

Socially intelligent task and motion planning for human-robot interaction

Andrea Frank and Laurel D. Riek
Computer Science and Engineering, UC San Diego

I. INTRODUCTION

Robots are entering human social environments (HSEs) such as hospitals, homes, and factories [1]. As social beings, much human behavior is predicated on *social context* - the ambient social state that describes cultural norms, social signals, and individual preferences [2–4]. Cues such as facial expressions, body posture, and group behavior encode social affordances that are critical for understanding and predicting human behavior [5–8]. Furthermore, humans consistently view, communicate with, and react to robots as social actors [9–13]. Many tasks can only be completed by understanding, respecting, and interacting with social context. For example, even in the simple task of buying an ice cream, one must understand why people line up to order, wait in line without cutting, and execute an interactive dialogue before receiving the treat. Thus, in order to act both appropriately and effectively in HSEs, it is crucial for robots to understand social context and how to leverage social interactions to complete tasks.

Endowing robots with social agency—the understanding of oneself as a social actor—is an open challenge that must be addressed before robots can plan safe and reliable behavior in HSEs [1–4, 14]. Socially-aware planning seeks to model social context in order to generate appropriate policies for HSEs [13–16]. Many existing socially-aware planners use proxemics, the social affordance of space [5], to represent social context as an objective function over the environment [17–19]. These planners consider humans as *social obstacles*, i.e. obstacles with an additional social occupancy that encodes proxemic affordances such as personal space [20] and visibility preferences [21]. Social obstacle planners then generate human-aware paths on these costmaps using lightweight graph-search [19], potential field [22], or sampling-based [23] planners. However, these path planners do not consider high-level decisions about subtasks beyond motion and obstacle avoidance, and are therefore insufficient for planning the complex, interactive policies necessary in HSEs.

Joint task and motion planning (TMP) is a central area of research in robotics and artificial intelligence that seeks to address this issue by sharing information between task and motion planners. Many TMP algorithms interleave task and motion planning, replanning at the task level (i.e. reordering or choosing new subtasks) when the high-level symbolic plan violates low-level geometric constraints (e.g. the robot is un-

able to reach the desired object without collision) [24–27]. The majority of this work focuses on manipulation tasks, where the interactions between the robot and the world are simple and well-defined. Few studies have addressed joint TMP in HSEs [28, 29], and these approaches still largely incorporate human models for the purpose of avoiding collisions during task execution. While these methods are a great start to a difficult problem, one key gap is the consideration of social context. Without this, robots are at risk of both distressing those around it by violating social norms and of failing to find the optimal task plan by ignoring possible social interactions.

It is nontrivial to integrate social awareness with modern TMP. Social constraints can be considered relaxed constraints, i.e. constraints that can be only partially satisfied when needed to meet a stronger constraint such as collision avoidance. While some researchers incorporated relaxed constraints into robot planners [30–32], the majority only do so in the context of replanning low-level obstacle avoidance policies to satisfy a static task. The few works that explore relaxed constraints in dynamic TMP either only replan subtask order [33, 34] or consider only simple, well-defined subtasks (e.g. recording information, picking up a cup) [33, 35, 36].

Furthermore, most existing work does not weight the degree of strictness of different constraints. While [35] and [36] propose ways to weight constraints, these methods rely on a human expert to designate their importance. In HSEs, the relative importance of constraints varies with social context, and since social context is also influenced by each robot action, it is inadmissible to assume static, externally-supplied weights. To our knowledge, there exists no joint task and motion planner that is able to reason about complex high-level task replanning with dynamically-varying, relaxed constraints.

II. PROPOSED ALGORITHM

In this work, we propose a socially-aware TMP algorithm that leverages an understanding of social context to generate appropriate and effective policy in HSEs. The key strength of our algorithm is that it explicitly models how potential actions affect not only objective cost (e.g. path length), but transform social context (e.g. interrupting conversations via an alarm). We also inform the planner of the relative importance or urgency of its current task goal, which it uses along with its own calculation of social constraints to determine when it is and is not appropriate to violate social expectations to optimize the objective function [3]. This social awareness allows a robot to understand a fundamental aspect of society: just because

Some research reported in this article is based upon work supported by the National Science Foundation under Grant IIS-1527759.

something makes your job easier does not make it the right thing to do.

We make several assumptions in this work. We assume knowledge of the location of the robot, static obstacles, and people. We further assume that there is a finite, known set of available interactions the robot can initiate with a person or group of people. We do not assume perfect knowledge of social context, and instead estimate its state from observing human activity and social cues. From our previous work in real-world hospitals and factories, we found it is feasible to employ wearable sensors to accurately track people and equipment and recognize important activities [37, 38].

In this work, we adopt a classic TMP architecture with a high-level task planner and a low-level motion planner. We separate our robot’s social interactions into two domains: active (e.g. asking for help) and passive (e.g. accommodating personal space). We incorporate passive social interactions into our low-level planner in a lightweight social obstacle planner in order to leverage the benefits of existing work. Our task planner models active social interactions based on a partially-observable Markov Decision Process (POMDP). Here, social context consists of hidden states (e.g. each person’s current activity) that are approximated from features such as proxemics [18] and group behavior [39].

