Inferring Strategies of Avoidance: Towards Socially Competent Navigation in Human Environments

Christoforos I. Mavrogiannis[†] Valts Blukis[‡] and Ross A. Knepper[‡]
[†] Sibley School of Mechanical & Aerospace Engineering, Cornell University, Ithaca NY
[‡] Department of Computer Science, Cornell University, Ithaca NY
Email: {cm694,vb295}@cornell.edu, rak@cs.cornell.edu

***Abstract*—We present a framework for online navigation planning in multi-agent environments, where no explicit communication takes place among agents, such as pedestrian scenes. Inspired by pedestrian navigation, our approach encodes the concept of coordination into agents' decision making through an inference mechanism about joint strategies of avoidance. Strategies of avoidance represent avoidance protocols that agents engage in to avoid colliding with each other throughout the scene. In this work, we model such strategies topologically, by employing the formalism of braids. This model allows us to characterize the collective behavior of a set of agents but also to enumerate future scene outcomes for a multi-agent scene. Inspired by the mechanisms of human action interpretation, we design an inference mechanism that enables an agent to infer future strategies of avoidance, by observing agents' past behaviors. We integrate this mechanism into the agent's decision making towards generating intent-expressive and socially compliant behaviors that reduce the planning effort for everyone in the scene. Results of statistical significance, generated upon extensive simulation evaluations, indicate faster average uncertainty reduction and faster average destination arrival, compared to purely efficient agents.**

I. Introduction

Over the past decade, the importance of nonverbal communication in human-robot interaction has become increasingly appreciated. Several works have proposed frameworks for intent-expressive motion generation, aiming at achieving increased productivity and a higher degree of integration of robots into human environments [4]. To this end, roboticists have often been following the insights of studies on the mechanisms underlying human inference, which appear to be *teleological* in nature, with humans attributing *goals* to observed *actions* [3].

Human navigation is not an explicitly collaborative activity; however, the insights of studies on pedestrian behavior indicate the existence of context-specific mechanisms of cooperation (e.g. [16]). Inspired by these, we develop a topological model of collective behavior in navigation, based on braids [2], which allows us to enumerate distinct classes of collective navigation strategies of avoidance. This model constitutes the basis for the design of an inference mechanism, according to models of human inference, supported by Csibra and Gergely [3]. This mechanism is incorporated in a policy for intent-expressive motion planning for crowded environments. Our policy leads to the generation of socially competent robot behaviors, by compromising between a cost representing individual efficiency and a cost representing the state of consensus regarding a strategy of avoidance among all agents.

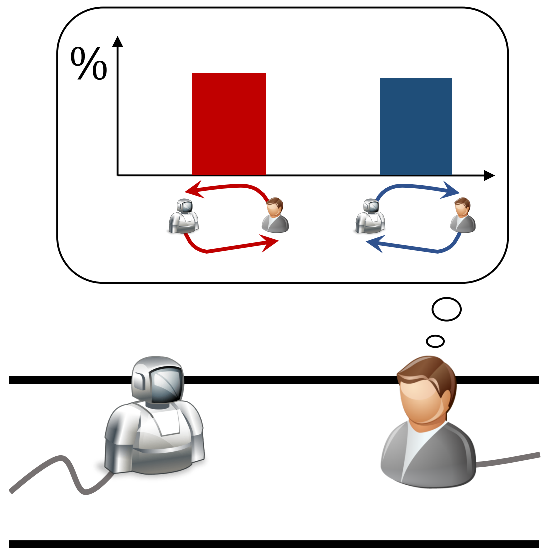

Fig. 1: A human and a robot navigating a hallway. From the perspective of the human, the behavior of the robot so far makes it unclear how the agents are going to avoid each other, yielding a high-entropy probability distribution.

This paper contains key results from Mavrogiannis and Knepper [10], Mavrogiannis et al. [11] but also newer results from a work in progress. Specifically, we study the effect of our policy in multi-agent environments under discrete and continuous settings.

