

Decentralized Multi-Agent Navigation Planning with Braids

Christoforos I. Mavrogiannis[†] and Ross A. Knepper[‡]

[†]Sibley School of Mechanical & Aerospace Engineering, Cornell University
cm694@cornell.edu

[‡]Department of Computer Science, Cornell University
rak@cs.cornell.edu

Abstract. We present a novel planning framework for navigation in dynamic, multi-agent environments with no explicit communication among agents, such as pedestrian scenes. Inspired by the collaborative nature of human navigation, our approach treats the problem as a coordination game, in which players coordinate to avoid each other as they move towards their destinations. We explicitly encode the concept of coordination into the agents' decision making process through a novel inference mechanism about future joint strategies of avoidance. We represent joint strategies as equivalence classes of topological trajectory patterns using the formalism of braids. This topological representation naturally generalizes to any number of agents and provides the advantage of adaptability to different environments, in contrast to the majority of existing approaches. At every round, the agents simultaneously decide on their next action that contributes collision-free progress towards their destination but also towards a global joint strategy that appears to be in compliance with all agents' preferences, as inferred from their past behaviors. This policy leads to a smooth and rapid uncertainty decrease regarding the emerging joint strategy that is promising for real world scenarios. Simulation results highlight the importance of reasoning about joint strategies and demonstrate the efficacy of our approach.

Keywords: motion planning, navigation, topology, braids

1 Introduction

Pedestrian environments, such as crowded hallways, sidewalks, and rooms are often characterized by un-structured motion, imposed by the lack of formal rules to control traffic and the lack of explicit communication among pedestrians. Nonetheless, humans are capable of traversing such environments with remarkable efficiency and without hindering each other's paths. Human navigation not only achieves collision avoidance; it does so while respecting several social considerations, such as the passing preference of others and their personal space, ensuring smooth co-navigation. Cooperation has been identified as the key for the generation of such a complex behavior (e.g. [23]). In the absence of explicit communication, cooperation relies on intention inference, which in turn is based on trust. Pedestrians infer others' intentions and preferences by observing their motion, while communicating their own, essentially negotiating a joint strategy of avoidance. Past experiences of cooperatively resolved pedestrian encounters build and

reinforce a form of trust among pedestrians that allows them to relax uncertainty and agree on a joint strategy of avoidance.

Inspired by the efficiency of humans in resolving pedestrian conflicts, we explicitly employ the concept of cooperation in the design of an online navigation algorithm for multi-agent environments. Our approach treats navigation as a cooperative game in which agents make decisions by compromising between a notion of personal efficiency and a concept of joint efficiency. The concept of joint efficiency concerns the joint strategy that all of the agents will follow to reach their destinations, while avoiding others. We model joint strategies as topological patterns of agents' trajectories using *braids* [3].

In contrast to the majority of existing approaches that are either too myopic, only focusing on local collision avoidance resolution, or too specific, reproducing demonstrated behaviors in specific contexts, we contribute: (1) a topological model of a multi-agent scene, based on braids, that allows us to reduce the problem of planning joint strategies to a graph search problem that can be efficiently solved with existing techniques; (2) a framework for collaborative motion planning that generalizes across environments and numbers of agents; (3) a probabilistic intent inference mechanism for cooperative navigation that accelerates the rate of convergence among agents' plans; and (4) simulation results demonstrating the importance of incorporating a collective, global topological understanding in the planning process. Our framework was designed according to the insights of sociology studies on pedestrian behavior and psychology studies on action interpretation, reflecting our goal to employ it on a mobile robot platform navigating in crowded human environments. The topological structure that our model offers to the motion planning process is expected to reduce the emergence of undesired situations such as deadlocks and livelocks that are frequently observed in human-robot pedestrian encounters.

2 Related Work

2.1 Navigation

Navigation has been the focus of various diverse scientific communities, ranging from sociology and cognitive science to computer vision and robotics, aiming at understanding and simulating human navigation but also at reproducing robotic navigation.

Several works have proposed models for crowd dynamics that have been validated in simulation of various scenarios in different contexts. The social force model [9] and its variants [12, 17, 21] introduced a physics-inspired way of modeling pedestrian interactions: pedestrians are attracted to their destination and repulsed by obstacles or other agents. Hoogendoorn and Bovy [11] modeled the problem as a differential game in which the agents are cost minimizing predictive controllers. Moussaïd et al. [15] looked at the problem from a cognitive science perspective, proposing a set of behavioral heuristics that guide human walking behaviors. Bonneaud and Warren [4] proposed a decomposition of locomotion into a set of elementary behaviors, each modeled as an experimentally tuned nonlinear dynamical system. Finally, Zhou et al. [25] presented a data-driven approach for learning macroscopic collective crowd behaviors.

Designing artificial agents, capable of seamlessly navigating dynamic human environments typically requires a predictive framework and a planning framework. Over

the past two decades, the robotics community has made a number of significant contributions related to both components.

In human motion prediction, roboticists have employed learning techniques to derive models of human behavior. Bennewitz et al. [2] clustered human behavior into typical motion patterns which they used to perform online trajectory predictions on a mobile robot platform. Ziebart et al. [26] and Henry et al. [10] presented data-driven approaches, based on Inverse Reinforcement Learning for learning context-specific humanlike navigation behaviors for static and dynamic environments respectively.

