Optimization-based Motion Planning with Human in The Loop for Non-Prehensile Manipulation

Rafael Papallas, Anthony G. Cohn, Mehmet R. Dogar

***Abstract*— We are interested in the reaching through clutter problem where a number of robots are trying to reach for a goal objects from the back of cluttered shelves. We investigate the performance increase that can be achieved by using a human-in-the-loop to guide these robots. The Reaching Through Clutter problems are difficult for fully autonomous planners as they have to search for a solution in a high-dimensional space. Furthermore, physics simulators suffer from the uncertainty problem where a valid trajectory in simulation can be invalid when executing the trajectory in the real-world. We propose an online-replanning method with human-in-the-loop to tackle these problems. This system enables a robot to plan and execute a trajectory autonomously, but also to seek high-level input from a human operator if needed. This method aims to minimize the human effort required, thereby increasing the number of robots that can be guided in parallel by a single human operator.**

I. INTRODUCTION

Consider the scenario where 6 robots are working in a warehouse, reaching for items on cluttered shelves in parallel with a single human operator providing guidance when required. Initially, all 6 robots try to generate a plan, using trajectory optimization. Some of the robots quickly generate a feasible trajectory and start autonomous execution without requiring any human help. Since there is uncertainty in how objects move, the robots perform online replanning (similar to model predictive control), where they re-optimize and execute the trajectory at each time step. The second robot however, decides to ask for human help early on and prompts the user. The human engages with the second robot, quickly inspects the scene, uses an interface to provide high-level input, and disengages. After a while the second robot tries to generate a trajectory again, this time making use of the human provided input, and then proceeds with autonomous execution using online replanning. In the meantime, after a duration of autonomous execution, the objects in the forth robot's environment move very differently from the planner's expectations, resulting in robot 4 requiring human help. The human is prompted for input, and execution proceeds.

Such a system has certain advantages. One advantage is the availability of human help in *planning* non-prehensile reaching through clutter motions. A variety of autonomous planners have recently been proposed to solve this problem

This research has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Sklodowska-Curie grants agreement No. 746143, and from the UK Engineering and Physical Sciences Research Council under grant EP/N509681/1, EP/P019560/1 and EP/R031193/1, and a Turing Fellowship to the second author.

Authors are with the School of Computing, University of Leeds, United Kingdom {r.papallas, a.c.cohn, m.r.dogar}@leeds.ac.uk

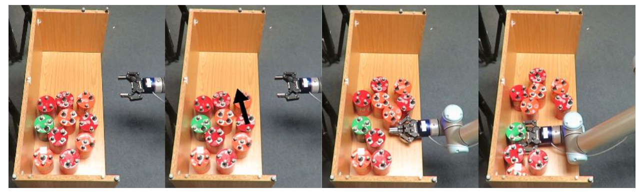

Fig. 1: Robot tries to reach for the goal object (green). Arrow indicates human input.

[1]–[8], though difficult instances still result in low success rates or long planning times in the order of tens of seconds or minutes. In this work, we investigate the performance increase that can be achieved by using a human-in-the-loop providing guidance, particularly by minimizing the human effort required, thereby increasing the number of robots that can be guided in parallel by a single human. Such high-level guidance is usually easy and natural for a human to provide and can dramatically accelerate the performance of planners in difficult scenes.

Another advantage of the system described above is the use of *online replanning*. When a robot executes a non-prehensile plan, objects in the real world move differently to the model's predictions, which makes it necessary to update the plan. Trajectory optimization based planning approaches are particularly effective in such settings, because a new optimization cycle can be warm-started with the previous solution, and convergence can be achieved in few optimization iterations.

Our previous work uses human guidance to solve reaching through clutter problems [9], but that approach requires human input at all times and uses a sampling-based planning approach which we expect to perform poorly in our online-replanning framework.

II. ADAPTIVE OR-HITL FRAMEWORK

We use an optimization-based approach that integrates human input to solve the problem of reaching through clutter. Our system starts tackling the problems fully autonomously and decides to ask for human help only when needed.

Optimization: In line 5 using the OPTIMIZE function we pass to the optimizer the initial state x^{world}, the current trajectory τ and the current cost function, $\mathcal{C}$. The optimizer will optimize for some duration and then return a *result*. The result can either be "human input required" or "success". If the optimizer returns "success", it also updates τ with the new trajectory. If the optimizer decides that human input

Algorithm 1: OR-HITL Framework

1: **function** OR-HITL(τ, $\mathcal{C}$)
2: $x^{world} \leftarrow$ observe current real-world state
3: **if** x^{world} reached the goal **then** ***stop***
4: **do**
5: *result* $\leftarrow$ OPTIMIZE(x^{world}, τ, $\mathcal{C}$)
6: **if** *result* is "human input required" **then**
7: *input* $\leftarrow$ obtain input from human
8: update cost function $\mathcal{C}$ based on *input*
9: update τ based on *input*
10: **while** *result* is not "success"
11: execute $\tau_{[0]}$ in real-world
12: $\tau \leftarrow \tau_{[1,n-1]}$ and expanded with u_{togoal}
13: **return** OR-HITL (τ, $\mathcal{C}$)

is required (we describe how this decision is taken later), then on line 7 we obtain a high-level input from a human operator. This high-level input includes information to update the cost function and to instantiate a new initial trajectory τ. We repeat these steps until *result* is "success" (line 10). Once the optimization is successful, we proceed with the *execution* part of the framework.