A. Related Work

A significant amount of research has been devoted to the development of models of crowd dynamics that may generate realistic simulations. To this end, the seminal Social Force model of Helbing and Molnár [5] and its variants have simulated pedestrian flows in various environments. Besides, several works have proposed planning frameworks towards generating smooth robot motion in human environments. A class of works directly incorporates measures of human comfort into the motion planning (e.g. Sisbot et al. [13]), whereas other approaches have focused on the design of mechanisms for learning to predict human behaviors and agent interactions (e.g. see [9, 15]). Finally, several works incorporate models of multi-agent interactions and achieve oscillation-free collision avoidance (e.g., [7]).

Our framework differs from the aforementioned frameworks by (1) incorporating a model of collective behavior that allows us to predict classes of trajectories instead of actual trajectories and (2) incorporating a mechanism for intent-expressive decision making in a multi-agent context and (3) incorporating a model of observer uncertainty into agents' decision making.

II. Foundations

We consider the problem of navigation planning in a crowded workspace where multiple agents navigate towards their destinations. The agents do not explicitly exchange any kind of information with each other but are assumed to be acting rationally, which in our context means that (1) they aim at making progress towards their destinations and (2) they have no motive for acting adversarially against other agents (e.g. blocking their paths or colliding with them). Inspired by studies on human behavior, we design a planning framework that enables a planning agent to understand the dynamics of interaction among all agents in a scene, towards acting in a *socially competent* fashion. In the following subsections we provide a definition for social competence in the context of navigation in multi-agent environments and describe our model, that constitutes an abstraction of collective navigation behaviors.

A. Socially Competent Navigation

Based on observations of pedestrian behavior, Wolfinger [16] concluded to a high-level protocol that appears to be the basis of the social order of human navigation, the *Pedestrian Bargain*. The Pedestrian Bargain comprises a set of social rules: (1) *people must behave like competent pedestrians* and (2) *people must trust copresent others to behave like competent pedestrians*.

Although Wolfinger did not explicitly define competence, from the examples included in his work, we may deduce that he refers to a notion of *Social Competence*. The concept of Social Competence has been extensively studied in the field of Psychology from different perspectives and for different scenarios (for an extensive review see [12]). In the context of navigation, we may define social competence as:

> *The ability of an agent to perceive the context*[1], *analyze it and pick an action that appears to be compatible with it, according to a pattern of behavior that the agent assumes observing agents expect from him/her by having observed and analyzed the context themselves.*

According to Csibra and Gergely [3], humans tend to attribute *goals* to observed *actions* in a given context. Therefore, socially competent navigation behaviors should be indicative of agents' intentions and compatible with the context. In other words, socially competent agents should be cognizant of the fact that their behaviors implicitly communicate their *intentions* to any observing agents. The importance of implicit communication for human-robot interaction applications has lately been increasingly appreciated [4, 8].

A problem that arises though, involves providing an appropriate, concrete definition of *intentions* in the context of multi-agent navigation. In our problem setup, we make the assumption of rational agents, aiming at reaching their destinations efficiently. Clearly, the pursuit of destination drives agents' behavior. However, ensuring collision avoidance and comfortable navigation often requires paths of complex shapes, comprising maneuvers and turns. Even in a densely crowded scene, multiple such paths may be possible for an agent. Thus, from the perspective of an observer, inferring *where* an agent is going is not sufficient to predict its future behavior; it is important to be able to infer *how* the agent is going to its destination. We argue that in the context of navigation in a workspace with multiple navigating intelligent agents, a proper model of intentions should incorporate information for both *where* and *how* but also contain a form of structure, i.e., a configuration space over the possible interactions of agents. In this work, we develop such a model of collective intentions, based on the topological formalism of *braids* [2]. In the following subsections, we provide a model of multi-agent navigation, introduce braids and describe how we use them to model collective intentions in navigation.