In the area of planning and control, emphasis has been given to the design of strategies that would enable robots to integrate smoothly in human environments. To this end, Sisbot et al. [18] presented a cost-based planner that incorporates considerations of human comfort and context-specific social conventions, whereas Park et al. [16] proposed an online model-predictive control framework that generates locally optimal collision-free smooth trajectories for autonomous robotic wheelchairs. Another class of works have focused on modeling the interactions among multiple agents. The reactive multi-robot planning framework of van den Berg et al. [22] made explicit use of the assumption that the responsibility for collision avoidance is shared among interacting agents. Under the same assumption, Knepper and Rus [13], inspired by human navigation, contributed a sampling-based planner that also incorporates predictions about other agents' trajectories in the planning process. Kuderer et al. [14] and Trautman et al. [20] presented learning frameworks for predicting the trajectories of interacting pedestrians, which they used to plan socially compliant robot motion.

2.2 Human Behavior

One of the central principles guiding human decision making in pedestrian environments appears to be *cooperation*. Wolfinger [23] concluded to a concise, high-level protocol that captures the essence of the cooperative nature of human navigation: the *pedestrian bargain*. The pedestrian bargain is a set of foundational social rules that regulate pedestrian cooperation: (1) *people must behave like competent pedestrians* and (2) *people must trust copresent others to behave like competent pedestrians*. Pedestrians' trust to the rules of the bargain constitutes the basis of smooth co-navigation in shared environments.

In the absence of explicit communication, pedestrians rely on inference mechanisms for both prediction of others' behaviors and for generation of their own behaviors. As a result, building an autonomous system capable of seamlessly navigating human environments requires the design of a realistic, human-like inference mechanism. To this end, under the assumption of rational action, we design a goal-driven probabilistic model for action understanding and generation, that is in line with the insights of Csibra and Gergely [5, 6] regarding teleological action interpretation of humans. The concept of teleological reasoning, describing the tendency of humans to attribute potential *goals* to observed *actions*, has recently been employed in human robot-interaction by Dragan and Srinivasa [8], who formalized a framework for intent-expressive robot motion.

2.3 Topology

Finally, the foundational inspiration for this work is the topological concept of braids. The formalism of braids, first formulated by Artin [1] and extensively studied by Bir-

man [3] has been an inspiration for applications in various disciplines, including robotics. Diaz-Mercado and Egerstedt [7] were the first to develop a framework for centralized multi-robot mixing, in which the agents are assigned trajectories that contribute to a specified topological pattern corresponding to a given braid. Although we are also making use of braids to model multi-robot behaviors, the scope of our approach is inherently different, since our target application concerns navigation in dynamic environments where no explicit communication takes place. In our case the agents do not follow a pre-specified braid pattern, but rather employ a braid-based probabilistic reasoning to reach a topological consensus that best complies with everyone’s intentions or preferences. For our purposes, braids provide a basis for reasoning about uncertainty in a principled fashion, as their dual geometric and algebraic representation enables us to symbolically enumerate diverse distinct topological scene evolutions. As a result, our algorithm generates *socially competent* behaviors, i.e., behaviors that explicitly take into consideration the *social welfare* of the whole system of agents.

3 Foundations

Consider n agents navigating a workspace \mathcal{W} . Each agent i starts from an initial configuration $q_i \in \mathcal{W}$ and aims at reaching a destination $d_i \in \mathcal{W}$ by following a trajectory $\zeta_i : I \rightarrow \mathbb{R}^2$, with $I = [0, 1]$ being a normalized time parametrization. The agents do not explicitly exchange information regarding their planned paths and are assumed to be acting rationally, which in our context means that (1) they always aim at making progress towards their destinations and (2) they have no motive for acting adversarially against other agents (e.g. blocking others’ paths or colliding with them). The notion of rationality is in line with the concept of *competence* as described by Wolfinger [23] in his definition of the *Pedestrian Bargain*.

3.1 Game-Theoretic Formulation

Inspired by Wolfinger’s observations on the cooperative nature of human navigation, we approach the problem of robotic navigation in multi-agent, dynamic environments as a finitely repeated coordination game of imperfect information and perfect recall. The game is repeated a finite number of rounds (until all agents reach their destinations). At every round t , each agent decides on an action a_i from a set of available¹ actions A_i . All 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).

Based on our assumption of rationality, we can model agents’ decision making to be the result of optimizing a utility function u_i . Agents generally aim at reaching their destination by spending low energy. However, in a multi-agent, uncertain environment, each agent’s decisions are not independent of the decisions of others. Many decisions might lead to collisions, whereas others, seemingly greedily serving one agent’s own interests might not be sufficient to guarantee long term efficiency in such a complex context. In particular, the latter might actually lead to undesired outcomes such as longer paths, antisocial hindering of others’ paths or even deadlocks and livelocks. For this

¹ Actions that could probabilistically lead to collisions or actions that violate the agent’s dynamics are excluded.

reason, it is important that each agent’s utility function incorporates a term reflecting the *social welfare*, i.e., the “common good”, besides its own efficiency.

