

E-Graphs: Bootstrapping Planning with Experience Graphs

Mike Phillips

mlphilli@andrew.cmu.edu
Carnegie Mellon University

Benjamin Cohen

bcohen@seas.upenn.edu
University of Pennsylvania

Sachin Chitta

sachinc@willowgarage.com
Willow Garage

Maxim Likhachev

maxim@cs.cmu.edu
Carnegie Mellon University

Abstract—Human environments possess a significant amount of underlying structure that is under-utilized in motion planning and mobile manipulation. In domestic environments for example, walls and shelves are static, large objects such as furniture and kitchen appliances most of the time do not move and do not change, and objects are typically placed on a limited number of support surfaces such as tables, countertops or shelves. Motion planning for robots operating in such environments should be able to exploit this structure to improve its performance with each execution of a task. In this paper, we develop an online motion planning approach which learns from its planning episodes (experiences) a graph, an *Experience Graph*. This graph represents the underlying connectivity of the space required for the execution of the mundane tasks performed by the robot. The planner uses the *Experience graph* to accelerate its planning efforts whenever possible. On the theoretical side, we show that planning with Experience graphs is complete and provides bounds on sub-optimality with respect to the graph that represents the original planning problem. On the experimental side, we show in simulations and on a physical robot that our approach is particularly suitable for higher-dimensional motion planning tasks such as planning for single-arm manipulation and two armed mobile manipulation. The approach provides significant speedups over planning from scratch and generates predictable motion plans: motions planned from start positions that are close to each other to goal positions that are also close to each other tend to be similar. In addition, we show how the Experience graphs can incorporate solutions from other approaches such as human demonstrations, providing an easy way of bootstrapping motion planning for complex tasks.

I. INTRODUCTION

Motion planning is essential for robots operating in dynamic human environments. Tasks like picking and placing objects and mobile manipulation of larger boxes require the robot to approach and pick up (or place) objects with minimal collision with the rest of the environment. Fast performance of motion planning algorithms in such tasks is critical, to account for the speed of operation expected by humans and to account for sudden changes in the environment. This is especially true of tasks involving higher-dimensional configuration spaces, e.g. for a two-armed mobile manipulation system (Figure 1).

At the same time, many of mundane manipulation tasks such as picking and placing various objects in a kitchen are highly repetitive. It is therefore expected that robots should be capable of learning and improving their performance with every execution of these repetitive tasks. In particular, robots



Fig. 1. Motion planning is often used to compute motions for repetitive tasks such as dual-arm mobile manipulation in a kitchen.

should be capable of exploiting learned knowledge about the underlying geometric structure in tasks and human environments. While human environments can be very dynamic, e.g. with people walking around, large parts of the environment are still static for significant periods of time. Similarly, tasks tend to have some form of spatial structure, e.g. objects are often found on support surfaces like tables and desks.

This work focuses on learning from experience for motion planning. Our approach relies on a graph-search method for planning that builds an *Experience Graph* online to represent the high-level connectivity of the free space used for the encountered planning tasks. New motion planning requests reuse this graph as much as possible, accelerating the planning process significantly by eliminating the need for searching large portions of the search-space. While previously encountered motion planning problems can speed up the planner dramatically, in their absence, it gracefully falls back to searching the original search-space, adding the newly generated motion to the Experience graph. Planning with Experience graphs is therefore complete. Furthermore, we show that it provides bounds on sub-optimality with respect to the graph that represents the original planning problem. Planning with Experience graphs leads to consistent and predictable solutions for motion plans requested in similar (but not the same) scenarios, e.g. when the goal states of the robot are close to each other. Our approach is particularly useful when the tasks are somewhat repeatable spatially, e.g. in moving a set of dishes off a particular counter into a dishwasher. Although the start and goal states would be different for each motion plan, the general motion of moving a dish from the counter to the washer would essentially be the same each time.

We provide experimental results demonstrating the use of our approach both in simulation and on a real robot. A full-body planner for the PR2 robot was developed for dual-arm mobile manipulation tasks using a search-based planner.

We thank Willow Garage for their support of this work. This research was also sponsored by ARL, under the Robotics CTA program grant W911NF-10-2-0016.

We show how the use of Experience graphs improves the performance of the planner in this high-dimensional state space. We also present results comparing our planner against sampling-based planning approaches. Finally, we show a preliminary application of Experience graphs to learning by demonstration.

II. RELATED WORK

Initial approaches to motion planning focused on planning from scratch, i.e. there was no reuse of the information from previous plans. Recently, there has been more work on reuse of previous information for motion planning, especially in the context of performance optimization for motion planning in realtime dynamic environments. Lien et. al. [12] presented an approach that involved constructing roadmaps for obstacles, storing them in a database, and reusing them during motion planning. Bruce et. al. [3] extended the traditional RRT approach to reuse cached plans and bias the search towards waypoints from older plans. Extensions of this approach can be found in [20, 6].

