

Online Multi-Agent Trajectory Generation for Adaptive Navigation Planning

Christoforos I. Mavrogiannis
Sibley School of Mechanical & Aerospace Engineering
Cornell University
cm694@cornell.edu

I. INTRODUCTION

One of the most ambitious visions of robotics for the near future is the smooth integration of robots into human populated spaces, such as crowded pedestrian environments. Naturally, the problem of generating socially compliant, humanlike robot motion that ensures human comfort has attracted significant attention over the past few decades. This problem has been proven to be particularly challenging, mainly due to the lack of formal rules regulating navigation in unstructured environments, the lack of explicit communication among agents and the complexity of the environment.

II. RELATED WORK

In an effort to relax this problem from the aforementioned complications, researchers have been drawing inspiration from human navigation. Several approaches have focused on modeling social rules and imbuing robots with an understanding of them [8, 19, 5]. Others, observing the cooperative nature of human navigation (as highlighted for example by Wolfinger [20]) have proposed planning algorithms that distribute the responsibility for collision avoidance across the navigating agents [9, 18, 11]. Finally, a few works, leveraging the existence of sophisticated mechanisms of implicit communication in humans [6], have focused on the generation of intent-expressive robot behaviors [12, 10, 16] which have been shown to be of particular importance for various areas of human-robot interaction (e.g. [7, 4]).

Although these works have captured different elements of what constitutes competent pedestrian behavior, they make use of specific context assumptions that prevent them from being deployed widely in different environments and under different settings. These assumptions are introduced for example when (1) deciding on a training set, (2) employing techniques that aim at imitating observed human behavior, (3) adopting context-specific models of human behavior, (4) engineering specific classes of robot behavior or (5) ignoring the complex dynamics of interaction among agents.

This work proposes a planning framework for the generation of socially competent robot motion that aims at leveraging the topological, spatiotemporal structure of the multi-agent collision avoidance problem towards achieving a greater generalization across diverse environments and settings and ensuring adaptation to unexpected events such as the appearance of heterogeneous agents or agents with changing intentions. Our

approach combines the processes of prediction and planning into a joint decision making scheme that simultaneously reasons about the behavior of other agents and about the role of the robot in the crowd. Our scheme favors actions that exhibit a consistent inclination towards a mutually acceptable strategy of avoidance (e.g. pass from the right side). This strategy reinforces a consensus among agents regarding a collective, mutually beneficial strategy of avoidance which constitutes the basis of competent pedestrian behavior, in line with the insights of Wolfinger [20].

III. APPROACH

This section introduces the foundational components of our approach and describes our online planning architecture.

A. A Symbolic Representation of Multi-Agent Trajectories

In past work [13, 14], we proposed an abstraction that maps a multi-agent trajectory in the form of Cartesian coordinates into a symbolic representation in the form of a topological braid word [1, 3]. Noticing that agents' navigation strategies over the course of the scene are reflected in the entanglement of their trajectories, the formalism of braids serves as a data structure that enables an artificial agent to enumerate a set of classes of multi-agent behaviors that could emerge. By observing agents' past behaviors, an agent may approach the problem of predicting the future behaviors of multiple agents towards making an informed decision about its own navigation strategy. We have showed how a predictive mechanism of this form may allow an agent to act in a consistently intent-expressive and socially compliant way [13, 15, 14]. However, in some cases, the braid representation may constitute an overspecification that is not required by the structure of the problem. For example, a braid may be prescribing a specific sequence of agent crossings that is not crucial for the prediction problem. Considering such redundant constraints in the prediction problem may result in high computational costs. For this reason, in this work, we employ a different data structure that makes use of the topological invariant of the *Winding Number* to prescribe topological properties to agents' trajectories.