To choose relevant features for a specific application, we plan to consult experts in the area, extract features via learning algorithms like neural networks, and perform extensive cross-validation of our selection. We will collect data on social interactions in real-world manufacturing and emergency department settings to accurately model how human-human and human-robot interactions transform social context in these HSEs. We will then develop interaction models for a set of possible human-robot interactions, which our task planner will use to plan optimal high-level policy.

The high complexity of HSEs poses concerns for POMDPs. In order to remain tractable, our algorithm will employ a policy-switching architecture that balances the benefits of sophisticated social reasoning with the speed of a socially-naive motion planner. In this model, an executor manages the cooperation between the complex task-level POMDP and the lightweight social obstacle motion planner. The executor uses a lazy approach that activates the high-level planner only where there is significant potential for active social interaction. Assuming that interaction is only possible in a small subset of states (often those proximal to a human), our algorithm avoids high-level replanning for the majority of states. Furthermore, this constraint limits the effective state space of the POMDP to this subset. We anticipate these strategies will enable our algorithm to remain tractable in real-world HSEs.

III. EVALUATION

To validate our algorithm, we will simulate a crowded emergency department in which a robot must deliver medicine to a patient’s room. We will investigate three scenarios: 1) social interaction is helpful, but unnecessary to complete the task; 2) social interaction is necessary; and 3) social interaction

is helpful, but socially unacceptable (e.g. if it would interrupt a clinician treating a patient).

We will perform our study at a medical simulation center at our institution. We will recruit clinicians and clinical learners to realistically populate the simulation, who will be assigned tasks to simulate various social states. For each trial, we will record the executed plan’s path length, duration, and “social burden” as estimated by a post-trial questionnaire. We will compare our algorithm against two competing state-of-the-art methods: a task-naive social motion planner, such as a pure social obstacle model, and a socially-naive TMP, such as *Ffrob* [26]. We anticipate that our algorithm will generate lower cost plans than the task-naive social obstacle planner by leveraging social interaction. In addition, we expect that the task-naive planner will fail when social interaction is necessary, as in (2) and (3), while our algorithm will find a viable plan. We also expect that our algorithm will guarantee socially acceptable policy, where the socially-naive TMP may fail in (3).

IV. DISCUSSION

We anticipate this work will offer the following contributions: First, introduce a socially intelligent joint task and motion planner that generates appropriate and effective policies for robots in HSEs. Second, we present a framework for dividing social interactions into active and passive domains to limit computational complexity and leverage existing work in social obstacle planning. Third, we present a socially-aware, lazy policy-switching architecture that selectively activates high-level replanners only when necessary in order to remain effective and tractable in complex HSEs.

To our knowledge, the proposed work is the first TMP algorithm that enables a robot to act as an intelligent social agent in HSEs. By explicitly modeling social context while remaining tractable in complex environments, we anticipate our approach will be applicable in a variety of real-world applications. Through this work, robots will be able to understand, respect, and leverage social context to produce acceptable and functional policy in HSEs.

REFERENCES

- [1] L. D. Riek, “Healthcare robotics,” *Communications of the ACM*, 2017.
- [2] L. D. Riek, “The social co-robotics problem space: Six key challenges,” *RSS, Robotics Challenges and Visions*, 2014.
- [3] A. Nigam and L. D. Riek, “Social context perception for mobile robots,” in *IROS*. IEEE, 2015.
- [4] M. F. O’Connor and L. D. Riek, “Detecting social context: A method for social event classification using naturalistic multimodal data,” in *11th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (F&G)*. IEEE, 2015.
- [5] E. T. Hall *et al.*, “Proxemics [and comments and replies],” *Current anthropology*, 1968.
- [6] K. Hogan, *Can’t get through: eight barriers to communication*. Pelican Publishing, 2003.