In [15], an evolutionary algorithm approach was used to bias RRT search for replanning in dynamic environments towards the edges of the explored areas, intended to reduce the time spent on searching parts of the space that have already been explored. In [20], workspace probability distributions were automatically learned for certain classes of motion planning problems. The distributions attempted to capture a locally-optimal weighting of workspace features.

Trajectory libraries have seen use for adapting policies for new situations [18], especially for control of underactuated systems and high-dimensional systems. In [1], new trajectories in a state space were generated by combining nearby trajectories, including information about local value function estimates at waypoints. In [14], a trajectory library was used in combination with an optimal control method for generating a balance controller for a two link robot. Transfer of policies across tasks, where policies designed for a particular task were adapted and reused in a new task, were discussed in [19].

A learning based approach to reuse information from previous motion plans, the environment, and the types of obstacles was presented in Jetchev [7]. Here, a high-dimensional feature vector was used to capture information about the proximity of the robot to obstacles. After a dimensionality reduction in the feature space, several learning methods including nearest neighbor, locally weighted regression and clustering are used to predict a good path from the database in a new situation.

The work in [8] is most similar to our work and involves the use of a database of older motion plans. The approach uses a bi-directional RRT and tries to draw the search towards a path which is most similar to the new motion planning problem (based on distances to the start, goal and obstacles). Some more recent work [2] also attempts to repair paths from a database of past paths using sampling-based planners. It's future work to compare against this approach.

The use of a database of motion plans is a key feature of our approach. However, we differ from other approaches

in that we attempt to use all the information from previous searches instead of attempting to pick the most similar or best path. Our approach can reuse parts of the Experience Graph even when the start and goal states change. Contrary to other approaches, our approach also gracefully degenerates to planning from scratch (with no reused information) to deal with new scenarios which might be completely different from the ones in the database. Our approach also does not rely on object or shape recognition and is thus agnostic to the representation of the environment, e.g. as a voxel grid or individual objects represented using meshes. Although our approach is also comparable to Probabilistic Roadmaps [11], a crucial difference is that Experience Graphs are generated from task-based requests instead of sampling the whole space. We are also able to provide a bound on the quality of the returned solution, with respect to the discretization of the action and state spaces. Finally, our approach tends to generate consistent solutions as we confirm in our experimental analysis.

III. ALGORITHM

A. Overview

An *Experience Graph* or E-Graph is a graph formed from the solutions found by the planner for previous planning queries or from demonstrations. We will abbreviate this graph as G^E . The graph G^E is incomparably smaller than graph G used to represent the original planning problem. At the same time, it is representative of the connectivity of the space exercised by the previously found motions. The key idea of planning with G^E is to bias the search efforts, using a specially constructed heuristic function, towards finding a way to get onto the graph G^E and to remain searching G^E rather than G as much as possible. This avoids exploring large portions of the original graph G . In the following we explain how to do this in a way that guarantees completeness and bounds on solution quality with respect to the original graph G .

B. Definitions and Assumptions

First, we will list some definitions and notations that will help explain our algorithm. We assume the problem is represented as a graph where a start and goal state are provided (s_{start}, s_{goal}) and the desired output is a path (sequence of edges) that connect the start to the goal.

- $G(V^G, E^G)$ is a graph modeling the original motion planning problem, where V^G is the set of vertices and E^G is the set of edges connecting pairs of vertices in V^G .
- $G^E(V^E, E^E)$ is the E-Graph that our algorithm builds over time ($G^E \subseteq G$).
- $c(u, v)$ is the cost of the edge from vertex u to vertex v
- $c^E(u, v)$ is the cost of the edge from vertex u to vertex v in graph G^E and is always equal to $c(u, v)$

Edge costs in the graph are allowed to change over time (including edges being removed and added which happens when new obstacles appear or old obstacles disappear). The more static the graph is, the more benefits our algorithm provides. The algorithm is based on heuristic search and is

therefore assumed to take in a heuristic function $h^G(u, v)$ estimating the cost from u to v ($u, v \in V^G$). We assume $h^G(u, v)$ is admissible and consistent for any pair of states $u, v \in V^G$. An admissible heuristic never overestimates the minimum cost from any state to any other state. A consistent heuristic $h^G(u, v)$ is one that satisfies the triangle inequality, $h^G(u, v) \leq c(u, s) + h^G(s, v)$ and $h^G(u, u) = 0, \forall u, v, s \in V^G$ and $\forall(u, s) \in E^G$.