Execution: To cope with physics uncertainty when executing a trajectory in the real-world, we propose an Online-Replanning approach. In line 11 we execute the first control of the trajectory in the real-world. We then update our trajectory τ in line 12 to be the remaining τ trajectory, expanded with a control towards the goal (u_{togoal}).

Decide if human help is required: To decide if human help is required we check if the optimizer is stuck at a local minimum. To decide if the optimizer is stuck at a local minimum, we look at the absolute difference between the previous iteration's cost and at the current iteration's cost. If this difference is lower than a threshold for a number of consecutive iterations, then we say that we are stuck at a local minimum and a signal to request human help is sent.

Human-Input: A user's high-level action suggests a particular object o_i to be pushed to particular point on the plane. We formalize this high-level action with the triple (o_i, x_i, y_i), where $o_i \in \mathcal{O}$ is an object, and (x_i, y_i) is a point of on the plane that o_i needs to be pushed to.

III. Results

We evaluate Fixed 5, Fixed 20 (planners time-out every 5 and 20 seconds respectively to ask for human help) and Adaptive (our proposed approach, Sec. II) and we compare them with an Autonomous planner in simulation. For this experiment, we generated 30 scenes for *each* planner (120 different scenes) because we use the same user in this experiment and we would like to avoid the chance that the user memorizes some pattern from the problems. For the Autonomous planner, since there is no learning, we evaluate it over all 120 scenes. These results are presented in Table I.

In a real-world setting we evaluated Adaptive with Online Replanning with Adaptive that is not using Online Replanning (OL-Adaptive). OL-Adaptive simply executes the trajectory in an open-loop manner and checks the state after the trajectory execution. We performed 30 real-world experiments, 15 for each planner in 15 different scenes. The robot was asked to reach for the green object in small shelf among other 9 obstacles (e.g., Fig. 1). The results are presented in Table II.

TABLE I: Simulation Results

	Success (%)	Planning Time (s) 95% CI	Guidance Time (s) 95% CI
Fixed 5	90.0	38.1 ± 15.6	9.6 ± 4.1
Fixed 20	93.3	44.2 ± 13.8	7.0 ± 1.8
Adaptive	**96.6**	$\mathbf{31.0 \pm 12.8}$	$\mathbf{2.5 \pm 0.9}$
Autonomous	74.6	79.8 ± 11.2	–

TABLE II: Real-world results.

Real-world (15 scenes, 30 experiments in total)		
	Adaptive	OL Adaptive
Success	13	8
Optimization Failures	1	1
Execution Failures	1	6

Both experiments suggest that our framework with the human-in-the-loop and online replanning increases the success rate of the robots in these experiments especially in the real-world and reduces the planning time considerably compared to the autonomous planner. Using our adaptive approach the human engagement is minimal (2.5 seconds on average) and this suggests that our framework can be used with multiple robots and a single human operator to guide these robots effectively simultaneously.

References

[1] M. Dogar, K. Hsiao, M. Ciocarlie, and S. Srinivasa, "Physics-based grasp planning through clutter," in *Robotics: Science and Systems*, 2012.

[2] J. A. Haustein, J. King, S. S. Srinivasa, and T. Asfour, "Kinodynamic randomized rearrangement planning via dynamic transitions between statically stable states," in *Robotics and Automation (ICRA), 2015 IEEE International Conference on*. IEEE, 2015, pp. 3075–3082.

[3] M. ud din, M. Moll, L. Kavraki, J. Rosell *et al.*, "Randomized physics-based motion planning for grasping in cluttered and uncertain environments," *IEEE Robotics and Automation Letters*, vol. 3, no. 2, pp. 712–719, 2018.

[4] W. Bejjani, R. Papallas, M. Leonetti, and M. R. Dogar, "Planning with a receding horizon for manipulation in clutter using a learned value function," in *2018 IEEE-RAS 18th International Conference on Humanoid Robots (Humanoids)*. IEEE, 2018, pp. 1–9.

[5] C. Nam, J. Lee, Y. Cho, J. Lee, D. H. Kim, and C. Kim, "Planning for target retrieval using a robotic manipulator in cluttered and occluded environments," *arXiv preprint arXiv:1907.03956*, 2019.

[6] J. E. King, M. Cognetti, and S. S. Srinivasa, "Rearrangement planning using object-centric and robot-centric action spaces," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 3940–3947.

[7] E. Huang, Z. Jia, and M. T. Mason, "Large-scale multi-object rearrangement," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 211–218.

[8] K. Kim, J. Lee, C. Kim, and C. Nam, "Retrieving objects from clutter using a mobile robotic manipulator," in *2019 16th International Conference on Ubiquitous Robots (UR)*. IEEE, 2019, pp. 44–48.

[9] R. Papallas and M. R. Dogar, "Non-prehensile manipulation in clutter with human-in-the-loop," in *2020 International Conference on Robotics and Automation (ICRA)*. IEEE, 2020.