informed Consent**: This article does not contain any studies with human or animal subjects performed by any of the authors.

## References


[1] Ian F Akyildiz, Weilian Su, Yogesh Sankarasubramaniam, and Erdal Cayirci. Wireless sensor networks: a survey. *Computer networks*, 38(4):393–422, 2002.

[2] Alexandre Alahi, Raphael Ortiz, and Pierre Vandergheynst. Freak: Fast retina keypoint. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 510–517. Ieee, 2012.

[3] Karen Anderson and Kevin J Gaston. Lightweight unmanned aerial vehicles will revolutionize spatial ecology. *Frontiers in Ecology and the Environment*, 11(3):138–146, 2013.

[4] Shi Bai, Fanfei Chen, and Brendan Englot. Toward autonomous mapping and exploration for mobile robots through deep supervised learning. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2379–2384. IEEE, 2017.

[5] Shi Bai, Jinkun Wang, Fanfei Chen, and Brendan Englot. Information-theoretic exploration with bayesian optimization. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2016, Daejeon, South Korea, October 9-14, 2016*, pages 1816–1822, 2016.

[6] Shi Bai, Jinkun Wang, Fanfei Chen, and Brendan Englot. Information-theoretic exploration with bayesian optimization. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1816–1822. IEEE, 2016.

[7] Shi Bai, Jinkun Wang, Kevin Doherty, and Brendan Englot. Inference-enabled information-theoretic exploration of continuous action spaces. In *Robotics Research*, pages 419–433. Springer, 2018.

[8] Tirthankar Bandyopadhyay, Kok Sung Won, Emilio Frazzoli, David Hsu, Wee Sun Lee, and Daniela Rus. Intention-aware motion planning. In *Algorithmic foundations of robotics X*, pages 475–491. Springer, 2013.

[9] Peter W Battaglia, Jessica B Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, et al. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018.




 

[10] Marco Bazzi, Francisco Blasques, Siem Jan Koopman, and André Lucas. Time varying transition probabilities for markov regime switching models. Tinbergen Institute Discussion Papers 14-072/III, Tinbergen Institute, 2014.

[11] Dimitir P Bertsekas and Steven Shreve. *Stochastic optimal control: the discrete-time case.* 2004.

[12] Jonathan Binney, Andreas Krause, and Gaurav S. Sukhatme. Informative path planning for an autonomous underwater vehicle. In *International Conference on Robotics and Automation*, pages 4791–4796, 2010.

[13] Jonathan Binney, Andreas Krause, and Gaurav S. Sukhatme. Optimizing waypoints for monitoring spatiotemporal phenomena. *International Journal on Robotics Research (IJRR)*, 32(8):873–888, 2013.

[14] Shaunak D Bopardikar, Brendan Englot, Alberto Speranzon, and Jur van den Berg. Robust Belief Space Planning Under Intermittent Sensing via a Maximum Eigenvalue-based Bound. *The International Journal of Robotics Research*, 35(13):1609–1626, 2016.

[15] Frederic Bourgault, Alexei A Makarenko, Stefan B Williams, Ben Grocholsky, and Hugh F Durrant-Whyte. Information based adaptive robotic exploration. In *IEEE/RSJ international conference on intelligent robots and systems*, volume 1, pages 540–545. IEEE, 2002.

[16] Craig Boutilier, Thomas Dean, and Steve Hanks. Decision-theoretic planning: Structural assumptions and computational leverage. *Journal of Artificial Intelligence Research*, 11:1–94, 1999.

[17] Alex Brooks, Alexei Makarenko, and Ben Upcroft. Gaussian process models for indoor and outdoor sensor-centric robot localization. *IEEE Transactions on Robotics*, 24(6):1341–1351, 2008.

[18] Jake Bruce, Niko Sünderhauf, Piotr Mirowski, Raia Hadsell, and Michael Milford. One-shot reinforcement learning for robot navigation with interactive replay. *Advances in neural information processing systems*, 2017.

[19] Adam Bry and Nicholas Roy. Rapidly-exploring Random Belief Trees for Motion Planning Under Uncertainty. In *IEEE International Conference on Robotics and Automation*, pages 723–730, 2011.