C. Algorithm Detail

The planner maintains two graphs, G and G^E . At the high-level, every time the planner receives a new planning request the *findPath* function is called. It first updates G^E to account for edge cost changes and perform some precomputations. Then it calls the *computePath* function, which produces a path π . This path is then added to G^E . The *updateEGraph* function works by updating any edge costs in G^E that have changed. If any edges are invalid (e.g. they are now blocked by obstacles) they are put into a disabled list. Conversely, if an edge in the disabled list now has finite cost it is re-enabled. At this point, the graph G^E should only contain finite edges. A *precomputeShortcuts* function is then called which can be used to compute shortcut edges before the search begins. Ways to compute shortcuts are discussed in Section V. Finally, our heuristic h^E , which encourages path reuse, is computed[†].

findPath(s_{start}, s_{goal})

- 1: *updateEGraph*(s_{goal})
 - 2: $\pi = \text{computePath}(s_{start}, s_{goal})$
 - 3: $G^E = G^E \cup \pi$
-

updateEGraph(s_{goal})

- 1: *updateChangedCosts*()
 - 2: disable edges that are now invalid
 - 3: re-enable disabled edges that are now valid
 - 4: *precomputeShortcuts*()
 - 5: compute heuristic h^E according to Equation 1
-

Our algorithm’s speed-up comes from being able to reuse parts of old paths and avoid searching large portions of graph G . To accomplish this we introduce a heuristic which intelligently guides the search toward G^E when it looks like following parts of paths in G^E will help the search get close to the goal. We define a new heuristic h^E in terms of the given heuristic h^G and edges in G^E for any state $s_0 \in V^G$.

$$h^E(s_0) = \min_{\pi} \sum_{i=0}^{N-1} \min\{\varepsilon^E h^G(s_i, s_{i+1}), c^E(s_i, s_{i+1})\} \quad (1)$$

where π is a path $\langle s_0 \dots s_{N-1} \rangle$ and $s_{N-1} = s_{goal}$ and ε^E is a scalar ≥ 1 .

Equation 1 returns the cost of the minimal path from the the queried state s_0 to the goal where the path consists of

[†] This can be done with one search prior to planning or can be done on the fly for states generated by the planner. The first option is feasible whenever the heuristic is computed for lower-dimensional projections of the original problem, such as 3D space for the end-effector of a high-dimensional arm.




Fig. 2. Effect of ε^E . The dark solid lines are paths in G^E while the dark dotted lines are best paths from s_0 to s_{goal} according to h^E . Note as ε^E increases, the heuristic prefers to travel on G^E . The light gray circles and lines show the graph G and the filled in gray circles show the expanded states when planning with E-Graphs. The dark gray arrow shows the returned path.

an arbitrary number of two kinds of segments. The first type of segment corresponds to an instantaneous jump between s_i and s_{i+1} at a cost equal to the original heuristic inflated by ε^E (this is equivalent to saying all states can reach all of the other states according to the original heuristic but inflated by ε^E). The second kind of segment is an edge from G^E and it uses its actual cost. As the penalty term ε^E increases, the heuristic path from s_0 to the goal will go farther out of its way to travel toward the goal using E-Graph edges.

The larger ε^E is, the more the actual search avoids exploring G and focuses on traveling on paths in G^E . Figure 2 demonstrates how this works. As ε^E increases, the heuristic guides the search to expand states along parts of G^E as shown in Figure 2. In Figure 2a, the heuristic causes the search to ignore the graph G^E because without inflating h^G at all ($\varepsilon^E = 1$), the heuristic will never favor following edges in G^E . This figure also shows how during the search, by following G^E paths, we can avoid obstacles and have far fewer expansions. The expanded states are shown as filled in gray circles, which change based on ε^E .

The *computePath* function (shown below) runs weighted A* without re-expansions [13]. Weighted A* uses a parameter $\varepsilon > 1$ to inflate the heuristic used by A*. The solution cost is guaranteed to be no worse than ε times the cost of the optimal solution and in practice it runs dramatically faster than A*. The main modification to Weighted A*, is that in addition to using the edges that G already provides (*getSuccessors*), we add two additional types of successors: *shortcuts* and *snap motions* (line 9). The only other change is that instead of using the heuristic h^G , we use our new heuristic h^E (lines 4 and 16).

Shortcut successors are generated when expanding a state $s \in G^E$. A shortcut successor uses G^E to jump to a place much closer to s_{goal} (closer according to the heuristic h^G). This shortcut may use many edges from various previous paths. The shortcuts allow the planner to quickly get near the goal without having to re-generate paths in G^E . Possible shortcuts are discussed in Section V.

Finally, for environments that can support it, we introduce *snap motions*. Sometimes, the heuristic may lead the search to a local minimum at the “closest” point to G^E with respect to the heuristic, h^G , but it may not be a state on G^E . For example, in (x, y, θ) navigation, a 2D (x, y) heuristic will create a minimum for 2 states with the same x,y but different

```

computePath( $s_{start}, s_{goal}$ )
1:  $OPEN = \emptyset$ 
2:  $CLOSED = \emptyset$ 
3:  $g(s_{start}) = 0$ 
4:  $f(s_{start}) = \varepsilon h^E(s_{start})$ 
5: insert  $s_{start}$  into  $OPEN$  with  $f(s_{start})$ 
6: while  $s_{goal}$  is not expanded do
7:   remove  $s$  with the smallest  $f$ -value from  $OPEN$ 
8:   insert  $s$  in  $CLOSED$ 
9:    $S = getSuccessors(s) \cup shortcuts(s) \cup snap(s)$ 
10:  for all  $s' \in S$  do
11:    if  $s'$  was not visited before then
12:       $f(s') = g(s') = \infty$ 
13:    end if
14:    if  $g(s') > g(s) + c(s, s')$  and  $s' \notin CLOSED$  then
15:       $g(s') = g(s) + c(s, s')$ 
16:       $f(s') = g(s') + \varepsilon h^E(s')$ 
17:      insert  $s'$  into  $OPEN$  with  $f(s')$ 
18:    end if
19:  end for
20: end while

