Analyzing Multiagent Interactions in Traffic Scenes via Topological Braids

Christoforos Mavrogiannis[1], Jonathan DeCastro[2], and Siddhartha S. Srinivasa[1]

***Abstract*— We focus on the problem of analyzing multiagent interactions in traffic domains. Understanding the space of behavior of real-world traffic may offer significant advantages for algorithmic design, data-driven methodologies, and benchmarking. However, the high dimensionality of the space and the stochasticity of human behavior may hinder the identification of important interaction patterns. Our key insight is that traffic environments feature significant geometric and temporal structure, leading to highly organized collective behaviors, often drawn from a small set of dominant modes. In this work, we propose a representation based on the formalism of topological braids that can summarize arbitrarily complex multiagent behavior into a compact object of dual geometric and symbolic nature, capturing critical events of interaction. This representation allows us to formally enumerate the space of outcomes in a traffic scene and characterize their complexity. We illustrate the value of the proposed representation in summarizing critical aspects of real-world traffic behavior through a case study on recent driving datasets. We show that despite the density of real-world traffic, observed behavior tends to follow highly organized patterns of low interaction. Our framework may be a valuable tool for evaluating the richness of driving datasets, but also for synthetically designing balanced training datasets or benchmarks.**

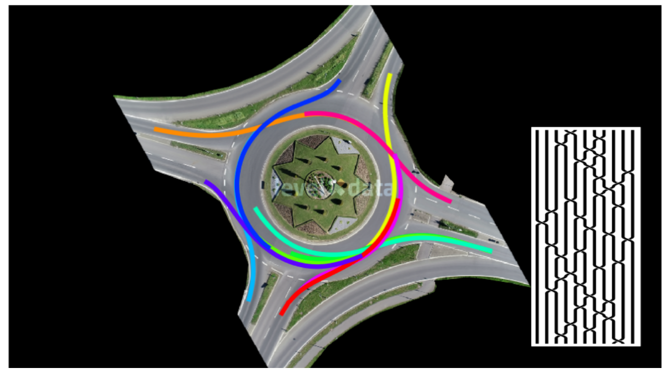

Fig. 1: This work proposes a formal framework for the characterization of multiagent behavior in driving domains. Complex multiagent interactions encountered in real-world driving domains such as a roundabout can be compactly represented as topological braids (right).

I. Introduction

Traffic scenes pose unique challenges for prediction and planning [33, 13, 32, 9, 22] due to their high dimensionality, the complexity of modeling human behavior, and the performance standards motivated by the safety-critical nature of the domain. Despite these complications, real-world traffic scenes often feature significant structure. Vehicles follow designated lanes, and traffic is regulated through traffic signals and signs or coordinated via turn indicators. Driver behavior can often be modeled as rational, characterized by risk aversion and efficiency-seeking objectives. Recent work has leveraged these observations in the design of data-driven models for behavior prediction and planning [41, 5, 32, 15, 34]. To measure and transfer their performance to the real world, such models require large, balanced datasets, containing a diversity of behavior that is representative of real-world traffic. This requires an understanding of the space of behavior that typically unfolds in different scenes. While there exist tools for characterizing various aspects of behavior [42, 20, 9, 22], summarizing multiagent interaction in an intuitive and formal way is hindered by the high dimensionality and the stochasticity of human behavior.

Our key insight is that the behavior encountered in traffic scenes often exhibits prevalent topological structure: common events like overtaking, merging or traversing a roundabout (Fig. 1) give rise to *interactions* of distinct topological signatures [2]. In this work, we abstract multiagent behavior in a traffic scene into a topological braid [1], a compact and interpretable topological object with symbolic and geometric descriptions. Building on past work on the use of braids for multiagent navigation [25], we make the following contributions. First, we adapt the representation of Mavrogiannis and Knepper [25] to structured domains like driving environments through a rigorous mathematical presentation. We then study its computational properties, and discuss how a measure of complexity based on braids [11] may capture the interactivity of a traffic scene. We show that our framework is applicable to complex scenes through a case study on real-world intersections and roundabouts [4, 18]. We cluster the behaviors exhibited in these scenes into braids, and characterize their complexity. We find that in the majority of scenes, a few simple braids dominate, indicating a low degree of interaction despite the high traffic density. Our methodology can be valuable for the analysis and design of road networks, the design and benchmarking of data-driven frameworks for prediction and planning, the evaluation and generation of driving datasets.

This work was (partially) funded by the National Science Foundation IIS (#2007011), National Science Foundation DMS (#1839371), the Office of Naval Research, US Army Research Laboratory CCDC, Amazon, and Honda Research Institute USA.

[1]C. Mavrogiannis and S. S. Srinivasa are with the Paul G. Allen School of Computer Science & Engineering, University of Washington, Seattle, WA 98102, USA. E-mail: {cmavro, siddh}@cs.washington.edu.