[20] Cesar Cadena, Luca Carlone, Henry Carrillo, Yasir Latif, Davide Scaramuzza, José Neira, Ian Reid, and John J Leonard. Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age. *IEEE Transactions on robotics*, 32(6):1309–1332, 2016.

[21] Jeffrey A Caley, Nicholas RJ Lawrance, and Geoffrey A Hollinger. Deep learning of structured environments for robot search. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3987–3992. IEEE, 2016.

[22] Nannan Cao, Kian Hsiang Low, and John M. Dolan. Multi-robot informative path planning for active sensing of environmental phenomena: A tale of two algorithms. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems*, pages 7–14, 2013.



 

[23] Michal Cáp and Javier Alonso-Mora. Multi-objective analysis of ridesharing in automated mobility-on-demand. In *Robotics: Science and Systems*, 2018.

[24] Benjamin Charrow, Gregory Kahn, Sachin Patil, Sikang Liu, Ken Goldberg, Pieter Abbeel, Nathan Michael, and Vijay Kumar. Information-theoretic planning with trajectory optimization for dense 3d mapping. In *Robotics: Science and Systems*, volume 11, 2015.

[25] Fanfei Chen, Shi Bai, Tixiao Shan, and Brendan Englot. Self-learning exploration and mapping for mobile robots via deep reinforcement learning. In *AIAA Scitech 2019 Forum*, page 0396, 2019.

[26] Fanfei Chen, John D Martin, Yewei Huang, Jinkun Wang, and Brendan Englot. Autonomous exploration under uncertainty via deep reinforcement learning on graphs. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020.

[27] Fanfei Chen, Jinkun Wang, Tixiao Shan, and Brendan Englot. Autonomous exploration under uncertainty via graph convolutional networks. In *International Symposium on Robotics Research*. IEEE, 2019.

[28] Kevin Chen, Juan Pablo de Vicente, Gabriel Sepulveda, Fei Xia, Alvaro Soto, Marynel Vázquez, and Silvio Savarese. A behavioral approach to visual navigation with graph localization networks. In *Robotics: Science and Systems*, 2019.

[29] Weizhe Chen and Lantao Liu. Long-term autonomous ocean monitoring with streaming samples. In *OCEANS 2019 MTS/IEEE SEATTLE*, pages 1–8. IEEE, 2019.

[30] Weizhe Chen and Lantao Liu. Multi-objective and model-predictive tree search for spatiotemporal informative planning. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, pages 5716–5722. IEEE, 2019.

[31] Weizhe Chen and Lantao Liu. Pareto monte carlo tree search for multi-objective informative planning. In *Robotics: Science and Systems*, 2019.

[32] Sanjiban Choudhury, Ashish Kapoor, Gireeja Ranade, and Debadeepta Dey. Learning to gather information via imitation. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 908–915. IEEE, 2017.

[33] Shushman Choudhury, Christopher M Dellin, and Siddhartha S Srinivasa. Pareto-optimal search over configuration space beliefs for anytime motion planning. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3742–3749. IEEE, 2016.

[34] Nello Cristianini, Andre Elisseeff, John Shawe-Taylor, and Jaz Kandola. On kernel-target alignment. *Advances in neural information processing systems*, 2001.

[35] Lehel Csató and Manfred Opper. Sparse on-line gaussian processes. *Neural computation*, 14(3):641–668, 2002.

[36] Rongxin Cui, Yang Li, and Weisheng Yan. Mutual information-based multi-auv path planning for scalar field sampling using multidimensional rrt. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 46(7):993–1004, 2016.



 

[37] Alexander G Cunningham, Enric Galceran, Ryan M Eustice, and Edwin Olson. Mpdm: Multipolicy decision-making in dynamic, uncertain environments for autonomous driving. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1670–1677. IEEE, 2015.

[38] Mark Cutler and Jonathan P How. Autonomous drifting using simulation-aided reinforcement learning. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5442–5448. IEEE, 2016.

[39] Shoubhik Debnath, Lantao Liu, and Gaurav Sukhatme. Solving markov decision processes with reachability characterization from mean first passage times. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7063–7070. IEEE, 2018.

[40] Noel E Du Toit and Joel W Burdick. Robot motion planning in dynamic, uncertain environments. *IEEE Transactions on Robotics*, 28(1):101–115, 2011.

