

Implicit Multiagent Coordination at Uncontrolled Intersections via Topological Braids

Christoforos Mavrogiannis¹, Jonathan DeCastro², and Siddhartha Srinivasa¹

¹ Paul G. Allen School of Computer Science & Engineering,
University of Washington, Seattle, WA 98195, USA
`{cmavro,siddh}@cs.washington.edu`

² Toyota Research Institute, Cambridge, MA 02139, USA
`jonathan.decastro@tri.global`

Abstract. We focus on decentralized navigation among multiple non-communicating agents at uncontrolled street intersections. Avoiding collisions under such settings demands nuanced *implicit* coordination. This is challenging to accomplish; the high dimensionality of the space of possible behavior and the lack of explicit communication among agents complicate prediction and planning. However, the structure of these domains often collapses the space of possible collective behavior into a finite set of *modes*. Our key insight is that enabling agents to reason about modes may enable them to coordinate *implicitly* via intent signals encoded in their actions. In this paper, we represent modes as low-dimensional multiagent motion primitives in a compact and interpretable fashion using the formalism of topological braids. Based on this representation, we derive a probabilistic model that maps past behavior of multiple agents to a future mode. Using this model, we design a decentralized control algorithm that treats navigation as uncertainty minimization over the space of modes. This algorithm enables agents to collectively reject unsafe intersection crossing strategies in a distributed fashion. We demonstrate our approach in a simulated four-way uncontrolled intersection. Our model is shown to reduce the frequency of collisions by over 65% against baselines explicitly reasoning in the space of trajectories, while maintaining comparable time efficiency in challenging scenarios.

Keywords: Multi-Robot Systems; Distributed Control; Topological Braids

1 Introduction

Street environments such as intersections often feature significant spatial structure through crosswalks, sidewalks or dedicated lanes. However, due to driver-to-driver variability, local customs and inconsistencies in the placement of signs and traffic lights, they do not always feature concrete mechanisms for organizing traffic flows temporally. For instance, street intersections lacking traffic lights and signage is a situation most drivers have encountered, and is, in fact, prevalent in developing countries [24]. While standard means of signaling such as turn

signals, horns, or even gaze, gestures and verbal communication could be useful under such circumstances, relying solely on these cues poses significant risk.

The effects of wrongly responding at intersections vary from inefficient intersection crossings (traffic backups) to catastrophic collisions. For reference, in the United States, during the year 2018, 43.7% of motor vehicle crashes occurred at intersections (2,943,717 out of 6,734,416 incidents). Out of these, 8,245 incidents involved fatalities, representing the 24.5% of all fatal crashes for the same year [23]. While the circumstances of each accident may differ, we view the high uncertainty and lack of timely coordination as major contributing factors to this sad reality.

Motivated by these observations, in this paper we consider the model setup of an uncontrolled, four-way street intersection [22] where multiple rational¹, non-communicating but perfectly-observing agents navigate in close proximity. The lack of explicit communication among agents results in high uncertainty which complicates decision making. However, the spatial structure of the environment and the assumed rationality of agents tend to collapse multiagent behavior to a discrete set of *modes*. Our key insight is that instead of attempting to predict agents’ future behavior in the form of Cartesian trajectories—which comes with high requirements in sample or runtime complexity—it might often be sufficient to imbue agents with a model of mode prediction. For a example, in a simple scenario involving two vehicles, an agent could reason about the likelihood passing before or after the other.

In prior work [17], we formalized modes as topological *braids* [1] and showed that this representation can be used to analyze traffic in complex, real-world scenes [16]. Building upon that work, in this paper, we develop a probabilistic inference mechanism that predicts future multiagent behavior, represented as a braid, given a history of agents’ Cartesian trajectories. We integrate this mechanism into an optimization-based reactive control algorithm that selects actions balancing uncertainty reduction over a mode, and safety (see Fig. 5). Through a simulation study, we compare our algorithm against a set of baselines that reason directly over the space of trajectories (Sec. 6). We demonstrate that our framework enables *multiple* non-communicating agents to coordinate implicitly and follow significantly safer paths (at least 65% fewer collisions) across a series of challenging intersection-crossing scenarios. Our findings suggest that incorporating topological features in the decision making process of non-communicating

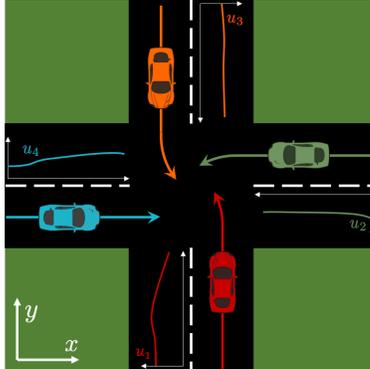


Fig. 1: We study a minimalistic setup of decentralized navigation at uncontrolled intersections. By judiciously adjusting their speeds, agents may reach consensus over safe intersection traversals. We enable agents to visualize possible types of traversal using the representation of topological braids.

¹ Agents that intend to follow time-efficient and collision-free paths.

agents enables effective coordination even in the absence of explicit communication or sophisticated trajectory reconstruction.

2 Related Work

In recent years, the robotics community has paid considerable attention to problems in multiagent driving domains.

The complexity of real-world traffic has motivated foundational work on abstractions of multiagent behavior. Pierson et al. [25] propose a congestion cost function that enables agents to plan lane changes within desired risk levels. Some works have contributed discrete, semantic representations of traffic. Wang et al. [34] classify discrete driving styles using multi-dimensional time series analysis and Gadepally et al. [7] use Hidden Markov Models to estimate long-term driver behaviors. Other approaches extract symbolic representations of traffic behavior from data [30] or leveraging properties like topological invariance [15, 26]. Finally, a body of work focuses on developing tools for testing and validation tailored to autonomous driving applications. Tian et al. [33] propose a verification testbed for navigation algorithms using tools from game theory whereas Liebenwein et al. [14] apply compositional and contract-based principles to driving controller verification. In prior work, we proposed a framework for characterizing the complexity of traffic datasets using topological braids [16].