```

θ . A problem then arises because there isn't a useful heuristic gradient to follow, and therefore, many states will be expanded blindly. We borrow the idea of *adaptive motion primitives* [4] to generate a new action which can snap to a state on G^E whenever states s_i, s_j have $h^G(s_i, s_j) = 0$ and $s_j \in G^E$ and $s_i \notin G^E$. The action is only used if it is valid with respect to the current planning problem (e.g. doesn't collide with obstacles). As with any other action, it has a cost that is taken into account during the search.

IV. THEORETICAL ANALYSIS

Our planner provides a guarantee on completeness with respect to G (the original graph representation of the problem).

Theorem 1. *For a finite graph G , our planner terminates and finds a path in G that connects s_{start} to s_{goal} if one exists.*

Since no edges are removed from the graph (we only add) and we are searching the graph with Weighted A* (a complete planner), if a solution exists on the original graph, our algorithm will find it.

Our planner provides a bound on the sub-optimality of the solution cost. The proof for this bound depends on our heuristic function h^E being ε^E -consistent.

Lemma 1. *If the original heuristic function $h^G(u, v)$ is consistent, then the heuristic function h^E is ε^E -consistent.*

From Equation 1, for any $s, s' \in V^G, (s, s') \in E^G$.

$$h^E(s) \leq \min\{\varepsilon^E h^G(s, s'), c^E(s, s')\} + h^E(s')$$

$$h^E(s) \leq \varepsilon^E h^G(s, s') + h^E(s')$$

$$h^E(s) \leq \varepsilon^E c(s, s') + h^E(s')$$

The argument for the first line comes from Equation 1 by contradiction. Suppose the line is not true. Then, during the computation of $h^E(s)$, a shorter path π could have been found by traveling to s' and connecting to s with $\min\{\varepsilon^E h^G(s, s'), c^E(s, s')\}$. The last step follows from h^G being admissible. Therefore, h^E is ε^E -consistent.

Theorem 2. *For a finite graph G , the planner terminates, and the solution it returns is guaranteed to be no worse than $\varepsilon \cdot \varepsilon^E$ times the optimal solution cost in graph G .*

Consider $h'(s) = h^E(s)/\varepsilon^E$. $h'(s)$ is clearly consistent. Then, $\varepsilon h^E(s) = \varepsilon \cdot \varepsilon^E h'(s)$. The proof that $\varepsilon \cdot \varepsilon^E h'(s)$ leads to Weighted A* (without re-expansions) returning paths bounded by $\varepsilon \cdot \varepsilon^E$ times the optimal solution cost follows from [13].

V. IMPLEMENTATION DETAIL

In this section we discuss how various parts of the algorithm could be implemented.

A. Heuristic

Some heuristics h^G are derived using dynamic programming in a lower dimensional state space, such as a 2D Dijkstra search for an (x, y, θ) navigation problem. Alternatively, some heuristics h^G can be computed in $O(1)$ upon request (e.g. euclidean distance).

In the first case, we can compute h^E by running a Dijkstra search on the low-dimensional projection of G , with additional edges from G^E connecting the low-dimensional projection of the states in G^E . This can be computed with similar efficiency to the original heuristic h^G , so it doesn't hurt the planning times (this is what we used in our experiments).

In the second case, h^E can be computed by constructing a fully connected graph on V^E and s_{goal} (using both $\varepsilon^E h^G$ and c^E). A Dijkstra search is used to find the heuristic h^E from each of these states to the goal. Then during the search for any state s , $h^E(s) = \min_{s' \in \{V^E \cup s_{goal}\}} (h^G(s, s') + h^E(s', s_{goal}))$.

In our implementation, we compute the heuristic h^E in an on-demand fashion. Our computation runs Dijkstra's algorithm until the heuristic of the requested state is computed and then suspends until another un-computed heuristic is requested.

B. Shortcuts

Shortcuts accelerate the search by allowing the search to bypass retracing an old path (re-expanding the states on it) in G^E . The algorithm works with or without the shortcuts. Basically, the shortcuts are pre-computed edges that connect all states in G^E to a very small set of states in G^E . Shortcut successors can only be generated when expanding a state $s \in G^E$. There are several obvious choices for this subset. For example, it can contain all states s in G^E that are closest to the goal within each connected component of G^E . The closeness can be defined by h^G or h^E . In our experiments we used h^G . Other ways can also be used to compute this subset of states. It is future work to explore these options.

