Decentralized Navigation Planning Using Multi-Agent Trajectory Prediction Governed by Hamiltonian Dynamics

Christoforos I. Mavrogiannis† and Ross A. Knepper‡

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 [5, 8, 20]. Others, observing the cooperative nature of human navigation (as highlighted for example by Wolfinger [21]) have proposed planning algorithms that distribute the responsibility for collision avoidance across the navigating agents [9, 11, 19]. 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 [10, 12, 17] which have been shown to be of particular importance for various areas of human-robot interaction (e.g. [4, 7]).

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 smooth, consistently intent-expressive and adaptive behaviors in multi-agent domains that aims at approaching a greater level of generalization across different environments, settings and types of agents. In order to do so, we leverage the underlying topological structure of multi-agent collision avoidance. We introduce a novel abstraction that maps a multi-agent trajectory in the form of Cartesian coordinates to a word, representing the joint strategy of avoidance that agents followed throughout the course of the scene. This abstraction serves as a data structure containing a specification for the generation of a multi-agent trajectory. We propose a trajectory planner that generates multi-agent trajectories from topological specifications in the form of a word. This planner is used to make online multi-agent trajectory predictions, essentially combining efficiently the processes of motion planning and prediction. Benefits of this approach include adaptation to unexpected events such as the emergence of heterogeneous agents or agents with changing intentions and acceleration of a consensus towards a mutually beneficial strategy of collision avoidance despite the complication of no explicit communication.

III. Approach

In this section, we introduce our model for abstracting a multi-agent trajectory from a cartesian representation to a symbolic representation and a method for generating a multi-agent trajectory that satisfies a topological specification, specified in the form of a symbol. Finally, we describe how we employ this method to design an online, decentralized algorithm for navigation in multi-agent environments.

A. *From Trajectories to Symbols*

In past work [13, 15], 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 a scene are reflected in the entanglement of their trajectories, the formalism of braids serves as a data structure that enables an agent to reason about the emerging joint behavior in a principled fashion. In particular, by observing everyone's past trajectories, an agent may infer the braid describing the future joint behavior of multiple agents towards making an informed decision about its own navigation strategy. We have shown how a predictive mechanism of this form may allow an agent to plan consistently intent-expressive and socially compliant actions towards facilitating the convergence to a consensus over a joint strategy of avoidance [13, 15, 16].

The braid representation may encode any multi-agent navigation scenario with any number of agents. As such, it is a very useful tool for analyzing the topological properties of an execution (see e.g. [17]). However, it is not always a practical tool for online inference, as for any given

†C. Mavrogiannis is with Sibley School of Mechanical & Aerospace Engineering, Cornell University, Ithaca NY, 14853, USA. Email: cm694@cornell.edu

‡R. Knepper is with the Department of Computer Science, Cornell University, Ithaca NY, 14853, USA. Email: rak@cs.cornell.edu

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

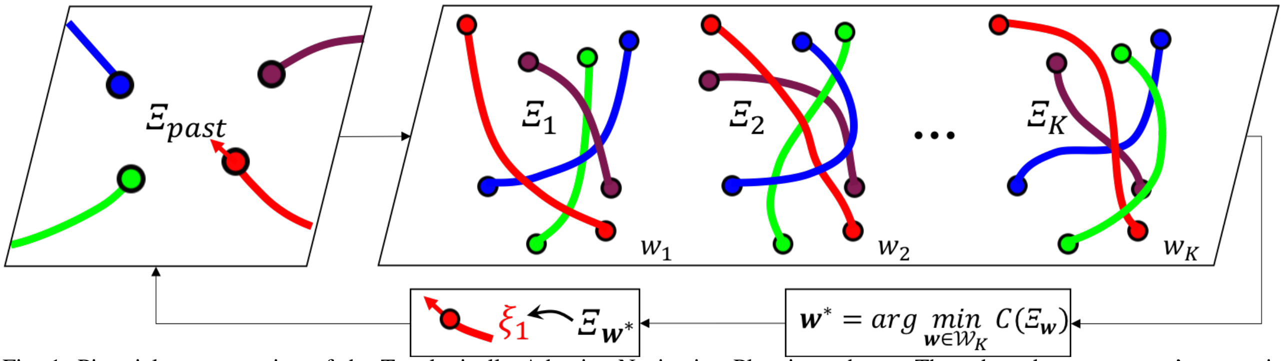

Fig. 1: Pictorial representation of the Topologically Adaptive Navigation Planning scheme. The robot observes agents' past trajectories ζ, generates a set of m possible scene evolutions $w_1, \ldots, w_m$, derives geometric representations of them $\xi_1, \ldots, \xi_m$ and picks the next action assigned to it from the trajectory of the lowest cost, ξ^*.

environment the set of braids is infinite and determining a finite set of likely and realistic candidates is not trivial. Learning techniques may help guide the search but their efficacy is highly dependent on the selected dataset and their applicability may vary depending on the context of the scene in consideration. Motivated by these issues, in this work, we introduce a different data structure that makes use of the topological invariant of the *Winding Number* towards characterizing a multi-agent trajectory with respect to its topological properties.

For a pair of agent trajectories $a, b : [0, 1] \to \mathbb{R}^2$, the *Winding Number* is defined as:

$$w_{ab} = \frac{1}{2\pi} \int_0^1 d\theta, \tag{1}$$

where $\theta(t) = \tan^{-1}(b(t) - a(t))$ is the angle between agents a and b at time $t \in [0, 1]$. This quantity represents the number of times the two agents revolved around each other throughout their motion from $t = 0$ to $t = 1$. For our application, where agents aim at moving efficiently to their destinations while avoiding others, the exact quantity w_{ab} is not important; it is its sign that carries the topological property of the avoidance. 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 from the left hand side.