[2]J. A. DeCastro is with the Toyota Research Institute, Cambridge, MA 02139, USA. E-mail: jonathan.decastro@tri.global.

II. Related Work

Recent work on behavior prediction and decision making for autonomous driving applications has leveraged discrete, semantic representations of traffic behavior. For instance, Wang et al. [43] classify discrete driving styles using a variant of hidden Markov models (HMM). Gadepally et al. [13] also use HMM to estimate long-term driver behaviors from a sequence of discrete decisions. Others, such as Konidaris et al. [17] and Shalev-Shwartz et al. [36], propose using learned symbolic representations for high-level planning and collision avoidance, via a hierarchical options model. Similarly, Shu et al. [37] learn a latent representation of interactions. While these works uncover discrete representations of driving behavior, they either require large datasets to learn discrete modes or specify them manually without harnessing the rich geometric structure of the environment.

Another body of work has focused on tools for testing and validation in realistic settings, leveraging a semantic-level understanding of interactions. Tian et al. [41] model traffic at unsignalized intersections using tools from game theory and propose a verification testbed for navigation algorithms. Liebenwein et al. [20] propose a framework for safety verification of driving controllers based on compositional and contract-based principles. Hsu et al. [15] investigate how velocity signals generated by Markov decision processes affect interaction dynamics at intersections. DeCastro et al. [9] construct a representation of multi-vehicle interaction outcomes based on latent parameters using a generative model. Tolstaya et al. [42] propose an *Interactivity* score that enables the identification of interesting interactive scenarios for training generative models. Our work is similar in spirit and complementary to this latter line of work. We also approach a notion of interactivity between agents. However, instead of investigating statistical properties of distributions, we focus on the aspect of the representation, through the introduction of tools from low-dimensional topology.

Recently, roboticists have been increasingly making use of topological representations to model the rich structure of real-world domains. These include knitting [21], untangling [14] or knot planning [44], aircraft conflict resolution [16] or multiagent navigation [10]. Some works leverage insights from homotopy theory [7, 3], persistent homology [30] and fiber bundles [28]. Some other works make use of topological abstractions such as invariants [24, 23, 32] and braids [26, 25] as representatives of multiagent motion primitives for prediction and planning. In this paper, we are following up on this latter body of work work by employing topological braids as an abstraction of traffic behavior. While past work considered simplified simulation domains [26, 25], in this paper, we adapt the braid presentation to accommodate rich traffic environments such as real-world intersections or roundabouts. To the best of our knowledge this paper is the first to investigate the applicability of braids as primitives for multiagent behavior in realistic real-world environments.

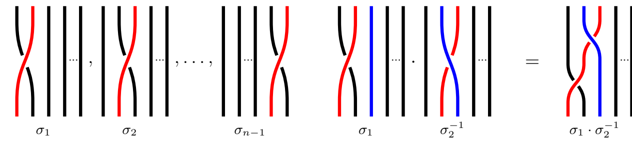

(a) Generators of B_n. (b) Example of composition.

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

III. Abstracting Driving Interactions as Topological Braids

We introduce a representation based on topological braids [1], that captures critical interaction events in street environments (e.g., overtaking, merging, crossing). This representation describes such interactions as sequences of symbols describing topological relationships between agents; any possible interaction manifests as a unique symbolic representation of their trajectories. Our representation adapts the presentation of Mavrogiannis and Knepper [25] to real-world traffic domains through theoretical developments.

A. Domain

Consider a structured domain $\mathcal{Q} \subseteq \mathbb{R}^2$ where $n > 1$ agents are navigating from time $t = 0$ to a finite final time t_∞. Define by $q_i \in \mathcal{Q}$ the position of agent $i \in \mathcal{N} = \{1, \ldots, n\}$ with respect to a fixed reference frame. Agent i follows a trajectory $\xi_i : [0, t_\infty] \to \mathcal{Q}$. Collectively, agents follow a system trajectory $\Xi = (\xi_1, \ldots, \xi_n)$. This trajectory is a detailed representation of the collective strategy that agents followed to avoid each other while following their paths. Their strategy can be summarized as a set of discrete relationships, such as the passing sides or crossing order of agents. These relationships are formed as a result of the geometric structure of the environment, traffic regulations, and agents' policies. In this paper, we show that such relationships feature topological properties that can be succinctly captured by the representation of topological braids [1].

