

MPTP: Motion-planning-aware task planning for navigation in belief space

Antony Thomas*, Fulvio Mastrogiovanni, Marco Baglietto

Department of Informatics, Bioengineering, Robotics, and Systems Engineering, University of Genoa, Via All'Opera Pia 13, 16145 Genoa, Italy



ARTICLE INFO

Article history:

Received 31 August 2019
Received in revised form 8 August 2020
Accepted 8 April 2021
Available online 16 April 2021

Keywords:

Task-motion planning
Belief space planning
Autonomous navigation

ABSTRACT

We present an integrated Task-Motion Planning (TMP) framework for navigation in large-scale environments. Of late, TMP for manipulation has attracted significant interest resulting in a proliferation of different approaches. In contrast, TMP for navigation has received considerably less attention. Autonomous robots operating in real-world complex scenarios require planning in the discrete (task) space and the continuous (motion) space. In knowledge-intensive domains, on the one hand, a robot has to reason at the highest-level, for example, the objects to procure, the regions to navigate to in order to acquire them; on the other hand, the feasibility of the respective navigation tasks have to be checked at the execution level. This presents a need for motion-planning-aware task planners. In this paper, we discuss a probabilistically complete approach that leverages this task-motion interaction for navigating in large knowledge-intensive domains, returning a plan that is optimal at the task-level. The framework is intended for motion planning under motion and sensing uncertainty, which is formally known as belief space planning. The underlying methodology is validated in simulation, in an office environment and its scalability is tested in the larger Willow Garage world. A reasonable comparison with a work that is closest to our approach is also provided. We also demonstrate the adaptability of our approach by considering a building floor navigation domain. Finally, we also discuss the limitations of our approach and put forward suggestions for improvements and future work.

© 2021 Elsevier B.V. All rights reserved.

1. Introduction

Autonomous robots operating in complex real world scenarios require different levels of planning to execute the assigned tasks. High-level (task) planning helps break down a given set of tasks into a sequence of sub-tasks. Actual execution of each of these sub-tasks would require low-level control actions to generate appropriate robot motions. In fact, the dependency between logical and geometrical aspects is pervasive in both task planning and execution. Hence, planning should be performed in the task-motion or the discrete-continuous space [1].

In recent years, combining high-level task planning with low-level motion planning has been a subject of great interest among the Robotics and Artificial Intelligence (AI) communities. Traditionally, task planning and motion planning have evolved as two independent fields. AI planning frameworks such as the Planning Domain Definition Language (PDDL) [2] mainly focus on high-level task planning supposing that the geometric preconditions (e.g., grasping poses for a pick-up task [3]) for the robot motion

to carry out these tasks are achievable. In reality, such an assumption can be catastrophic as an action or sequence of actions generated by the task planner might turn out to be unfeasible at the controller execution level.

Over the past few years, Task-Motion Planning (TMP) for manipulation has received considerable interest among the research community [3–7]. Robot-based manipulation domain calls for discrete and continuous reasoning to execute the required action reliably. For example, a simple table top domain requires the robot to reason at the discrete level to decide the objects to be picked up and also the order of these high-level actions. The execution of these discrete actions require continuous reasoning in the configuration space of the robot to generate appropriate motions. Yet, a discrete action might turn out to be unfeasible due to the end-effector's reachability workspace. This might be due to the availability of a partial map leading to unmodeled objects or occlusions leading to unobserved objects or simply because the robot is too close the target object, rendering a grasp action impossible. This presents the need for a tight coupling between task planning and motion planning, enabling an interface for efficient interaction between the symbolic and geometric layers. TMP for navigation presents different challenges in comparison to TMP for manipulation. As such, TMP for navigation has not yet received much attention and therefore lacks sufficient literature.

* Corresponding author.

E-mail addresses: antony.thomas@dibris.unige.it (A. Thomas), fulvio.mastrogiovanni@unige.it (F. Mastrogiovanni), marco.baglietto@unige.it (M. Baglietto).

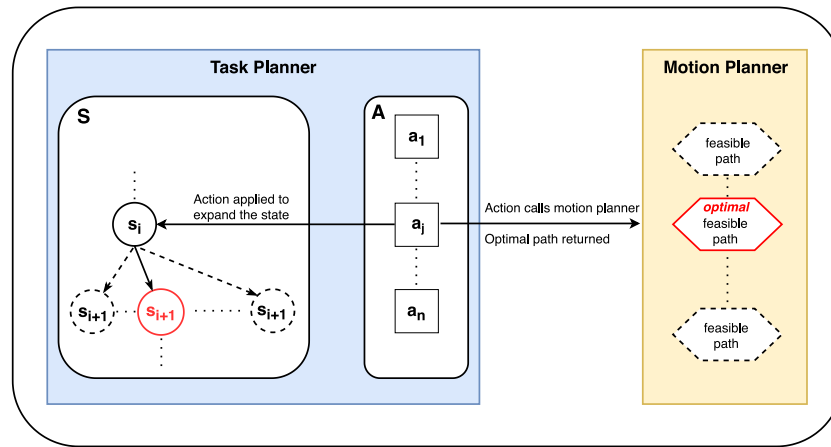


Fig. 1. The discrete actions available to the planner are denoted by $A = \{a_1, a_2, a_3, \dots, a_n\}$. Different motion plans are generated for the action that requires appropriate robot motion via an external module. This module is essentially a motion planner. The optimal path among the feasible motion plans is then selected, returning the optimal cost to the task planner. The corresponding action and the optimal path is the task-motion plan for changing the task state of the robot from s_i to s_{i+1} .

TMP for navigation essentially involves reasoning about different objects and their properties, deciding which objects to procure, selecting high-level actions that satisfy the low-level continuous motion constraints to navigate to the objects or other locations of interest, and finally procuring the objects and delivering it to the respective goal locations subject to task and motion constraints. For example, consider a robot in an office environment where it needs to deliver documents for evaluation to the respective project managers. At the task level, it is required that the robot first identifies the project in order to navigate to the respective sections, collect the documents and then deliver them to the project manager. A task planner computes a plan in terms of these symbolic actions, subject to minimizing a certain metric. This metric, for example, might correspond to different types of action costs or the number of actions. Since we are concerned with navigation, in this paper we associate the symbolic actions to their associated motion costs. Certain symbolic actions may not require robot motions. For example, for collecting a document, the robot may have to stay at a particular location for a given amount of time waiting for a human to place the document. Such actions are assigned a fixed cost. Selecting the best set of discrete actions for a given objective requires computing the navigation costs (and other fixed costs) for each of these actions. Hence motion planning should be interleaved with task planning to compute the motion costs for each of the respective discrete actions. Though it can be argued that the motion costs can be approximated a priori and fed to the task planner, in large knowledge-intensive domains such an assumption can be harder to justify, especially in the presence of localization and map uncertainty. Moreover, real-world scenarios often induce uncertainties. Such uncertainties arise due to insufficient knowledge about the environment, inexact robot motion or imperfect sensing. In such scenarios, the robot poses or other variables of interest can only be dealt with, in terms of probabilities. Planning is therefore done in the *belief* space, which corresponds to the probability distributions over possible robot states. Consequently, for efficient planning and decision making, it is required to reason about future belief distributions due to candidate actions and the corresponding expected observations. Such a problem falls under the category of Partially Observable Markov Decision Processes (POMDPs) [8]. Our motion planner is therefore equipped to perform planning in partially-observable state-spaces with motion and sensing uncertainty.

This paper extends the work presented in [9] providing a comprehensive discussion and an extensive performance evaluation. Specifically, this paper contributes to the literature with a

Motion-Planning-aware Task Planning (MPTP) approach providing an interface between task and motion planning for navigating in large knowledge-intensive domains. Such domains require a robot to reason about different objects and locations to navigate to, subject to minimizing (or maximizing) the navigation cost (objective function). Our task-motion interface layer facilitates this reasoning by communicating the motion feasibility and the corresponding planned motion costs to the task planner, synthesizing an optimal plan. To this end, we develop a probabilistically complete Task-Motion Planning (TMP) framework for mobile robot navigation under partial-observability, embedding a motion planner within a task planner through an interface layer. We would like to stress the fact that our implementation is independent of any particular form of cost function. In this paper, we use a standard cost function (see Section 4) as the MPTP cost and compare it with different cost functions in Section 5.

An overview of our MPTP approach is shown in Fig. 1. We define $A = \{a_1, \dots, a_n\}$ as the finite set of symbolic/discrete actions available to the task planner. For example, let us again consider an office setting where a robot is tasked with collecting and delivering documents. In such a setting, some of the actions include, *collect_document*– which might correspond to a human placing the document on the robot and therefore the robot waiting at a specific location for a certain duration, *deliver_document*– similar to *collect_document* action but a human picks up the document, *goto_region*– corresponds to navigating through the environment. Once an action that require appropriate robot motions to be generated is expanded by the task planner, a call to an external library is triggered. The symbolic parameters are then converted to their corresponding geometric instantiations. For example, for an action that takes the robot to a particular cubicle/region, the instantiations would be the different sampled poses in that cubicle. Once the map of the environment is obtained, the geometric instantiations can be pre-sampled. The instantiations give rise to different motion plans and the best among them is chosen according to a certain metric. The cost of the selected motion plan cost is then returned to the task planner as the cost of the corresponding action. The task-motion plan for changing the task state of the robot from the state s_i to s_{i+1} is the ordered tuple of the action a_i and the corresponding optimal path. For instance, in the office setting where a robot navigates from one cubicle (s_i) to another (s_{i+1}), the tuple is $\{\text{goto_region}, \tau_i\}$. Here, *goto_region* is the task-level action a_i and τ_i is the planned trajectory for achieving this high-level action. This tuple is appended for all the task-level

actions to generate the complete task-motion plan. While our approach is applicable to any domain that require task-motion interaction, we establish the key ideas in Section 4 through two different navigation domains and further validate our approach in Section 5 using the same.

2. Related work

TMP has emerged as an active research area in the recent past, with particular focus on robot-based manipulation. Manipulation tasks are often rendered infeasible due to the end-effector's reachability workspace. This calls for an integrated TMP approach to ensure geometric feasibility of high-level tasks.

The genesis of TMP can be credited to Fikes and Nilsson for their work on STRIPS [10] which further led to the Shakey project [11]. Initial works on TMP performed task planning first, synthesizing a sequence of actions to be executed later by a motion planner. Shakey's planner performed a logical search first, assuming that the resulting robot motion plans can be formulated. This assumption limits the capability of the robot as the high-level actions may turn out to be non executable due to geometric limitations of the environment or the robot or both. [12] interleaves task and motion planning by checking individual high-level action feasibility using *semantic attachments*. [13] perform a combined search in the logical and geometric spaces using a state composed of both the symbolic and geometric paths. The aSyMov planner described in [13] adopts a combination of Metric-FF [14] and a sampling-based motion planner. In contrast, we use a temporal task planner, POPF-TIF [15] with roadmap-based sampling, incorporating robot state uncertainty. Srivastava et al. [3] implicitly incorporate geometric variables, performing symbolic-geometric mapping using a planner-independent interface layer. Erdem et al. [16] leverage a boolean satisfiability (SAT) solver, computing a task-level plan and then refining it until a feasible motion plan is found.

Kaelbling and Lozano-Péres [17] propose a hierarchical approach that tightly integrates logical and geometric planning. The complexities arising out of long-horizon¹ planning are tackled to the extent that planning is done at different levels of abstraction, thereby reducing the long-horizons to a number of feasible sub-plans of shorter horizon. This regression²-based planner assumes that the actions are reversible while backtracking. This work is extended in [5] to consider the current state uncertainty, modeling the planning problem in the belief space. The hierarchical approach is also employed in [20,21] to compute discrete actions with unbounded continuous variables. A geometric backtrack search is used to instantiate the symbolic actions in [22]. They also prune certain geometric instantiations, reducing the complexity. FFRob [4] performs task planning by performing search over a sampled finite set of poses, grasps and configurations. The authors of [4] extend the FF heuristics, incorporating geometric and kinematic planning constraints that provide a tight estimate of the distance to the goal. Our approach is similar to FFRob in the sense that we also pre-sample robot configurations and then plans with them, incorporating motion constraints.

Toussaint [23] performs optimization over an objective function based on the final geometric configuration (and the cost thereby), finding approximately locally optimal solutions by minimizing the objective function. The planning problem is modeled as a constraint satisfaction problem with symbolic states used to define the constraints in the optimization. This *logic-geometric programming* is applied to a four manipulator setting in [24].

¹ Large environments require a robot to perform many actions to reach the goal, resulting in a long planning horizon [18].

² Goal regression is the process of planning backwards from the goal [19].

Lozano-Péres and Kaelbling [25] model the motion planning as a constraint satisfaction problem over a subset of the configuration space. Iteratively Deepened Task and Motion Planning (IDTMP) is a constraint-based task planning approach that incorporates geometric information to account for the motion feasibility at the task planning level [6]. In our architecture, the motion costs are returned to the task planner, similar to the motion planner information that guides the IDTMP task planner. IDTMP performs task-motion interaction using abstraction and refinement functions whereas we use *semantic attachments* [26].

Though the approaches discussed above fall under the category of TMP for manipulation, the scope of TMP is not limited to manipulation problems alone. TMP for navigation is pervasive in most real world scenarios. For example, a mobile office robot may be tasked with collecting documents and delivering them across multiple floors. Yet, TMP for robot navigation has received less attention in the past. Real-world planning problems in large scale environments often require solving several sub-problems. For example, while navigating to a goal, the robot might have to visit other places of interests. Visiting these places of interest are high-level tasks that can be addressed using traditional task planners. Yet, these symbolic planners cannot compute the exact motion costs for these tasks, let alone perform navigation and path planning. This calls for task plans that are motion planning aware, in terms of motion costs and its feasibility.