The set of actions selected by all players at a round t , $A = \{a_1, \dots, a_n\}$ constitute a *strategy profile*. The sequence of strategy profiles of all rounds from the beginning to the end of the game form a global *joint strategy* s . In this paper, we make use of the concept of joint strategies to imbue artificial agents with an understanding of how their own actions affect the actions of others.

3.2 Modeling Joint Strategies

As the agents move from their initial configurations $Q = \langle q_1, \dots, q_n \rangle$ to their intended destinations $D = \langle d_1, \dots, d_n \rangle$, their trajectories $Z : I \rightarrow (\mathbb{R}^2)^n$ form a 3-dimensional pattern in space-time. This pattern is the result of the global *joint strategy* s that the agents engaged in, to avoid each other and reach their destinations, from the beginning to the end of the game. Its topological properties are particularly interesting, as they can provide a qualitative characterization of the strategy and hence of the agents’ interactions. For this reason, we represent joint strategies by their corresponding topological pattern. Thus, a joint strategy s is an equivalence class of trajectory patterns that belongs to a set of classes \mathcal{S} . In this paper, we model this set, using braid groups [3].




Fig. 1: Alternative braid representations.

3.2.1 Background on Braids

Consider the finite set $N = \{1, 2, \dots, n\}$, where $n \in \mathbb{N}^+$ and denote by $Perm(N)$ the set of permutations on N . A permutation in $Perm(N)$ is a bijection $b : N \rightarrow N$ that can be written² as:

$$b = \begin{pmatrix} 1 & 2 & \dots & n \\ b(1) & b(2) & \dots & b(n) \end{pmatrix}, \quad (1)$$

where $b(i)$, $i \in N$, is the image of element i , through the permutation b .

From a geometric perspective, a braid on $n \geq 1$ strands can be described as a system of n curves in \mathbb{R}^3 , called the strands of the braid such that each strand i connects the

² Adopting Cauchy’s *two-line notation*

point $(i, 0, 0)$ with the point $(b(i), 0, 1)$ and intersects each plane $\mathbb{R}^2 \times t$ exactly once for any $t \in I$ (see Fig. 1a). A braid is usually represented with a *braid diagram*, a projection of the braid to the plane $\mathbb{R} \times 0 \times I$ with indications of the strand crossings (see Figs. 1b).




Fig. 2: The generators of the Braid Group \mathcal{B}_n .

The set of all braids on n strands, along with the composition operation, form a group³ \mathcal{B}_n that can be generated from a set of $n - 1$ elementary braids (shown in Fig. 2), called the generators of \mathcal{B}_n that satisfy the following *relations*:

$$\sigma_j \sigma_k = \sigma_k \sigma_j, \quad |j - k| > 1; \quad \sigma_j \sigma_k \sigma_j = \sigma_k \sigma_j \sigma_k, \quad |j - k| = 1. \quad (2)$$

Intuitively, a generator σ_i , $i \in \{1, 2, \dots, n - 1\}$, can be described as the pattern that emerges upon exchanging the i th strand (counted from left to right) with the $(i + 1)$ th strand, such that the left strand passes over the right, whereas the inverse element σ_i^{-1} implements the same strand exchange, with the difference that the left strand passes under the right. A trivial element is a braid where no strand exchanges occur. The group operation can be described as a concatenation of braids: given two braids $B_1, B_2 \in \mathcal{B}_n$, the product $B_1 \cdot B_2$ results in B_2 being placed at the top of B_1 , by attaching the top endpoints of B_1 to the bottom endpoints of B_2 and shrinking each braid by a factor of 2, along the t axis (e.g. see Fig. 1b). Any braid can be written as a product of generators and their inverses. This product is commonly referred to as *braid word* (Fig. 1c).

3.2.2 Representing a Trajectory Collection as a Braid

We use the aforementioned braid representations to characterize topologically a collection of trajectories, based on the method of Thiffeault [19]. Given a trajectory collection Z and a line $\epsilon \in \mathcal{W}$, we can extract the braid word that corresponds to the scene evolution by (1) projecting the states of all agents at every point in time, $Z(t)$ to ϵ , (2) labeling any emerging projected trajectory intersections as generators (or their inverses) according to the intersection pattern, and (3) arranging them into temporal order. This word describes the topological properties of agents' trajectories Z , from the beginning to the end of time. Fig. 3 depicts an example of transitioning from a collection of trajectories to a braid diagram and finally a braid word (considering ϵ to be the x axis).

³ They satisfy the four group axioms of *closure*, *associativity*, existence of *identity* element and existence of *inverse* element.




Fig. 3: Projecting a trajectory collection (left) to the x axis to derive a braid diagram and a braid word (right).

3.2.3 Using Braids to Represent Joint Strategies

In this paper, we make use of the braid group \mathcal{B}_n to represent the set of classes of *joint strategies* \mathcal{S} that a set of n agents could follow in a common workspace to avoid each other and reach their destinations. Instead of explicitly planning geometric representations of agents' potential future behaviors, i.e., predicting their future trajectories Z' , the agents are reasoning symbolically about possible emerging collective topologies—braids. A planning agent i , headed towards a destination d_i , determines a set of joint strategies \mathcal{S} by considering an appropriate braid definition, i.e., selecting an appropriate projection line ϵ to define braids with respect to. It should be noted that a change of projection line only changes the braid by conjugation [19]. Therefore, without loss of generality, we assume that all agents have access to the same set of joint strategies, although they may look at them from different perspectives.