B. Topological Braids

A braid is a tuple $b_f = (f_1, \ldots, f_n)$ of functions $f_i : I \to \mathbb{R}^2 \times I$, $i \in \mathcal{N}$, 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 $p_f : \mathcal{N} \to \mathcal{N}$ 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, $b_f = (f_1, \ldots, f_n)$, $b_g = (g_1, \ldots, g_n)$, can be composed through a *composition* operation (Fig. 2b): their composition, $b_h = b_f \cdot b_g$, is also a braid $b_h = (h_1, \ldots, h_n)$, comprising a set of n curves, defined as:

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

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

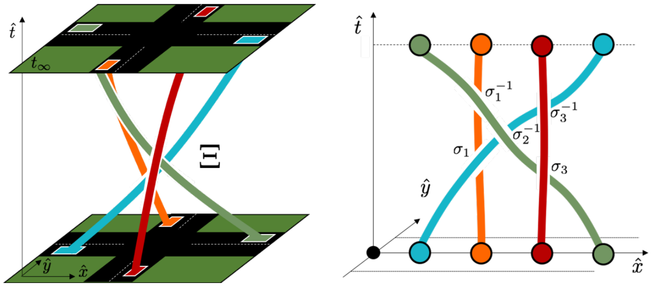

(a) Trajectories of four agents as they navigate an intersection, plotted in spacetime.

(b) Braid $\sigma_3\,\sigma_1\,\sigma_2^{-1}\sigma_3^{-1}\sigma_1^{-1}$ capturing the topological entanglement of agents' trajectories.

Fig. 3: Transition from Cartesian trajectories (a) to topological braids (b) via eq. (4) assuming a x-t projection.

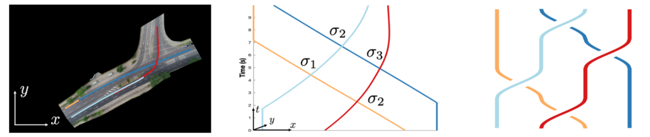

(a) Top view.

(b) Side projection.

(c) Extracted braid.

Fig. 4: Transitioning from a real-world episode to a braid. The trajectories of (a) are first projected on the plane x-t (b) and then the braid $\sigma_2\sigma_1\sigma_3\sigma_2$ is reconstructed (c).

Braid group on n strands. Following Artin's presentation [1], the braid group B_n can be generated from $n-1$ primitive braids $\sigma_1, ..., \sigma_{n-1}$ (see Fig. 2a), called generators, and their inverses, via composition.

A **generator** is a braid $\sigma_i = (\upsigma_1, \ldots, \upsigma_n)$, $i \in \mathcal{N} \setminus \{n\}$ for which: (a) $\upsigma_i(0) = (1, 0, 0)$, and $\upsigma_i(1) = (p_i(i), 0, 1)$, where $p_i : \mathcal{N} \to \mathcal{N}$ is an adjacent transposition swapping i and $i+1$; (b) there exists a unique $t_c \in [0, 1]$ such that $(\upsigma_i(t_c) - \upsigma_{i+1}(t_c)) \cdot \hat{x} = 0$ and $(\upsigma_i(t_c) - \upsigma_{i+1}(t_c)) \cdot \hat{y} > 0$.

The **inverse** of σ_i is the braid $\sigma_i^{-1} = (\upsigma_1^{-1}, \ldots, \upsigma_n^{-1})$, $i \in \mathcal{N} \backslash n$, for which: (a) $\upsigma_i^{-1}(0) = (1, 0, 0)$, and $\upsigma_i^{-1}(1) = (p_i(i), 0, 1)$; (b) there exists a unique $t_c \in [0, 1]$ such that $(\upsigma_i^{-1}(t_c) - \upsigma_{i+1}^{-1}(t_c)) \cdot \hat{x} = 0$ and $(\upsigma_i^{-1}(t_c) - \upsigma_{i+1}^{-1}(t_c)) \cdot \hat{y} < 0$.

The **identity** braid $\sigma_0 = (\upsigma_1^0, \ldots, \upsigma_n^0)$ implements no swap, i.e., $p_0(k) = k$ for any $k \in \mathcal{N} \setminus \{n\}$, yielding $\upsigma_k^0(0) = (k, 0, 0)$, $\upsigma_k^0(1) = (k, 0, 1)$ and it holds that $\nexists t_c \in [0, 1]$ such that $(\upsigma_k^0(t_c) - \upsigma_{k+1}^0(t_c)) \cdot \hat{x} = 0$.

Any braid can be written as a *word*, i.e., a product of generators and their inverses (Fig. 2b), satisfying the *relations*:

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

C. Transforming Traffic Trajectories into Braids

We will convert a system trajectory Ξ into a Cartesian object with the structure of a topological braid through a sequence of operations that retain the topological relationships among agents' trajectories.

We define by $\xi_i^x : [0, t_\infty] \to \mathbb{R}$ and $\xi_i^y : [0, t_\infty] \to \mathbb{R}$ the x and y projections of ξ_i. For $t = 0$, ranking agents in order of increasing $\xi_i^x(0)$, $i \in \mathcal{N}$ value defines a starting permutation $p_s : \mathcal{N} \to \mathcal{N}$, where $p_s(i)$ denotes the order of agent i. For $t = t_\infty$, ranking agents in order of increasing $\xi_i^x(t_\infty)$ value defines a final permutation $p_d : \mathcal{N} \to \mathcal{N}$, where $p_d(p_s(i))$ denotes the final ranking of agent i. Thus the execution in Ξ can be abstracted into a transition from p_s to p_d.