Task planning for robot Navigation Among Movable Obstacles (NAMO) is introduced in [27], where each object is displaced at most once throughout the plan. Van Den Berg et al. [28] provide a probabilistically complete algorithm for the NAMO class of problems. However, the robot state is assumed to be known perfectly. In contrast, we plan in the belief space, computing an estimate of the robot state at each instant. Hauser and Latombe [29,30] consider *multi-model motion planning* for manipulation and legged locomotion, wherein the space of feasible configurations consists of intersecting spaces of different dimensions. In [31] a TMP approach is presented in the context of Human-Robot Interaction (HRI). They integrate probabilistic reasoning with symbolic reasoning by implementing a spoken dialog system, enabling the robots to ask intelligent queries. Their task planner is based on Answer Set Programming (ASP) [32]. Jiang et al. [33] focus exclusively on task planning in robotics, assuming that a feasible motion plan exists for the synthesized task plan. They provide a comparison between ASP-based and PDDL-based task planners using different benchmark domains and conclude that PDDL-based planners perform better on tasks with long solutions, and ASP-based planners tend to perform better on shorter tasks. In this paper, we employ a PDDL-based task planner. UP2TA [34] develops a unified path planning and task planning framework for mobile robot navigation. In this approach, the robot is required to perform a series of tasks at different locations before returning back to the initial location. An interesting feature of UP2TA is its task planner heuristic, which is a combination of the FF heuristic [14] and the Euclidean distance between the waypoints associated with locations. The path planning layer computes the optimal path between each waypoint with the help of a Digital Terrain Model (DTM). Wong et al. [35] develop a task planning approach that takes into account the optimal traversal costs³ to synthesize a plan. Similar to UP2TA, they define tasks that are to be performed at different waypoints. However, the path planner pre-computes an optimal path for all pairs of waypoints, which are then passed to the task planner to find the optimal sequence of tasks. In contrast, we consider a general approach where the robot has to reason at a high-level about different objects or

³ The costs are defined in terms of mechanical work and the objective is to find the path with optimal mechanical work. For more details, refer to [35].

locations or regions to navigate to. The objects/locations/regions are instantiated to their geometric counterpart, by considering a set of sampled poses. For example, if a robot has to reach a location close to a chair, the geometric instantiations of this symbolic goal would correspond to a set of poses around the chair.

Jiang et al. [36] introduced a framework that integrates TMP with reinforcement learning that is robust to changes in the environment. The inner loop of their dual layer architecture is a TMP planner that generates task-motion plans to be sent to the outer loop. The outer loop executes the generated plans to learn from rewards. In contrast MPTP is a purely planning approach. Lo et al. [37] introduced PETLON, a purely planning approach for navigation that is task-level optimal and is the work closest to our approach. The inner loop in [36] uses a TMP planner that is similar to PETLON. However, in PETLON, the action costs returned by the motion planner is the trajectory length and complete observability is assumed. In contrast, our framework is more general, since we additionally consider the cost due to motion and sensing uncertainty and the distance to the goal. It is to be noted that our approach is not limited to any particular cost function and can be easily adapted to support any general cost formulation. In Section 5, we benchmark the scalability of our approach and provide a comparison with PETLON by considering a motion planner that evaluates the geometric-level cost of navigation. In this way we compare MPTP to PETLON by adapting our cost function to incorporate only the geometric-level cost of traversing from one location to another. Further, PETLON first compute a task plan using an admissible heuristic which is then sent to the motion planner for cost evaluation. This updates the heuristic and a refinement process repeats until the optimal plan is found. In contrast, MPTP evaluates the motion cost as each action is expanded by the task planner and hence the plan returned is optimal and needs no refinement.

3. Preliminaries and definitions

We begin by formally defining the notions of task and motion planning. Then, we state the TMP problem that we discuss in this paper. The notations and formalism correspond to that of a state-transition system [19].

3.1. Task planning

Task planning or classical planning can be defined as the process of finding a discrete sequence of actions from the current state to a desired goal state [19].

Definition 1. A task domain Ω can be represented as a state transition system and is a tuple $\Omega = (S, A, \gamma, s_0, S_g)$ where:

- S is a finite set of states, each state is a conjunction of propositions⁴;
- A is a finite set of actions;
- $\gamma : S \times A \rightarrow S$ is the state transition function such that $s' = \gamma(s, a)$;
- $s_0 \in S$ is the start state;
- $S_g \subseteq S$ is the set of goal states.

Definition 2. The task plan for a task domain Ω is the sequence of actions a_0, \dots, a_n such that $s_{i+1} = \gamma(s_i, a_i)$, for $i = 0, \dots, n$ and s_{n+1} satisfies S_g .

⁴ A proposition is represented by a tuple of elements, which may be constants or variables, and can be negated [38].

The Planning Domain Definition Language (PDDL) [2] being the de facto standard syntax for task planning, we resort to the same for modeling our task domain. PDDL is an action-centered language, where each action a_i is described as a tuple $a_i = (pre_{a_i}, eff_{a_i})$, where pre_{a_i} (a set of preconditions for a_i) is a conjunction of propositions with either positive or negative terms that must hold for action execution and eff_{a_i} (the set of effects of a_i) is a conjunction of positive ($eff_{a_i}^+$) and negative ($eff_{a_i}^-$) propositions that are added or deleted upon action application, thereby changing the system state. The set of positive effects $eff_{a_i}^+$ contains propositions that become true upon the execution of action a_i and the set of negative effects $eff_{a_i}^-$ contains propositions that evaluates to false upon action execution. An action a_i is said to be applicable to a state s_i if each proposition of the preconditions holds in s_i , that is, $pre_{a_i} \subseteq s_i$. If an action a_i is applicable in state s_i , the corresponding successor state s_{i+1} is obtained as, $s_{i+1} = \gamma(s_i, a_i)$, where $s_{i+1} = (s_i \setminus eff_{a_i}^-) \cup eff_{a_i}^+$. A valid plan is a sequence of actions that when executed from s_0 results in S_g .

A planning problem with PDDL is created by providing a domain description that describes the predicates and action schemas with free variables, and a problem description that specifies the objects, initial state and the goal condition. The objects are used to instantiate the predicates and action schemas, through a process called grounding. Grounding is the process by which every combination of objects is used to replace the free variables in predicates and action schemas to obtain propositions and ground actions respectively. In this paper, we use an extension of PDDL [39] that supports durative actions and numeric-valued fluents. Temporal planning introduces the possibility of computing concurrent plans. A temporal task domain can be defined by extending the task domain in Definition 1 as follows

Definition 3. A temporal task domain Ω can be represented as state transition system and is a tuple $\Omega = (S, A, \gamma, s_0, S_g)$ where:

- S is a finite set of states;
- V is a set of real valued variables;
- A is a finite set of actions;
- $\gamma : S \times A \rightarrow S$ is the state transition function such that $s' = \gamma(s, a)$;
- $s_0 \in S \cup V$ is the start state;
- $S_g \subseteq S \cup V$ is the set of goal states.

A durative action is a tuple $a_i = (pre_{a_i}, eff_{a_i}, dur_{a_i})$, where pre_{a_i} and eff_{a_i} are temporally annotated by specifying conditions/effects that holds at the *start*, *end* or during the *entire* action interval and are expressed using the constructs *at start*, *at end* and *over all* respectively. Note that these constructs are specific to PDDL formalism. dur_{a_i} corresponds to the duration of action a_i .

3.2. Motion planning

Motion planning finds a sequence of collision free poses from a given initial/start pose (position and orientation) to a desired goal pose [40].

Definition 4. A motion planning problem is a tuple $M = (C, f, q_0, G)$ where:

- C is the configuration space or the space of possible robot poses;
- $f = \{0, 1\}$ determines if a configuration/pose is in collision ($f = 0$) or not (C_{free} with $f = 1$). C_{free} denotes the set of all poses that are not in collision;
- q_0 is the initial configuration;
- G is the set of goal configurations.

Definition 5. A motion plan for M finds a valid trajectory in C from q_0 to $q_n \in G$ such that f evaluates to true for q_0, \dots, q_n .

In addition to the sequential form of the definition above, a motion plan can also be defined by a continuous trajectory

Definition 6. A motion plan for M is a function of the form $\tau : [0, 1] \rightarrow C_{free}$ such that $\tau(0) = q_0$ and $\tau(1) \in G$.

We will use a combination of the two to define the TMP problem and use roadmap based motion planner to generate collision free configurations.

3.3. Task-motion planning

TMP essentially involves combining discrete and continuous decision-making to facilitate efficient interaction between the two domains. Starting from an initial state, TMP synthesizes a plan to a goal state by a concurrent or interleaved set of discrete actions and continuous collision-free motions. Below we define the TMP problem formally.

Definition 7. A task-motion planning is a tuple $\Psi = (C, \Omega, \phi, \xi, q_0)$ where:

- $\phi : S \rightarrow 2^C$, is a function mapping states to the configuration space. For example, if s represents the task state — the robot is in a corridor, then $\phi(s)$ corresponds to all configurations such that the robot is in the corridor;
- $\xi : A \rightarrow 2^C$, is a function mapping actions to motion plans. We recall here that motion planning is essentially computing collision free poses in C .

Definition 8. The TMP problem for the TMP domain Ψ is to find a sequence of actions a_0, \dots, a_n such that $s_{i+1} = \gamma(s_i, a_i)$, $s_{n+1} \in S_g$ and to find a sequence of motion plans τ_0, \dots, τ_n such that for $i = 0, \dots, n$, it holds that

$$\tau_i(0) \in \phi(s_i) \text{ and } \tau_i(1) \in \phi(s_{i+1}) \quad (1)$$

$$\tau_{i+1}(0) = \tau_i(1) \quad (2)$$

$$\tau_i \in \xi(a_i) \quad (3)$$

3.4. Problem definition

In this paper, we consider the TMP problem for a mobile robot operating in a partially-observable environment. The map of the environment is either known a priori or is built using a standard Simultaneous Localization and Mapping (SLAM) algorithm.⁵ At any time k , we denote the robot pose (or configuration q_k) by $x_k \doteq (x, y, \theta)$, the acquired measurement is denoted by z_k and the applied control action is denoted as u_k . We consider a standard motion model with Gaussian noise

$$x_{k+1} = f(x_k, u_k, w_k), \quad w_k \sim \mathcal{N}(0, W_k) \quad (4)$$

where w_k is the random unobservable noise, modeled as a zero mean Gaussian. To process the landmarks in the environment we measure the range and the bearing of each landmark relative to the robot's local coordinate frame. In general, we consider the observation model with Gaussian noise

$$z_k = h(x_k) + v_k, \quad v_k \sim \mathcal{N}(0, Q_k) \quad (5)$$

It is to be noted that we assume data association as solved and hence given a measurement we know the corresponding

landmark that generated it. This is not a limitation and our approach can be extended to incorporate reasoning regarding data association, as shown recently in [41]. The motion (4) and observation (5) models can be written probabilistically as $p(x_{k+1}|x_k, u_k)$ and $p(z_k|x_k)$ respectively. Given an initial distribution $p(x_0)$, and the motion and observation models, the posterior probability distribution at time k can be written as

$$p(X_{0:k}|Z_{0:k}, U_{0:k-1}) = p(x_0) \prod_{i=1}^k p(x_i|x_{i-1}, u_{i-1}) p(z_i|x_i) \quad (6)$$

where $X_{0:k} \doteq \{x_0, \dots, x_k\}$, $Z_{0:k} \doteq \{z_0, \dots, z_k\}$ and $U_{0:k-1} \doteq \{u_0, \dots, u_{k-1}\}$. This posterior probability distribution is the *belief* at time k , denoted by $b[X_k] \sim \mathcal{N}(\mu_k, \Sigma_k)$. Similarly, given an action u_k , the propagated belief can be written as

$$b[X_{k+1}] = p(X_{0:k}|Z_{0:k}, U_{0:k-1}) p(x_{k+1}|x_k, u_k) \quad (7)$$

Given the current belief $b[X_k]$ and the control u_k , the propagated belief parameters can be computed using the standard Extended Kalman Filter (EKF) [42] prediction as

$$\bar{\mu}_{k+1} = f(\mu_k, u_k) \quad (8)$$

$$\bar{\Sigma}_{k+1} = F_k \Sigma_k F_k^T + V_k W_k V_k^T$$

where F_k is the Jacobian of $f(\cdot)$ with respect to x_k and V_k is the Jacobian of $f(\cdot)$ with respect to u_k . For brevity, the linearized process noise will be denoted as $R_k = V_k W_k V_k^T$. Upon receiving a measurement z_k , the posterior belief $b[X_{k+1}]$ is computed using the EKF update equations

$$\begin{aligned} K_k &= \bar{\Sigma}_{k+1} H_k^T (H_k \bar{\Sigma}_{k+1} H_k^T + Q_k)^{-1} \\ \mu_{k+1} &= \bar{\mu}_{k+1} + K_k (z_{k+1} - h(\bar{\mu}_{k+1})) \\ \Sigma_{k+1} &= (I - K_k H_k) \bar{\Sigma}_{k+1} \end{aligned} \quad (9)$$

where H_k is the Jacobian of $h(\cdot)$ with respect to x_k , K_k is the Kalman gain and $I \in \mathbb{R}^{3 \times 3}$ is the identity matrix.