[41] Yan Duan, John Schulman, Xi Chen, Peter L Bartlett, Ilya Sutskever, and Pieter Abbeel. Rl $^2$: Fast reinforcement learning via slow reinforcement learning. *arXiv preprint arXiv:1611.02779*, 2016.

[42] M. Dunbabin and L. Marques. Robotics for environmental monitoring: Significant advancements and applications. *IEEE Robot. Autom. Mag.*, 19(1):24 – 39, 2012.

[43] Alberto Elfes. Robot navigation: Integrating perception, environmental constraints and task execution within a probabilistic framework. In *International Workshop on Reasoning with Uncertainty in Robotics*, pages 91–130. Springer, 1995.

[44] Brendan Englot, Tixiao Shan, Shaunak D Bopardikar, and Alberto Speranzon. Sampling-based Min-Max Uncertainty Path Planning. In *55th Conference on Decision and Control (CDC)*, pages 6863–6870, 2016.

[45] Philip Erickson, Michael Cline, Nishith Tirpankar, and Tom Henderson. Gaussian processes for multi-sensor environmental monitoring. In *2015 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI)*, pages 208–213, 2015.

[46] Jay A Farrell, Shuo Pang, and Wei Li. Chemical plume tracing via an autonomous underwater vehicle. *IEEE Journal of Oceanic Engineering*, 30(2):428–442, 2005.

[47] Wendell H Fleming and Raymond W Rishel. *Deterministic and stochastic optimal control*, volume 1. Springer Science & Business Media, 2012.

[48] Christian Forster, Matia Pizzoli, and Davide Scaramuzza. Svo: Fast semi-direct monocular visual odometry. In *2014 IEEE international conference on robotics and automation (ICRA)*, pages 15–22. IEEE, 2014.

[49] Christian Forster, Zichao Zhang, Michael Gassner, Manuel Werlberger, and Davide Scaramuzza. Svo: Semidirect visual odometry for monocular and multicamera systems. *IEEE Transactions on Robotics*, 33(2):249–265, 2016.

[50] Friedrich Fraundorfer, Davide Scaramuzza, and Marc Pollefeys. A constricted bundle adjustment parameterization for relative scale estimation in visual odometry. In *2010 IEEE International Conference on Robotics and Automation*, pages 1899–1904. IEEE, 2010.




 

[51] Nicholas Fung, John Rogers, Carlos Nieto, Henrik I Christensen, Stephanie Kemna, and Gaurav Sukhatme. Coordinating multi-robot systems through environment partitioning for adaptive informative sampling. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 3231–3237. IEEE, 2019.

[52] Enric Galceran, Alexander G Cunningham, Ryan M Eustice, and Edwin Olson. Multipolicy decision-making for autonomous driving via changepoint-based behavior prediction. In *Robotics: Science and Systems*, volume 1, 2015.

[53] Sahika Genc, Sunil Mallya, Sravan Bodapati, Tao Sun, and Yunzhe Tao. Zero-shot reinforcement learning with deep attention convolutional neural networks. *Advances in neural information processing systems*, 2019.

[54] Robert Ghrist, Jason M O'Kane, and Steven M LaValle. Pareto optimal coordination on roadmaps. In *Algorithmic foundations of robotics VI*, pages 171–186. Springer, 2004.

[55] Shixiang Gu, Timothy Lillicrap, Ilya Sutskever, and Sergey Levine. Continuous deep q-learning with model-based acceleration. In *International Conference on Machine Learning*, pages 2829–2838, 2016.

[56] Carlos Guestrin, Andreas Krause, and Ajit Paul Singh. Near-optimal sensor placements in gaussian processes. In *Proceedings of the 22nd international conference on Machine learning*, pages 265–272. ACM, 2005.

[57] Carlos Ernesto Guestrin. *Planning Under Uncertainty in Complex Structured Environments*. PhD thesis, Stanford, CA, USA, 2003. AAI3104233.

[58] Saurabh Gupta, James Davidson, Sergey Levine, Rahul Sukthankar, and Jitendra Malik. Cognitive mapping and planning for visual navigation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2616–2625, 2017.