B. Modeling Multi-Agent Navigation

Consider a set of $n \geq 2$ agents $N = \{1, \ldots, n\}$ navigating a workspace $\mathcal{Q} \subset \mathbb{R}^2$. Denote by $q_i \in \mathcal{Q}$ the configuration of agent $i \in N$. Agent i starts from an initial configuration $q_i^s \in \mathcal{Q}$ at time $t = 0$ and reaches a final configuration q_i^d at time $t = T_i$. The final configuration q_i^d corresponds to a landmark d_i from a set of landmarks $D \subset Q$ (we assume that $d_i \neq d_j$ for any two agents $i, j \in N$). The path that agent i follows to reach its destination is a function $\xi_i : [0, T_i] \to \mathcal{Q}$.

Let us collect the state of the system of all agents in a tuple $Q = (q_1, \ldots, q_n) \in \mathcal{Q}^n$. The system state evolves from a starting configuration $Q^s = (q_1^s, \ldots, q_n^s)$ to a final configuration $Q^d = \left(q_1^d, \ldots, q_n^d\right)$, by following a path $\Xi : [0, T] \to \mathcal{Q}^n$, from the space of system paths $\mathcal{Z}$, starting from Q^s and ending at Q^d. The system path is a function $\Xi : [0, T] \to \mathcal{Q}^n \backslash \Delta$, where $\Delta = \{Q = (q_1, q_2, \ldots, q_n) \in \mathcal{Q}^n : q_i = q_j \text{ for some } i \neq j \in N\}$ is the set of all system states with agents in collision and $T = \max_{i \in N} T_i$ (it is assumed that agents remain at their destinations until everyone reaches their own). Naturally Δ partitions the space of system paths $\mathcal{Z}$ into a set of classes of homotopically equivalent system paths. Each such class has distinct topological properties which indicate a distinct joint strategy behavior that the agents followed to reach their destinations, while avoiding collisions with each other. To enumerate such classes of behavior but also to label system paths, we develop a model of behavior using the concept of braids [2].

C. Background on Braids

Braids are topological objects with algebraic and geometric presentations. We first introduce them as geometrical entities, following a presentation based on Artin [1] and continue with

[1]By context, we refer to information that is publicly available (e.g. the map), information that may be directly acquirable through sensing (e.g. agent trajectories) and information that may be acquired through standard inference processes (e.g. agent groups).

a discussion of their algebraic presentation and their group formation.

Denote by x, y, z the cartesian coordinates of a Euclidean space $\mathbb{R}^2 \times I$, where $I = [0, 1]$. A braid on n-strings or n-braid is a set of n, monotonically increasing curves $X_i(z) : I \rightarrow \mathbb{R}^2$, $i \in N = \{1, \ldots, n\}$ for which:

1) $X_i(z) \neq X_j(z)$, for $i \neq j \ \forall z \in \mathbb{R}$
2) $X(0) = (i, 0)$ and $X(1) = (p(i), 0)$,

where $p(i)$ is the image of an element $i \in N$, through a permutation $p : N \rightarrow N$ from the set of permutations of N, $Perm(N)$, defined as:

$$p = \begin{pmatrix} 1 & 2 & \ldots & n \\ p(1) & p(2) & \ldots & p(n) \end{pmatrix}. \qquad (1)$$

This geometric representation of a braid is commonly referred to as a *geometric braid*. More formally, a geometric braid is often represented with a *braid diagram*, a projection of the braid onto the plane $\mathbb{R} \times 0 \times I$ (see e.g. Fig. 2).