4. Approach

PDDL-based planning frameworks are limited, as they are incapable of handling rigorous numerical calculations.⁶ Most approaches perform such calculations via external modules or *semantic attachments*, e.g. [26]. The term semantic attachment was coined by Weyhrauch [44] to describe the association of algorithms to function and predicate symbols via external procedures. However, the effects returned by these semantic attachments are not exploited in identifying *helpful actions* during search and hence do not provide any heuristic guidance, deeming the task unsolvable most often [45]. An action is considered *helpful* if it achieves at least one of the lowest level goals in the relaxed plan to the state at hand [14]. Recently, Bernardini et al. [45] developed a PDDL-based temporal planner to implicitly trigger such external calls via a specialized semantic attachments called *external advisors*. They classify variables into direct (V^{dir}), indirect (V^{ind}) and free (V^{free}). V^{dir} and V^{free} variables are the normal PDDL function variables whose values are changed in the action effects, in accordance with PDDL semantics. V^{ind} variables are affected by the changes in the V^{dir} variables. A change in a V^{dir} variable invokes the external advisor which in turn computes the V^{ind} variables. The Temporal Relaxed Plan Graph (TRPG) [46] construction stage of the planner incorporates the indirect variable values for heuristic calculation, thereby synthesizing an efficient goal-directed search. We employ this semantic attachment based approach for the task-motion interface. The overall procedure and the interface layer are discussed in detail in the remainder of this Section.

⁶ PDDL+ [43], an extension of PDDL supports mixed discrete and continuous non-linear changes.

⁵ http://wiki.ros.org/slam_gmapping/.

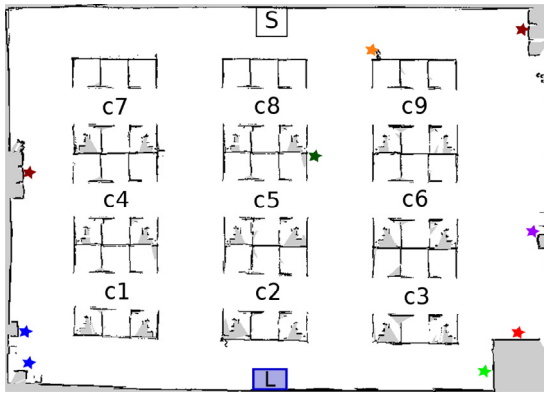


Fig. 2. Map of the office environment obtained after a SLAM session.

4.1. Task planning

TMP for navigation requires that the task planner takes into account the motion feasibility and the corresponding motion costs while synthesizing a plan. As opposed to the manipulation domain, where the motion feasibility is corroborated with the end-effector's reachability workspace, in navigation domains this is often validated against the cost constraints, for example, a robot navigating in a corridor with a bound on the pose covariance to avoid collisions. As such, any task planner customized to enable the task-motion interface can be employed for our approach. In our tests, PDDL is used to define the task domain.

Below, we elucidate the PDDL formalism for two different navigation domains that we have considered. It is to be noted that the semantic attachment procedure is domain independent and remains the same in both the domains. But the PDDL domain and problem description differ, as the two domains are different in nature. In the first domain, the underlying roadmap for motion planning does not change during plan computation. However, in the second domain, the roadmap is updated during plan computation. Description of the two domains are detailed below.

4.1.1. Office domain

We consider a robot navigating in an office environment to collect and deliver documents. The map of the environment following a SLAM session is shown in Fig. 2 (snapshot of the environment can be seen in Fig. 6). The regions c_1, \dots, c_9 are cubicles and L denotes a lift. The robot, starting from region S has to visit certain cubicles to receive documents. Navigating to cubicles/regions is encoded using a single high-level action `goto_region`. Once a robot reaches a cubicle from which a document is to be collected, we assume that a human places the requisite document. Thus, the robot needs to wait at the specific location for a fixed duration of time in which the human places the required document on the robot. This is encoded using a high-level action `collect_document`. These documents then have to be delivered to another floor, which implies using the lift L . Navigating to the lift is modeled using a different high-level action `goto_lift`. This is because, unlike the action `goto_region`, `goto_lift` is to be performed only if the robot has collected all the necessary documents to be delivered. The stars with different colors represent certain unique features assumed to be known and modeled like, printer, trash can, lounge, that aids the robot in better localization. Hence, once the robot knows the regions to visit, then it suffices to perform `goto_region` actions and collect the documents from these regions. However, to synthesize an optimal plan it is necessary to sequence these actions in an order that minimizes the cost function. It is therefore inevitable

to obtain the motion costs of these `goto_region` actions, so as to accurately synthesize the optimal plan.

A fragment of the PDDL domain is shown in Fig. 3. The PDDL domain dynamics is specified through a set of durative actions (`:durative-action`). We use the following parameters to model these actions: $?v$ is the name of the robot, $?from$ is the cubicle the robot is currently at and $?to$ is the cubicle to which the robot needs to move, $?r$ corresponds to the different regions or cubicles in the environment. Each action is described using `:condition` and `:effect`, as defined in Section 3, and defines the conditions and effects that holds at the start (at start), end (at end) or during the entire action interval (overall), respectively. The predicate `robot_in` checks if the robot is in a particular region. The function triggered encodes the fact that the robot is moving from one cubicle (from) to another (to). The functions `get` and `collected` model the cubicles from which the document is collected and whether it has been collected. Finally, `act-cost` stores the cost associated with the actions and `goal-trace` keeps the robot state uncertainty bounded. The actions `goto_region` and `goto_lift` invoke the external module call once the facts (`(increase (act-cost) (external))` and (`(increase (goal-trace) (bound))`) are encountered. Here, `act-cost`, `goal-trace` are the direct variables in V^{dir} and `external`, `bound` are the indirect variables V^{ind} . The function (triggered $?from ?to$) is assigned the numerical value 1 each time the actions are expanded and re-initialized to 0 once the action duration is completed. In this way, the grounded variables `from` (start) and `to` (goal) are communicated to the motion planner. The variables `external` and `bound` returns the motion cost and the goal covariance trace respectively, which are computed by the external module. The action `collect_document` does not invoke the motion planner. In the problem description, the function (`get ?r`), where r is a free variable denoting cubicles, is initialized to 1 for the cubicles from which the documents are to be collected and to 0 for the remaining.

4.1.2. Corridor domain

We consider a navigation domain, similar to the one in [33], wherein a robot navigates through a building floor that consists of several rooms connected to one another through a corridor. These rooms have doors, which can either be closed or open, connecting them to the corridor. In addition, some of the rooms are also accessible from each other, through doors in between them. The robot can navigate through the entire building by opening these doors. We assume that once the robot is near to a closed door that directly connects a room to the corridor, a human opens the door to allow the robot to pass through. Navigating to rooms can hence be encoded using a single high-level action `goto_room`. However, the doors between any two rooms are automatic, that opens only when the robot is directly in front of the door. This requires the robot to navigate to the door and is encoded using the high-level action `goto_door`. Upon reaching the goal, since the robot is uncertain about its pose, the robot can be anywhere within its current belief distribution. Taking this into account, on reaching the door it is open only if the trace of the pose covariance is within a certain bound η . If the trace is within the bound, an edge is added to the Probabilistic Roadmap (PRM) [47] graph between the current node and the nearest node in the next room to which the robot can navigate via the door. Once the robot traverses the door to reach the next room, the newly added edge is removed from the roadmap. This process is illustrated in Fig. 4. The addition and deletion of edges is performed by the external module.

A fragment of the corridor PDDL domain is shown in Fig. 5. Similar to the office domain, we use the following parameters: $?from$ is the room the robot is currently at and $?to$ is the room

```

(:durative-action goto_region
:parameters (?v - robot ?from ?to - region)
:duration (= ?duration 100)
:condition (at start (robot_in ?v ?from))
:effect (and (at start (not (robot_in ?v ?from)))
(at start (increase (triggered ?from ?to) 1))
(at end (robot_in ?v ?to)) (at end (assign (triggered ?from ?to) 0))
(at end (increase (act-cost) (external)))
(at end (increase (goal-trace) (bound))))

(:durative-action collect_document
:parameters (?v - robot ?r - region)
:duration (= ?duration 20)
:condition (and (at start (robot_in ?v ?r)) (at start (> (get ?r) 0))
(over all (robot_in ?v ?r)))
:effect (and (at end (collected ?r))(at end (increase (act-cost) 4))))

(:durative-action goto_lift
:parameters (?v - robot ?from ?to - region)
:duration (= ?duration 100)
:condition (at start (robot_in ?v ?from))
:effect (and (at start (not (robot_in ?v ?from)))
(at start (increase (triggered ?from ?to) 1))
(at end (reached ?to)) (at end (assign (triggered ?from ?to) 0))
(at end (increase (act-cost) (external))))

```

Fig. 3. A fragment of the PDDL office domain.

which the robot needs to visit, $?d$ is any door. The predicate `visited_in` checks if the robot has visited a room, `hasdoor` checks if the room has a door that opens to another room, and expanded model the change in the roadmap. Similar to the previous domain, the actions `goto_room` and `goto_door` invoke the external module call once the fact `(increase (act-cost) (external))` is encountered. Here, `act-cost` is the direct variable in V^{dir} variable and `external` is the indirect variable in V^{ind} . The function `(triggered ?from ?to)` and `(expanded ?r ?d)` are assigned the value of 1 each time the actions are expanded and re-initialized to 0 once the action duration is completed. This is performed so that the grounded variables `from` (start) and `to` (goal) as well as `r` (start) and `d` (goal) are communicated to the motion planner. The variables `from`, `to` and `r` are used to denote the rooms and the variable `d` represents the doors available. This can be seen in the `parameters` definition of the actions. The variable `external` returns the motion cost computed by the external module.

4.2. Motion planning

Independently of the domain, we use a sampling based PRM to instantiate robot poses for the task actions. To begin with, the initial mean and covariance of the robot pose is assumed to be known. This means that the initial state s_0 corresponds to a single pose instantiation q_0 . The regions to be navigated to are also instantiated into poses, by sampling from the pose space within each region. Once an action a_i is expanded by the task planner, the corresponding start and goal states, that is s_i and s_{i+1} are communicated to the motion planner. This is facilitated by the

functions `triggered` and `expanded`, as detailed in the previous section. For example, the task state s_i might specify that the robot is in cubicle c_2 and the goal state s_{i+1} can be for the robot to reach cubicle c_4 . In this scenario $\phi(s_i)$ and $\phi(s_{i+1})$, that is, the mapping from states to configurations, correspond to all possible poses such that the robot is in cubicles c_2 and c_4 respectively. Since the set of possible poses is infinite, we randomly sample a set of poses corresponding to each task state s_i . It is to be noted that this sampling is an independent problem and this set is incorporated while building the entire roadmap. For each region s_i , the number of pose instantiations will be denoted by s_i^n and a particular instantiation by $s_i^{n_k}$. With the pose instantiation of s_i as the start node, for each pose instantiation of s_{i+1} , we simulate a sequence of controls along each edge starting from $s_i^{n_k}$ and ending in $s_{i+1}^{n_j}$, estimating the beliefs at the each of these nodes using (8)–(9). The $s_{i+1}^{n_j}$ that corresponds to the minimum cost is then selected as the goal pose to reach, for the state s_{i+1} . Thereafter, this instantiation becomes the start node when an expansion is attempted from state s_{i+1} . It is true that PRM is in the configuration space and not in the belief space, but the basic problem remains the same since we are essentially finding a sequence of actions that minimizes the objective function which is a function of the resulting beliefs. Our PRM approach is similar to the Belief Roadmap (BRM) [48] approach and differs in the way one-step belief updates are performed. Moreover, BRM assume maximum likelihood observations but we do not.