In parallel, there has been rich interest in algorithmic approaches for multiagent driving domains. Many works study mechanisms of coordination. For example, Sadigh et al. [27] plan intent-expressive maneuvers that reinforce safe and efficient coordination in mixed traffic scenarios whereas Lazar et al. [13] plan optimal lane changes that reinforce prosocial behaviors such as platooning to increase capacity in congested highways. Some works focus on centrally managed intersections. Buckman et al. [5] plan vehicle rearrangements using a social psychology metric to reduce system delays in centrally managed intersections whereas Miculescu and Karaman [20] develop a centralized control framework with safety and efficiency guarantees for continuous car flows at an unsignalized intersection. Many works apply tools from belief space planning to the problem of safe lane merging [2, 4, 8, 9, 29]. Finally, some approaches focus on dealing with occlusions and faulty perception using probabilistic modeling [19] or deep reinforcement learning [11].

Building upon recent work on the use of topological braids [3, 16, 17] as a symbolic abstraction of multiagent behavior in driving domains, we derive a probabilistic framework for decision making in complex multiagent scenes with no explicit communication among agents. Unlike prior work [17] which considered simplified, obstacle-free domains and holonomic agents [18], in this paper we consider a continuous workspace with a realistic street intersection structure and nonholonomically-constrained agents. Using first principles of probability theory, we develop a Bayesian framework that enables an agent to reason about how its own local decisions impact global interactions with others in tight driving environments like four-way intersections. We then treat decentralized multiagent

navigation as uncertainty reduction over the space of collective navigation strategies, demonstrating safer intersection crossing behaviors compared to baselines performing inference in the trajectory space.

3 Problem Statement

Consider the uncontrolled street intersection of Fig. 1 where $n > 1$ non communicating agents with car-like kinematics are navigating. Denote by $q_i = (x_i, y_i, \theta_i) \in \mathcal{Q} \subseteq SE(2)$ the state of agent $i \in \mathcal{N} = \{1, \dots, n\}$ with respect to (wrt) a fixed reference frame, defined by a basis $(\hat{x}, \hat{y}, \hat{t})$. Each agent i starts from an initial state $s_i \in \mathcal{Q}$, lying on a side of the intersection, and moves towards a final –unknown to others– state $d_i \in \mathcal{Q}$ lying on a different side. They do so by tracking a path $\tau_i : I \rightarrow \mathcal{Q}$, for which it holds that $\tau_i(0) = s_i$ and $\tau_i(1) = d_i$, where $I = [0, 1]$ is a path parametrization. Observing the complete system state $Q = (q_1, \dots, q_n) \in \mathcal{Q}^n$, agent i tracks τ_i by executing a policy $\pi_i : \mathcal{Q} \rightarrow \mathcal{U}$, generating actions $u_i \in \mathcal{U}$ (speed and steering angle), where $\mathcal{U} \subseteq \mathbb{R} \times \mathbb{S}$ is a space of controls. Agent i is not aware of the intended path τ_j , the destination d_j or the exact policy π_j of any other agent $j \neq i \in \mathcal{N}$ but is able to perfectly observe their state at every timestep.

In this paper, we design a policy π_i with the goal of enabling agents to fluently coordinate collision-free intersection crossings while following time-efficient trajectories under uncertainty in a distributed fashion.

4 Preliminaries

Our approach is based on coupling individual agents’ behaviors into a joint representation. Collectively, agents track a system path $T = (\tau_1, \dots, \tau_n)$ by executing a system trajectory $\Xi = (\xi_1, \dots, \xi_n)$, where $\xi_i : [0, t_\infty] \rightarrow \mathcal{Q}$ is the trajectory of agent i , and t_∞ is a terminal time corresponding to when the last agent reaches its destination. The trajectory Ξ can be abstracted into a topological braid [3] using our method from prior work [16]. We quickly recap braids and the process for transitioning from a Cartesian trajectory into a braid.



(a) Generators of B_n .

(b) Example of composition.

Fig. 2: Presentation of the braid group, B_n .

4.1 Topological Braids

A braid is a tuple $b_f = (f_1, \dots, f_n)$ of functions $f_i : I \rightarrow \mathbb{R}^2 \times I$, $i \in \mathcal{N}$, defined on a domain $I = [0, 1]$ and embedded in a Cartesian space $(\hat{x}, \hat{y}, \hat{t})$. These functions, called *strands*, are monotonically increasing along the \hat{t} direction, satisfying the properties: (a) $f_i(0) = (i, 0, 0)$, and $f_i(1) = (p_f(i), 0, 1)$, where $p_f : \mathcal{N} \rightarrow \mathcal{N}$ is a permutation in the symmetric group N_n ; (b) $f_i(t) \neq f_j(t) \forall t \in I, j \neq i \in \mathcal{N}$. Two braids, $b_f = (f_1, \dots, f_n)$, $b_g = (g_1, \dots, g_n)$, can be composed through a *composition* operation (Fig. 2b): their composition, $b_h = b_f \cdot b_g$, is also a braid $b_h = (h_1, \dots, h_n)$, comprising a set of n curves, defined as:

$$h_i(t) = \begin{cases} f_i(2t), & t \in \left[0, \frac{1}{2}\right] \\ g_j(2t - 1), & t \in \left[\frac{1}{2}, 1\right] \end{cases}, \quad (1)$$

where $j = p_f(i)$. The set of all braids on n strands, along with the composition operation form a group, B_n , called the Braid group on n strands. Following Artin's presentation [1], the braid group B_n can be generated from $n - 1$ primitive braids $\sigma_1, \dots, \sigma_{n-1}$ (see Fig. 2a), called generators, and their inverses, via composition. Any braid can be written as a *word*, i.e., a product of generators and their inverses (Fig. 2b), satisfying the *relations*:

$$\begin{aligned} \sigma_i \sigma_j &= \sigma_j \sigma_i, \quad |j - i| > 1, \\ \sigma_i \sigma_{i+1} \sigma_i &= \sigma_{i+1} \sigma_i \sigma_{i+1}, \quad \forall i. \end{aligned} \quad (2)$$