[59] Hadi Hajieghrary, Alex Fabregat Tom, M Ani Hsieh, et al. An information theoretic source seeking strategy for plume tracking in 3d turbulent fields. In *2015 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*, pages 1–8. IEEE, 2015.

[60] Tomislav Hengl. *A practical guide to geostatistical mapping*, volume 52. Hengl, 2009.

[61] Geoffrey A Hollinger, Arvind A Pereira, Jonathan Binney, Thane Somers, and Gaurav S Sukhatme. Learning uncertainty in ocean current predictions for safe and reliable navigation of underwater vehicles. *Journal of Field Robotics*, 33(1):47–66, 2016.

[62] Geoffrey A Hollinger, Arvind A Pereira, and Gaurav S Sukhatme. Learning uncertainty models for reliable operation of autonomous underwater vehicles. In *Robotics and Automation (ICRA), 2013 IEEE International Conference on*, pages 5593–5599. IEEE, 2013.

[63] Max Jaderberg, Volodymyr Mnih, Wojciech Marian Czarnecki, Tom Schaul, Joel Z Leibo, David Silver, and Koray Kavukcuoglu. Reinforcement learning with unsupervised auxiliary tasks. *arXiv preprint arXiv:1611.05397*, 2016.

[64] Leslie Pack Kaelbling, Michael L. Littman, and Anthony R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101:99–134, 1998.




 

[65] Stephanie Kemna, John G Rogers, Carlos Nieto-Granda, Stuart Young, and Gaurav S Sukhatme. Multi-robot coordination through dynamic voronoi partitioning for informative adaptive sampling in communication-constrained environments. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2124–2130. IEEE, 2017.

[66] Thomas Kollar and Nicholas Roy. Efficient optimization of information-theoretic exploration in slam. In *AAAI*, volume 8, pages 1369–1375, 2008.

[67] Thomas Kollar and Nicholas Roy. Trajectory optimization using reinforcement learning for map exploration. *The International Journal of Robotics Research*, 27(2):175–196, 2008.

[68] Hanna Kurniawati, David Hsu, and Wee Sun Lee. Sarsop: Efficient point-based pomdp planning by approximating optimally reachable belief spaces. In *Robotics: Science and systems*, volume 2008. Zurich, Switzerland., 2008.

[69] Guillaume Lample and Devendra Singh Chaplot. Playing fps games with deep reinforcement learning. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, pages 2140–2146. AAAI Press, 2017.

[70] Steven M LaValle and Seth A Hutchinson. Optimal motion planning for multiple robots having independent goals. *IEEE Transactions on Robotics and Automation*, 14(6):912–925, 1998.

[71] Cedric Le Gentil, Teresa Vidal-Calleja, and Shoudong Huang. IN2LAMA: Inertial lidar localisation and MApping. In *IEEE International Conference on Robotics and Automation*, pages 6388–6394, 2019.

[72] Jeongseok Lee, Daqing Yi, and Siddhartha S Srinivasa. Sampling of pareto-optimal trajectories using progressive objective evaluation in multi-objective motion planning. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1–9. IEEE, 2018.

[73] Alberto Quattrini Li, Adem Coskun, Sean M Doherty, Shervin Ghasemlou, Apoorv S Jagtap, Md Modasshir, Sharmin Rahman, A Singh, Marios Xanthidis, Jason M O'Kane, et al. Experimental comparison of open source vision-based state estimation algorithms. In *International Symposium on Experimental Robotics*, pages 775–786. Springer, 2016.

[74] Andreas Lichtenstern. Kriging methods in spatial statistics. *Technische Universität München*, 2013.

[75] Tony Lindeberg. Scale invariant feature transform. 2012.

[76] Lantao Liu and Gaurav S Sukhatme. A solution to time-varying markov decision processes. *IEEE Robotics and Automation Letters*, 3(3):1631–1638, 2018.

[77] Xiuming Liu, Teng Xi, and Edith Ngai. Data modelling with gaussian process in sensor networks for urban environmental monitoring. In *Proc. 24th International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems*, pages 457–462. IEEE Computer Society, 2016.