Since we plan in the belief space of the robot state, given the mean and covariance of the starting node we propagate the belief along the edges of the PRM as the roadmap is expanded during the search. Belief update is performed upon reaching a node if

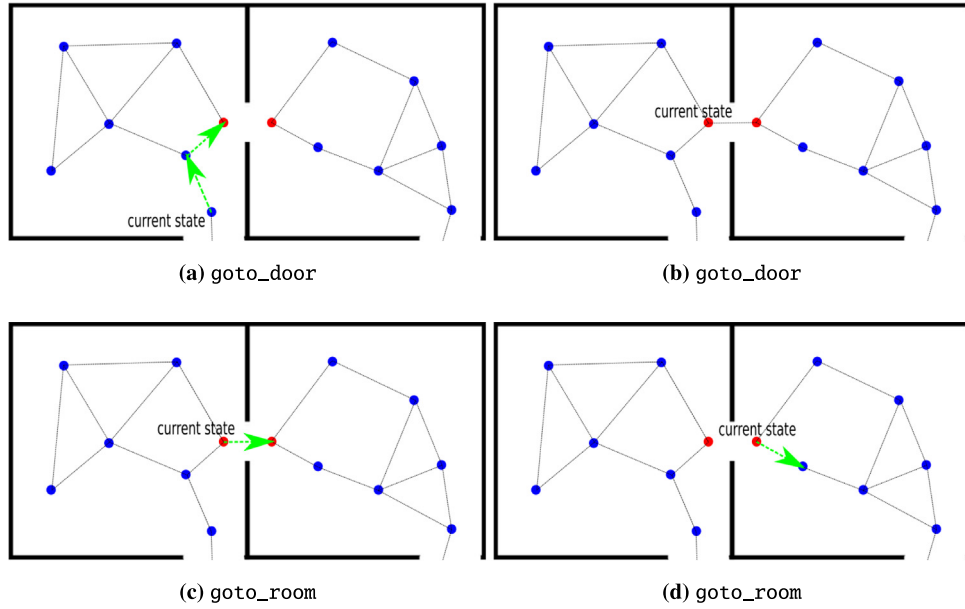


Fig. 4. The addition and deletion of an edge to the PRM graph. The red nodes are the ones that are close to the door. (a) Shows a possible path in green, when the `goto_door` action is expanded. Note that there is no edge between the two red colored nodes. (b) Upon satisfying the trace constraint, an edge added between the two nodes close to the door. (c) The `goto_room` action takes the robot to the next room. (d) As the robot navigates towards the first node (red colored node) in the new room, the edge connecting it to the room from which the robot traversed is removed from the roadmap. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

```
(:durative-action goto_room
:parameters (?from ?to - room)
:duration (= ?duration 100)
:condition (and (at start (robot_in ?from)) (at start
(connected ?from ?to)))
:effect (and (at start (not (robot_in ?from)))
(at start (increase (triggered ?from ?to) 1))
(at end (robot_in ?to)) (at end (assign (triggered ?from ?to) 0))
(at end (increase (act-cost) (external))) (at end (visited ?to))))

(:durative-action goto_door
:parameters (?r - room ?d - door)
:duration (= ?duration 40)
:condition (and (at start (robot_in ?r)) (at start (hasdoor ?r ?d))
(over all (robot_in ?r)))
:effect (and (at start (increase (expanded ?r ?d) 1))
(at end (assign (expanded ?r ?d) 0))
(at end (increase (act-cost) (external))))))
```

Fig. 5. A fragment of the PDDL corridor domain.

a landmark is successfully detected by the robot's perception system. Since we are in the planning phase and yet to obtain observations, we simulate future observations z_{k+1} given the propagated belief $b[X_{k+1}]$, the set of landmarks $L_{\mathbb{N}} = l_1, \dots, l_n$ and the measurement model (5). In this work, we model landmarks using AprilTags [49] which are placed on the objects of interest. Given a pose $x \in b[X_{k+1}]$, the nominal observation $\hat{z} = h(x, l_i)$ is corrupted with noise to obtain z_{k+1} , which is then used to compute the posterior belief.

4.3. Task-motion planning for navigation

In our approach, the interface between task and motion planning occurs through semantic attachments. Formally, semantic attachment can be defined as

Definition 9. Semantic attachments is a functional mapping from the set of direct variables to the set of indirect variables, that is, $\chi : V^{dir} \rightarrow V^{ind}$.

Algorithm 1 TMP for Navigation in Belief Space

Input: $\Psi = (C, \Omega, \phi, \xi, q_0)$: Task-Motion domain, η : Uncertainty budget

```

1: while true do
2:    $a_i \leftarrow \text{task planning}(\Omega)$ 
    $\triangleright a_i$  is an action selected to expand the next state
3:    $\pi^* \leftarrow \emptyset$  // Task-Motion Plan
4:   if  $a_i \in A_s$  then
5:     External module  $\leftarrow V^{dir}$ 
      $\triangleright V^{dir} = \{\text{act} - \text{cost}, \text{goal} - \text{trace}\}$ 
6:     current task state  $\leftarrow s_i$ , next task state  $\leftarrow s_{i+1}$ 
7:      $c \leftarrow \emptyset, T \leftarrow \emptyset$ 
8:     current task state  $\leftarrow \phi(s_i)$ , next task state  $\leftarrow \phi(s_{i+1})$ 
9:     for each  $s_{i+1}^{n_j} \in \phi(s_{i+1})$  do
10:      start node  $\leftarrow s_i^{n_k}$ , goal node  $s_{i+1}^{n_j}$ 
11:      Belief space search from start node to goal node.
12:       $c \leftarrow c^j, T \leftarrow \tau_i^j$ 
13:    end for
14:     $j^* = \arg \min c$ 
15:     $\tau_i \leftarrow \tau_i^{j^*}$ 
      $\triangleright \tau_i$  is the selected motion plan to arrive at the task
    state  $s_{i+1}$ .
16:     $V^{ind} \leftarrow \text{External module}$ 
17:     $\pi^* \leftarrow \text{append}(\pi^*, (a_i, \tau_i))$ 
18:  end if
19: end while
20: return  $\pi^*$ 

```

Algorithm 2 Belief space search

Input: Roadmap (sampled poses and edges), start node n with belief (μ_n, Σ_n) corresponding to start state s_i , goal node $(\phi(s_{i+1}))$

```

1:  $\tau_i \leftarrow n$ 
2: while  $\phi(s_{i+1})$  not reached do
3:   for each edge from  $n$  to  $n'$  do
4:     Propagate the belief (8)
5:     if Landmark within sensing range then
6:       Compute posterior belief (9).
7:     end if
8:     Select  $n'$  with minimum cost.
9:      $c \leftarrow \text{minimum cost}, \tau_i \leftarrow \text{append}(\tau_i, n')$ 
10:     $n' = n$ 
11:  end for
12: end while
13: return  $c, \tau_i$ 

```

We recall here that for the *office domain* $V^{dir} = \{\text{act-cost}, \text{goal-trace}\}$ and $V^{ind} = \{\text{external}, \text{bound}\}$. For the *corridor domain*, we have $V^{dir} = \{\text{act-cost}\}$ and $V^{ind} = \{\text{external}\}$. The planner receives as input the PDDL domain, problem description, the shared library and other input parameters. The input parameter specifies the regions/rooms and the corresponding pose instantiations. For the office domain, these pose instantiations are the poses that lie inside the cubicles and for the corridor domain they are the poses that lie inside the rooms. These poses are sampled once the map of the environment is available as described in the previous section.

An overview of our TMP approach is presented in Algorithm 1. The external module computes the V^{ind} values and is invoked only when a change occurs in V^{dir} variables due to the action effects. The PDDL keyword *increase* is overloaded to refer to an encapsulated object [15] and the external module is called if the

PDDL action to be expanded has an *effect* of the form (increase (v_i^{dir}) (v_j^{ind})), where $v_i^{dir} \in V^{dir}$ and $v_j^{ind} \in V^{ind}$. We denote the set of such actions by A_s . It is to be noted that the elements of this set can vary depending on the requirements of a particular domain. However, the process for achieving the semantic attachments of the external module remains the same. In this paper, the set $A_s = \{\text{goto_region}, \text{goto_lift}, \text{goto_room}, \text{goto_door}\}$. Every time a v_i^{dir} is changed due to the direct effects of an action $a_i \in A_s$, the values of the respective v_j^{ind} is calculated by the external module, attaching the computed value to the indirect variable v_j^{ind} , thereby updating the state. Once an action a_i is expanded by the task planner, the corresponding start (s_i) and goal (s_{i+1}) task states are communicated to the motion planner through the function (*triggered ?from ?to*) (line 6). For the task state s_i , the robot pose $\tau_i(0) = \phi(s_i)$ is known since it is the mean of the current belief distribution. For the task state s_{i+1} , each pose instantiation $s_{i+1}^{n_j} \in \phi(s_{i+1})$ is considered as a goal node (line 9). With $\tau_i(0)$ as the start node, a motion plan is attempted to each of the goal node $s_{i+1}^{n_j}$. The set of feasible motion plans is obtained by performing a search over the roadmap. Along each edge of the roadmap, the belief at s_i is propagated to $s_{i+1}^{n_j}$ by simulating the sequence of controls and observations. We use EKF to compute the appropriate matrices for belief computation as shown in (8). The posterior belief is computed at each node if a landmark is detected by the robot's sensor. This belief search process is shown in Algorithm 2. The motion costs and the corresponding feasible motion plans are populated to the sets c and T respectively (line 12). The motion plan that corresponds to minimum cost is then computed as $\tau_i^{j^*}$ (lines 14–15). The computed values by the external module is then passed to the respective indirect variables V^{ind} (line 16), achieving semantic attachments. The corresponding motion plan τ_i and the goal node $s_{i+1}^{n_j^*}$ are stored and this goal node subsequently becomes the start node for the roadmap search from s_{i+1} . Consequently, the belief estimates returned by the semantic attachments guide the TRPG in identifying the *helpful actions*, besides providing an efficient heuristic evaluation for the task plan.

For the *office domain*, the feasibility of the motion plan $\tau_i^{j^*}$ is checked by accounting for the trace of the covariance matrix upon reaching a cubicle associated with s_{i+1} , that is, $\text{trace}(\Sigma_{s_{i+1}}^{n_j^*})$. Since the cubicle doors are of specific length, we bound the trace by a constant η . However, the failure of an action a_i to find a feasible motion plan during the current expansion does not mean that it has to be discarded. Feasibility also depends on the sequence of actions performed earlier. A different action sequence prior to a_i can render a_i feasible. Hence infeasible actions are not discarded and are set aside for reattempting later. Consequently the feasibility check is performed for the returned optimal plan π^* . The plan is feasible if for each $a_i \in \pi^*$, the $\text{trace}(\Sigma_{s_{i+1}}^{n_j^*}) < \eta$; else there is no feasible plan.

4.3.1. Cost function

So far we have been agnostic about the cost function used while selecting the nodes for expansion. Though our formulation can be adapted to any generic cost functions we use a standard cost function [50]

$$c \doteq M_u c_u + M_G c_G + M_\Sigma c_\Sigma \quad (10)$$

where c_u is the control usage, c_G is the distance to goal and c_Σ is the cost due to uncertainty, defined as $\text{trace}(\Sigma)$, where Σ is the state covariance associated with the robot belief. M_u , M_G and M_Σ are user-defined weights. For the current node n that is considered for expansion, the cost c is computed for each of the nodes that shares an edge with n . The node with the minimum c is selected as the next node n^* for expansion. As such, this can

be extended to non-myopic planning in a trivial manner, but it is not the current focus of this paper. It is to be noted that n^* is considered only if it is not already in the expanded path with the n being the last node added to the path. So if n^* leads to a cycle, the next best node n^{**} is selected.

As mentioned in the previous section, in case of the *office domain* we add the condition $c_{\Sigma_g} < \eta$, where Σ_g is the trace of the goal state covariance and η is a constant. The cubicle doors have a width of $2m$ and considering maximum uncertainty along the door width we fix $\eta = 3m^2$ as the maximum upper limit and discard the motion plans with $c_{\Sigma_g} > 3$ (see lines 19–24, Algorithm 1). For the *corridor domain*, since the automatic doors are of $1m$ in length, we set an upper bound of $\eta = 0.75m^2$, which corresponds to an uncertainty budget of $0.5m$ in each of the pose component. This check is performed when the robot is at a node directly in front of the door as a result of executing the action `goto_door`. If the estimated covariance is within the uncertainty budget an edge is added between the current node and the nearest node in the next room to which the robot can navigate via the door. Once the robot traverses the door to reach the next room by executing the action `goto_room`, the newly added edge is removed from the roadmap. The process of addition and deletion of an edge occur within the external module as a consequence of the `goto_door` and `goto_room` actions.

4.3.2. Optimality

For a given roadmap, the plan synthesized by our approach is optimal at the task-level. This means that the task plan cost returned by our approach (c^*) is lower than any of the other possible task plan costs (c). Let us denote the optimal plan corresponding to c^* as π^* . Suppose that there exists a plan π with associated cost c such that $c < c^*$. If π and π^* have the same sequence of actions, this is not possible since the action costs are evaluated by the motion planner and for a given roadmap, the motion cost returned is the optimal for each action, giving $c^* \leq c$. If π and π^* have a different sequence of actions, the task planner ensures that the returned plan is optimal, giving $c^* \leq c$. Therefore, in both the case, we have $c^* \leq c$.

4.3.3. Completeness

We provide a sufficient condition under which the probability of our approach returning a plan approaches one exponentially with the number of samples used in the construction of the roadmap. A task planning problem, $\Omega = (S, A, \gamma, s_0, S_g)$ is complete if it does contain any dead-ends [51], that is there are no states from which goal states cannot be reached. The PRM motion planner is probabilistically complete [52], that is the probability of failure decays to zero exponentially with the number of samples used in the construction of the roadmap. Therefore, if the motion planner terminates each time it is invoked then probability of finding a plan, if it exists, approaches one.

On the one hand our approach is probabilistically complete; on the other hand, it is also resolution complete since the motion plan feasibility depends on the parameter η . Nevertheless, given a fixed value of η , the probability that the planner fails to return a solution, if one exists, tends to zero as the number of samples approaches infinity. In this sense the best that we can guarantee is probabilistic completeness.

5. Implementation and experimental results

In this Section, we validate our approach in two different robot navigation domains, namely *office domain* and *corridor domain* as described in Sections 4.1.1 and 4.1.2. We use the temporal POPF-TIF [45] as our task planner by customizing it to achieve semantic attachments of an external module. The external module

performs a PRM-based planning in the belief space and is implemented as a dynamically loaded shared library that is passed as an input to the planner. The enumeration into direct variables V^{dir} and indirect variables V^{ind} are listed in the external module. The performance are evaluated on an Intel® Core i7-6500U under Ubuntu 16.04 LTS.

First, we present the motion and sensor models used in our experiments.⁷ Then, we discuss the metrics devised to evaluate the usefulness and validity of our approach. Finally, we present the evaluation of our approach in the two navigation domains using the devised metrics.