C. Pre-Computations

Some of the computations on G^E can be done before the goal is known. In particular, $shortcuts(s)$ (line 9 of *computePath*) need to know the costs of least-cost paths between pairs of states in G^E . If state $u \in G^E$ is being expanded and has a shortcut to the state v on the same component in G^E then we need to assign an edge cost $c(u, v)$. In order to do that we need to know the cost of a least-cost path on G^E from u to v .

These costs can be computed before knowing the goal by using an all-pairs shortest path algorithm like Floyd-Warshall. This can be done in a separate thread between planning queries (as well as adding the path from the previous query into G^E). To make Floyd-Warshall run faster and to save memory, we can also exploit the fact that most of the paths in G^E don't intersect each other in many places. We can therefore compress it into a much smaller graph containing only vertices of degree $\neq 2$ and run Floyd-Warshall on it. Then, the cost of a path between any pair $x, y \in G^E$ is given by $\min_{x_i, y_i} \{c(x, x_i) + c(x_i, y_i) + c(y_i, y)\}$. Where $x_i \in \{x_1, x_2\}$ and $y_i \in \{y_1, y_2\}$. x_1 and x_2 are states with degree $\neq 2$ that contain the path on which x resides. y_1 and y_2 are defined similarly.


VI. EXPERIMENTAL RESULTS

A variety of experiments were run in multiple domains to verify the performance gains of our algorithm. We compared our approach against Weighted A* without using G^E as well as with a variety of sampling-based planners. The test domains include planning for a 7 degree of freedom arm as well as full-body planning for the PR2 robot (two arms, the telescoping spine and the navigation of the base).

A. Full Body Planning

Planning in higher-dimensional spaces can be challenging for search-based planners. A full-body planning scenario for a PR2 robot is thus a good test of the capabilities developed in this work. Our test involves the PR2 carrying objects in a large environment. We restrict the objects to be upright in orientation, a constraint that often presents itself in real-world tasks like carrying large objects, trays, or liquid containers. We assume that the two end-effectors are rigidly attached to the object. The state space is 10 dimensional. Four define the position and yaw of the object in a frame attached to the robot. The planner operates directly in this workspace (instead of joint space) and uses inverse kinematics to map movements back onto the joint angles of both arms. To restrict the number of inverse kinematics solutions to one, the redundant degree of freedom in each arm (upper arm roll joint) are part of our state. The robot's position and yaw in the map frame and the height of the telescoping spine are the last four dimensions.

A 3D Dijkstra heuristic is used to plan (backwards) for a sphere inscribed in the carried object from its goal position to the start position (this is the low-dimensional projection for the heuristic). The heuristic is useful in that it accounts for collisions between the object and obstacles. However, it does not handle the complex kinematic constraints on the motion of the object due to the arms, spine, and base and does not account for collisions between the body of the robot and the environment. In all experiments, $\varepsilon = 2$ and $\varepsilon^E = 10$ resulting in a sub-optimality bound of 20. We chose these values for the parameters (manually) as they provided a good combination of speed-up and sub-optimality bound. In future work, we will look into ways to automatically reduce the sub-optimality bound as planning time allows. Setting ε^E to 10 greatly encourages the search to go to G^E if it looks like



(a) Bootstrap goals (b) One of the test sets

Fig. 3. Full-body planning in a warehouse

TABLE I
WAREHOUSE ENVIRONMENT: E-GRAPH PLANNING TIME

successes(of 500)	mean time(s)	std dev(s)	max(s)
500	0.30	0.43	3.37

following it can get it closer to the goal. Setting ε to 2 inflates the whole heuristic including the part of the h^E using paths in G^E . This encourages the planner to use shortcuts as soon as they become available, preventing it from re-expanding an old path.

The results were compared against regular Weighted A* with $\varepsilon = 20$ so that both approaches would have the same sub-optimality bound. We also compared against several sampling techniques from the OMPL (Open Motion Planning Library), including RRT-Connect (a bi-directional RRT), PRM, and RRT* [5, 9, 11, 10]. Since the PRM is a multi-query approach, we allowed it to keep its accumulated roadmap between queries like our method does. Since RRT* is an anytime method that improves as more time is allowed we provide results for the first solution found as well as the solution after our timeout (we gave all planners up to 2 minutes to plan). We post-process all sampling planner paths using OMPL's short-cutter (E-Graph and Weighted A* paths have no post-processing).

1) *Warehouse scenario*: Our first scenario is modeled on an industrial warehouse where the robot must pick up objects off a pallet and move them onto a shelving unit (Figure 3). The goals alternate between pallet and shelves. Since our planner improves in performance with repeated attempts in a particular spatial region of space, we first bootstrap it with 45 uniformly distributed goals (split between the pallet and the shelves). The bootstrap goals and the resultant G^E (projected onto 3D end effector position) after processing them are shown in Figure 3a. We bootstrapped the PRM with the same 45 goals.

A set of 100 random goals (with varying positions and yaws of the object) alternating between the pallet and the shelves were then specified to the planner. This entire process was repeated 5 times with G^E cleared before running each new set of 100 goals. The planning times for E-Graphs for all 500 random goals are shown in Table I. On average, 94% of the edges on a path produced by our planner were recycled from G^E . The mean time to update G^E between queries (add the new path and compute Floyd-Warshall) was 0.31 seconds.