We denote by $\uptau : I \to [0, t_\infty]$ a time normalization function, uniformly mapping $I = [0, 1]$ to the execution time in the range $[0, t_\infty]$. We then define the trajectory bounds as $x_{min} = \min_{i,t} \xi_i^x(t)$, $x_{max} = \max_{i,t} \xi_i^x(t)$, and $y_{min} = \min_{i,t} \xi_i^y(t)$, $y_{max} = \max_{i,t} \xi_i^y(t)$. Assuming $x_{max} \neq x_{min}$, $y_{max} \neq y_{min}$, we define the ratio functions:

$$r_i^x(t) = \frac{\xi_i^x(t) - x_{min}}{x_{max} - x_{min}}, \qquad r_i^y(t) = \frac{\xi_i^y(t) - y_{min}}{y_{max} - y_{min}}. \tag{3}$$

Finally, we define a set of functions $(f_1, \ldots, f_n)$, with $f_j : I \to \mathbb{R}^2 \times I$, $j \in \mathcal{N}$, such that:

$$f_j(a) = \begin{cases} (j, 0, 0)\,, & a = 0 \\ (f_j^x(a), f_j^y(a), a), & a \in (0, 1)\,, \\ (p_d(j), 0, 1)\,, & a = 1 \end{cases} \tag{4}$$

where

$$\begin{aligned} f_j^x &= 1 + r_j^x(\uptau(a))(n-1), \\ f_j^y &= -1 + 2r_j^y(\uptau(a)), \end{aligned} \tag{5}$$

and $j = p_s(i)$, $i \in \mathcal{N}$. For $a \in (0, 1)$, the expressions of (4) scale x-coordinates to lie within $[1, n-1]$ and the y-coordinates to lie within $[-1, 1]$ in a way that preserves topological relationships among trajectories. The set of functions $(f_1, \ldots, f_n)$ is a topological *braid* β following the definition of Sec. III-B. The braid β is topologically equivalent – ambient-isotopic [27]—to the system trajectory Ξ.

D. Braids as Modes of Traffic Coordination

The transformation of Sec. III-C enables summarization of a traffic episode into a braid capturing multiagent collision-avoidance relationships. This braid can be written as a word, similarly to how Thiffeault [39] converted particle motion in a fluid to a braid (Fig. 3): a) we label any trajectory crossings that emerge within the x-t projection as braid generators by identifying *under* or *over* crossings (Fig. 3b); b) we arrange these generators in temporal order into a *braid word*. Note that alternative reference frames can be employed; we selected the $\hat{x}$-$\hat{t}$ plane projection for convenience.

In Fig. 3, four agents cross an intersection. The braid $\sigma_3\,\sigma_1\,\sigma_2^{-1}\sigma_3^{-1}\sigma_1^{-1} \in B_4$ is a description of how agents coordinated to avoid each other. The group B_4 contains all ways in which these four agents could possibly avoid each other. In a scene with n agents, a braid represents a *mode* of coordination from the set of possible modes in B_n.

E. Computational Properties of the Representation

To highlight the possible computational benefits arising from the summarization of traffic episodes into braids, we study the runtime of enumerating modes of coordination as topological braids in comparison to enumerating Cartesian

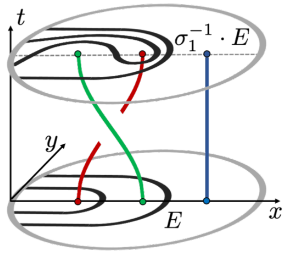

(a) Curve diagram for σ_1^{-1}.

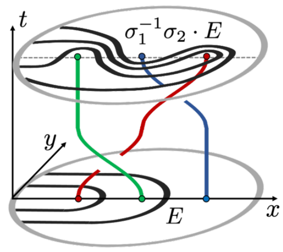

(b) Curve diagram for $\sigma_1^{-1}\sigma_2$.

Fig. 5: Curve diagrams for braids of different complexity. The braid $\sigma_1^{-1}\sigma_2$ (b) is more complex ($TC = 2$) than the braid σ_1^{-1} ($TC = 1.585$) shown in (a). This is reflected in the higher number of intersections between the curve diagram $\sigma_1^{-1}\sigma_2 \cdot E$ and the x-axis (dotted line).

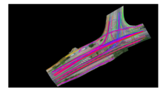

(a) inD 1.

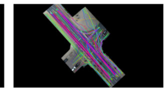

(b) inD 2.

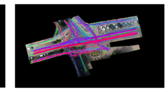

(c) inD 3.

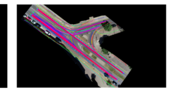

(d) inD 4.

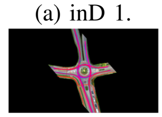

(e) rounD 1.

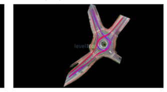

(f) rounD 2.

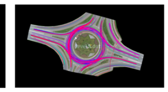

(g) rounD 3.