The set of all braids on n strings, along with the composition operation, form a group B_n. The group may be generated from a set of $n-1$ elementary braids $\sigma_1, \sigma_2, \ldots, \sigma_{n-1}$ (see Fig. 2), called the generators of B_n, that satisfy the following *relations*:

$$\sigma_j \sigma_k = \sigma_k \sigma_j, \qquad |j - k| > 1, \qquad (2)$$

$$\sigma_j \sigma_k \sigma_j = \sigma_k \sigma_j \sigma_k, \qquad |j - k| = 1. \qquad (3)$$

A generator σ_i, $i \in \{1, 2, ..., n-1\}$ can be described as the crossing pattern that emerges upon exchanging the ith string (counted from left to right) with the $(i+1)$th string, such that the initially left string passes *over* the initially right one, whereas the inverse element, σ_i^{-1}, implements the same string exchange, with the difference that the left string passes *under* the right (see Fig. 3). An identity element, e, is a braid with no string exchanges.

Two braids $b_1, b_2 \in B_n$ may be composed through the composition operation $(\cdot)$, which is algebraically denoted as a product $b_1 \cdot b_2$. Geometrically, this composition results in the pattern that emerges upon attaching the lower endpoints of b_2 to the upper endpoints of b_1 and shrinking each braid by a factor of 2, along the z axis (see Fig. 2,Fig. 3). Any braid can be written as a product of generators and generator inverses. This form of representation is commonly referred to as an *algebraic braid* or a *braid word* (Fig. 3).

D. Abstracting Collective Navigation Behaviors Using Braids

Denote by $f_x : \mathcal{Q}^n \rightarrow Perm(N)$ a function that takes as input the system state $Q \in \mathcal{Q}^n$ and outputs a permutation $p \in Perm(N)$ corresponding to the arrangement of all agents in order of increasing x-coordinates.

As the agents move towards their destinations, they employ navigation strategies – maneuvers to avoid collisions. These result in a system path Ξ, which corresponds to a path of permutations $\pi : [0, T] \rightarrow Perm(N)$ that may be extracted by evaluating f_x throughout the whole path Ξ. This path can be represented by a sequence of permutations of minimal length $\pi^* = (p_0, \ldots, p_K)$, i.e., $p_{j-1} \neq p_j$, $\forall\, j = \{1, \ldots, K\}$

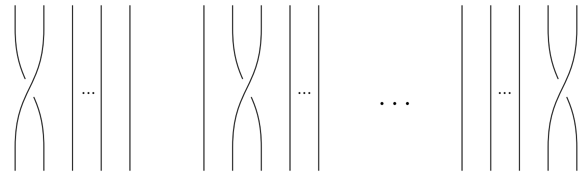
(a) σ_1 (b) σ_2 (c) σ_{n-1}

Fig. 2: The generators of the Braid Group B_n.

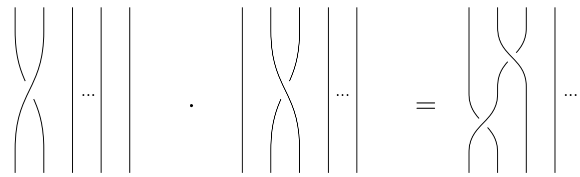
σ_1 σ_2^{-1} $\sigma_1 \cdot \sigma_2^{-1}$

Fig. 3: Example of the Composition operation $\sigma_1 \cdot \sigma_2^{-1}$ for $\sigma_1, \sigma_2^{-1} \in B_n$.

and consecutive waypoints are adjacent transpositions[2], i.e., permutations that differ by exactly one swap of adjacent elements.