[78] Jaime Lloret, Miguel Garcia, Diana Bri, and Sandra Sendra. A wireless sensor network deployment for rural and forest fire detection and verification. *sensors*, 9(11):8722–8747, 2009.



 

[79] Kian Hsiang Low. *Multi-robot Adaptive Exploration and Mapping for Environmental Sensing Applications*. PhD thesis, Carnegie Mellon University, Pittsburgh, PA, USA, 2009.

[80] Kian Hsiang Low, John M Dolan, and Pradeep Khosla. Active markov information-theoretic path planning for robotic environmental sensing. In *Proceedings of the 10th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS-11)*, pages 753–760, May 2011.

[81] Brandon D Luders, Sertac Karaman, and Jonathan P How. Robust Sampling-based Motion Planning with Asymptotic Optimality Guarantees. In *AIAA Guidance, Navigation, and Control (GNC) Conference*, page 5097, 2013.

[82] Kai-Chieh Ma, Lantao Liu, Hordur K Heidarsson, and Gaurav S Sukhatme. Data-driven learning and planning for environmental sampling. *Journal of Field Robotics*, 35(5):643–661, 2018.

[83] Kai-Chieh Ma, Lantao Liu, and Gaurav S. Sukhatme. An information-driven and disturbance-aware planning method for long-term ocean monitoring. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2016.

[84] Kai-Chieh Ma, Lantao Liu, and Gaurav S. Sukhatme. Multi-robot informative and adaptive planning for persistent environmental monitoring. In *International Symposium on Distributed Autonomous Robotic Systems (DARS)*, 2016.

[85] Kai-Chieh Ma, Lantao Liu, and Gaurav S. Sukhatme. Informative planning and online learning with sparse gaussian processes. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2017.

[86] Dhanvin Mehta, Gonzalo Ferrer, and Edwin Olson. Backprop-mpdm: Faster risk-aware policy evaluation through efficient gradient optimization. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1740–1746. IEEE, 2018.

[87] Alexandra Meliou, Andreas Krause, Carlos Guestrin, and Joseph M. Hellerstein. Non-myopic informative path planning in spatio-temporal models. In *Proceedings of National Conference on Artificial Intelligence (AAAI)*, pages 602–607, 2007.

[88] Igor Mordatch, Kendall Lowrey, and Emanuel Todorov. Ensemble-cio: Full-body dynamic motion planning that transfers to physical humanoids. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5307–5314. IEEE, 2015.

[89] Bruno Morier and Vladimir Teles. A time-varying markov-switching model for economic growth. Textos para discussão 305, Escola de Economia de São Paulo, Getulio Vargas Foundation (Brazil), 2011.

[90] Raul Mur-Artal, Jose Maria Martinez Montiel, and Juan D Tardos. Orb-slam: a versatile and accurate monocular slam system. *IEEE transactions on robotics*, 31(5):1147–1163, 2015.

[91] Erik Nelson and Nathan Michael. Information-theoretic occupancy grid compression for high-speed information-based exploration. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4976–4982. IEEE.

[92] Duy Nguyen-tuong and Jan Peters. Local gaussian process regression for real time online model learning and control. In *In Advances in Neural Information Processing Systems 22 (NIPS*, 2008.



 

[93] Farzad Niroui, Kaicheng Zhang, Zendai Kashino, and Goldie Nejat. Deep reinforcement learning robot for search and rescue applications: Exploration in unknown cluttered environments. *IEEE Robotics and Automation Letters*, 4(2):610–617, 2019.

[94] Stephen Nuske, Supreeth Achar, Terry Bates, Srinivasa Narasimhan, and Sanjiv Singh. Yield estimation in vineyards by visual grape detection. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2352–2358. IEEE, 2011.

[95] Ruofei Ouyang, Kian Hsiang Low, Jie Chen, and Patrick Jaillet. Multi-robot active sensing of non-stationary gaussian process-based environmental phenomena. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems*, pages 573–580, 2014.

[96] Lerrel Pinto and Abhinav Gupta. Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours. In *2016 IEEE international conference on robotics and automation (ICRA)*, pages 3406–3413. IEEE, 2016.

[97] Samuel Prentice and Nicholas Roy. The Belief Roadmap: Efficient Planning in Belief Space by Factoring the Covariance. *The International Journal of Robotics Research*, 28(11-12):1448–1465, 2009.