4.2 Transforming Traffic Trajectories into Braids

In prior work [16], we showed that any collection of trajectories in a traffic scene can be represented as a topological braid. We also described a method for converting traffic trajectories into braids, inspired by the technique of Thiffeault [31]. In brief, the technique involved: a) projecting trajectories onto a selected spacetime plane; b) labeling any trajectory crossings that emerge in the projection as braid generators and their inverses by identifying *under* or *over* crossings (Fig. 3d); c) forming a *braid word* by arranging them in temporal order.

This technique enables the summarization of a traffic episode as an object with symbolic and geometric

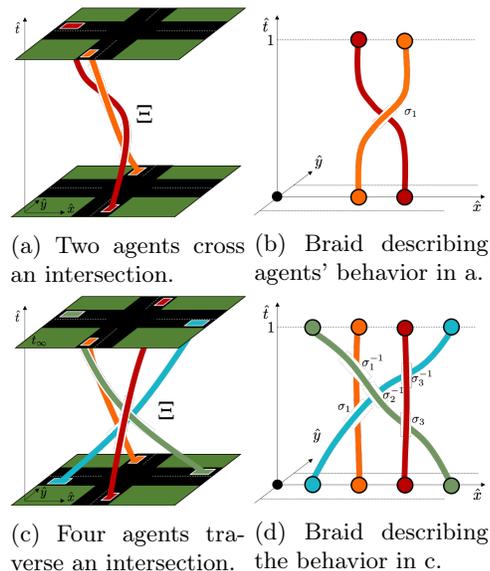


Fig. 3: Vehicle trajectories (left) are summarized as topological braids (right).

descriptions. For example, in Fig. 3d, the braid $\sigma_3\sigma_1\sigma_2^{-1}\sigma_3^{-1}\sigma_1^{-1} \in B_4$ is a summary of how agents coordinated to avoid each other throughout the episode. This braid represents a *mode* of behavior involving the four agents. In general, in a scene with n agents, the braid group B_n represents the set of all modes of multiagent behavior that are likely. As shown [16], the braid representation *compresses* multiagent behavior into a low-dimensional compact object that retains salient properties of interaction.

5 Decentralized Navigation as Braid Prediction

We describe a probabilistic model that links past agents’ trajectories to a braid representing the spatiotemporal entanglement of their future trajectories at an intersection domain. Fig. 4 illustrates the setup of the proposed model in a four-agent scenario. Based on this mechanism, we build an optimization-based control scheme for decentralized navigation at uncontrolled intersections.

5.1 Reasoning about Braids of Multiagent Interaction

At time $t \in [0, t_\infty]$, agent i , having access to the complete system state history so far, Ξ , maintains a belief $bel_i = P(\beta_i|\Xi)$ over the braid $\beta_i \in B_n$ that describes the topology of the emerging (future) system trajectory $\Xi' = \Xi_{t \rightarrow \infty}$ from the perspective³ of agent i . The braid β_i depends on agents’ intended system path T . To capture this dependency, we marginalize over \mathcal{T}_i , the subset of the set of system paths, \mathcal{T} , for which agent i (the ego agent) follows its intended path:

$$\begin{aligned} bel_i &= P(\beta_i|\Xi) \\ &= \sum_{T \in \mathcal{T}_i} P(\beta_i|\Xi, T)P(T|\Xi). \end{aligned} \quad (3)$$

For a given system path T , different braids could possibly emerge, depending on the path tracking behavior of agents. To capture this dependency, we marginalize the probability $P(\beta_i|\Xi, T)$ over the control profile $U \in \mathcal{U}^n$ that could be taken by agents at the current time step:

$$P(\beta_i|\Xi, T) = \sum_{U \in \mathcal{U}^n} P(\beta_i|\Xi, U, T)P(U|\Xi, T). \quad (4)$$

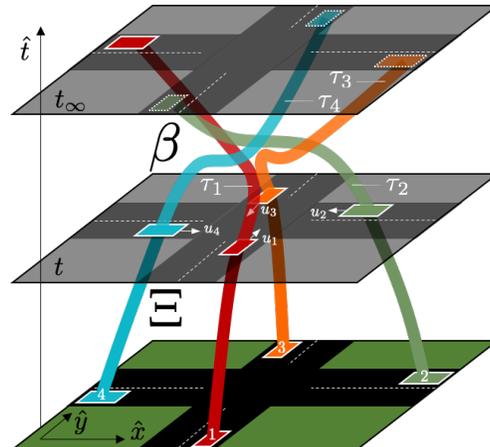


Fig. 4: Topological inference. At time t , given state history Ξ , the ego agent (red), following path τ_1 , predicts the topology β of the unfolding multiagent interaction.

³ Each agent uses a distinct projection plane to define their own braid set.

Substituting to eq. (3), we get:

$$bel_i = \sum_{\tau_i} \left\{ \sum_{U^n} P(\beta_i|\Xi, U, T)P(U|\Xi, T) \right\} P(T|\Xi). \quad (5)$$

Eq. (5) combines a local action selection model $P(U|\Xi, T)$ with a model of intent inference $P(T|\Xi)$ and a global behavior prediction model $P(\beta_i|\Xi, U, T)$.