Therefore, due to continuity, a transition from the $(j-1)$th permutation, p_{j-1}, to the jth permutation, p_j, implies the occurrence of an *event* τ_j, which may be described as *the intersection of the x-projections of the paths of two agents that were adjacent in the permutation* p_{j-1}. The event τ_j may be represented as an elementary braid $\tau_j \in \sigma_i^{\pm 1}, i \in \{1, ..., n-1\}$, where i corresponds to the index of the leftmost swapping agent in permutation p_{j-1}. Therefore the whole execution from $t = 0$ to $t = T$ may be abstracted into the braid that corresponds to the temporal sequence of events:

$$\tau = \tau_1 \tau_2 \ldots \tau_K \in B_n. \qquad (4)$$

This braid word not only constitutes a topological characterization of the system path (see Fig. 4 for an example of characterizing a system path as a braid) but it also represents a topological class of system paths that are homotopy-equivalent with the system path in consideration. In the remainder of this paper, we will be referring to the sequence τ as the *joint strategy* or the *entanglement* of the system path. Essentially, we model the space of joint strategies $\mathcal{T}$ as the braid group, i.e., $\mathcal{T} := B_n$.

Remark 1. *It should be noted that since braids are defined with respect to a specific frame of reference, the model of joint strategies is agent-specific. Furthermore, different projection planes could be selected.*

III. Socially Competent Navigation Planning

Generating socially competent navigation behaviors in our context involves understanding the collective intention of the system of agents regarding a future strategy of avoidance and acting in a compliant fashion. Doing so requires (1) a

[2]A transposition can be described as a permutation involving exactly one swap of a pair of elements. An adjacent transposition is a transposition involving an exchange of two adjacent elements.

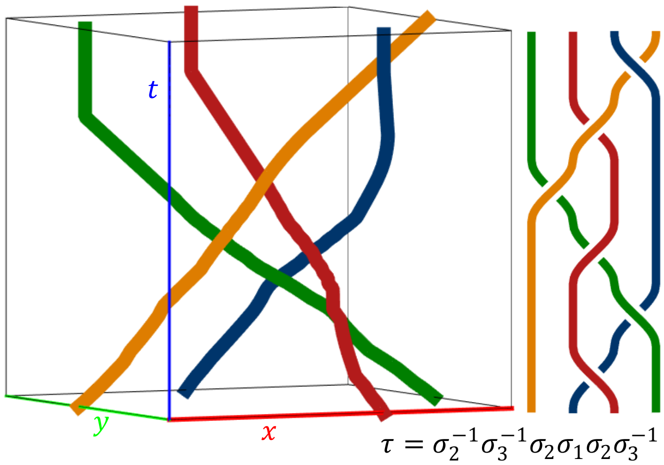

Fig. 4: A space-time representation of a system path in a workspace with 4 agent (left) along with its corresponding braid diagram (right) and braid word (bottom), defined with respect to the path's x-projection. The visualization of the braid diagram and the extraction of the braid word was done using BraidLab [14].

proper inference mechanism and (2) a corresponding policy for motion generation. In the following subsections, we describe our approach to designing them.

A. Inferring Avoidance Strategies from Context

Following the insights of Csibra and Gergely [3] regarding the goal-directed mechanisms of human action-interpretation, we design an inference mechanism of the form $P(\tau|\Xi_t, M_t)$, corresponding to a belief over an emerging strategy $\tau \in \mathcal{T}$, given a partial system path Ξ_t and the state of the context at time t, M_t. The strategy τ represents a collective *intention*, whereas the system path represents a collective *action*. By context we refer to publicly available information, such as a model of the static environment (e.g. a map, obstacles, points of interest etc) but also information extracted through processing, e.g. by employing secondary inference mechanisms regarding group formations, identification of reactive agents etc.

From (4), the belief $P(\tau|\Xi_t, M_t)$ may be expanded as:

$$P(\tau|\Xi_t, M_t) = P(\tau_1, \dots, \tau_K|\Xi_t, M_t), \tag{5}$$

which, by applying the chain rule, may be factored as:

$$P(\tau|\Xi_t, M_t) = \prod_{k=1}^{K} P(\tau_k|\bigcap_{j=1}^{k-1} \tau_j, \Xi_t, M_t). \tag{6}$$

This belief quantifies the likelihood of a sequence of events $\tau_1, \dots, \tau_K$ given observation of agents' past behaviors and the context. Essentially, this corresponds to predicting the emerging sequence of permutations π^* but also the quality of the physical transitions between consecutive permutation waypoints, i.e., the passing protocol (right/left) followed.