[98] Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., 1st edition, 1994.

[99] Chao Qin, Haoyang Ye, Christian E Pranata, Jun Han, Shuyang Zhang, and Ming Liu. LINS: A Lidar-Inertial State Estimator for Robust and Efficient Navigation. In *IEEE International Conference on Robotics and Automation*, pages 8899–8906, 2020.

[100] S.Joe Qin and Thomas A. Badgwell. A survey of industrial model predictive control technology. *Control Engineering Practice*, 11(7):733 – 764, 2003.

[101] Ananth Ranganathan, Ming-Hsuan Yang, and Jeffrey Ho. Online Sparse Gaussian Process Regression and Its Applications. *IEEE Transactions on Image Processing*, 20(2):391–404, February 2011.

[102] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 2005.

[103] Diederik M Roijers, Peter Vamplew, Shimon Whiteson, and Richard Dazeley. A survey of multi-objective sequential decision-making. *Journal of Artificial Intelligence Research*, 48:67–113, 2013.

[104] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.

[105] Alvaro Sanchez-Gonzalez, Nicolas Heess, Jost Tobias Springenberg, Josh Merel, Martin Riedmiller, Raia Hadsell, and Peter Battaglia. Graph networks as learnable physics engines for inference and control. In *International Conference on Machine Learning*, pages 4470–4479, 2018.

[106] Davide Scaramuzza and Friedrich Fraundorfer. Visual odometry [tutorial]. *IEEE robotics & automation magazine*, 18(4):80–92, 2011.



 

[107] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2008.

[108] Suresh Sethi and Gerhard Sorger. A theory of rolling horizon decision making. *Ann. Oper. Res.*, 29(1-4):387–416, 1991.

[109] Tixiao Shan and Brendan Englot. Belief Roadmap Search: Advances in Optimal and Efficient Planning Under Uncertainty. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5318–5325, 2017.

[110] Tixiao Shan and Brendan Englot. LeGO-LOAM: Lightweight and Ground-Optimized Lidar Odometry and Mapping on Variable Terrain. *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4758–4765, 2018.

[111] Tixiao Shan, Brendan Englot, Drew Meyers, Wei Wang, Carlo Ratti, and Daniela Rus. LIO-SAM: Tightly-coupled Lidar Inertial Odometry via Smoothing and Mapping. *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5135–5142, 2020.

[112] Amarjeet Singh, Andreas Krause, Carlos Guestrin, William Kaiser, and Maxim Batalin. Efficient planning of informative paths for multiple robots. In *Proceedings of the 20th International Joint Conference on Artifical Intelligence*, IJCAI'07, pages 2204–2211, 2007.

[113] Ryan N Smith, Mac Schwager, Stephen L Smith, Burton H Jones, Daniela Rus, and Gaurav S Sukhatme. Persistent ocean monitoring with underwater gliders: Adapting sampling resolution. *Journal of Field Robotics*, 28(5):714–741, 2011.

[114] Daniel E. Soltero, Mac Schwager, and Daniela Rus. Generating informative paths for persistent sensing in unknown environments. In *IROS*, pages 2172–2179, 2012.

[115] Cyrill Stachniss, Giorgio Grisetti, and Wolfram Burgard. Information gain-based exploration using rao-blackwellized particle filters. In *Robotics: Science and Systems*, volume 2, pages 65–72, 2005.

[116] Cyrill Stachniss, Christian Plagemann, Achim J Lilienthal, and Wolfram Burgard. Gas distribution modeling using sparse gaussian process mixture models. In *Robotics: Science and Systems*, volume 3, 2008.

[117] Hauke Strasdat, José MM Montiel, and Andrew J Davison. Visual slam: why filter? *Image and Vision Computing*, 30(2):65–77, 2012.

[118] Lei Tai, Shaohua Li, and Ming Liu. A deep-network solution towards model-less obstacle avoidance. In *2016 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pages 2759–2764. IEEE, 2016.

[119] Lei Tai and Ming Liu. Towards cognitive exploration through deep reinforcement learning for mobile robots. *arXiv preprint arXiv:1610.01733*, 2016.