For a pair of agent trajectories $a, b : [0, 1] \rightarrow \mathbb{R}$, the *Winding Number* is defined as:

$$\lambda_{ab} = \frac{1}{2\pi} \int_0^1 d\theta, \quad (1)$$




Fig. 1: Pictorial representation of the Topologically Adaptive Navigation Planning scheme. The robot observes agents’ past trajectories ζ , generates a set of m possible scene evolutions $\Lambda_1, \dots, \Lambda_m$, derives geometric representations of them ξ_1, \dots, ξ_m and picks the next action assigned to it from the trajectory of the lowest cost, ξ^* .

where $\theta(t) = \tan^{-1}(b(t) - a(t))$ is the angle between agents a and b at time $t \in [0, 1]$ and counts the number of times the two agents revolved around each other. A positive winding number indicates a collision avoidance involving two agents passing each other from the right hand side whereas a negative winding number represents collision avoidance by passing from the left hand side.

For a set of n agents navigating on a planar workspace, we may synthesize global specifications for collision avoidance among all agents by specifying a combination of signs for all pairwise winding numbers. We thus represent a global strategy of collision avoidance in the form of a tuple of winding numbers:

$$\Lambda = (\lambda_{12}, \lambda_{13}, \dots, \lambda_{(n-1)n}). \quad (2)$$

Under the assumption that all agents are moving from starting positions towards destinations, i.e., agents are not looping around each other, the exact value of the winding numbers is not important and we only need their signs.

B. From Symbols to Trajectories

In order to make use of this representation in a planning problem, we need a way to transition from the symbolic representation Λ to a trajectory representation ξ . To do so, we employ the method of Berger [2], who make use of Hamiltonian vortex dynamics to grow braided trajectories with topological properties specified in the form of topological invariants. This method allows us to plan a multi-agent trajectory that drives a group of agents from a set of initial conditions to a set of destinations, while satisfying a set of topological specifications, formulated in the form of pairwise winding numbers. This approach constitutes a computationally efficient approach to multi-agent trajectory planning as it allows us to plan the motion of multiple agents by growing them from initial conditions with a rule-based decision making scheme.

C. TANP: Topologically Adaptive Navigation Planning

We propose an online navigation planner, called TANP (Topologically Adaptive Navigation Planning), based on our outlined method for multi-agent trajectory generation. The planner runs in replanning cycles. During each cycle it (1) **predicts** a set of possible scene evolutions in the form of combinations of winding number signs $\Lambda_1, \dots, \Lambda_m$, (2) **computes**

the probability of each of the braids considered given observations of agents’ past trajectories ξ_{past} with a model of form $P(\Lambda|\xi_{past})$, (3) **plans** corresponding geometric representations ξ_1, \dots, ξ_m with the method of sec. III-B for a subset of the most likely of them, (4) **scores** the the generated trajectory representations with respect to a cost comprising trajectory quality measures such as efficiency, acceleration and distance from other agents and (5) executes the first action a^* from the trajectory of the lowest cost. A schematic representation of the proposed planning architecture is depicted in Fig. 1. This architecture is not tied to the selection of the aforementioned quality criteria. Different cost functions could be employed to introduce a variety of costs such as *Legibility* [7, 15] and dimensions of human-awareness [17].

IV. RESULTS AND FUTURE WORK

Simulation results demonstrating the effectiveness of the described inference mechanism [15] and of preliminary versions of the presented planning scheme [13, 14] have been presented. An online user study has also demonstrated the importance of incorporating a topological reasoning for socially aware navigation [16]. The main novelty of the present approach involves (1) the combination of winding numbers for the specification of global topological properties of multi-agent trajectories, (2) the multi-agent trajectory generation from topological specification that makes use of pairwise Hamiltonian vortex dynamics and (3) the online adaptation through the consideration of a diverse set of topologically distinct trajectory predictions. The main benefit of this approach is that it enables an artificial agent to understand the qualitatively distinct options that it has access to in a changing environment. Ongoing work involves extensive testing of the proposed motion planner in simulation in different types of environments and settings and a user study to measure the effects of the planner on the behaviors of human subjects in a controlled lab environment.

ACKNOWLEDGMENT

This material is based upon work supported by the National Science Foundation under Grants IIS-1526035 and IIS-1563705. We are grateful for this support.