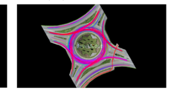

(h) rounD 4.

Fig. 6: Top view of the 8 scenes from the inD and rounD datasets that we analyzed using topological tools. All trajectories are overlayed on top of the street structures.

trajectories. Consider a traffic episode of H timesteps, involving n agents. Each agent has $\mathcal{T}$ options of routes to follow and $\mathcal{U}$ actions to take at every timestep. We assume that there is at most one agent per lane, i.e., $n \leq \mathcal{T}$. The horizon of the execution is long, and thus $n \ll H$. Further, $\mathcal{U}$ is a realistically rich space of controls, and thus $n \ll \mathcal{U}$, $H \ll \mathcal{U}$ and $\mathcal{T} \ll \mathcal{U}$. Finally, we assume that agents are goal-driven for the horizon of each episode, and thus they will cross paths with each other at most once.

The number of possible Cartesian trajectories in this domain is $N_c = |\mathcal{T}|^n(|\mathcal{U}|^n)^H$. Enumerating these trajectories runs in time $O(2^{nH\log\mathcal{U}})$. For the same scene, the number of possible braids generally depends on the structure of the road network. However, we can bound the number of possible outcomes as $N_b \leq 3^{\binom{n}{2}}$, where the exponent is the binomial coefficient representing the number of all pairs of agents and the base represents the 3 types of possible interactions per pair that could be represented by a braid, i.e., "over-crossing", "under-crossing", or no crossing. This enumeration runs in time $O(2^{n^2})$.

Theorem III.1. The runtime of enumerating braids is lower than the runtime of enumerating Cartesian trajectories for the class of driving problems considered above.

Proof. We want to show that $2^{n^2} < 2^{nH\log\mathcal{U}}$. This inequality is equivalent to $n < H\log\mathcal{U}$. We assumed that $n \ll H$, $n \ll \mathcal{U}$, therefore it should also hold that $n \ll H\log\mathcal{U}$. Thus the initial inequality holds and supports the statement that the runtime of enumerating braids is significantly lower than the runtime of enumerating Cartesian trajectories. ■

F. Complexity of Braid Entanglement

The entanglement of the trajectories described by a braid is indicative of the complexity of the interaction between agents. We quantify braid complexity using the *Topological Complexity* index (TC) of Dynnikov and Wiest [11] for which we provide an informal definition below.

Assume that a braid $\beta \in B_n$ represents the collective motion of n agents from initial locations $\beta(0)$ to final locations $\beta(1)$. Denote by D^2 a closed disk surrounding agents' initial positions, $\beta(0)$. Define by E a set of $n-1$ disjoint arcs, anchored on the disk, and separating the agents for $t = 0$, defining $n-1$ distinct regions in the disk (see Fig. 5). Assume that these regions are rigidly attached on the agents. As the agents follow the motion described by β from $t = 0$ to $t = 1$, the regions dynamically deform. The image $D = \beta \cdot E$ representing the shape of the regions obtained upon applying the motion described by β on E is called a *curve diagram*. The *norm* of curve diagram D is defined as the number of intersections of D with the x axis. Based on the above definitions, we can define the TC index of a braid $\beta \in B_n$ as:

$$TC(\beta) = \log_2(||\beta \cdot E||) - \log_2(||E||). \qquad (6)$$

This expression is equivalent to the logarithm of the *gain* of intersections with the x-axis, upon application of a braid. Fig. 5 depicts curve diagrams acquired upon inducing motion of two different braids on the canonical curve diagram E.

IV. A Case Study on Traffic Datasets

We demonstrate how braids may abstract traffic episodes through a case study on real-world datasets.

A. Datasets

We consider the inD [4] and rounD [18] datasets, containing trajectory data (of vehicles, pedestrians and bicycles) recorded respectively in four intersection scenes and four roundabout scenes of the German road network. Both datasets were extracted from drone footage in 25 fps via computer vision techniques, yielding an estimated positional error in the order of 10 cm. A top view of the eight scenes is shown in Fig. 6, and their dimensions are listed in Table I.

B. Methodology

We split each scene into a set of sequential episodes, sweeping the whole duration of the recording. Each episode has a fixed duration of $\Delta T = 10\,\mathrm{s}$, containing trajectories of simultaneously navigating agents. From qualitative inspection, we observed that the most interesting interactions across all scenes involved vehicle traffic; to highlight dynamic vehicle traffic, we filtered out agents moving with a speed lower than $14\,\mathrm{m\,s^{-1}}$ and agents that are too far from each other (agents that kept a minimum distance greater than $d_{min} = 10\,\mathrm{m}$ throughout the episode). This resulted in a set of episodes summarized in Table I. Using the framework

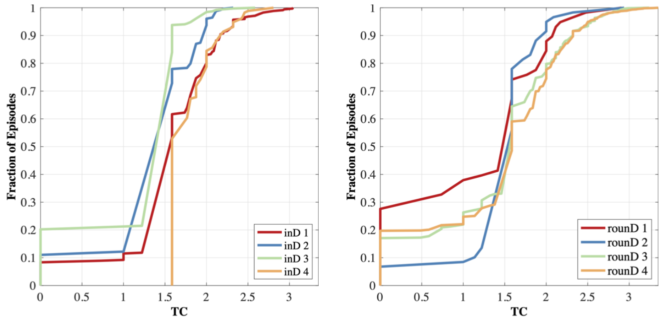

(a) inD scenes.