The intention of agent $j \neq i$ over a path τ_j is conditionally independent of the intention of any other agent, given the past system trajectory Ξ . The probability over the path intention of agent j does not depend on the trajectories of others. Thus, we simplify the computation of the system path prediction as:

$$P(T|\Xi) = \prod_{j \in \mathcal{N} \setminus i} P(\tau_j|\xi_j), \quad (6)$$

where the product only considers the probabilities over the paths of others, since agent i is certain about their own path.

Similarly, since agents select a control input independently, without having access to the policies of others, we decompose the computation of the control profile prediction as:

$$P(U|\Xi, T) = \prod_{i=1}^n P(u_i|\Xi, T), \quad (7)$$

where the distribution $P(u_i|\Xi, T)$ represents the control input that agent i executes to make progress along its path τ_i , incorporating considerations such as preferred navigation velocity and a local tracking controller class.

The model of inference of eq. (5) focuses on topology prediction, without considerations of collision avoidance. To filter out unsafe braids, we redefine eq. (5) by incorporating a model of collision prediction. Define by c a boolean random variable representing the event that Ξ' , the emerging future trajectory contains collisions (**true** for a collision, **false** for no-collision). Denote by $\tilde{\beta} = (\beta_i, \neg c)$ the joint event that Ξ' is both topologically equivalent, i.e., ambient-isotopic [21] to a braid $\beta_i \in B_n$, and not in collision, i.e., c is **false**. Then the belief, \tilde{bel}_i , of agent i that $\tilde{\beta}_i$ is true can be computed as:

$$\tilde{bel}_i = P(\tilde{\beta}_i|\Xi) = \sum_{\tau} \left\{ \sum_U P(\tilde{\beta}_i|\Xi, U, T)P(U|\Xi, T) \right\} P(T|\Xi). \quad (8)$$

The occurrence of a collision is conditionally independent of the emerging braid—the braid only describes the topological pattern of the trajectories, ignoring any geometric intersections among the volumes of the vehicles. Thus, we may compute their joint distribution as:

$$\begin{aligned} P(\tilde{\beta}_i|\Xi, U, T) &= P(\tilde{\beta}_i, \neg c|\Xi, U, T) \\ &= P(\neg c|\Xi, U, T)P(\beta_i|\Xi, U, T) \\ &= (1 - P(c|\Xi, U, T))P(\beta_i|\Xi, U, T). \end{aligned} \quad (9)$$

5.2 Decision Making

An outcome $\tilde{\beta}_i$ represents a class of trajectories Ξ_{β_i} that are topologically equivalent to the braid β_i and not in collision. During execution, the distribution

$P(\tilde{\beta}_i|\Xi, U, T)$ is reshaped as a result of agents' decisions. Our approach contributes towards a minimum-entropy shape of $P(\tilde{\beta}_i|\Xi, U, T)$, which corresponds to a state of consensus over a braid $\tilde{\beta}_i$ from the perspective of agent i . We do so through the following receding-horizon control scheme:

$$u_i^* = \arg \min_{u_i \in \mathcal{U}} H(\tilde{\beta}_i), \quad (10)$$

where

$$H(\tilde{\beta}_i) = - \sum_{B_n} P(\tilde{\beta}_i|\Xi) \log P(\tilde{\beta}_i|\Xi), \quad (11)$$

is the information entropy of $P(\tilde{\beta}_i|\Xi)$, representing agent i 's uncertainty over a solution $\tilde{\beta}_i$ where $P(\tilde{\beta}_i|\Xi)$ is recovered using eq. (8). This optimization scheme contributes uncertainty-reducing actions over the emerging outcome $\tilde{\beta}_i$. Note that the use of the information entropy as a cost reflects the insight that multiple elements of B_n could be valid solutions to the collision avoidance problem at a given instance. The ego agent behavior could still contribute to collision-free navigation even when there is not a unique winner within B_n , a strategy that has been successfully applied to domains like shared control [12].

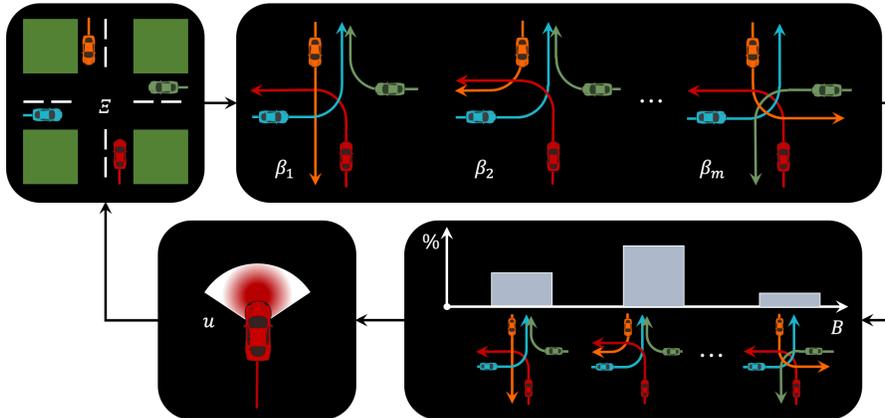


Fig. 5: Decision-making scheme. At every cycle, the ego agent forward simulates a set of distinct futures, classifies them into topological outcomes, and selects the action that minimizes the uncertainty over such outcomes.

6 Application

We deploy our decision-making mechanism in a simulated study in an uncontrolled intersection with multiple (n) vehicles.

6.1 Setup

Our setup is the 4-way symmetric intersection of Fig. 1. Each lane at the intersection is $50m$ long and $3.6m$ wide, whereas each car is represented as a rectangle

with a length of $4.7m$ and a width of $1.7m$. We assume that any side a is connected to any other side $b \neq a$ with a unique, publicly known legal path τ_{ab} , lying along the middle of the lane. We assume that any agent $i \in \mathcal{N}$ that attempts to reach side b from side a will attempt to track this path, τ_{ab} .