B. Generating Socially Competent Motion

The outlined inference mechanism is used to develop a decision making policy that compromises between achieving personal efficiency and reinforcing a consensus over a joint strategy among agents.

Let $u_i : \mathcal{A}_i \rightarrow \mathbb{R}$ be a cost function evaluating the quality of an action $a_i \in \mathcal{A}_i$ that agent i selects from its action space $\mathcal{A}_i$. We design this cost to comprise two terms: (1) E_i, which represents the action's *Efficiency* with respect to the agent's destination and (2) C_i, which represents the state of *Consensus* over a joint strategy among agents, from the perspective of agent i, upon taking an action $a_i \in \mathcal{A}_i$:

$$u_i(a_i) = \lambda E_i(a_i) + (1 - \lambda)C_i(a_i), \tag{7}$$

where λ is a weighting factor, expressing the compromise between efficiency and consensus.

We define the personal efficiency term E_i, to be the length of the shortest path to the agent's destination, whereas C_i is modeled as the Information Entropy of the belief distribution over joint strategies $P(\tau|\Xi_t, M_t)$, from the perspective of agent i, i.e.,

$$C_i(a_i) = -\sum_{\tau \in \mathcal{T}} P(\tau|a_i, \Xi_t, M_t) \log_2 P(\tau|a_i, \Xi_t, M_t), \tag{8}$$

Formally, the decision making policy for socially competent navigation (SCN) may be described as a minimization of eq. (7):

$$a_i^* = \arg\min_{a_i \in \mathcal{A}_i} u_i(a_i). \tag{9}$$

Overall, this policy enables an agent to make decisions that not only contribute progress towards its destination but also towards a mutually beneficial consensus over a joint strategy. The faster such a consensus is established, the lower the uncertainty will be for all agents throughout the remainder of the execution. The Efficiency term represents agents' intention of reaching their destinations by spending low energy and is in line with the principle of rational action as highlighted in the definitions of the pedestrian bargain [16] and the teleological reasoning [3]. The Consensus term scores the current state of the global consensus among agents regarding the joint strategy to be followed and therefore, it directly incorporates a form of social understanding into the agent's decision making policy. The lower the entropy, the lower the uncertainty regarding the emerging joint strategy. Thus, by consistently picking actions that contribute to entropy reduction, an agent communicates its intention of complying with a subset of scene outcomes that appear to be preferable by everyone according to the model $P(\tau|\Xi, M)$. As a result, the agents are expected to reach a consensus over τ faster, avoiding ambiguous situations such as livelocks or deadlocks and reach their destinations with lower planning effort.

IV. Applications

So far, we have employed the SCN framework on two different setups: a discrete setup, in which agents navigate a discretized board, and a continuous setup, in which agents navigate a continuous workspace. For both cases, we demonstrate

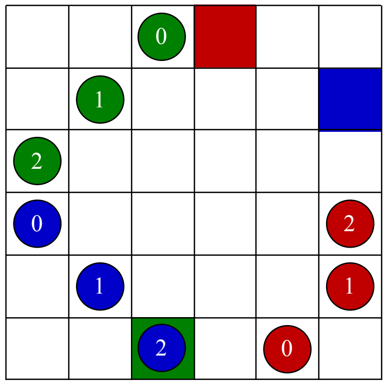

(a) Partial execution from a game with 3 agents on a square board.

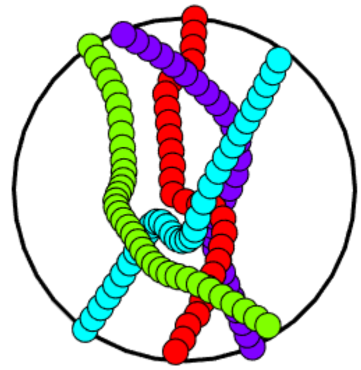

(b) Swept volumes of 4 agents navigating a circular workspace.