5.1. Motion and sensor model

The robot dynamics is modeled using the following non-linear model [53]

$$\begin{aligned} x_{k+1}(1) &= x_k(1) + \delta_{trans} \cdot \cos(x_k(3) + \delta_{rot1}) \\ x_{k+1}(2) &= x_k(2) + \delta_{trans} \cdot \sin(x_k(3) + \delta_{rot1}) \\ x_{k+1}(3) &= x_k(3) + \delta_{rot1} + \delta_{rot2} \end{aligned} \quad (11)$$

where $x_k \doteq (x, y, \theta)$, is the robot pose at time k with $x_k(1) = x$, $x_k(2) = y$ and $x_k(3) = \theta$ and $u_k \doteq (\delta_{rot1}, \delta_{trans}, \delta_{rot2})$ is the applied control. The model in (11) assumes that the robot ideally implements the following commands in order: rotation by an angle of δ_{rot1} , translation of δ_{trans} and a final rotation of δ_{rot2} orienting the robot in the required direction.⁸ It is to be noted that the robot accrue translational and rotational errors while executing u_k .

In the EKF, the Jacobian of the state transition model with respect to the state x_k denoted by F_k (see (8) and (9)) is obtained by linearizing the state transition function about the mean state at x_k and is given by

$$\begin{aligned} F_k &= \begin{bmatrix} \frac{\partial f}{\partial x_k(1)} & \frac{\partial f}{\partial x_k(2)} & \frac{\partial f}{\partial x_k(3)} \\ \frac{\partial f}{\partial x_k(1)} & \frac{\partial f}{\partial x_k(2)} & \frac{\partial f}{\partial x_k(3)} \\ \frac{\partial f}{\partial x_k(1)} & \frac{\partial f}{\partial x_k(2)} & \frac{\partial f}{\partial x_k(3)} \end{bmatrix} \\ &= \begin{bmatrix} 1 & 0 & -\delta_{trans} \cdot \sin(x_k(3) + \delta_{rot1}) \\ 0 & 1 & \delta_{trans} \cdot \cos(x_k(3) + \delta_{rot1}) \\ 0 & 0 & 1 \end{bmatrix} \end{aligned} \quad (12)$$

Similarly, the linearized process noise, $R_k = V_k W_k V_k^T$, is obtained by computing the Jacobian of V_k

$$\begin{aligned} V_k &= \begin{bmatrix} \frac{\partial f}{\partial \delta_{rot1}} & \frac{\partial f}{\partial \delta_{trans}} & \frac{\partial f}{\partial \delta_{rot2}} \\ \frac{\partial f}{\partial \delta_{rot1}} & \frac{\partial f}{\partial \delta_{trans}} & \frac{\partial f}{\partial \delta_{rot2}} \\ \frac{\partial f}{\partial \delta_{rot1}} & \frac{\partial f}{\partial \delta_{trans}} & \frac{\partial f}{\partial \delta_{rot2}} \end{bmatrix} \\ &= \begin{bmatrix} -\delta_{trans} \cdot \sin(x_k(3) + \delta_{rot1}) & \cos(x_k(3) + \delta_{rot1}) & 0 \\ \delta_{trans} \cdot \cos(x_k(3) + \delta_{rot1}) & \sin(x_k(3) + \delta_{rot1}) & 0 \\ 1 & 0 & 1 \end{bmatrix} \end{aligned} \quad (13)$$

The noise covariance matrix W_k is formulated as below with α_1 to α_4 being the robot-specific error parameters [53] modeling the accuracy of the robot motion

$$W_k = \begin{bmatrix} \alpha_1 \cdot \delta_{rot1}^2 + \alpha_2 \cdot \delta_{trans}^2 & 0 & 0 \\ 0 & \alpha_3 \cdot \delta_{trans}^2 + \alpha_4 \cdot (\delta_{rot1}^2 + \delta_{rot2}^2) & 0 \\ 0 & 0 & \alpha_2 \cdot \delta_{trans}^2 + \alpha_1 \cdot \delta_{rot2}^2 \end{bmatrix} \quad (14)$$

⁷ To simplify the notation, most variables are presented without time indexes.

⁸ The state transition model form of (11) is given in (4).

As for the sensor model, we use a landmark-base model

$$z_k = \begin{bmatrix} r = \sqrt{(l_i(1) - x_k(1))^2 + (l_i(2) - x_k(2))^2} \\ \phi = \arctan\left(\frac{l_i(2) - x_k(2)}{l_i(1) - x_k(1)}\right) - x_k(3) \end{bmatrix} + v_k, \quad v_k \sim \mathcal{N}(0, Q_k) \quad (15)$$

where r and ϕ are the range and bearing of the i th landmark l_i relative to the robot frame. The sensor model is linearized to obtain the Jacobian H_k , which is the partial derivative of the measurement function with respect to the robot state.⁹

$$H_k = \begin{bmatrix} \frac{\partial r}{\partial x_k(1)} & \frac{\partial r}{\partial x_k(2)} & \frac{\partial r}{\partial x_k(3)} \\ \frac{\partial \phi}{\partial x_k(1)} & \frac{\partial \phi}{\partial x_k(2)} & \frac{\partial \phi}{\partial x_k(3)} \end{bmatrix} = \begin{bmatrix} -\frac{(l_i(1) - x_k(1))}{r} & -\frac{(l_i(2) - x_k(2))}{r} & 0 \\ \frac{(l_i(2) - x_k(2))}{r^2} & -\frac{(l_i(1) - x_k(1))}{r^2} & -1 \end{bmatrix} \quad (16)$$

We would like to reiterate the fact that since we are in the planning phase, the nominal observation $\hat{z} = h(x, l_i)$ is corrupted with noise to simulate future observations.

5.2. Plan metrics

To benchmark our approach we consider four different cost formulations that differ in their motion cost computation and thereby the task-level action costs. Though our formulation can be adapted to any general cost function (see Section 4.3.1), we choose the following four cost functions to demonstrate the efficiency of our approach:

- *Euclidean cost*: The motion planner is never called and the task cost are evaluated computing the Euclidean distance c_{euc} between the geometric instantiations of s_i and s_{i+1} , that is, between $\tau_i^j(0)$ and $\tau_i^j(1)$. Here $c \doteq c_{euc}$.
- σ -*Euclidean cost*: This configuration evaluates the motion cost as the sum of Euclidean distance between $\tau_i(0)$ and $\tau_i(1)$ and the cost due to uncertainty, defined as $c_\Sigma = \text{trace}(\Sigma)$, where Σ is the covariance at each node of τ_i . The general form of this cost function is $c \doteq M_{euc}c_{euc} + M_\Sigma c_\Sigma$.
- *PETLON cost*: In this configuration, the motion planner returns the trajectory length or the geometric-level cost of traversing from s_i to s_{i+1} , that is, from $\tau_i^j(0) \in \phi(s_i)$ to $\tau_i^j(1) \in \phi(s_{i+1})$. The general form of the cost for this configuration is $c \doteq M_u c_u + M_G c_G$, where c_u is the control usage and c_G is the distance to goal. Since we assume straight line path between two sampled poses, the applied control for translation, that is δ_{trans} represents the trajectory length. We note here that the motion planner in PETLON [37] computes the geometric-level cost of traversing from one state to another and hence this configuration will be used to compare MPTP with PETLON.
- *MPTP cost*: In this configuration, we use the cost function as defined in Section 4.3.1, that is, $c \doteq M_u c_u + M_G c_G + M_\Sigma c_\Sigma$, where c_u is the control usage, c_G is the distance to goal and c_Σ is the cost due to uncertainty. It is noteworthy that *PETLON cost* is subsumed in *MPTP cost* since *MPTP cost* is fundamentally *PETLON cost* added with the cost due to uncertainty.

5.3. Office domain

This domain is simulated in Gazebo [54] by constructing an office environment of 36 m \times 25 m; top view of the simulated

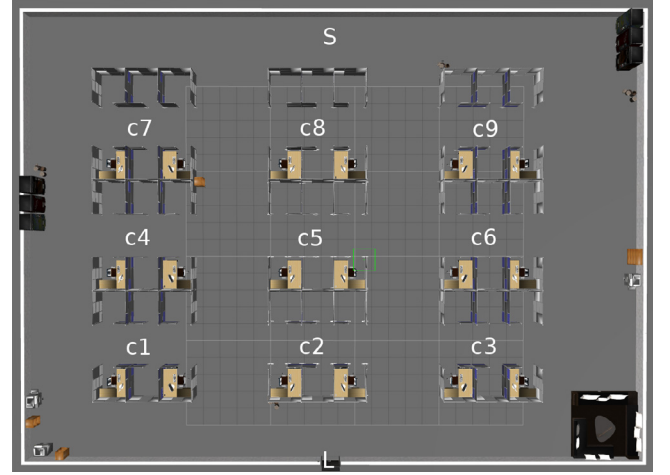


Fig. 6. Top view of the simulated environment in Gazebo. See office domain in Section 4 for a detailed description.

environment is shown in Fig. 6. We note here that the landmarks considered in this domain are the objects outside the cubicles like printers, trash cans, lounge, vending machines and book-shelves. The robot is required to collect documents from different cubicles, and the documents are then taken to the next floor via the lift L .

5.3.1. Validation

We first demonstrate the need for a combined TMP for navigation. Unless otherwise stated, the planning times presented is an average for 25 different planning sessions. Consider the following scenario in which the robot is required to collect documents from the cubicles $c3$, $c4$, $c6$ and $c9$. We first run the planner with *Euclidean cost* to synthesize the task plan. We remind that in this configuration the motion planner is never called and the action costs are evaluated by considering the Euclidean distance between the start and goal regions. The plan synthesized is $S \rightarrow c3 \rightarrow c4 \rightarrow c6 \rightarrow c9 \rightarrow L$. This plan is then given to the motion planner, to compute the corresponding cost due to uncertainty c_Σ which is the trace of the robot state covariance. The task planning cost and the motion planning cost are added to estimate the overall planning cost, which equated to 298.84. The addition of the two costs is possible because we first compute the task plan which is then passed to the motion planner to compute the cost due to uncertainty. Therefore the overall planning cost is the task planning cost combined with motion planning cost. In the same way, the overall planning time was computed to be 0.94 (± 0.09) seconds by adding the time for task planning and motion planning, respectively. Next, we ran the planner with σ -*Euclidean cost*, returning the plan $S \rightarrow c4 \rightarrow c9 \rightarrow c6 \rightarrow c3 \rightarrow L$, in 1.28 (± 0.06) seconds with a total cost of 90.89. This configuration evaluates the motion cost as the sum of Euclidean distance and the cost due to uncertainty. It is seen that there is a significant difference in the plan quality as the cost is improved by a factor of 3 for σ -*Euclidean cost*. This difference in cost is attributed to the different task sequence synthesized. Essentially, *Euclidean cost* corresponds to planners that pre-compute motion costs of all task-level actions or use an admissible heuristic for the same (for example, the approach in [35]). The task plan is then given to the motion planner for execution, assuming that such a motion plan exists. In contrast, σ -*Euclidean cost* checks for the motion feasibility and estimates the motion costs while expanding each task-level action and thus corresponds to an integrated TMP approach as discussed in this paper. The difference in plan quality between *Euclidean cost*

⁹ The measurement function form of (15) is given in (5).

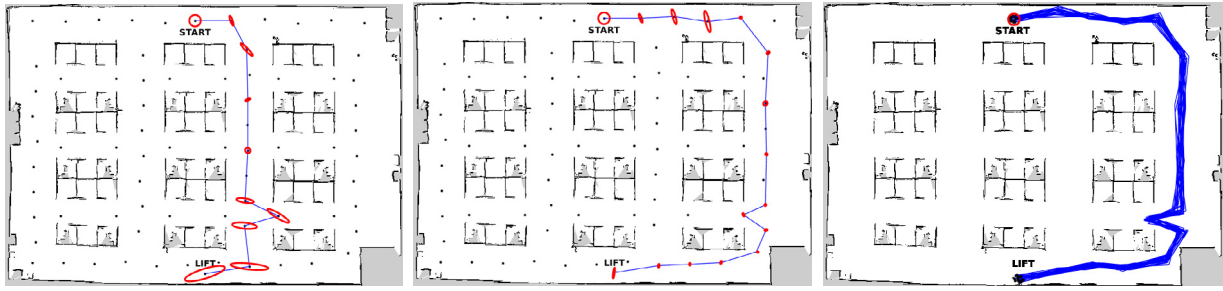


Fig. 7. (left and center) The propagated belief distributions along the planned paths for *PETLON* cost and *MPTP* cost. The belief estimates for a single planning instantiation corresponding to a unique set of simulated observations are shown. Black dots represent the sampled poses. (left) Shortest path route that corresponds to *PETLON* cost. (center) Belief space planning corresponding to *MPTP* cost, returning an information rich route. (right) Traces of robot's true state while starting from the initial belief – run with *MPTP* cost.

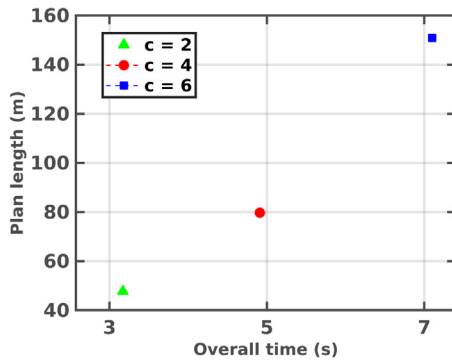


Fig. 8. Plan length with overall planning time. MPTP is run with *PETLON* cost and a sampling density of $d = 1.5$.

and σ -Euclidean cost clearly demonstrates the efficiency of a combined TMP approach as opposed to performing task planning and motion planning separately. Though our considered scenario is much less knowledge-intensive than real-world scenarios, the above example conveys the need for a combined task-motion planner.