3.3 Modeling Agents' Inference Mechanism

The braid group \mathcal{B}_n is infinite; therefore infinitely many alternative joint strategies could be mathematically possible. However, under a context M and observations of past collective behaviors Z , agents may form a belief over the set of emerging joint strategies \mathcal{S} . For a planning agent, this inference mechanism serves as a form of context-specific social understanding, in the sense that it enables the agent to understand how its decision over a navigation strategy is coupled with the decisions of others.

In this paper, we design agents' inference mechanism in a human-inspired fashion, reflecting our goal to employ our framework on an autonomous social robot. In particular, we follow the main insight of [5, 6] regarding the teleological nature of human action interpretation: humans tend to attribute potential context-specific *goals* to observed *actions*. In our framework, a joint strategy $s \in \mathcal{S}$ is a goal, whereas agents' state history Z is the action and M is a variable that models the context of the scene (encoding the understanding of the static environment, such as obstacles, points of interest but also secondary inferences such as predictions about the destinations of others, agents' groupings etc). Formally, we model an agent i 's inference as a belief distribution over a set of future joint strategies \mathcal{S} , given past trajectories Z and the context M :

$$\text{bel}_i(s) = P(s|Z, M). \quad (3)$$

3.4 Modeling Agents' Utilities

We model the interests of an agent i with a utility function $u_i : \mathcal{A}_i \rightarrow \mathbb{R}$ that maps an action $a_i \in \mathcal{A}_i$ to a real number, representing a reward for selecting it (higher rewards are better), whereas the *Social Welfare* is the mean of the individual utilities:

$$W(A) = \frac{1}{n} \sum_{i=1}^n u_i(a_i). \quad (4)$$

As we discussed earlier, in a multi-agent environment where no explicit communication takes place among agents, it is important that agents take into consideration their scene understanding in their decision making process. For this reason, we model u_i to be a weighted sum, compromising between personal efficiency and social compliance:

$$u_i(a_i) = \lambda E_i(a_i) - (1 - \lambda) H_i(a_i). \quad (5)$$

$E_i = \exp(-C_i(a_i))$ represents the efficiency of an action $a_i \in \mathcal{A}_i$ with respect to a geometric cost to destination $C_i : \mathcal{A}_i \rightarrow \mathbb{R}$, whereas H_i is the information entropy of the agent's belief regarding the emerging strategy s , defined as:

$$H_i(a_i) = - \sum_{s \in \mathcal{S}} P(s|Z^+, M) \log_2 P(s|Z^+, M) \quad (6)$$

where Z^+ is Z (the state history so far), augmented with the action in consideration a_i and λ is a weighting factor.

The Efficiency 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 [23] and the teleological reasoning [5]. The entropy reflects the state of the global consensus among pedestrians regarding the joint strategy to be followed and therefore, it directly incorporates a form of social understanding in an 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 joint strategies that appear to be more likely or "social" according to the (assumed shared) model $P(s|Z, M)$. Another interpretation of the functionality of this policy is that it implicitly biases others towards complying with the same strategy. As a result, the agents are expected to reach a consensus over s easier and faster, avoiding ambiguous situations such as livelocks or deadlocks and reach their destinations more *comfortably*; not necessarily faster or with less energy, but with a higher degree of cooperation, requiring lower planning cognitive load.

The superposition of the specifications for Efficiency and Entropy reduction, represent what -to our interpretation of the pedestrian bargain [23]- constitutes *competent* behavior in the pedestrian context: (1) rationality and (2) social understanding.

4 Planning Global Joint Strategies with Braids

As discussed in Sec. 3.3, the braid group \mathcal{B}_n is infinite; however, in practice, only a subset of joint strategies-braids are meaningful under the context of a scene and given observations of agents' past behaviors. In particular, given predictions of agents' destinations D , we can determine a set \mathcal{S} , comprising only strategies that take agents from their current configurations to their predicted destinations. Making use of the observation that all braids in \mathcal{B}_n describe transitions between permutations of the set $N = \{1, 2, \dots, n\}$,




Fig. 4: A multi-agent scene from the perspective the planning agent (blue color). The robot arranges all agents according to the projections of their current configurations on the line ϵ (coincident to the x-axis of its body frame) and derives a corresponding permutation b_0 . Given agents' past trajectories (dashed lines) and the context it also makes a prediction of their intended destinations (denoted with crosses) and derives a corresponding permutation b_1 . Transitioning from b_0 to b_1 is implemented with a joint strategy $s \in \mathcal{S}$.

we convert the problem of planning joint strategies to a graph search in a permutation graph.

Fig. 4 illustrates the concept of our method. Assume that at planning time, the agents have already followed trajectories denoted with thick dashed lines and are located at positions contained in the tuple Q (circles in vivid colors, thick lines). The planning agent (blue color) has predicted that they are aiming for the destinations contained in the tuple D (denoted with crosses). From its perspective, Q and D correspond to the permutations b_0 and b_1 respectively, derived upon their projection on the line ϵ (parallel to the x-axis of its body frame).