Each agent follows a path out of three options (left, right, or straight). Thus, agent i needs to consider $|\mathcal{T}_i| = 3^{n-1}$ possible system paths, extracted upon iterating over all possible combinations for other agents' paths. We consider a trial to be split in two phases: (a) the *negotiation* phase, which corresponds to the initial straight-path part of the intersection (denoted as \mathcal{Q}_i^{neg} for agent i), within which the agent attempts to reach a consensus with others wrt a joint strategy of collision avoidance; (b) the *execution* phase, which corresponds to the rest of the path (denoted as \mathcal{Q}_i^{exec} for agent i), within which the agent tracks the remainder of its path, by maintaining a constant speed. This decision emphasizes the importance of proactive negotiation during the first portion, and provides a natural metric of quality: the count of collisions during the execution part.

6.2 Models

We describe models for the components of eq. (11).

Intention prediction. We assume that agent i has no knowledge of the path τ_j of any other agent $j \neq i \in \mathcal{N}$ while j is in the negotiation stage. However, we assume that τ_j becomes immediately obvious to agent i when agent j enters the intersection:

$$P(\tau_j | \xi_j(t)) = P(\tau_j | q_j) = \begin{cases} 1/m & \text{for } q_j \in \mathcal{Q}_j^{neg} \\ 1 & \text{for } q_j \in \mathcal{Q}_j^{exec}, \end{cases} \quad (12)$$

where $q_j = \xi_j(t)$ is agent j 's current state, and $m = 3$ is the number of paths that agent j selects from.

Behavior prediction. Here, $P(u_j | \Xi, T)$ expresses the belief of agent i that agent $j \in \mathcal{N}$ will execute a control input u_j in the next timestep. We assume that all agents employ the same PID controller converting a desired speed ν_j , drawn from a set \mathcal{V}_j into a control input u_j towards the next waypoint along their path, τ_j . We also assume that \mathcal{V}_j contains two speeds: ν_j^{low} and ν_j^{high} . Agents generally prefer the high speed, which we encode in a distribution $P(\nu_j | \xi_j, \tau_j) = P(\nu_j)$. For each agent, $P(\nu_j = \nu_j^{high})$ is sampled at random from the range $[0.6, 0.8]$ and remains fixed throughout the execution. Note that a speed ν_j is deterministically mapped to a control input u_j through the low-level controller, thus $P(u_j | \xi_j, \tau_j) = P(\nu_j)$. Each agent i has a noisy estimate about the speed preferences of others: it assumes that others have the same exact preferences.

Topology prediction. To extract the probability of a topological outcome, agent i forward simulates each path set $T \in \mathcal{T}_i$ a total of 2^n times, each corresponding to a unique assignment of speeds (drawn from \mathcal{V}) to agents that are held constant throughout execution. For each path set / speed combination, agent i extracts a future system trajectory Ξ' and extracts a corresponding braid word β_i by projecting Ξ' onto a selected plane, as described in Sec. 4.2. For convenience and generality, each agent uses a distinct projection plane defined by a

local x axis (see Fig. 4), and the time axis. This process results in a set $B \subset B_n$ containing the set of all braids that could be realizable in the remainder of the execution. We model the probability that a future trajectory is equivalent to a braid $\beta^* \in B$ as follows:

$$P(\beta_i = \beta^* | \Xi, U, T) = \frac{1}{z} \sum \mathbb{1}(\Xi', \beta^*), \quad (13)$$

where the indicator function $\mathbb{1}(\Xi, \beta^*) = 1$ if a trajectory Ξ' is topologically equivalent to the braid β^* and 0 otherwise, and z is a normalizer across B . We perform the above computations using the Braidlab [32] package.

Collision prediction. During the forward simulations detailed above, for each trajectory Ξ' we compute a minimum inter-agent distance d_{min} . We model the probability of a collision $P(c | \Xi, U, T)$ as follows:

$$P(c | \Xi, U, T) = \frac{1}{1 + e^{a(d_{min} - \delta)}}, \quad (14)$$

where a (set to 10) controls the rate of change of the function and δ (set to $15m$) denotes a threshold distance beyond which collision is imminent. According to this model, the smaller d_{min} is, it is exponentially more likely to have a collision.

6.3 Experiment Design

We define three scenarios, involving 2, 3, and 4 agents respectively, designed to give rise to challenging interactions among agents. For each scenario, we generate a set of randomized experiments by varying agents' speed preferences. We execute each experiment under 5 conditions, each corresponding to a distinct variation of the proposed algorithm, executed by all agents.

Scenarios

S2: Two agents, starting from the bottom / right sides of the intersection, are moving straight towards the top and left sides respectively. They both draw speeds from \mathcal{U}_{s2} containing 12 evenly spaced speeds within $[5, 10]$ (m/s). We generate 144 experiments corresponding to the Cartesian product \mathcal{U}_{s1}^2 .

S3: Three agents, starting from the bottom, right and top are moving straight towards the top, left and bottom sides respectively. They draw speeds from \mathcal{U}_{s3} , containing 5 evenly spaced speeds within the range $[5, 10]$ (m/s). We generate 125 experiments corresponding to \mathcal{U}_{s2}^3 .

S4: Four agents, starting from the bottom, right, top, and left, are moving straight towards the top, left, bottom, and right sides respectively. They draw speeds from \mathcal{U}_{s4} , containing 3 evenly spaced speeds within the range $[5, 10]$ (m/s). We generate 81 experiments corresponding to \mathcal{U}_{s3}^4 .

Conditions

C1: Agents track their desired paths with their desired speeds, without accounting for avoiding collisions with others. This condition serves as a reference of the intensity of interactions for each scenario.

C2: Our complete proposed algorithm from eq. 10.