In Table II we compare E-Graphs to several other methods. We show the average speed-up by computing the average E-

TABLE II
WAREHOUSE ENVIRONMENT: PLANNING TIME COMPARISON

method	successes(of 500)	mean speed-up	std dev speed-up	max speed-up
Weighted A*	497	15.80	54.05	625.54
RRT-Connect	500	0.59	0.42	3.08
PRM	500	2.16	14.79	226.51
RRT* (first solution)	440	11.90	45.57	594.45

TABLE III
WAREHOUSE ENVIRONMENT: PATH QUALITY COMPARISON

method	object XYZ path length ratio	std dev ratio	base XY path length ratio	std dev ratio
Weighted A*	0.70	0.15	1.36	2.89
RRT-Connect	0.94	0.59	1.55	3.31
PRM	0.78	0.33	0.64	0.39
RRT* (first solution)	0.86	0.37	1.47	3.46
RRT* (final solution)	0.87	0.43	1.41	2.84

Graph to other method ratio (across the cases where both methods found a solution to the query). We can see that generally E-Graphs are significantly faster than other methods (except for bi-directional RRT). In Table III we can see results comparing path quality on two metrics: the length of the path the carried object travels and the length of the path the base travels. The results in this table are also ratios and we can see that all methods produce shorter paths than E-Graphs. This is expected from Weighted A* since E-Graph paths sometimes go out of their way a bit to find solutions quicker. Sampling planners generally have poor path quality but the shortcutting post-process step works well on this relatively simple experiment where most paths are very close to straight lines in configuration space. Our next experiment is a more difficult kitchen scenario where the more complicated paths cause shortcutting to be less effective.


2) *Kitchen Environment*: A second set of tests was run in a simulated kitchen environment. 50 goals were chosen in locations where object are often found (e.g. tables, countertops, cabinets, refrigerator, dishwasher). 10 representative goals were chosen to bootstrap our planner, which was then tested against the remaining 40 goals. Figure 4 shows G^E both after bootstrapping and after all the goals have been processed.

Table IV shows the planning times for E-Graphs. On average, 80% of the edges on a path produced by our planner were recycled from G^E . The mean time to update G^E between queries (add the new path and compute Floyd-Warshall) was 0.34 seconds. In Table V we can see that E-Graphs provide a significant speed-up over all the other methods, generally over 30 times (though only a factor of 2 for bi-directional RRT). The PRM and RRT* also failed to solve many of the queries within the 2 minute limit.


In Table VI we can see that E-Graphs provide significantly shorter paths than bi-directional RRT (Figure 5 shows an example). While RRT* path quality does improve with time, after 2 minutes, it still is no better than E-Graphs. The PRM still does better than E-Graphs in the kitchen scenario,

TABLE IV
KITCHEN ENVIRONMENT: E-GRAPH PLANNING TIME

successes(of 40)	mean time(s)	std dev(s)	max(s)
40	0.33	0.25	1.00



(a) G^E after bootstrap goals (b) G^E after bootstrap goals



(c) G^E after test goals (d) G^E after test goals


Fig. 4. Full-body planning in a kitchen scenario

TABLE V
KITCHEN ENVIRONMENT: PLANNING TIME COMPARISON

method	successes(of 40)	mean speed-up	std dev speed-up	max speed-up
Weighted A*	37	34.62	87.74	506.78
RRT-Connect	40	1.97	2.35	11.32
PRM	25	16.52	74.25	372.90
RRT* (first solution)	23	50.99	141.35	613.54

however, PRMs only solved 25 of the 40 cases and therefore is only compared against E-Graphs on easier queries where shortcutting still works well.

We ran an additional experiment in the kitchen scenario to show the consistency of E-Graph solutions. Consistency measures how similar output of a planner is, given similar inputs (start and goal). In many domains, this kind of path



(a) RRT-Connect (b) E-Graph

Fig. 5. An example comparing an RRT-Connect path to an E-Graph path (path waypoints move from black to white). The E-Graph path is shorter and more goal directed.

TABLE VI
KITCHEN ENVIRONMENT: PATH QUALITY COMPARISON

method	object XYZ path length ratio	std dev ratio	base XY path length ratio	std dev ratio
Weighted A*	0.91	0.68	1.14	1.40
RRT-Connect	2.54	4.67	3.45	9.67
PRM	0.85	0.32	0.88	0.48
RRT* (first solution)	1.08	0.60	1.39	1.79
RRT* (final solution)	1.03	0.48	1.36	1.96




Fig. 6. Consistency experiment in a kitchen scenario. 58 goals were split into two groups and the various methods were asked to plan between them.