Next, we run the planner with *PETLON* cost and *MPTP* cost to demonstrate the advantage of planning in belief space, that is using our MPTP approach. We recall here that similar to PETLON [37], with *PETLON* cost, the motion planner evaluates the geometric-level cost of traversing $\tau_i(0)$ to $\tau_i(1)$, whereas with *MPTP* cost, in addition to considering the geometric-level cost of traversing, the cost due to uncertainty is also incorporated. We consider a scenario in which the robot has to collect a document from cubicle c_3 . The planned trajectories in both the scenarios with the corresponding covariance estimated at each node (only the (x, y) portion is shown) is shown in Fig. 7. Clearly, the belief space task-motion planner (*MPTP* cost) returns a route which is rich in sensor information (see Fig. 7 in the mid), enabling effective localization. *PETLON* cost returns the shortest path trajectory but with an increased robot state uncertainty. Fig. 7 on the right hand side shows the traces of true robot state for 25 different simulations while running on *MPTP* cost—the initial state being sampled from the known initial belief.

5.3.2. Scalability

We test the scalability of our approach by increasing the task-level complexity. We run our planner on three different scenarios where 2, 4, 6 number of cubicles ($c = 2, 4, 6$) are to be visited to collect the corresponding number of documents. This results in evaluating more task-level actions, escalating the task level complexity. We also test these scenarios on varying levels of sample densities. We choose $d = 1, 1.5, 2$, where $d = i$ corresponds

to an average of i samples per square meter. The tests are run using *MPTP* cost and *PETLON* cost. The overall planning time and the returned cost can be seen in Table 1. While we ran with the *MPTP* cost, for $d = 1$ and $c = 6$, no feasible motion plan is found since the condition $\eta < 1$ is violated. However, for higher sample densities, a feasible motion plan is found. The plan quality is increased with increase in d , but at the expense of exponentially increasing computation time. It is clearly seen that for our considered scenario $d = 1.5$ can be chosen, without much loss of plan quality.

In [37], TMP for navigation is performed by evaluating the geometric cost of traversing. They consider a scenario in which nine objects are placed at different locations. Two objects from among them are to be collected and delivered to a person such that the geometric cost of traversing is minimum. They report a total planning time of about 15 s with a plan length of 37 m. Though the environment considered in [37] is larger than ours, to provide a comparison with PETLON, we run our task-motion planner with *PETLON* cost and evaluate the planning time with respect to the plan length. In comparison, MPTP with *PETLON* cost fares superiorly with respect to increased task-level complexity. To demonstrate this, we first consider three scenarios where 2, 4, and, 6 documents are to be collected to be delivered to the next floor. The results can be seen in Table 1 under the *PETLON* cost section. We note here that for $d = 1.5$ and collecting 6 documents ($c = 6$) MPTP with *PETLON* cost took only about 7 (± 0.34) seconds with a plan length of about 150 m (see Fig. 8). To provide a better comparison, we also evaluate our approach by considering a much larger environment, the Willow Garage world of 58 m \times 45 m as shown in Fig. 9(a). In this example, the robot (at start) needs to collect any two objects from among nine different objects (location of objects marked as green blobs), and deliver it to a person at the goal location (shown in red). We ran our planner with *PETLON* cost, returning an optimal plan of length 53.94 m in 3.69 (± 0.09) seconds. We recall here that for the same scenario, PETLON [37] report a planning time of about 15 s for a plan length of 37 m. In contrast, MPTP with *PETLON* cost is almost three faster. This clearly elucidates the superiority of our approach. PETLON first computes a task plan using an admissible heuristic which is then sent to the motion planner for actual cost evaluation. This cost refinement process is iterated until the optimal plan is found. MPTP does not require such an iteration since it evaluates the motion cost using semantic attachments as the action is expanded by the task planner. The scalability to increasing task complexity is tested by varying the number of objects to be collected (see Fig. 9(b)). The task in which four objects are to be collected was completed in only about 25 (± 1.64) seconds. Therefore MPTP reveals to be much faster than PETLON and is robust to the increasing number of objects and map size.

Table 1

Overall planning time and cost returned while running the task-motion planner with *MPTP cost* and *PETLON cost*. The average number of samples per square meter is denoted by d . $c = 2, 4$ and 6 denotes the number of cubicles to be visited, increasing the task-level complexity. '-' denotes the fact that no plan is found as the condition $\eta < 1$ is violated.

	d	Overall time (s)			Cost		
		$c = 2$	$c = 4$	$c = 6$	$c = 2$	$c = 4$	$c = 6$
<i>MPTP cost</i>	1	1.34 ± 0.05	2.24 ± 0.15	-	83.84	90.27	-
	1.5	3.41 ± 0.08	7.16 ± 0.12	14.04 ± 0.09	88.18	101.01	237.59
	2	9.11 ± 1.17	28.48 ± 1.19	46.15 ± 2.23	92.32	126.96	260.092
<i>PETLON cost</i>	1	0.47 ± 0.02	0.77 ± 0.04	1.77 ± 0.01	47.80	84.88	161.47
	1.5	3.17 ± 0.03	4.91 ± 0.02	7.10 ± 0.10	55.77	95.74	174.90
	2	6.08 ± 0.11	9.86 ± 0.17	15.14 ± 1.09	56.19	95.77	181.06

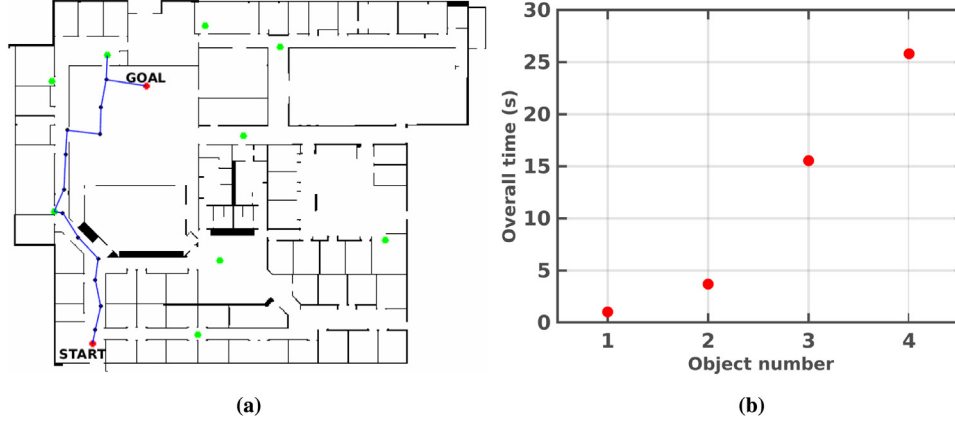


Fig. 9. (a) Willow Garage world with nine objects whose instances are marked as green blobs. The optimal path when two objects are to be collected is shown in blue. The planner is run with *PETLON cost*. (b) Overall planning time with increasing number of objects to be collected for delivering. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

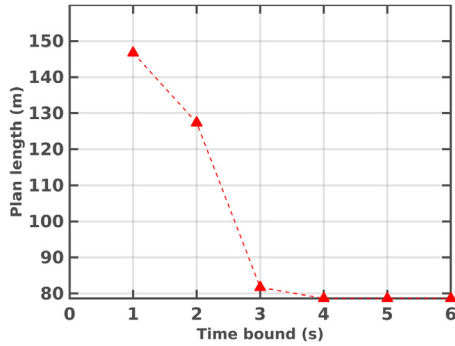


Fig. 10. Anytime property of MPTP. Valid solutions are returned even when strict bounds are placed on the planning time.

POPF-TIF supports *anytime* planning which means that the planner searches for improved solutions until it has exhausted the search space or is interrupted. Specifically, POPF-TIF is run with a $-n$ flag to activate anytime search. A time bound may be specified with the flag $-tx$, where x is the time bound in seconds and is used in situations with strict time bounds where optimality is sacrificed. We demonstrate this by considering the Willow Garage world in which the robot needs to collect any three documents from among the nine objects and deliver it to a person. We start with a time bound of 1 s and increment it by a second until an optimal solution is found. The result is plotted in Fig. 10. As the time bound is incremented, the plan quality is increased and for a planning time bound of 4 s, the optimal plan length of 78.63 m is returned.

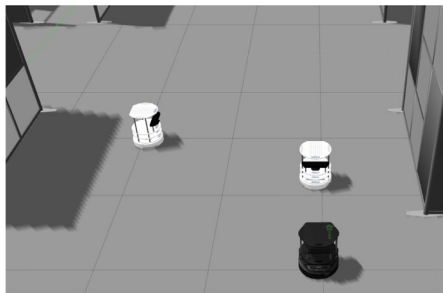
We stress here the fact that in this paper we are mainly concerned with planning and the synthesized plans are given to the robot for execution. Thus, any such execution approach may be employed. In this work, the generated plans are executed with a TurtleBot robot in the simulated Gazebo environment. We use AprilTags [49] to identify objects like printers, trash cans, as landmarks. TurtleBot robot in front of one such landmark is seen in Fig. 11. A ROS-based architecture has been developed to implement the approach. Belief estimation is carried out using EKF. We note here that presently we consider static obstacles while planning and therefore the planned trajectories are collision-free. However, to be robust to dynamic obstacles, the plan execution is trivially extended to employ any collision avoidance approach in dynamic environments [55,56]. Snapshots of dynamic obstacle avoidance during the execution of a plan can be seen in Fig. 12. As seen in the figure, dynamic obstacles are simulated using TurtleBot robots (white in figure). We now report here the execution time for the scenario discussed in Section 5.3.2. When 2, 4, 6 number of cubicles are to be visited to collect the corresponding number of documents, the execution times are 140.21s (± 3.11 s), 366.40s (± 4.99 s), and 664.71s (± 16.28 s), respectively. We note here that the execution time varies with robot and its control limits.

5.4. Corridor domain

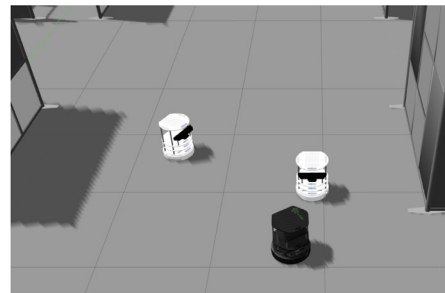
Our *corridor domain* (see Section 4.1.2 for a detailed description) is a variant of the robot navigation domain in [33]. However, they treat it as a task planning problem assuming that feasible motion plans exist for the synthesized task plans. In contrast, we perform task-motion planning. In this domain of 12 m \times 25 m,



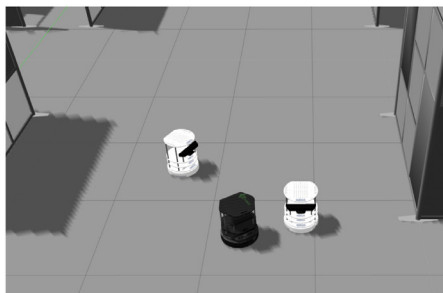
Fig. 11. A robot in front of AprilTags which provide the transformation between the robot pose and the landmark pose.



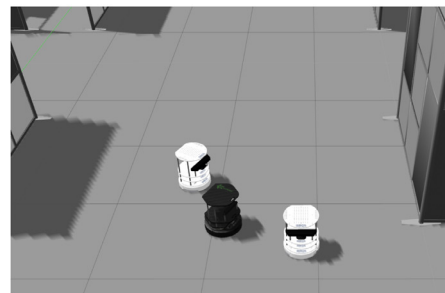
(a)



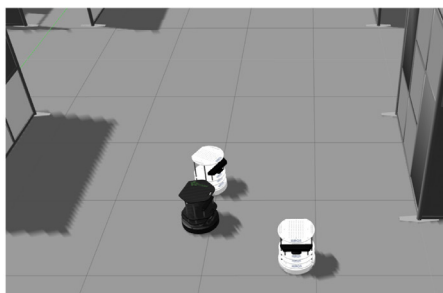
(b)



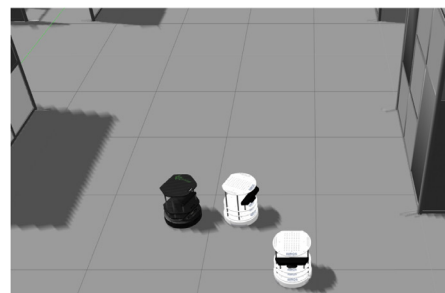
(c)



(d)



(e)



(f)

Fig. 12. A robot avoiding a couple of dynamic obstacles (white TurtleBot robots) during execution. Our approach is not restrictive to any particular execution strategy and any approach that employs dynamic obstacle avoidance may be used.