(b) rounD scenes.

Fig. 7: Cumulative density of TC (Topological Complexity index) in intersections (a) and roundabouts (b).

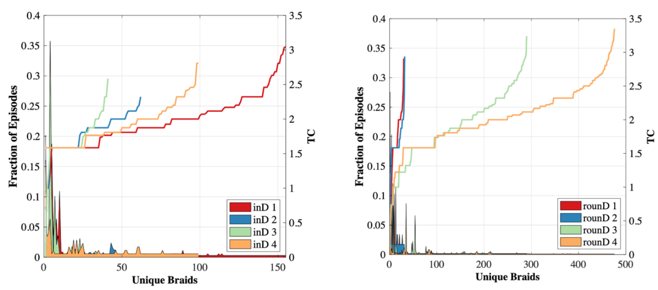

(a) inD scenes.

(b) rounD scenes.

Fig. 8: Frequency of unique braids in intersections (a) and roundabouts (b), arranged in order of increasing TC (Topological Complexity index).

of Sec. III-B, we abstracted the trajectory of each episode into a braid, leveraging the braid relations of eq. (2). Finally, we computed the TC score for each braid. We performed all computations using the Braidlab package [40].

C. Analysis

The behavior in each scene is clustered into a small number of unique braids, describing vehicles' interaction patterns (see Table I). This highlights that real-world traffic tends to collapse to a small set of outcomes. The extracted braids are mapped onto the TC values on the right. Fig. 9 depicts episodes of varying TC, drawn from the two datasets, along with their braid representatives and TC scores. We see that complex interactions get mapped onto higher TC values.

Fig. 7 shows the empirical cumulative density of TC across the inD and rounD dataset scenes. We see that each scene has a distinct complexity growth pattern but in both datasets, about 60% of episodes are concentrated below $TC = 1.5$. This is highlighted in Fig. 8, which shows the relative frequency of unique braids per scene, organized in order of increasing TC. We see that the mass of the frequency is concentrated on the left side for both plots, suggesting that the majority of episodes feature a relatively low degree of interaction. This indicates that despite the dense traffic exhibited in the datasets (Table I), the vast majority of episodes involve traffic that is orderly and well organized.

TABLE I: Scene details and interaction statistics.

Scene	Dimensions (m^2)	Episodes	Agents/Episode (M, SD)	Unique braids	TC (M, SD)
inD 1	131×110	347	3.62 ± 1.76	155	1.62 ± 0.59
inD 2	59×64	254	2.82 ± 1.00	62	1.48 ± 0.55
inD 3	85×45	386	2.62 ± 0.90	41	1.28 ± 0.66
inD 4	79×67	174	4.10 ± 1.51	99	1.79 ± 0.28
rounD 1	99×143	58	3.16 ± 1.45	30	1.20 ± 0.84
rounD 2	99×122	59	3.85 ± 1.75	32	1.54 ± 0.50
rounD 3	127×69	574	4.36 ± 2.28	290	1.43 ± 0.79
rounD 4	92×98	1050	4.07 ± 2.00	476	1.46 ± 0.83

This is an artifact of the underlying spatiotemporal structure (geometry, traffic rules, driving styles).

D. Discussion

Our representation enables enumeration of the classes of multiagent interaction patterns that are *theoretically* possible in a traffic scene in a compact and interpretable form. Further, given a traffic dataset, it allows us to extract the subset of interaction patterns that are *empirically* relevant. This may inform algorithmic design, benchmarking and even road network design. Importantly, our framework can be valuable for characterizing a traffic dataset with respect to the support it provides over the space of theoretically possible behavior in a domain. Understanding the support of a dataset is crucial for data-driven approaches [32, 34, 22] but also for guiding the process of synthetically generating simulated scenarios to produce diverse datasets.

Our framework is complementary to alternative approaches for characterizing interaction, such as the interactivity score [42] and distribution-based KL-divergence. The Interactivity score may miss crucial interaction events: scores can be large when there is high correlation between two trajectories (e.g., one car following another), but small when trajectories are dissimilar (e.g., cars crossing an intersection). In contrast, TC will account for these situations through the consideration of the underlying topological structure. Further, our framework may be directly applicable to any traffic dataset [12, 6, 8] and even to alternative domains like pedestrian tracking [29] or sports analysis [38] without additional modifications. It may complement temporal logic approaches for trajectory labeling [31, 19] which often require involved and domain-specific mathematical treatment [35].