TABLE VII
KITCHEN ENVIRONMENT: PATH CONSISTENCY COMPARISON


method	dynamic time warping similarity
E-Graphs	407
RRT-Connect	1512
PRM	748
RRT* (first solution)	1432
RRT* (final solution)	1034

predictability is critical for people to be comfortable around robots. We tested this by placing two groups of goals in the kitchen environment as shown in Figure 6. There are 58 goals split evenly between a cabinet and a refrigerator shelf. We ran through the goals alternating between the two groups so that while no two queries have the same start and goal, the paths should be mostly the same for each. The results of the experiment are shown in Table VII. We used the dynamic time warping similarity metric [16] to compare the methods. Having a value closer to 0 means the paths are more similar. Since this method is for comparing pairs of paths, we performed an all-pairs comparison within each group and then took the average path similarity. We can see that E-Graphs by far have the most consistent paths due to the deterministic nature of the planner and reuse of previous path segments.

B. Learning by Demonstration


The high dimensionality of the full-body domain makes it easy for the planner to get stuck in large local minima. This is especially true when the heuristic is misleading, such as in Figure 7a. After picking up an object at the pallet and given a goal in the shelf, the heuristic guides the search around the right side of the pole where the robot’s body won’t fit through. The search without Experience Graphs then takes over 30 minutes to solve the problem.

However, if a path were demonstrated, and then added to G^E , our approach could harness this information to reach all the desired goals. We demonstrated such a path through teleoperation in the simulated world (Figure 7a). We added this demonstrated path to G^E by having each recorded waypoint be a vertex with a cost represented by the cost function used in our approach. After the demonstration, our approach had an average planning time of 1.08 seconds to all 12 goals as shown in Figure 7b. While learning by demonstration is not the focus of this work, these preliminary results show it as a promising application. It is also interesting to note that any path can be demonstrated regardless of how sub-optimal the



(a) A demonstrated path (b) G^E after 12 goals

Fig. 7. Learning by demonstration in a more difficult warehouse scenario



(a) Grasping pipeline setup (b) G^E partway through the experiment

Fig. 8. Tabletop manipulation experiments

quality is. Our planner will use parts of it if possible and return a solution with a bound on the resulting solution quality.

C. Single Arm Planning

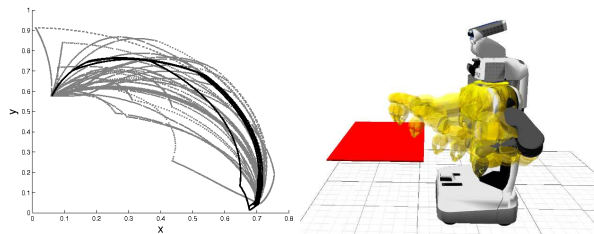
The planner’s performance was also tested for tabletop manipulation using the PR2 robot (Figure 8a). A search-based planner [4] generates safe paths for each of the PR2’s 7-DOF arms separately. The goals for the planner are specified as the position and orientation of the end-effector. Our implementation builds on the ROS grasping pipeline to pick up and put down objects on the table in front of the robot. Our approach is used during both the pick and place motions to plan paths for the robot’s arms (individually).

In the experiments, statistics for 411 planning requests were recorded using Weighted A*, our approach, and a randomized planner (SBL [17]). The results are shown in Table VIII and Figure 8b shows G^E part-way through the experiment. We set $\varepsilon = 2$ and $\varepsilon^E = 50$ for a sub-optimality bound of 100. We ran the regular Weighted A* planner with $\varepsilon = 100$.

Table VIII shows that we have a speed increase of about 2.5 over both methods. The heuristic computation time (0.1s) dominates the planning times for both our approach and the Weighted A* approach, resulting in a smaller speedup than expected by the ratio of expansions. On average, 95% of the edges on a path produced by our planner were recycled from

TABLE VIII
RESULTS ON TABLETOP MANIPULATION (411 GOALS)

	mean time(s)	std dev time(s)	mean expands	mean cost
E-Graphs (E)	0.13	0.07	4	117349
Weighted A* (W)	0.26	0.10	145	109297
SBL (S)	0.24	0.09	N/A	N/A
Ratio (W/E)	2.50	1.23	66	1.03
Ratio (S/E)	2.44	1.50	N/A	N/A



(a) 40 end effector trajectories with similar starts and goals. Our approach is in black, while the sampling approach is in gray. (b) One visualized path

Fig. 9. E-Graphs provide similar solutions to similar problems

TABLE IX
LENGTH OF 40 SIMILAR QUERIES IN TABLETOP MANIPULATION

	mean length (m)	std dev length (m)
E-Graphs	1.378	0.012
SBL	1.211	0.178

G^E . The mean time to update G^E (add the new path and compute Floyd-Warshall) was 0.12 seconds.

These results show that our approach is competitive with sampling-based method in terms of planning times. However, our approach provides explicit cost minimization and therefore, some notion of consistency (for similar goals we provide similar paths). Figure 9a shows the paths of the end effector for 40 paths that had similar starts and goals (within a 3x3x3cm cube). The gray paths are from the randomized approach while the black are from our approach. Notice that the randomized approach produces highly varying paths even after shortcutting. On the other hand our approach consistently provides almost the same path each time (we also applied a simple short-cutter to eliminate artifacts from the discretized environment). Table IX shows that our approach has only a slightly longer path length (for the end effector distance) but a significantly lower standard deviation. While our planner’s cost function is actually trying to minimize the change in joint angles, our average end effector path length is still relatively competitive with the sampling-based approach. Figure 9b shows one of these paths visualized in more detail.