For a system of n navigating agents, let us collect the pairwise winding numbers of all agents into the *winding* tuple:

$$\boldsymbol{w} = (w_{12}, w_{13}, \ldots). \tag{2}$$

This tuple is a symbolic description of the global, topological properties of a multi-agent trajectory.

B. From Symbols to Trajectories

We propose a computational framework that explicitly leverages the outlined topological structure of multi-agent collision avoidance towards informing the motion planning process of an agent navigating in a multi-agent environment. This design enables an agent to rapidly adapt to unexpected events, such as the emergence of heterogeneous agents or agents with changing intentions.

Our approach is based on a method for transitioning from a symbolic representation $\boldsymbol{w}$ to a trajectory representation ξ. It is inspired by the framework of Berger [2] which makes use of Hamiltonian dynamics as a driving force a set of particles to move along trajectories of desired topological specifications. 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 $\boldsymbol{w}$. 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. At each cycle, it (1) **predicts** a set of candidate windings $\mathcal{W}$, representing a set of distinct scene evolutions, (2) **computes** the probability of these windings given observations of agents' past trajectories ξ_{past} with a model of form $P(\boldsymbol{w}|\xi_{past})$, (3) **plans** corresponding geometric representations $\xi_1, \ldots, \xi_k$ with the method of sec. III-B for the k most likely among them, (4) **scores** the k 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 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, 16] and dimensions of human-awareness [18].

IV. Results

Fig. 2 depicts planning examples from scenarios involving 2, 3 and 4 agents respectively. The agent running our algorithm (red color) observes the past trajectories of others, maps them to a set of intended destinations (represented as landmark pointers of same color) and grows a set of topologically distinct, time-parametrized, collision-free trajectories (overlayed on top of each other, with varying levels of transparency) that drive all agents (including the red one) to their destinations. More details about the proposed algorithm and a more extensive validation will be available at our upcoming paper [14].

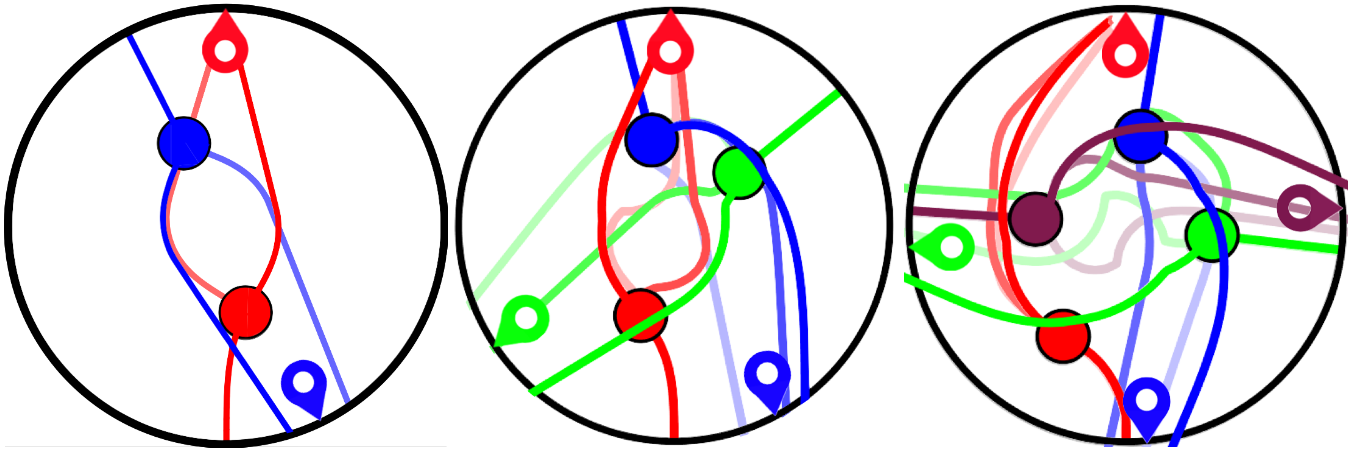

(a) Two agents. (b) Three agents. (c) Four agents.

Fig. 2: Trajectory prediction for examples with different numbers of agents. The TANP agent (red color) is moving towards the red destination. It first makes a coarse prediction about the destinations of others (colored pointers) and then grows a set of qualitatively distinct trajectory predictions, denoted with different color tones.

V. Conclusion

This work contributes: (1) a data structure, built around the topological invariant of the winding number to represent the topological properties of a multi-agent trajectory; (2) a planner that generates global, multi-agent trajectory predictions from symbolic, topological specifications; (3) an online algorithm that makes use of topologically distinct multi-agent trajectory predictions to adapt robustly to the navigation strategies of potentially heterogeneous agents in dynamic environments with no explicit communication. 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 trajectory prediction and generation with topological invariants enforced by hamiltonian dynamics. In *Proceedings of the International Workshop on the Algorithmic Foundations of Robotics*, 2018. [Submitted].

[15] 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.

[16] 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.

[17] 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.

[18] 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.

[19] 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.

[20] 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.

[21] Nicholas H. Wolfinger. Passing Moments: Some Social Dynamics of Pedestrian Interaction. *Journal of Contemporary Ethnography*, 24(3): 323–340, 1995.