V. Conclusion

Abstractions like braids highlight topological patterns of interaction like vehicles' crossings or overtaking maneuvers through projection transforms or simplification rules like eq. 2. They filter out geometric features, like the temporal spacing between vehicles or a driver's erratic maneuvers. These artifacts could be relevant to traffic analysis. Thus, the proposed framework is not meant to replace existing, geometry-focused tools but rather to complement them.

Our goal in this study was to demonstrate that tools from braid theory can be valuable for the analysis of multiagent behavior in traffic scenes. Thus, some of the parameters chosen, e.g., the braid projection plane, the episode duration, agents' distance and speed thresholds were not optimally selected. These parameters could be adapted to the specific context of a scene or the scope of an investigation.

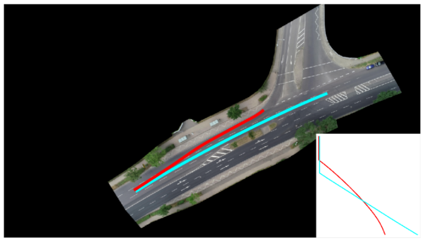

(a) inD 1, $TC = 0$.

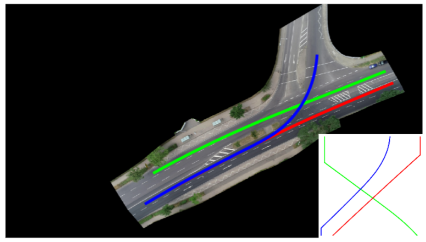

(b) inD 1, $TC = 1.5850$.

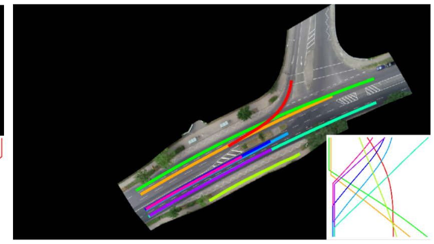

(c) inD 1, $TC = 3.0444$.

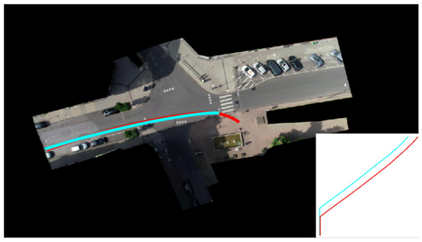

(d) inD 3, $TC = 0$.

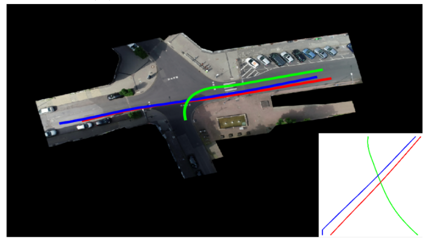

(e) inD 3, $TC = 1.5850$.

(f) inD 3, $TC = 2.5850$.

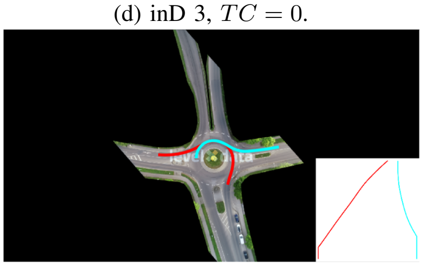

(g) rounD 1, $TC = 0$.

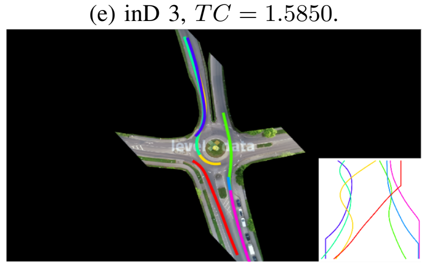

(h) rounD 1, $TC = 1.4150$.

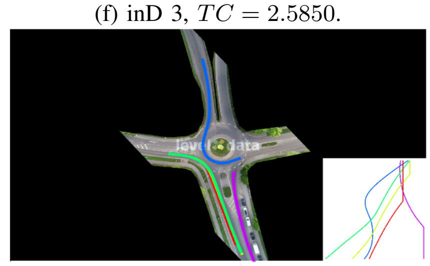

(i) rounD 1, $TC = 2.9069$.

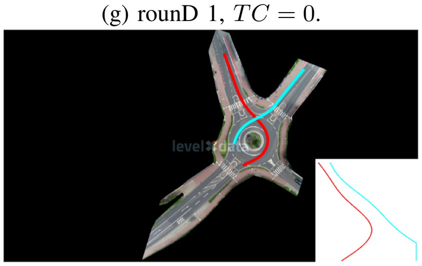

(j) rounD 2, $TC = 0$.