[120] Aviv Tamar, Yi Wu, Garrett Thomas, Sergey Levine, and Pieter Abbeel. Value iteration networks. *Advances in neural information processing systems*, 29:2154–2162, 2016.

[121] Pratap Tokekar, Deepak Bhadauria, Andrew Studenski, and Volkan Isler. A robotic system for monitoring carp in minnesota lakes. *Journal of Field Robotics*, 27(6):779–789, 2010.



 

[122] Rafael Valencia, Jaime Valls Miró, Gamini Dissanayake, and Juan Andrade-Cetto. Active pose slam. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1885–1891. IEEE, 2012.

[123] Jinkun Wang and Brendan Englot. Autonomous exploration with expectation-maximization. In *International Symposium on Robotics Research*. IEEE, 2017.

[124] Tingwu Wang, Renjie Liao, Jimmy Ba, and Sanja Fidler. Nervenet: Learning structured policy with graph neural networks. In *International Conference on Learning Representations*, 2018.

[125] Geoff Werner-Allen, Konrad Lorincz, Jeff Johnson, Jonathan Lees, and Matt Welsh. Fidelity and yield in a volcano monitoring sensor network. In *Proceedings of the 7th symposium on Operating systems design and implementation*, pages 381–396. USENIX Association, 2006.

[126] Peter Whaite and Frank P Ferrie. Autonomous exploration: Driven by uncertainty. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(3):193–205, 1997.

[127] Kai M Wurm, Daniel Hennes, Dirk Holz, Radu B Rusu, Cyrill Stachniss, Kurt Konolige, and Wolfram Burgard. Hierarchies of octrees for efficient 3d mapping. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4249–4255. IEEE, 2011.

[128] Junhong Xu, Kai Yin, and Lantao Liu. Reachable space characterization of markov decision processes with time variability. In *Proceedings of Robotics: Science and Systems*, FreiburgimBreisgau, Germany, June 2019.

> This paper attempts to explore the time variability property of the planning stochasticity and investigate the state reachability, based on which an efficient iterative method is developed to offer a good trade-off between solution optimality and time complexity.

[129] Junhong Xu, Kai Yin, and Lantao Liu. State-continuity approximation of markov decision processes via finite element analysis for autonomous system planning. *arXiv preprint arXiv:1903.00948*, 2019.

[130] Brian Yamauchi. Frontier-based exploration using multiple robots. In *Proceedings of the second international conference on Autonomous agents*, pages 47–53, 1998.

[131] Chenghai Yang. A high-resolution airborne four-camera imaging system for agricultural remote sensing. *Computers and electronics in agriculture*, 88:13–24, 2012.

[132] Kwangjin Yang, Seng Keat Gan, and Salah Sukkarieh. A gaussian process-based rrt planner for the exploration of an unknown and cluttered environment with a uav. *Advanced Robotics*, 27(6):431–443, 2013.

[133] Haoyang Ye, Yuying Chen, and Ming Liu. Tightly Coupled 3d Lidar Inertial Odometry and Mapping. In *IEEE International Conference on Robotics and Automation*, pages 3144–3150, 2019.

[134] Jingjin Yu, Mac Schwager, and Daniela Rus. Correlated orienteering problem and its application to informative path planning for persistent monitoring tasks. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2014.



 

Outstanding importance. This is one of the earliest informative planning methods that investigate realistic, time-varying, and spatially correlated scalar field.

[135] Ji Zhang and Sanjiv Singh. Low-drift and Real-time Lidar Odometry and Mapping. *Autonomous Robots*, 41(2):401–416, 2017.

[136] Tianhao Zhang, Gregory Kahn, Sergey Levine, and Pieter Abbeel. Learning deep control policies for autonomous aerial vehicles with mpc-guided policy search. In *2016 IEEE international conference on robotics and automation (ICRA)*, pages 528–535. IEEE, 2016.

[137] Yuke Zhu, Roozbeh Mottaghi, Eric Kolve, Joseph J Lim, Abhinav Gupta, Li Fei-Fei, and Ali Farhadi. Target-driven visual navigation in indoor scenes using deep reinforcement learning. In *2017 IEEE international conference on robotics and automation (ICRA)*, pages 3357–3364. IEEE, 2017.



 