C3: A modification of our proposed algorithm that assumes knowledge of the paths that other agents are following, i.e, they replace eq. (8) with

$$b\tilde{e}l_i = \sum_U P(\tilde{\beta}_i|\Xi, U, T)P(U|\Xi, T). \quad (15)$$

C4: A variation of C2 that does not use braids for clustering trajectory sets. Specifically, agents reason about the emerging collision-free system trajectory $\tilde{\Xi}_i$ (instead of $\tilde{\beta}_i$), replacing eq. (8) with

$$b\tilde{e}l_i = P(\tilde{\Xi}_i|\Xi, U, T)P(U|\Xi, T)P(T|\Xi). \quad (16)$$

C5: A modification of C4 that assumes knowledge of the paths that others are following, i.e, C5 replaces eq. (16) with

$$b\tilde{e}l_i = P(\tilde{\Xi}_i|\Xi, U, T)P(U|\Xi, T). \quad (17)$$

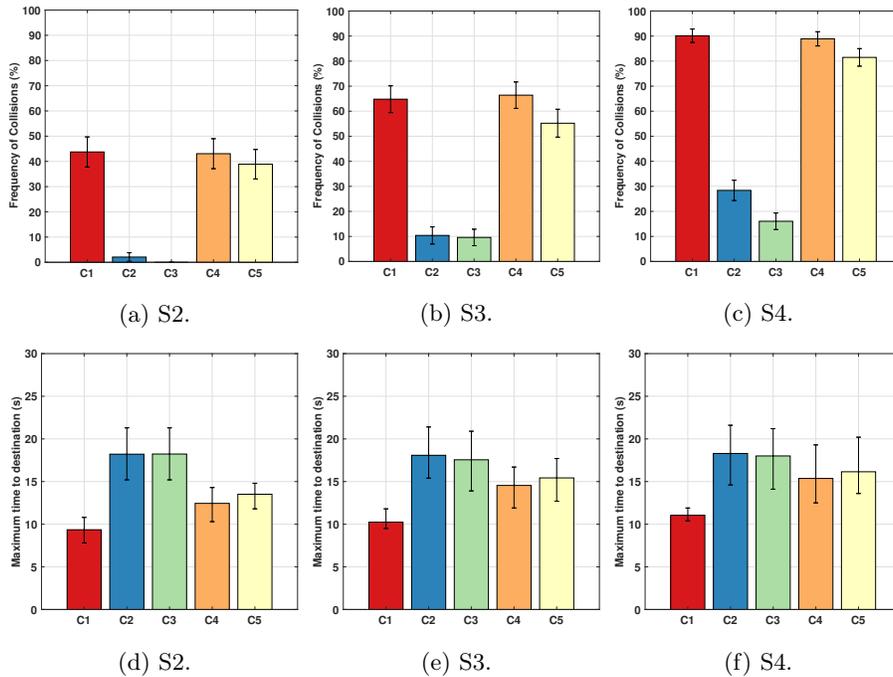


Fig. 6: Performance under homogeneous settings across conditions and scenarios. S2, S3, and S4 denote scenarios with 2, 3, and 4 agents respectively. Bars represent means; errorbars indicate standard deviations and 25/75 percentiles for the collision frequency (top) and time charts (bottom) respectively.

6.4 Results

Fig. 6 depicts performance under the 5 conditions across the 3 scenarios considered. As a reference, C1 scores the highest collision frequency and lowest time

to destination, serving as a characterization of the hardness of the selected scenarios. Overall, we observe that as the number of agents increases, the collision frequency increases as well. We also see that the topology-based approaches (C2, C3) generally outperform the trajectory-based ones (C4, C5) in terms of collision frequency. As expected, knowledge of agents’ intended paths generally reduces collisions.

Our approach (C2) achieves consistently low collision frequency across all scenarios. Compared to C4, C2 reduces collision frequency by: 95% in S2 (Fig. 6a); 65% in S3 (Fig. 6b); 66% in S4 (Fig. 6a). The price that C2, and C3 pay is the increased maximum time to destination. We note however that for the more complex scenarios (S3, S4), the time difference is not statistically significant (Fig. 6e, Fig. 6f).

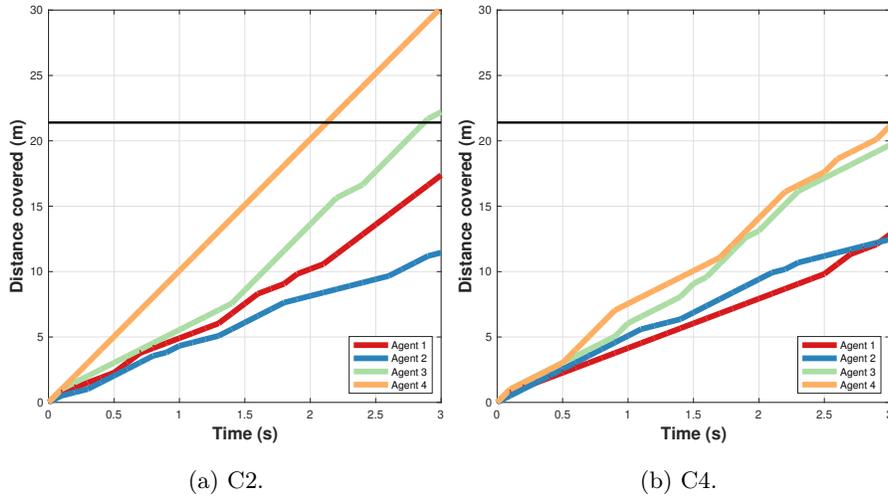


Fig. 7: Distance covered per agent over the first 3s of execution within a 4-agent experiment under C2 and C4. The black line indicates arrival at the intersection.