VII. CONCLUSION

In this paper we have presented planning with Experience Graphs, a general search-based planning method for reusing parts of previous paths in order to speed up future planning requests. Our approach is able to do this while still providing theoretical guarantees on completeness and a bound on the solution quality with respect to the graph representing the planning problem. The paths our planner uses can be fed back from each planning episode in an online fashion to allow the planner to get better over time. The paths can also be demonstrated by a human (or other planner) to allow our planner to improve (while still providing a bound on solution quality). We provided experiments in the robotics domains of planning for a 7 DoF arm in tabletop manipulation tasks and full-body planning for the PR2 robot performing mobile manipulation tasks in warehouse and kitchen scenarios. Our

comparison against the same search-based planning method without E-Graphs showed a speed-up of 1 to 2 orders of magnitude. Our comparison against a state of the art sampling based approach showed that we are competitive in terms of speed, but we yield more consistent paths.

As future work, we would like to use heuristics to prune G^E as it gets large, as well as taking a deeper look at applications in the field of learning by demonstration. We are also interested in making this into an anytime search so that the solution could approach optimality as more time is allowed.

REFERENCES

- [1] C. Atkeson and J. Morimoto. Nonparametric representation of policies and value functions: A trajectory-based approach. In *Advances in Neural Information Processing Systems (NIPS)*, 2003.
- [2] Dmitry Berenson, Pieter Abbeel, and Ken Goldberg. A robot path planning framework that learns from experience. In *ICRA*, 2012.
- [3] J. Bruce and M. Veloso. Real-time randomized path planning for robot navigation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2002.
- [4] B.J. Cohen, G. Subramanian, S. Chitta, and M. Likhachev. Planning for manipulation with adaptive motion primitives. In *Robotics and Automation (ICRA), 2011 IEEE International Conference on*, 2011.
- [5] Ioan A. Şucan, Mark Moll, and Lydia E. Kavraki. The Open Motion Planning Library. *IEEE Robotics & Automation Magazine*, 2012. URL <http://ompl.kavrakilab.org>. To appear.
- [6] D. Ferguson, N. Kalra, and A. T. Stenz. Replanning with rrt. In *IEEE International Conference on Robotics and Automation*, May 2006.
- [7] Nikolay Jetchev and Marc Toussaint. Trajectory prediction: Learning to map situations to robot trajectories. In *IEEE International Conference on Robotics and Automation*, 2010.
- [8] Xiaoxi Jiang and Marcelo Kallmann. Learning humanoid reaching tasks in dynamic environments. In *IEEE International Conference on Intelligent Robots and Systems*, 2007.
- [9] James J. Kuffner Jr. and Steven M. LaValle. Rrt-connect: An efficient approach to single-query path planning. In *ICRA*, 2000.
- [10] S. Karaman and E. Frazzoli. Incremental sampling-based algorithms for optimal motion planning. In *Robotics: Science and Systems (RSS)*, Zaragoza, Spain, June 2010.
- [11] L. E. Kavraki, P. Svestka, J.-C. Latombe, and M. H. Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation*, 12(4):566–580, 1996.
- [12] J. Lien and Y. Lu. Planning motion in environments with similar obstacles. In *Proceedings of the Robotics, Science and Systems Conference*, 2005.
- [13] M. Likhachev, G. Gordon, and S. Thrun. ARA*: Anytime A* with provable bounds on sub-optimality. In *Advances in Neural Information Processing Systems (NIPS) 16*. Cambridge, MA: MIT Press, 2003.
- [14] C. Liu and C. G. Atkeson. Standing balance control using a trajectory library. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2009.
- [15] S. Martin, S. Wright, and J. Sheppard. Offline and online evolutionary bi-directional rrt algorithms for efficient re-planning in environments with moving obstacles. In *IEEE Conference on Automation Science and Engineering*, 2007.
- [16] Hiroaki Sakoe and Seibi Chiba. Dynamic programming algorithm optimization for spoken word recognition. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 26, 1978.
- [17] Gildardo Sanchez and Jean claude Latombe. A single-query bi-directional probabilistic roadmap planner with lazy collision checking. In *International Symposium on Robotics Research*, 2001.
- [18] M. Stolle and C. Atkeson. Policies based on trajectory libraries. In *IEEE International Conference on Robotics and Automation*, 2006.
- [19] M. Stolle, H. Tappeiner, J. Chestnutt, and C. Atkeson. Transfer of policies based on trajectory libraries. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2007.
- [20] M. Zucker, J. Kuffner, and M. Branicky. Multipartite rrt for rapid replanning in dynamic environments. In *IEEE International Conference on Robotics and Automation*, 2007.