REFERENCES

- [1] E. Artin. Theory of braids. *Annals of Mathematics*, 48 (1):pp. 101–126, 1947.
- [2] Mitchell A Berger. Hamiltonian dynamics generated by vassiliev invariants. *Journal of Physics A: Mathematical and General*, 34(7):1363, 2001.
- [3] Joan S. Birman. *Braids Links And Mapping Class Groups*. Princeton University Press, 1975.
- [4] Daniel Carton, Wiktor Olszowy, and Dirk Wollherr. Measuring the effectiveness of readability for mobile robot locomotion. *International Journal of Social Robotics*, 8 (5):721–741, 2016.
- [5] Y. F. Chen, M. Everett, M. Liu, and J. P. How. Socially aware motion planning with deep reinforcement learning. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS '17)*, pages 1343–1350, 2017.
- [6] G. Csibra and G. Gergely. ‘Obsessed with goals’: Functions and mechanisms of teleological interpretation of actions in humans. *Acta Psychologica*, 124(1):60–78, 2007.
- [7] Anca D. Dragan and Siddhartha Srinivasa. Integrating human observer inferences into robot motion planning. *Autonomous Robots*, 37(4):351–368, 2014.
- [8] Beomjoon Kim and Joelle Pineau. Socially adaptive path planning in human environments using inverse reinforcement learning. *International Journal of Social Robotics*, 8(1):51–66, Jan 2016. ISSN 1875-4805.
- [9] Ross A. Knepper and Daniela Rus. Pedestrian-inspired sampling-based multi-robot collision avoidance. In *Proceedings of the International Symposium on Robot and Human Interactive Communication, RO-MAN '12*, pages 94–100, 2012.
- [10] Ross A. Knepper, Christoforos I. Mavrogiannis, Julia Proft, and Claire Liang. Implicit communication in a joint action. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction, HRI '17*, pages 283–292, 2017.
- [11] Henrik Kretzschmar, Markus Spies, Christoph Sprunk, and Wolfram Burgard. Socially compliant mobile robot navigation via inverse reinforcement learning. *The International Journal of Robotics Research*, 35(11):1289–1307, 2016.
- [12] Thibault Kruse, Patrizia Basili, Stefan Glasauer, and Alexandra Kirsch. Legible robot navigation in the proximity of moving humans. In *Proceedings of the 2012 IEEE Workshop on Advanced Robotics and its Social Impacts, ARSO '12*, pages 83–88, 2012.
- [13] Christoforos I. Mavrogiannis and Ross A. Knepper. Decentralized multi-agent navigation planning with braids. In *Proceedings of the 2016 International Workshop on the Algorithmic Foundations of Robotics (WAFR '16)*, 2016.
- [14] Christoforos I. Mavrogiannis and Ross A. Knepper. Multi-agent path topology in support of socially competent navigation planning. *The International Journal of Robotics Research*, 2018.
- [15] Christoforos I. Mavrogiannis, Valts Blukis, and Ross A. Knepper. Socially competent navigation planning by deep learning of multi-agent path topologies. In *Proceedings of the 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS '17)*, pages 6817–6824, 2017.
- [16] Christoforos I. Mavrogiannis, Wil B. Thomason, and Ross A. Knepper. Social momentum: A framework for legible navigation in dynamic multi-agent environments. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction, HRI '18*, pages 361–369, 2018.
- [17] Emrah Akin Sisbot, Luis Felipe Marin-Urias, Rachid Alami, and Thierry Siméon. A human aware mobile robot motion planner. *IEEE Transactions on Robotics*, 23(5):874–883, 2007.
- [18] Peter Trautman, Jeremy Ma, Richard M. Murray, and Andreas Krause. Robot navigation in dense human crowds: Statistical models and experimental studies of human-robot cooperation. *International Journal of Robotics Research*, 34(3):335–356, 2015.
- [19] X. T. Truong and T. D. Ngo. Toward socially aware robot navigation in dynamic and crowded environments: A proactive social motion model. *IEEE Transactions on Automation Science and Engineering*, 14(4):1743–1760, 2017.
- [20] Nicholas H. Wolfinger. Passing Moments: Some Social Dynamics of Pedestrian Interaction. *Journal of Contemporary Ethnography*, 24(3):323–340, 1995.