We attribute the performance gains of topology-based baselines to the incorporation of multiagent interaction modeling into agents’ decision making. Trajectory-based approaches ignore the existence of the domain’s structure which could be acquired through our proposed topological partitioning using braids. The braid group represents the set of distinct modes that could describe the collective motion of navigating agents. Explicitly reasoning about these modes enables a rational agent to anticipate the effect of its actions on system behavior. Our policy outputs local actions of global outlook that contribute towards reducing uncertainty over the emerging mode. Collectively, this results in implicit coordination, reflected in the reduced collision frequency of C2, C3. To illustrate this point, Fig. 7 depicts a comparative qualitative example of the behaviors generated by our policy. For the same experiment from S4, we observe that C2 agents (Fig. 7a) quickly converge to a clear order of intersection crossings as a

result of their proactive decision making. On the other hand, C4 agents (Fig. 7b), lacking the ability of modeling the complex multiagent dynamics, appear unable to coordinate their crossings and end up colliding.

7 Discussion

Our findings illustrate that the principled domain knowledge induced by the braid abstraction of multiagent interaction may enable superior performance than baselines reasoning directly in the vast space of Cartesian trajectories. Our setup is deliberately simplified to facilitate the extraction of foundational insights. Guided by the presented insights, we see the foundations of our approach as relevant to the areas of behavior prediction and decision making for multiagent navigation. For example, we anticipate that reasoning about the spatiotemporal topology of multiagent behavior could likely improve the performance of trajectory prediction models [28], belief-space approaches [4], reinforcement learning techniques [10] in traffic scenarios.

One interesting observation from our experiments is that our agents (C2) generally manage to coordinate collision-free intersection crossings even though they use different braid representations: each agent represents a space of modes by considering a different braid projection plane (their local \hat{x} - \hat{t} plane) which is unknown to others. This indicates that agents converge to the underlying multiagent behavior topology through our inference mechanism, even though they tend to use different *language* to represent the same topological events.

Another observation worth noting is that agents do not know the true speed preferences of others; they only maintain a noisy estimate. Despite that fact, they are able to avoid collisions to a significant extent by following the collaborative probabilistic approach of rejecting unsafe intersection traversals of eq. (10).

On a similar note, it is important to highlight that while each agent is using a distinct projection plane to represent braids internally, the collaborative filtering process induced through the iterative entropy reduction results in convergence to safer intersection crossings. In other words, even though agents might represent same coordination strategies with *different* braid words, they are still able to follow collision-free intersection crossings more often.

7.1 Limitations

The braid representation significantly compresses the space of outcomes from the continuum of Cartesian trajectories into a discrete set of topological modes. This compression holds the potential for computation speedups for inference and decision making. However, in this paper we did not explicitly demonstrate such computation gains, as our prototype implementation of eq. (10) involved forward simulations iterating over paths and controls. We envision that we could replace costly forward simulations with fast-inference models of the underlying probability distributions acquired through data-driven techniques [18, 26].

While deliberately minimalistic, our setup is an oversimplified abstraction of real-world uncontrolled intersections. Agents’ behavioral models could be replaced by more realistic ones, drawn from simulators [6] or real-world datasets. Additional considerations towards realism would involve asymmetrical intersections or different types of scenarios like roundabouts or merging highway lanes which, as we showed in prior work [16], can also be abstracted into braids. Finally, the implementations of the probabilistic models used by the framework (intention, behavior, topology, and collision prediction) could be replaced by higher-fidelity approximations through the use of data-driven techniques.

Acknowledgements his work was partially funded by the National Institute of Health R01 (#R01EB019335), National Science Foundation CPS (#1544797), National Science Foundation NRI (#1637748), the Office of Naval Research, the RTCA, Amazon, and Honda Research Institute USA.

Bibliography

- [1] Emil Artin. Theory of braids. *Annals of Mathematics*, 48(1):pp. 101–126, 1947.
- [2] Tirthankar Bandyopadhyay, Kok Sung Won, Emilio Frazzoli, David Hsu, Wee Sun Lee, and Daniela Rus. Intention-aware motion planning. In *Proceedings of the International Workshop on the Algorithmic Foundations of Robotics (WAFR)*, pages 475–491. Springer Berlin Heidelberg, 2013.
- [3] Joan S. Birman. *Braids Links And Mapping Class Groups*. Princeton University Press, 1975.
- [4] Maxime Bouton, Akansel Cosgun, and Mykel J. Kochenderfer. Belief state planning for autonomously navigating urban intersections. In *Proceedings of the IEEE Intelligent Vehicles Symposium (IV)*, pages 825–830, 2017.
- [5] Noam Buckman, Alyssa Pierson, Wilko Schwarting, Sertac Karaman, and Daniela Rus. Sharing is caring : Socially-compliant autonomous intersection negotiation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6136–6143, 2019.
- [6] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. CARLA: An open urban driving simulator. In *Proceedings of the Conference on Robot Learning (CoRL)*, volume 78 of *Proceedings of Machine Learning Research*, pages 1–16, 2017.
- [7] Vijay Gadeppally, Ashok Krishnamurthy, and Ümit Özgüner. A Framework for Estimating Long Term Driver Behavior. *Journal of Advanced Transportation*, 2017:1–11, 2017.
- [8] Ya-Chuan Hsu, Swaminathan Gopalswamy, Srikanth Saripalli, and Dylan A. Shell. An mdp model of vehicle-pedestrian interaction at an unsignalized intersection. In *Proceedings of the IEEE Vehicular Technology Conference (VTC)*, pages 1–6, 2018.