- [7] M. Moosaei, S. K. Das, D. O. Popa, and L. D. Riek, "Using facially expressive robots to calibrate clinical pain perception," in *HRI*. ACM, 2017.
- [8] A. Taylor, H. R. Lee, A. Kubota, and L. D. Riek, "Coordinating clinical teams: Using robots to empower nurses to stop the line," in *In review*, 2019.
- [9] B. Reeves and C. I. Nass, *The media equation: How people treat computers, television, and new media like real people and places*. Cambridge U. Press, 1996.
- [10] C. Breazeal, "Social interactions in hri: the robot view," *IEEE Trans. Syst., Man, Cybern. A, Syst. Humans, Part C (Applications and Reviews)*, 2004.
- [11] A. Sardar, M. Joosse, A. Weiss, and V. Evers, "Don't stand so close to me: users' attitudinal and behavioral responses to personal space invasion by robots," in *HRI*. ACM, 2012.
- [12] P. Baxter, P. Lightbody, and M. Hanheide, "Robots providing cognitive assistance in shared workspaces," in *Companion of HRI*. ACM, 2018.
- [13] A. G. Cunningham *et al.*, "Mpdm: Multi-policy decision-making from autonomous driving to social robot navigation," in *Control Strategies for Advanced Driver Assistance Systems and Autonomous Driving Functions*. Springer, 2019.
- [14] K. Charalampous, I. Kostavelis, and A. Gasteratos, "Recent trends in social aware robot navigation: A survey," *Robotics and Autonomous Systems*, 2017.
- [15] H. Khambhaita and R. Alami, "Assessing the social criteria for human-robot collaborative navigation: A comparison of human-aware navigation planners," in *2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 2017.
- [16] M. Rizwan, V. Patoglu, and E. Erdem, "Human-robot collaborative assembly planning using hybrid conditional planning," in *Proc. FAIM/ISCA Workshop on Artificial Intelligence for Multimodal HRI*, 2018.
- [17] J. Mumm and B. Mutlu, "Human-robot proxemics: physical and psychological distancing in human-robot interaction," in *HRI*. ACM, 2011.
- [18] J. Rios-Martinez, A. Spalanzani, and C. Laugier, "From proxemics theory to socially-aware navigation: A survey," *International Journal of Social Robotics*, 2015.
- [19] B. Okal and K. O. Arras, "Formalizing normative robot behavior," in *International Conference on Social Robotics*. Springer, 2016.
- [20] A. Bera, T. Randhavane, R. Prinja, and D. Manocha, "Sociosense: Robot navigation amongst pedestrians with social and psychological constraints," in *IROS*. IEEE, 2017.
- [21] R. Kirby, R. Simmons, and J. Forlizzi, "Companion: A constraint-optimizing method for person-acceptable navigation," in *RO-MAN*. IEEE, 2009.
- [22] K.-C. Huang, J.-Y. Li, and L.-C. Fu, "Human-oriented navigation for service providing in home environment," in *SICE*. IEEE, 2010.
- [23] M. Svenstrup, T. Bak, and H. J. Andersen, "Trajectory planning for robots in dynamic human environments," in *IROS*. IEEE, 2010.
- [24] L. P. Kaelbling and T. Lozano-Pérez, "Integrated task and motion planning in belief space," *The International Journal of Robotics Research*, 2013.
- [25] N. T. Dantam, Z. K. Kingston, S. Chaudhuri, and L. E. Kavraki, "Incremental task and motion planning: A constraint-based approach," in *RSS*, 2016.
- [26] C. R. Garrett, T. Lozano-Perez, and L. P. Kaelbling, "Ffrob: Leveraging symbolic planning for efficient task and motion planning," *The International Journal of Robotics Research*, 2018.
- [27] S. Srivastava *et al.*, "Combined task and motion planning through an extensible planner-independent interface layer," in *ICRA*. IEEE, 2014.
- [28] S. Pellegrinelli *et al.*, "A probabilistic approach to workspace sharing for human-robot cooperation in assembly tasks," *CIRP Annals*, 2016.
- [29] S. Pellegrinelli *et al.*, "Motion planning and scheduling for human and industrial-robot collaboration," *CIRP Annals*, 2017.
- [30] T. Kunz and M. Stilman, "Manipulation planning with soft task constraints," in *IROS*. IEEE, 2012.
- [31] M. Stilman, "Global manipulation planning in robot joint space with task constraints," *IEEE Trans. Robot.*, 2010.
- [32] J. Huh, B. Lee, and D. D. Lee, "Constrained sampling-based planning for grasping and manipulation," in *ICRA*. IEEE, 2018.
- [33] M. Lahijanian *et al.*, "Iterative temporal planning in uncertain environments with partial satisfaction guarantees," *IEEE Trans. Robot.*, 2016.
- [34] M. Guo and M. M. Zavlanos, "Probabilistic motion planning under temporal tasks and soft constraints," *IEEE Trans. Autom. Control*, 2018.
- [35] N. M. Ceriani *et al.*, "Reactive task adaptation based on hierarchical constraints classification for safe industrial robots," *IEEE/ASME T Mech*, 2015.
- [36] M. Guo and D. V. Dimarogonas, "Multi-agent plan reconfiguration under local ltl specifications," *The International Journal of Robotics Research*, 2015.
- [37] A. Kubota, T. Iqbal, J. A. Shah, and L. D. Riek, "Activity recognition in manufacturing: The roles of motion capture and semg+inertial wearables in detecting fine vs. gross motion," in *ICRA*. IEEE, 2019.
- [38] A. Frank, A. Kubota, and L. D. Riek, "Wearable activity recognition for robust human-robot teaming in safety-critical environments via hybrid neural networks," in *In review*. IEEE, 2019.
- [39] A. Taylor and L. D. Riek, "Robot perception of human groups in the real world: State of the art," in *2016 AAAI Fall Symposium Series*, 2016.