Fig. 5: Example executions from the discrete setup (left) and the continuous Setup (right).

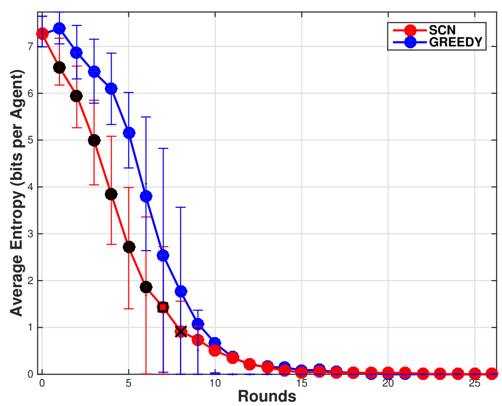

(a) Average Entropy profiles for the case of 4 agents navigating a discretized square board. Error bars indicate 25-75 percentiles, full black circles, circles and crosses indicate statistically different distributions with p-values $\leq 10^{-4}$, extracted by performing a T-test on every round.

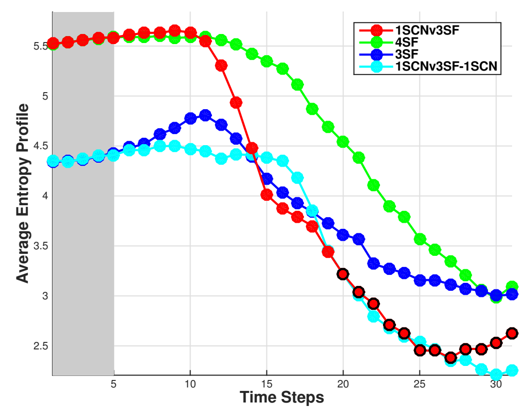

(b) Average Entropy profiles for the case of 4 agents navigating a continuous circular workspace. Different colors represent different setups, as shown in the legend. Circled datapoints correspond to timesteps at which 1SCNv3SF exhibited significantly lower entropy than all other setups (t-test, p-values ≤ 0.02).

Fig. 6: Average Entropy profiles for a discrete setup on a square board (left) and a continuous setup on a circular workspace (right), each derived upon running 100 experiments with the corresponding algorithm implementation.

the importance of reasoning about the collective behavior of multiple agents by comparing against *greedy* baselines that maximize agents' progress to destinations. In this section, we present corresponding models and main results.

A. Discrete Setup

Inspired by Wolfinger's observations on the cooperative nature of human navigation [16], we approach the problem of navigation in dynamic multi-agent environments, where no communication takes place among agents, as a finitely repeated coordination game of imperfect information and perfect recall [10]. The game is repeated a finite number of rounds M, which is unknown a priori and corresponds to the round at which the slowest agent reached its destination. At each round $m \in \{1, \ldots, M\}$, each agent i decides on an action a_i^k from a set of available actions (actions that could potentially lead to collisions and actions that violate the agent's dynamics are excluded) $\mathcal{A}_i^m$ by minimizing a cost function u_i. The agents are simultaneously selecting their actions and therefore they have no access to other agents' plans (imperfect information); we assume however that they maintain a history of all previous rounds (perfect recall).

To evaluate our approach, we constructed an analytical heuristic for approximating the distribution $P(\tau|\Xi_t, M_t)$ (for more details see Mavrogiannis and Knepper [10]). Fig. 5a represents an example execution with 3 agents, whereas Fig. 6a depicts average entropy profiles for 100 different scenarios involving 4 agents navigating the same workspace, by running our algorithm (red curve) or by only maximizing efficiency at every step (blue curve). It can be observed that agents running SCN achieve a faster entropy reduction, which represents a faster consensus over a joint strategy of avoidance.