- [9] Constantin Hubmann, Marvin Becker, Daniel Althoff, David Lenz, and Christoph Stiller. Decision making for autonomous driving considering interaction and uncertain prediction of surrounding vehicles. In *Proceedings of the IEEE Intelligent Vehicles Symposium (IV)*, pages 1671–1678, 2017.
- [10] David Isele, Akansel Cosgun, Kaushik Subramanian, and Kikuo Fujimura. Navigating intersections with autonomous vehicles using deep reinforcement learning. *CoRR*, abs/1705.01196, 2017.
- [11] David Isele, Reza Rahimi, Akansel Cosgun, Kaushik Subramanian, and Kikuo Fujimura. Navigating occluded intersections with autonomous vehicles using deep reinforcement learning. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 2034–2039, 2018.
- [12] Shervin Javdani, Henny Admoni, Stefania Pellegrinelli, Siddhartha S. Srinivasa, and J. Andrew Bagnell. Shared autonomy via hindsight optimization for teleoperation and teaming. *The International Journal of Robotics Research*, 37(7):717–742, 2018.
- [13] Daniel A. Lazar, Ramtin Pedarsani, Kabir Chandrasekher, and Dorsa Sadigh. Maximizing road capacity using cars that influence people. In *Proceedings of the IEEE Conference on Decision and Control (CDC)*, pages 1801–1808, 2018.
- [14] Lucas Liebenwein, Wilko Schwarting, Cristian-Ioan Vasile, Jonathan DeCastro, Javier Alonso-Mora, Sertac Karaman, and Daniela Rus. Compositional and contract-based verification for autonomous driving on road networks. In *Robotics Research*, pages 163–181, Cham, 2020. Springer International Publishing.
- [15] Christoforos Mavrogiannis and Ross A. Knepper. Hamiltonian coordination primitives for decentralized multiagent navigation. *The International Journal of Robotics Research*, 40(10-11):1234–1254, 2021.
- [16] Christoforos Mavrogiannis, Jonathan DeCastro, and Siddhartha S. Srinivasa. Analyzing multiagent interactions in traffic scenes via topological braids. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2022.
- [17] Christoforos I. Mavrogiannis and Ross A. Knepper. Multi-agent path topology in support of socially competent navigation planning. *The International Journal of Robotics Research*, 38(2-3):338–356, 2019.
- [18] 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 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6817–6824, 2017.
- [19] Stephen G. McGill, Guy Rosman, Teddy Ort, Alyssa Pierson, Igor Gilitschenski, Brandon Araki, Luke Fletcher, Sertac Karaman, Daniela Rus, and John J. Leonard. Probabilistic risk metrics for navigating occluded intersections. *IEEE Robotics and Automation Letters*, 4(4):4322–4329, 2019.
- [20] D. Miculescu and S. Karaman. Polling-systems-based autonomous vehicle coordination in traffic intersections with no traffic signals. *IEEE Transactions on Automatic Control*, 65(2):680–694, 2020.

- [21] Kunio Murasugi and Bohdan I. Kurpita. *A Study of Braids*. Mathematics and Its Applications. Springer Netherlands, 1999.
- [22] Wassim G. Najm, John D. Smith, and Mikio Yanagisawa. Pre-crash scenario typology for crash avoidance research. Technical Report DOT-HS-810 767, National Highway Transportation Safety Administration, 2007.
- [23] US Department of Transportation National Highway Traffic Safety Administration. Fatality analysis reporting system (fars) encyclopedia, 2018. URL <https://www-fars.nhtsa.dot.gov>. Retrieved: [01/24/2020].
- [24] Gopal R. Patil and Digvijay S. Pawar. Microscopic analysis of traffic behavior at unsignalized intersections in developing world. *Transportation Letters*, 8(3):158–166, 2016.
- [25] Alyssa Pierson, Wilko Schwarting, Sertac Karaman, and Daniela Rus. Navigating congested environments with risk level sets. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 5712–5719, 2018.
- [26] Junha Roh, Christoforos Mavrogiannis, Rishabh Madan, Dieter Fox, and Siddhartha S. Srinivasa. Multimodal trajectory prediction via topological invariance for navigation at uncontrolled intersections. In *Proceedings of the Conference on Robot Learning*, volume 155 of *Proceedings of Machine Learning Research*, pages 2216–2227, 2021.
- [27] Dorsa Sadigh, Nick Landolfi, Shankar S. Sastry, Sanjit A. Seshia, and Anca D. Dragan. Planning for cars that coordinate with people: leveraging effects on human actions for planning and active information gathering over human internal state. *Autonomous Robots*, 42(7):1405–1426, 2018.
- [28] Tim Salzmann, Boris Ivanovic, Punarjay Chakravarty, and Marco Pavone. Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 683–700, 2020.
- [29] Volkan Sezer, Tirthankar Bandyopadhyay, Daniela Rus, E. Frazzoli, and David Hsu. Towards autonomous navigation of unsignalized intersections under uncertainty of human driver intent. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3578–3585, 2015.
- [30] Shai Shalev-Shwartz, Shaked Shammah, and Amnon Shashua. Safe, Multi-Agent, Reinforcement Learning for Autonomous Driving. *arXiv:1610.03295*, 2016.
- [31] Jean-Luc Thiffeault. Braids of entangled particle trajectories. *Chaos*, 20(1), 2010.
- [32] Jean-Luc Thiffeault and Marko Budišić. Braidlab: A software package for braids and loops, 2013–2019. Version 3.2.4.
- [33] Ran Tian, Nan Li, Ilya Kolmanovsky, Yildiray Yildiz, and Anouck Girard. Game-theoretic Modeling of Traffic in Unsignalized Intersection Network for Autonomous Vehicle Control Verification and Validation. *IEEE Transactions on Intelligent Transportation Systems*, 23(3):2211–2226, 2022.
- [34] Wenshuo Wang, Weiyang Zhang, and Ding Zhao. Understanding V2V Driving Scenarios through Traffic Primitives. *IEEE Transactions on Intelligent Transportation Systems*, 23(1):610–619, 2022.