4.1 Permutation Graph Construction

The set of all permutations on N , $Perm(N)$, along with the composition operation, form the symmetric group S_n . S_n is a group of order $n!$, that can be generated by a set of adjacent transpositions $\beta_k = [k \ k + 1]$, with $1 \leq k < n - 1$ (i.e., the set of permutations that implement exactly one swap of a pair of adjacent elements in the set). It should be noted that these transpositions/generators just implement swaps of adjacent pairs, whereas *braid* generators, besides implementing adjacent swaps, also prescribe a swapping *quality* (which strand passes over or under, as discussed in Sec. 3.2.1).

We construct a permutation graph $G = (V, E)$, comprising a set of vertices that correspond to the elements of S_n . A pair of nodes $v_i, v_j \in V$ is connected iff there exists a permutation β_k (from the set of the generating transpositions described in the previous paragraph) that permutes v_i into v_j . Our graph can be graphically represented




Fig. 5: A permutation graph for a scene with four agents, represented as a permutohedron of order 4. Three alternative paths implementing the transition from the permutation 1234 to the permutation 3412 are depicted in different colors. Each path consists of a sequence of transitions, each of which can be implemented topologically with a braid generator or its inverse.

as a *permutohedron* [27]. Fig. 5 depicts a permutohedron of 4th order (for a scene with four agents) along with indications of braid transitions.

4.2 Searching in the Permutation Graph

Two vertices $v_i, v_j \in V$ corresponding to the permutations $b_i, b_j \in S_n$ respectively can be connected with a path of vertices p_{ij} in the graph G , containing the sequence of permutations that transitions b_i into b_j . Each edge in the graph corresponds to a generator of S_n that can only be implemented topologically with a generator of \mathcal{B}_n or its inverse. Therefore, given a path p_{ij} , a corresponding joint strategy (braid) s , transitioning b_i into b_j , can be derived by assigning a braid at each transition between consecutive way-points (vertices) in the path. For each transition, there are always two candidate braid generators, a σ_i^+ and a σ_i^- . Fig. 5 schematically demonstrates the assignment of braids to permutation transitions in three different paths implementing a symbolic plan with the same starting and ending permutations.

5 Algorithm Design

Algorithm 1 describes our online algorithm for socially competent navigation in multi-agent environments. The Algorithm starts by updating the context⁴ M with function `UpdateContext`. Next, the function `DetermineReactiveAgents` returns a subset of all observed agents that should be taken into consideration in the motion plan (e.g. ignoring agents that are behind the planning agent or agents that are too far ahead).

⁴ The context M comprises a static component (map, obstacles, points of interest in the scene etc) and a dynamic component that depends on agents' behaviors.

Subsequently, the algorithm determines the set of actions \mathcal{A} that are available to the planning agent, taking into consideration its dynamics and the positions and intentions of others (function `GetAvailableActions`). In case there are not other agents to which the planning agent should be reacting, the algorithm returns the most efficient action towards the agent’s destination (function `MaximizeEfficiency`). Otherwise, the algorithm continues with the function `GetStrategies` that derives a set of topological joint strategies/braids \mathcal{S} . Finally, the function `MaximizeUtility` returns a control command a that corresponds to the action that both makes progress towards the planning agent’s destination and communicates compliance with the most likely joint strategies at the given time. The algorithm runs until the planning agent reaches its destination, i.e., until the boolean variable `AtGoal` becomes 1.

Algorithm 1 `SCN($q, d, map, N, Z, AtGoal, a$)`

Input: q – agent’s current state; d – agent’s intended destination; `map`; Z – state history of all agents; `AtGoal` – boolean variable signifying arrival at agent’s destination; M – context

Output: a – action selected for execution

```

1: while  $\neg$ AtGoal do
2:    $M \leftarrow \text{UpdateContext}(Z, M)$ 
3:    $\mathcal{R} \leftarrow \text{DetermineReactiveAgents}(M)$ 
4:    $\mathcal{A} \leftarrow \text{GetAvailableActions}(M)$ 
5:   if  $\mathcal{R}$  then
6:      $\mathcal{S} \leftarrow \text{GetStrategies}(M, \mathcal{R})$ 
7:      $a \leftarrow \text{MaximizeUtility}(\mathcal{A}, \mathcal{S}, d, M, \mathcal{R})$ 
8:   else
9:      $a \rightarrow \text{MaximizeEfficiency}(\mathcal{A}, d)$ 
10: return  $a$ 

```

6 Application

We tested our algorithm in simulation in the following game. Consider a workspace \mathcal{W} , partitioned into a set of m polytopes (Fig. 6). A set of n agents navigate the workspace, each starting from an initial configuration $q_i \in \mathcal{W}$ and aiming at reaching a final configuration $d_i \in \mathcal{W}$. The game is played in rounds until all agents reach their destinations. At every round, the players simultaneously pick an action, i.e., a neighboring square. Forward, backward, left, right and diagonal, collision-free transitions are allowed. Since at planning time each agent has no access to others’ plans, in order to ensure collision avoidance, transitioning to a square that is adjacent to a square currently occupied by another agent is not allowed.