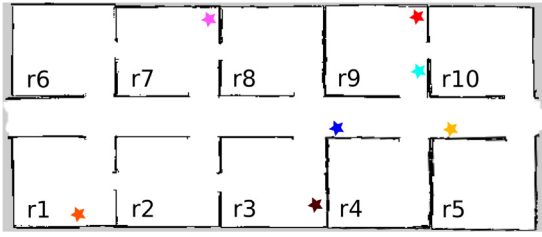


Fig. 13. Map of the building floor environment with half the rooms connected directly by doors. The stars with different colors represent landmarks that aids the robot in better localization. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

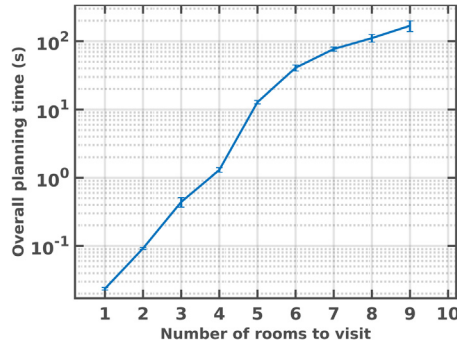


Fig. 14. Overall task-motion planning time for different number of rooms that need to be visited in log scale. Planning times are the average for 25 different runs.

a mobile robot, starting from a given room, navigates an office floor to visit a set of rooms that are selected randomly. The office floor has ten rooms and the robot is initially located in room 1. All the rooms are connected to each other through the central corridor. In addition, five rooms are directly connected with each other via doors which need to be opened by the robot. The goal is to visit a set of rooms R that are randomly selected for each run. Since these visits have to be carried out expending as less cost as possible, the robot needs to assess the accessibility between the rooms that are directly connected to each other via a door. This is facilitated through the `goto_door` action as discussed in Section 4.1.2. The map of the building floor is as shown in Fig. 13.

5.4.1. Validation and scalability

First, we run the planner with *MPTP cost*. For a fixed set cardinality $|R|$ (set elements are the rooms to be visited), 25 trials are performed, where the set elements are selected randomly for each trial. The average planning time for each of them is shown in Fig. 14. While the planning time does scale exponentially with $|R|$, the plan for $|R| = 9$ is computed in less than 3 min. The work in [33] evaluates the task planning performance on a similar domain randomly selecting the number of rooms to visit in each trial. Since MPTP performs task-motion planning, the overall MPTP planning times with increasing $|R|$ is greater than those reported in [33]. However the graph of $|R|$ with planning time (Fig. 14) follows a similar trend to that reported in [33]. It is noteworthy that for a given $|R|$, the difference in MPTP planning time and the planning time reported in [33] is significantly less.

Next, we run the planner with *PETLON cost* and *MPTP cost*. We consider a scenario in which the robot, starting from room $r1$, has to visit rooms $r2$ and $r3$. As seen in Fig. 13, rooms $r1$, $r2$ and $r2$, $r3$ are also connected by doors between them. Fig. 15 on top-left and top-right shows the planned trajectories in both the scenarios with the corresponding covariance estimated at each

node (only the (x, y) portion is shown). Note that the illustrations show a single planning instantiation corresponding to a unique set of simulated observations Z . Belief space planning (*MPTP cost*) enables effective localization by returning a route which is rich in sensor information (see Fig. 15 on top-right). Fig. 15 on the bottom-left, shows the traces of true robot states for 25 different simulations while running on *PETLON cost*. The initial poses are sampled from the known initial belief distribution. Out of the 25 trials, 20 lead to collision on the walls, giving a success rate of only 20%. The traces of true robot pose for 25 different simulations while running on *MPTP cost* is shown in Fig. 15 (bottom-right). Only 2 trials lead to collision, giving a success rate of 92%.

Finally, we test the scalability of our approach by running the planner with varying number of rooms that are directly connected by doors between them. We consider a scenario in which seven rooms are to be visited. We consider five different cases of this scenario, each of which has a fixed number of rooms that are directly connected by the doors. For each case, 25 trials are performed and for each trial, the rooms with doors between them are randomly selected. The overall planning time is seen in Fig. 16.

6. Discussion

In this section, we first discuss some limitations of our approach and later comment on the relation to multi-goal planning and traveling salesman problems.

MPTP has few limitations and assumptions and relaxing them would enhance the capability and robustness of our approach in challenging scenarios. First, we sample collision-free poses and therefore considering static obstacles, the planned trajectories are collision-free. In this sense, we employ a deterministic collision avoidance approach and do not compute the probability of collisions while computing a path during planning. It is a reasonable assumption for all practical purposes but is not the case in general while planning in narrow regions or corridors. The execution may be trivially extended to consider collision probabilities, making it robust to both static and dynamic obstacles. Second, we assume straight line path between two sampled poses. This might not fare well in some experimental domains and can lead to larger prediction uncertainties. Presently, as the number of samples vary, the search is performed again. It is our future direction to efficiently utilize the previous search results to reduce the computation time for increased samples. It is also an interesting future direction to extend the framework to an online real-time planning approach.

Multi-Goal Planning (MGP) [57], where a robot visits a sequence of goal configurations is a subset of the general class of TMP problems. Most existing MGP approaches [57–60] leverage the Traveling Salesman Problem (TSP) [61] solvers for task sequencing. A TSP problem finds a minimum cost path traversing a set of points such that every point is visited once. In an MGP problem these points correspond to the set of goal configurations the robot needs to visit. It can be argued that all MGP problems can be modeled as a TMP problem but not vice versa. For instance, consider the *office domain* presented in Section 4.1.1. In this scenario the robot not only has to visit regions of interest but execute actions such as collecting the documents, which is to be performed when visiting each cubicle ensuring that the action preconditions are met. Moreover, in certain scenarios cubicles may need to be visited multiple times violating the single visit constraint of traditional TSP solvers. The *corridor domain* (see Section 4.1.2) presents additional challenges for TSP solvers. If we consider that there are no doors between the rooms, then the problem reduces to just visiting different rooms and can be solved

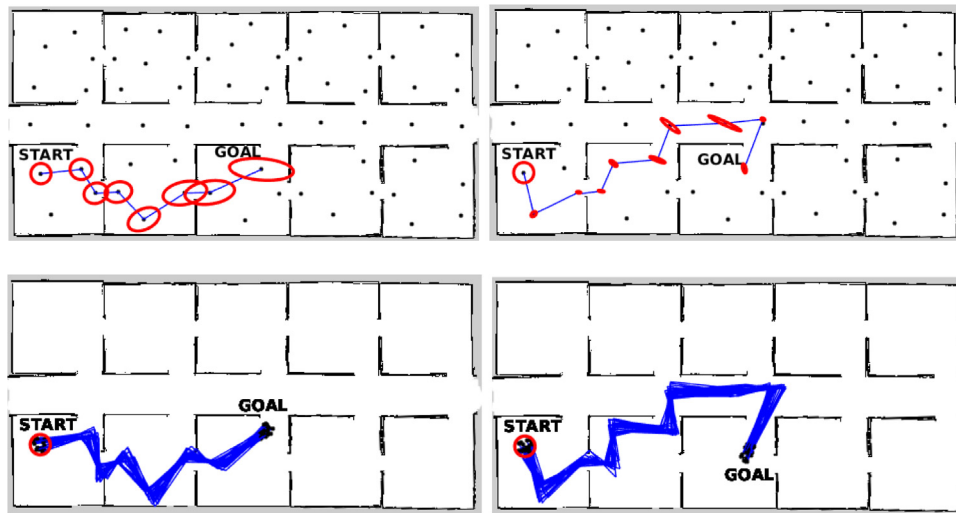


Fig. 15. (top-left and top-right) The propagated belief distributions along the planned paths while running MPTP with *PETLON* cost and *MPTP* cost. The belief estimates for a single planning instantiation corresponding to a unique set of simulated observations are shown. The black dots represent the sampled poses. (top-left) Shortest path route that corresponds to running the planner with *PETLON* cost. (top-right) Belief space planning corresponding to running the planner with *MPTP* cost, returning an information rich route. (bottom-left) Traces of robot's true state while starting from the initial belief and run on *PETLON* cost – 80% of the trajectories lead to collision. (bottom-right) Traces of robot's true state while starting from the initial belief and run on *PETLON* cost – only 8% of the trajectories lead to collision.

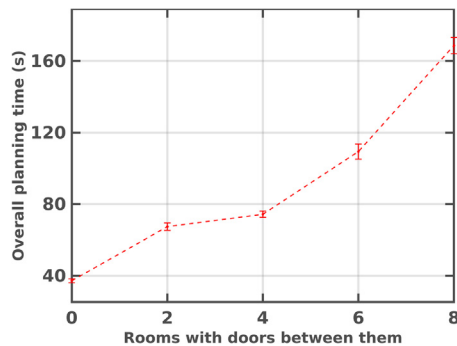


Fig. 16. Overall planning time for visiting 7 rooms when the number of rooms directly connected by doors are varying. Average time for 25 trails are plotted in each case.

using TSP solvers. However, in the considered scenario there are doors between certain rooms and the accessibility between the rooms that are directly connected to each other via a door needs to be assessed by the robot. This requires different levels of reasoning to verify the action preconditions such as, checking if a door exists, navigating to the door, checking if the trace of the robot pose covariance is within the uncertainty budget and if yes, then updating the roadmap. Moreover, if the robot passes through the door, the accomplishment of the action effect (in this case, closing the door corresponds to updating the roadmap) needs to be established. Thus MPTP is able to solve a larger class of problems than traditional TSP solvers.

7. Conclusions

This paper introduces an approach for task-motion planning under motion and sensing uncertainty. Task-motion interaction is facilitated by means of semantic attachments that return motion costs to the task planner. In this way, the action costs of the task plans are evaluated using a motion planner. The plan synthesized is optimal at the task-level since the overall action cost is less than that of other task plans generated for a given roadmap. It is to be noted that the action cost also encompasses

the motion cost. The proposed approach is probabilistically complete and we have validated the framework using a simulated office environment in Gazebo and a corridor environment. The approach has been evaluated with different configurations that correspond to different motion cost computation, illustrating the need for a combined TMP approach for navigation in belief space. Though we have validated MPTP in two different robot navigation domains, real-world scenarios often require large number of tasks to be performed. Real-world domains are much more knowledge-intensive, significantly increasing the task-level and motion-level complexity. The scalability results suggest that our approach fares well with respect to increased task-level complexity and plan length.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgment

We thank Chiara Piacentini for her valuable inputs on the POPF-TIF planner that were very helpful in our implementation.

References

- [1] F. Lagriffoul, N.T. Dantam, C. Garrett, A. Akbari, S. Srivastava, L.E. Kavraki, Platform-independent benchmarks for task and motion planning, *Robot. Autom. Lett.*
- [2] D. McDermott, M. Ghallab, A. Howe, C. Knoblock, A. Ram, M. Veloso, D. Weld, D. Wilkins, PDDL- The Planning Domain Definition Language, in: AIPS-98 Planning Competition Committee, 1998.
- [3] S. Srivastava, E. Fang, L. Riano, R. Chitnis, S. Russell, P. Abbeel, Combined task and motion planning through an extensible planner-independent interface layer, in: *Robotics and Automation (ICRA), IEEE International Conference on*, IEEE, 2014, pp. 639–646.
- [4] C.R. Garrett, T. Lozano-Perez, L.P. Kaelbling, Ffrob: Leveraging symbolic planning for efficient task and motion planning, *Int. J. Robot. Res.* 37 (1) (2018) 104–136.
- [5] L.P. Kaelbling, T. Lozano-Perez, Integrated task and motion planning in belief space, *Int. J. Robot. Res.* 32 (9–10) (2013) 1194–1227.