B. Continuous Setup

In order to apply our framework to more realistic settings, we employed a data-driven approach to learn a model of the inference mechanism $P(\tau|\Xi_t, M_t)$ from demonstrations of multi-agent behaviors [11]. To do so, we generated a dataset of system paths by employing the *Social Force* (SF) model [5] on different scenarios involving multiple agents navigating rectangular or circular workspaces. Each scenario was generated by sampling agents' initial and final configurations on the circumference of the workspace. Each agent traverses the workspace towards a diametrically opposed destination by executing control inputs generated by an individual instance of the social force model. The model parameters, as well as agents' initial positions and destinations are varied across agents and experiments according to gaussian distributions. The experiments were explicitly designed to enforce intense agent encounters in order to obtain interesting behaviors. Experiments on the circular workspace loosely simulate pedestrians crossing paths in free areas such as atriums or parking lots, whereas the rectangular workspace captures pedestrians passing in hallways with mostly parallel and orthogonal paths.

We extract models of the conditional probabilities of eq. (6) by employing a sequence to sequence encoder-decoder learning architecture. The encoder and decoder RNNs were implemented by following the Long Short-Term Memory (LSTM) paradigm [6] due to its effectiveness in capturing long-term sequence dependencies.

We evaluate our algorithm by comparing against the Social Force model in 4 different setups: (1) 1 agent runs SCN and 3 agents run SF, (2) 4 agents run SF, (3) 3 agents run SF, (4) case 1, upon removing the agent running SCN. Fig. 6b depicts corresponding average entropy profiles for 100 different scenarios. It can be observed that by putting 1 agent running SCN around a system of 3 agents running SF (red curve) results in significantly accelerated uncertainty decrease, compared to all other setups. Fig. 5b depicts an example execution with 4 agents.

V. Discussion and Future Work

We presented a planning framework for socially competent navigation in multi-agent environments with no explicit communication among agents. Simulation results demonstrated a faster decrease of uncertainty among agents in a scene that appears promising for application in human environments. Future work involves an experimental evaluation of our approach with a social robot in a crowded human workspace and a user study to get feedback from humans.

Acknowledgments

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

References

[1] E. Artin. Theory of braids. *Annals of Mathematics*, 48 (1):pp. 101–126, 1947.

[2] Joan S. Birman. *Braids Links And Mapping Class Groups*. Princeton University Press, 1975.

[3] 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.

[4] Anca D. Dragan and Siddhartha Srinivasa. Integrating human observer inferences into robot motion planning. *Auton. Robots*, 37(4):351–368, 2014.

[5] Dirk Helbing and Péter Molnár. Social force model for pedestrian dynamics. *Phys. Rev. E*, 51:4282–4286, 1995.

[6] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.

[7] Ross A. Knepper and Daniela Rus. Pedestrian-inspired sampling-based multi-robot collision avoidance. In *RO-MAN*, pages 94–100. IEEE, 2012.

[8] 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. ACM, 2017.

[9] 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.

[10] Christoforos I. Mavrogiannis and Ross A. Knepper. Decentralized multi-agent navigation planning with braids. In *Proceedings of the Workshop on the Algorithmic Foundations of Robotics (WAFR)*, 2016.

[11] Christoforos I. Mavrogiannis, Valts Blukis, and Ross A. Knepper. Socially competent navigation planning by deep learning of multi-agent path topologies. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2017*.

[12] Douglas W Nangle, David J Hansen, Cynthia A Erdley, and Peter J Norton. *Practitioner's guide to empirically based measures of social skills*. Springer Science & Business Media, 2009.

[13] 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.

[14] Jean-Luc Thiffeault and Marko Budišić. Braidlab: A software package for braids and loops, 2013–2016. URL http://arXiv.org/abs/1410.0849. Version 3.2.1.

[15] 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. *Int. J. of Robotics Res.*, 34(3): 335–356, 2015.

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