To demonstrate the importance of considering the emerging joint strategy in the decision making stage, we compare the performance of our algorithm against a simple baseline that only plans actions that seek to maximize the efficiency (the progress to the agent’s destination) at every round. This baseline is conceptually similar to the widely used *social force* algorithm [9]. We show that by explicitly reasoning about the emerging joint strategy when planning an action benefits everyone in the scene, as it leads

to a rapid uncertainty decrease that simplifies everyone’s decision making. This allows agents to avoid ambiguous situations that could lead to livelocks or deadlocks.

6.1 Implementation Details

For the simulations we made the assumption that the agents were aware of others’ destinations. The agents were starting from one side of the board and aiming at reaching a destination in the opposite side of the board. As a geometric cost C_i we selected the Manhattan distance to destination. For each agent, the projection line, with respect to which braids were defined, was selected to be constantly parallel to the line defining their starting board side (and coinciding with agent’s body frame x-axis). The belief over strategies $bel_i(s)$ was modeled as:

$$P(s|Z, M) = \frac{1}{\Lambda} \prod_{j=1:l} \exp(-(l-j)\Delta x_j) \quad (7)$$

where l is the length of the word representing a strategy s , Δx_j is the current distance along the x-axis between a pair of swapping agents, corresponding to the j th generator in the braid, and Λ is an appropriate normalizer. This distribution is a simplified approximation that cannot guarantee robust performance and generalization. We are using it in this paper to provide a proof of our concept. We plan on approximating it using human pedestrian data. Finally, for deriving a set of candidate paths in the permutation graph, we use the algorithm of Yen [24] for finding K shortest paths.

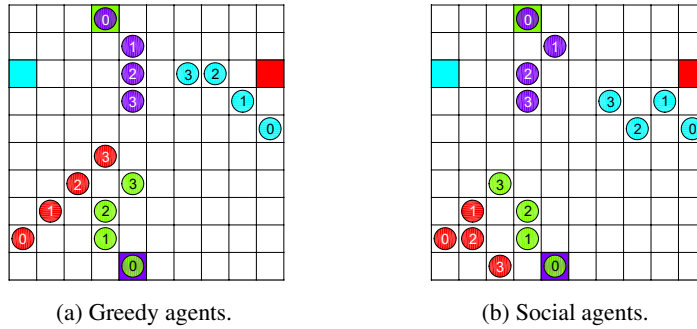


Fig. 6: Partial Execution: 4 agents play the same game (same initial configurations and destinations) running the baseline (Fig. 6a) and our algorithm (Fig. 6b). The destination of each agent corresponds to the square of the same color. The actions taken by each agent at every round so far are noted with a corresponding round number.

6.2 Simulation Results

Figs 6a and 6b depict partial executions of the same game, after three rounds, for the case of the baseline and our algorithm respectively. It can be observed that the agents running our algorithm (“social agents”) have achieved a better status, as all of their encounters are essentially resolved by the end of the third round. On the contrary, the agents running the baseline (“greedy agents”) are about to engage in an ambiguous encounter involving three of them aiming to pass from the same region of the board.

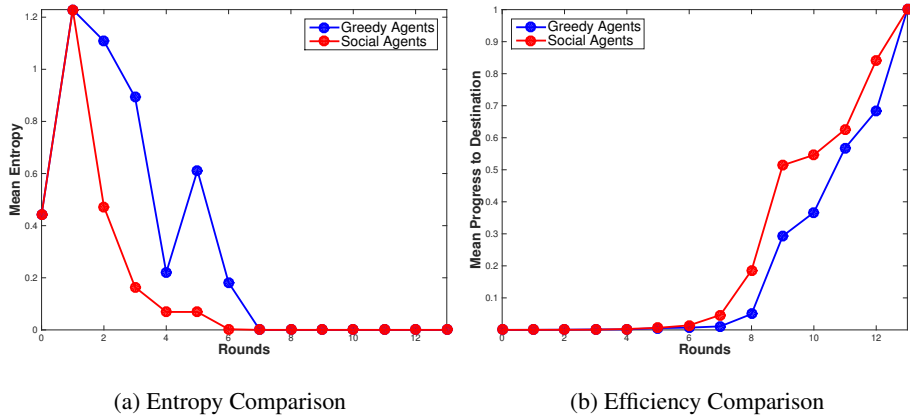


Fig. 7: Comparative Performance plots for a simple scenario involving four agents: Fig. 7a depicts the comparison of the mean entropy progression between the agents running our algorithm (red color) and the agents running the baseline (blue color), whereas Fig. 7b depicts the mean progress to destinations (mean efficiency) per round.

Fig. 7 depicts comparative performance diagrams for the game of Fig. 6. The progression of the quality of agents’ decision making is demonstrated by plots of the mean entropy⁵ (Fig. 7a) and the mean progress to the destination (Fig. 7b) per round of the game. Comparative plots of the complete trajectories of the game are shown in Fig. 8. It can be noticed that the *social agents* achieve a rapid decrease in uncertainty, expressed by the smooth convergence of the entropy, whereas in the case of the *greedy agents* the entropy fluctuates before the agents reach consensus, reflecting the ambiguity of agents’ actions. At the same time, it appears that the social agents’ actions are also ensuring faster progress to their destinations compared to the greedy ones.