- [6] N.T. Dantam, Z.K. Kingston, S. Chaudhuri, L.E. Kavraki, An incremental constraint-based framework for task and motion planning, in: Special Issue on the 2016 Robotics: Science and Systems Conference, *Int. J. Robot. Res.* 37 (10) (2018) 1134–1151.
- [7] C.R. Garrett, C. Paxton, T. Lozano-Pérez, L.P. Kaelbling, D. Fox, Online re-planning in belief space for partially observable task and motion problems, *arXiv preprint arXiv:1911.04577*.
- [8] L.P. Kaelbling, M.L. Littman, A.R. Cassandra, Planning and acting in partially observable stochastic domains, *Artificial Intelligence* 101 (1–2) (1998) 99–134.
- [9] A. Thomas, F. Mastrogiovanni, M. Baglietto, Task-motion planning for navigation in belief space, in: *The International Symposium on Robotics Research*, 2019.
- [10] R.E. Fikes, N.J. Nilsson, STRIPS: A new approach to the application of theorem proving to problem solving, *Artificial Intelligence* 2 (3–4) (1971) 189–208.
- [11] N.J. Nilsson, Shakey the robot, *Tech. Rep. 323*, Artificial Intelligence Center, SRI International, Menlo Park, California, 1984.
- [12] C. Dornhege, M. Gissler, M. Teschner, B. Nebel, Integrating symbolic and geometric planning for mobile manipulation, in: *Safety, Security & Rescue Robotics (SSRR)*, IEEE International Workshop on, IEEE, 2009, pp. 1–6.
- [13] S. Cambon, R. Alami, F. Gravot, A hybrid approach to intricate motion, manipulation and task planning, *Int. J. Robot. Res.* 28 (1) (2009) 104–126.
- [14] J. Hoffmann, The metric-FF planning system: Translating ignoring delete lists to numeric state variables, *J. Artificial Intelligence Res.* 20 (2003) 291–341.
- [15] C. Piacentini, V. Alimisis, M. Fox, D. Long, An extension of metric temporal planning with application to ac voltage control, *Artif. Intell.* 229 (2015) 210–245.
- [16] E. Erdem, K. Haspalamutgil, C. Palaz, V. Patoglu, T. Uras, Combining high-level causal reasoning with low-level geometric reasoning and motion planning for robotic manipulation, in: 2011 IEEE International Conference on Robotics and Automation, IEEE, 2011, pp. 4575–4581.
- [17] L.P. Kaelbling, T. Lozano-Pérez, Integrated robot task and motion planning in the now, *Tech. Rep. 2012-018*, Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, 2012.
- [18] H. Kurniawati, Y. Du, D. Hsu, W.S. Lee, Motion planning under uncertainty for robotic tasks with long time horizons, *Int. J. Robot. Res.* 30 (3) (2011) 308–323, <http://dx.doi.org/10.1177/0278364910386986>.
- [19] M. Ghallab, D. Nau, P. Traverso, *Automated Planning and Acting*, Cambridge University Press, 2016.
- [20] A.K. Pandey, J.-P. Saut, D. Sidobre, R. Alami, Towards planning human-robot interactive manipulation tasks: Task dependent and human oriented autonomous selection of grasp and placement, in: 2012 4th IEEE RAS & EMBS International Conference on Biomedical Robotics and Biomechatronics (BioRob), IEEE, 2012, pp. 1371–1376.
- [21] L. de Silva, A.K. Pandey, M. Gharbi, R. Alami, Towards combining htn planning and geometric task planning, in: *RSS Workshop on Combined Robot Motion Planning and AI Planning for Practical Applications*, 2013.
- [22] F. Lagriffoul, D. Dimitrov, J. Bidot, A. Saffiotti, L. Karlsson, Efficiently combining task and motion planning using geometric constraints, *Int. J. Robot. Res.* 33 (14) (2014) 1726–1747.
- [23] M. Toussaint, Logic-geometric programming: An optimization-based approach to combined task and motion planning, in: *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.
- [24] M. Toussaint, M. Lopes, Multi-bound tree search for logic-geometric programming in cooperative manipulation domains, in: 2017 IEEE International Conference on Robotics and Automation (ICRA), IEEE, 2017, pp. 4044–4051.
- [25] T. Lozano-Pérez, L.P. Kaelbling, A constraint-based method for solving sequential manipulation planning problems, in: *Intelligent Robots and Systems (IROS)*, IEEE/RSJ International Conference on, IEEE, 2014, pp. 3684–3691.
- [26] C. Dornhege, P. Eyerich, T. Keller, S. Trüg, M. Brenner, B. Nebel, Semantic attachments for domain-independent planning systems, in: *International Conference on Automated Planning and Scheduling (ICAPS)*, Thessaloniki, Greece, 2009, pp. 114–121.
- [27] M. Stilman, J. Kuffner, Planning among movable obstacles with artificial constraints, *Int. J. Robot. Res.* 27 (11–12) (2008) 1295–1307.
- [28] J. Van Den Berg, M. Stilman, J. Kuffner, M. Lin, D. Manocha, Path planning among movable obstacles: a probabilistically complete approach, in: *Workshop on the Algorithmic Foundations of Robotics VIII*, WAFR, Springer, Guanajuato, Mexico, 2009, pp. 599–614.
- [29] K. Hauser, J.-C. Latombe, Integrating task and PRM motion planning: Dealing with many infeasible motion planning queries, in: *ICAPS Workshop on Bridging the Gap between Task and Motion Planning*, 2009.
- [30] K. Hauser, J.-C. Latombe, Multi-modal motion planning in non-expansive spaces, *Int. J. Robot. Res.* 29 (7) (2010) 897–915.
- [31] P. Khandelwal, S. Zhang, J. Sinapov, M. Leonetti, J. Thomason, F. Yang, I. Gori, M. Svetlik, P. Khante, V. Lifschitz, et al., Bwibots: A platform for bridging the gap between ai and human-robot interaction research, *Int. J. Robot. Res.* 36 (5–7) (2017) 635–659.
- [32] V. Lifschitz, Answer set programming and plan generation, *Artificial Intelligence* 138 (1–2) (2002) 39–54.
- [33] Y.-q. Jiang, S.-q. Zhang, P. Khandelwal, P. Stone, Task planning in robotics: an empirical comparison of PDDL-and ASP-based systems, *Front. Inf. Technol. Electron. Eng.* 20 (3) (2019) 363–373.
- [34] P. Muñoz, M.D. R-Moreno, D.F. Barrero, Unified framework for path-planning and task-planning for autonomous robots, *Robot. Auton. Syst.* 82 (2016) 1–14.
- [35] C. Wong, E. Yang, X.-T. Yan, D. Gu, Optimal path planning based on a multi-tree T-RRT* approach for robotic task planning in continuous cost spaces, in: 2018 12th France-Japan and 10th Europe-Asia Congress on Mechatronics, IEEE, 2018, pp. 242–247.
- [36] Y. Jiang, F. Yang, S. Zhang, P. Stone, Task-motion planning with reinforcement learning for adaptable mobile service robots, in: *IROS*, 2019, pp. 7529–7534.
- [37] S.-Y. Lo, S. Zhang, P. Stone, Petlon: Planning efficiently for task-level-optimal navigation, in: *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, International Foundation for Autonomous Agents and Multiagent Systems, 2018, pp. 220–228.
- [38] T. Bylander, The computational complexity of propositional STRIPS planning, *Artificial Intelligence* 69 (1–2) (1994) 165–204.
- [39] M. Fox, D. Long, PDDL2. 1: An extension to PDDL for expressing temporal planning domains, *J. Artif. Intell. Res.* 20 (2003) 61–124.
- [40] J.-C. Latombe, *Robot Motion Planning*, Kluwer Academic Publishers, 1991.
- [41] S. Pathak, A. Thomas, V. Indelman, A unified framework for data association aware robust belief space planning and perception, *Int. J. Robot. Res.* 37 (2–3) (2018) 287–315, <http://dx.doi.org/10.1177/0278364918759606>.
- [42] R.E. Kalman, A new approach to linear filtering and prediction problems, *Trans. ASME D 82 (Series D)* (1960) 35–45.
- [43] M. Fox, D. Long, Modelling mixed discrete-continuous domains for planning, *J. Artificial Intelligence Res.* 27 (1) (2006) 235–297.
- [44] R.W. Weyhrauch, *Prolegomena to a theory of mechanized formal reasoning*, *Artif. Intell.* 13.
- [45] S. Bernardini, M. Fox, D. Long, C. Piacentini, Boosting search guidance in problems with semantic attachments, in: *International Conference on Automated Planning and Scheduling (ICAPS)*, Pittsburgh, PA, USA, 2017, pp. 29–37.
- [46] A.J. Coles, A.I. Coles, M. Fox, D. Long, Forward-chaining partial-order planning, in: *Twentieth International Conference on Automated Planning and Scheduling*, 2010.
- [47] L.E. Kavraki, P. Svestka, J.-C. Latombe, M.H. Overmars, Probabilistic roadmaps for path planning in high-dimensional configuration spaces, *IEEE Trans. Robot. Autom.* 12 (4) (1996) 566–580.
- [48] S. Prentice, N. Roy, The belief roadmap: Efficient planning in belief space by factoring the covariance, *Int. J. Robot. Res.* 28 (11–12) (2009) 1448–1465.
- [49] E. Olson, Apriltag: A robust and flexible visual fiducial system, in: *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2011, pp. 3400–3407.
- [50] V. Indelman, L. Carlone, F. Dellaert, Planning in the continuous domain: a generalized belief space approach for autonomous navigation in unknown environments, *Int. J. Robot. Res.* 34 (7) (2015) 849–882.
- [51] J. Hoffmann, B. Nebel, The FF planning system: Fast plan generation through heuristic search, *J. Artificial Intelligence Res.* 14 (2001) 253–302.
- [52] S. Karaman, E. Frazzoli, Sampling-based algorithms for optimal motion planning, *Int. J. Robot. Res.* 30 (7) (2011) 846–894.
- [53] S. Thrun, W. Burgard, D. Fox, *Probabilistic Robotics*, MIT press, 2005.
- [54] N. Koenig, A. Howard, Design and use paradigms for gazebo, an open-source multi-robot simulator, in: 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 3, IEEE, 2004, pp. 2149–2154.
- [55] C. Park, J.S. Park, D. Manocha, Fast and bounded probabilistic collision detection for high-DOF trajectory planning in dynamic environments, *IEEE Trans. Autom. Sci. Eng.* 15 (3) (2018) 980–991.
- [56] H. Zhu, J. Alonso-Mora, Chance-constrained collision avoidance for mavs in dynamic environments, *IEEE Robot. Autom. Lett.* 4 (2) (2019) 776–783.
- [57] M. Saha, G. Sánchez-Ante, J.-C. Latombe, Planning multi-goal tours for robot arms, in: 2003 IEEE International Conference on Robotics and Automation (Cat. No. 03CH37422), 3, IEEE, 2003, pp. 3797–3803.
- [58] F. Imeson, S.L. Smith, A language for robot path planning in discrete environments: The tsp with boolean satisfiability constraints, in: 2014 IEEE International Conference on Robotics and Automation (ICRA), IEEE, 2014, pp. 5772–5777.
- [59] S. Alartartsev, S. Stellmacher, F. Ortmeier, Robotic task sequencing problem: A survey, *J. Intell. Robot. Syst.* 80 (2) (2015) 279–298.
- [60] F. Imeson, S.L. Smith, An SMT-based approach to motion planning for multiple robots with complex constraints, *IEEE Trans. Robot.* 35 (3) (2019) 669–684.
- [61] D.L. Applegate, R.E. Bixby, V. Chvatal, W.J. Cook, *The Traveling Salesman Problem: A Computational Study*, Princeton university press, 2006.



Antony Thomas is currently a Ph.D. student at the Department of Informatics, Bioengineering, Robotics, and Systems Engineering (DIBRIS), University of Genoa, Italy. He received his M.Sc. degree in 2017 from the Technion-Israel Institute of Technology. In 2016, he was awarded the Leonard and Diane Sherman Interdisciplinary Graduate School Fellowship for his M.Sc. research. Antony also holds a dual degree (Bachelor of Technology and Master of Technology) in Aerospace Engineering, awarded by the Indian Institute of Technology (IIT) Madras in 2014. His current research

interests include integrated task-motion planning (TMP), belief space planning (BSP), and autonomous navigation.



Prof. Fulvio Mastrogiovanni, male, Ph.D., serves as Chief Research Officer at TESEO from 2019. Fulvio is Associate Professor at the University of Genova, Adjunct Professor at École Centrale de Nantes, and Affiliate Associate Professor at the Centre for Intelligent Robotics, Japan Advanced Institute of Science and Technology (JAIST). Fulvio has been Visiting Professor at the Asian Institute of Technology in Thailand, the Shanghai Jiao Tong University in China, the Karlsruhe Institute of Technology in Germany, Keio University and JAIST in Japan. He served as Project Coordinator,

Principal Investigator and co-Investigator, as well as Key Person in different national (SHELL), bilateral (WearAml, SSHH), FP7 (ROBOSKIN, CloPeMa), H2020 (CARESSES, InDex) and Erasmus+ (EMARO+, J-EMARO) projects. Fulvio served as EU Program Chair for IEEE RO-MAN 2013, IEEE RO-MAN 2015, and IEEE RO-MAN 2018, EU Program Chair for DARS 2014 and EU Program Chair for URAI 2014, Industry/Exhibitions co-Chair for IEEE/RSJ IROS 2016, Workshops/Tutorials co-Chair for IEEE RO-MAN 2016 and 2017, and for ERF 2017. Fulvio has served as Senior Editor of Intelligent Service Robotics (Springer) from 2013 to 2019, and he is Editor of Robotics and Autonomous Systems (Elsevier). Fulvio has published more than 140 papers in international journals and peer-reviewed international conferences, one edited book, and he is the inventor of three

patents. He is part of the Board of Directors of the Italian Association for Artificial Intelligence. His research interests focus on artificial intelligence techniques for Robotics with a special emphasis on human-robot interaction, knowledge-based representation and automatic reasoning. Fulvio received three Best Paper Awards in international conferences.



Marco Baglietto was born in Savona in 1970. He received the "Laurea" degree in Electronic Engineering in 1995 and the Ph.D. degree in Electronic Engineering and Computer Science in 1999, both from the University of Genoa. Since 1999, he has been an Assistant Professor and since 2011 he is Associate Professor of Automatic Control at the Dept. of Informatics, Bioengineering, Robotics and Systems Engineering-DIBRIS (formerly Dept. of Communications, Computer and Systems Science-DIST), University of Genoa. In the past years, he has been charged to hold the courses of

"Automatic Control", "Operations Research 2", "Systems Theory" and "Neural Networks for Classification and Identification" at the University of Genoa (Faculty of Engineering). At present, he is charged to hold the courses of "System Identification" and "Identification and Estimation of Dynamic systems." Prof. Baglietto is author of about 100 scientific papers appeared on international journals or conference proceedings. He is the co-recipient of the 2004 Outstanding Paper Award of the IEEE Transactions on Neural Networks. His research interests include: - State estimation for linear and nonlinear dynamic systems; in particular, in the so-called moving-horizon framework. - State estimation and control of switching systems. - Distinguishability for linear and nonlinear systems. - Information-based control and estimation of stochastic dynamic systems. - Control techniques applied to communication networks. - Neural networks for the solution of optimal control problems, with particular attention to informationally decentralized dynamic systems; Prof. Baglietto is currently an Associate Editor for the "Nonlinear Analysis, Hybrid Systems" and for the IEEE Control Systems Society Conference Editorial Board. "IEEE Trans. on Neural Networks and Learning Systems" and a member of the guest editorial team of the Special Issue of the "IEEE Transactions on Neural Network's on "Adaptive Learning Systems in Communication Networks."