Fig 9 depicts performance diagrams extracted from a similar scenario that besides four agents involves a static obstacle that blocks the agents’ way to their destinations. The obstacle is treated as an extra static agent. For our braid model an obstacle is not different than an agent, as it can be represented with a vertical strand. Fig. 9a demonstrates again an improved entropy progression, reflecting a faster consensus, while Fig. 9b shows that our algorithm stays quite close to the baseline in terms of efficiency. The complete trajectories for the two cases are depicted in Fig. 10.

7 Discussion and Future Work

We presented a framework for socially competent navigation in multi-agent environments with no explicit communication, inspired by the insights of recent studies on the cooperative nature of pedestrian behavior [23] and the goal-directed inference of humans [6]. Our framework explicitly incorporates the concept of cooperation by modeling multi-agent collective behaviors as topological global joint strategies, using the formalism of braids [3]. Our topological model forms the basis of an inference mech-

⁵ Entropy average across all agents.

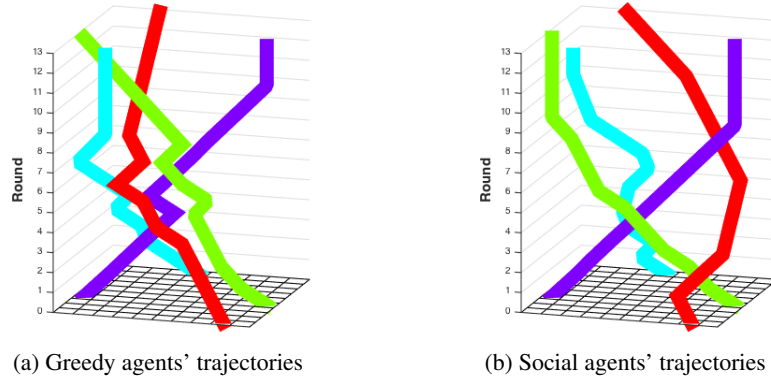


Fig. 8: Comparative plot of the trajectories followed by agents running the baseline (Fig. 8a) and our algorithm (Fig. 8b).

anism that associates observed behaviors with future collective topologies. In the decision making stage, each agent decides on an action that corresponds to a compromise between its personal efficiency (progress towards destination) and a form of joint efficiency (the status of a consensus on a joint strategy of avoidance). Simulation results on a discretized workspace demonstrated the benefit of incorporating this joint efficiency in the decision making stage, as opposed to picking actions that only contribute progress to one’s destination. Our approach was shown to lead to a rapid drop in uncertainty that allows agents to efficiently cooperate towards avoiding each other and reaching their destinations.

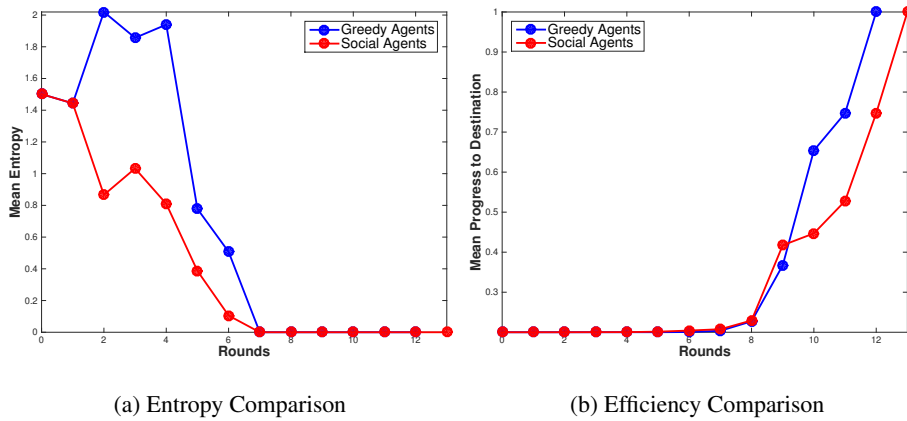


Fig. 9: Performance comparisons in a workspace with four agents and an obstacle.

Ongoing work involves the development of a continuous implementation of our algorithm and learning a distribution over topologies from human data. This will enable us to design a more realistic and accurate belief distribution over joint strategies and test our algorithm experimentally, in real world scenarios involving human agents in a variety of pedestrian environments. Finally, we plan on conducting a user study to get feedback and improve our design.

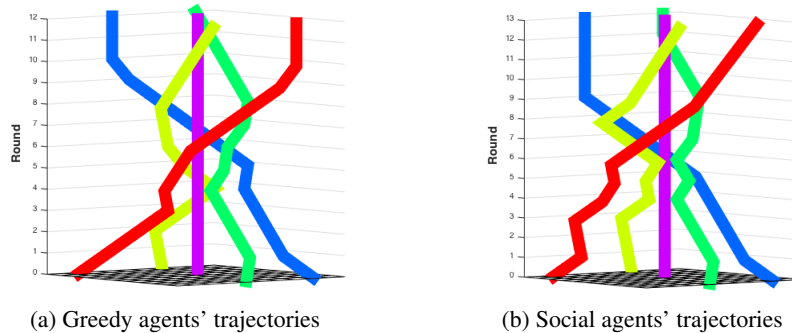


Fig. 10: Comparative plots of trajectories for a scenario involving an obstacle (purple trajectory) for greedy (Fig. 10a) and social agents (Fig. 10b).

Acknowledgement

This material is based upon work supported by the National Science Foundation under Grant No. 1526035. We are grateful for this support.

Bibliography

- [1] E. Artin. Theory of braids. *Annals of Mathematics*, 48(1):pp. 101–126, 1947.
- [2] M. Bennewitz, W. Burgard, G. Cielniak, and S. Thrun. Learning motion patterns of people for compliant robot motion. *Int. J. of Robotics Res.*, 24:31–48, 2005.
- [3] J. S. Birman. *Braids Links And Mapping Class Groups*. Princeton University Press, 1975.
- [4] S. Bonneaud and W. H. Warren. An empirically-grounded emergent approach to modeling pedestrian behavior. In *Pedestrian and Evacuation Dynamics 2012*, pages 625–638. Springer International Publishing, 2014.
- [5] G. Csibra and G. Gergely. The teleological origins of mentalistic action explanations: A developmental hypothesis. *Developmental Science*, 1(2):255–259, 1998.
- [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, Jan. 2007.
- [7] Y. Diaz-Mercado and M. Egerstedt. Interactions in multi-robot systems using braids. *Robotics, IEEE Transactions on*, 2016. Under Review.
- [8] A. D. Dragan and S. Srinivasa. Integrating human observer inferences into robot motion planning. *Auton. Robots*, 37(4):351–368, 2014.
- [9] D. Helbing and P. Molnár. Social force model for pedestrian dynamics. *Phys. Rev. E*, 51: 4282–4286, May 1995.
- [10] P. Henry, C. Vollmer, B. Ferris, and D. Fox. Learning to navigate through crowded environments. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2010.

- [11] S. Hoogendoorn and P. H. Bovy. Simulation of pedestrian flows by optimal control and differential games. *Optimal Control Applications and Methods*, 24(3):153–172, 2003.
- [12] I. Karamouzas, B. Skinner, and S. J. Guy. Universal power law governing pedestrian interactions. *Phys. Rev. Lett.*, 113:238701, Dec 2014.
- [13] R. A. Knepper and D. Rus. Pedestrian-inspired sampling-based multi-robot collision avoidance. In *RO-MAN*, pages 94–100. IEEE, 2012.
- [14] M. Kuderer, H. Kretschmar, C. Sprunk, and W. Burgard. Feature-based prediction of trajectories for socially compliant navigation. In *Proc. of Robotics: Science and Systems (RSS)*, Sydney, Australia, 2012.
- [15] M. Moussaïd, D. Helbing, and G. Theraulaz. How simple rules determine pedestrian behavior and crowd disasters. *Proceedings of the National Academy of Sciences*, 108(17):6884–6888, Apr. 2011.
- [16] J. J. Park, C. Johnson, and B. Kuipers. Robot navigation with model predictive equilibrium point control. In *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, pages 4945–4952. IEEE, 2012.
- [17] S. Pellegrini, A. Ess, K. Schindler, and L. J. V. Gool. You’ll never walk alone: Modeling social behavior for multi-target tracking. In *ICCV*, pages 261–268. IEEE Computer Society, 2009.
- [18] E. A. Sisbot, L. F. Marin-Urias, R. Alami, and T. Siméon. A human aware mobile robot motion planner. *IEEE Transactions on Robotics*, 23(5):874–883, 2007.
- [19] J.-L. Thiffeault. Braids of entangled particle trajectories. *Chaos*, 20(1), 2010.
- [20] P. Trautman, J. Ma, R. M. Murray, and A. 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.
- [21] A. Treuille, S. Cooper, and Z. Popović. Continuum crowds. *ACM Transactions on Graphics*, 25(3):1160–1168, July 2006.
- [22] J. van den Berg, S. J. Guy, M. C. Lin, and D. Manocha. Reciprocal n-body collision avoidance. In *Robotics Research - The 14th International Symposium, ISRR 2009, August 31 - September 3, 2009, Lucerne, Switzerland*, pages 3–19, 2009.
- [23] N. H. Wolfinger. Passing Moments: Some Social Dynamics of Pedestrian Interaction. *Journal of Contemporary Ethnography*, 24(3):323–340, 1995.
- [24] J. Yen. Finding the k shortest loopless paths in a network. *management Science*, pages 712–716, 1971.
- [25] B. Zhou, X. Tang, and X. Wang. Learning collective crowd behaviors with dynamic pedestrian-agents. *International Journal of Computer Vision*, 111(1):50–68, 2015.
- [26] B. D. Ziebart, N. Ratliff, G. Gallagher, C. Mertz, K. Peterson, J. A. Bagnell, M. Hebert, A. K. Dey, and S. Srinivasa. Planning-based prediction for pedestrians. In *Proc. of the International Conference on Intelligent Robots and Systems*, 2009.
- [27] G. M. Ziegler. *Lectures on polytopes*. Springer-Verlag, New York, 1995.