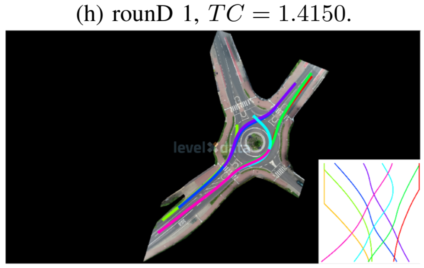

(k) rounD 2, $TC = 1.7162$.

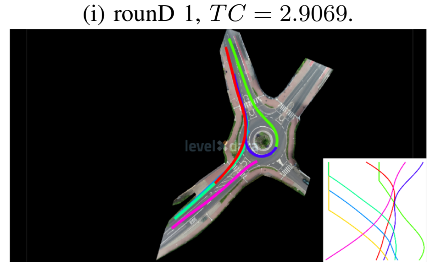

(l) rounD 2, $TC = 2.9386$.

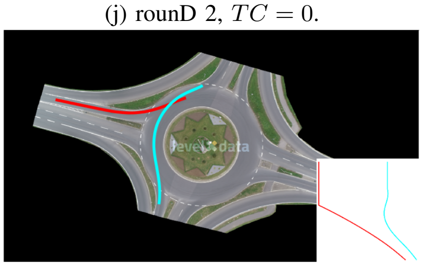

(m) rounD 3, $TC = 0$.

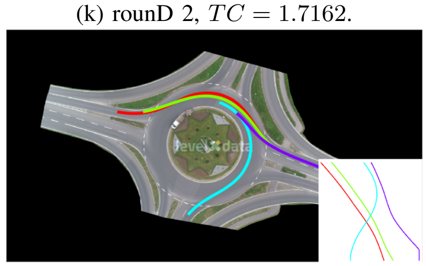

(n) rounD 3, $TC = 1.2224$.

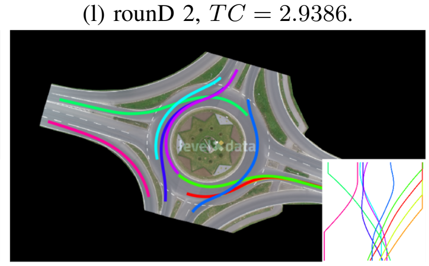

(o) rounD 3, $TC = 3.2395$.

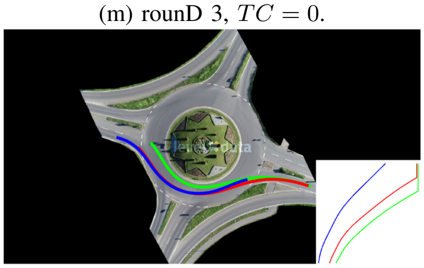

(p) rounD 4, $TC = 0$.

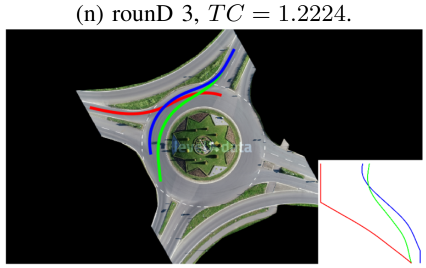

(q) rounD 4, $TC = 1$.

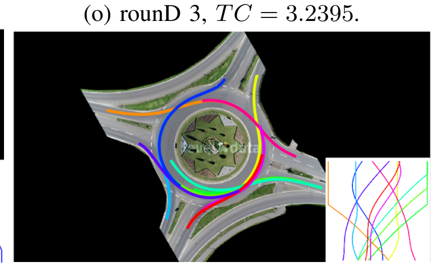

(r) rounD 4, $TC = 3.3505$.

Fig. 9: Episodes with different Topological Complexity (TC). Each row depicts three episodes yielding distinct braids in the same scene. At the bottom right of each figure, the braid formed by the data through a x-t side projection of the episode is plotted. The episodes on each row are organized from left to the right in order of increasing TC. In all scenes, the agents are following the right-hand traffic convention.

References

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

[2] M. A. Berger. Hamiltonian dynamics generated by vassiliev invariants. *Journal of Physics A: Mathematical and General*, 34(7):1363, 2001.

[3] S. Bhattacharya and R. Ghrist. Path homotopy invariants and their application to optimal trajectory planning. *Annals of Mathematics and Artificial Intelligence*, 84: 139–160, 2018.

[4] J. Bock, R. Krajewski, T. Moers, S. Runde, L. Vater, and L. Eckstein. The ind dataset: A drone dataset of naturalistic road user trajectories at german intersections. In *Proceedings of the IEEE Intelligent Vehicles Symposium (IV)*, pages 1929–1934, 2020.

[5] M. Bouton, A. Cosgun, and M. J. Kochenderfer. Belief state planning for autonomously navigating urban intersections. In *Proceedings of the IEEE Intelligent Vehicles Symposium (IV)*, pages 825–830, 2017.

[6] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom. nuscenes: A multimodal dataset for autonomous driving. *arXiv preprint arXiv:1903.11027*, 2019.

[7] C. Cao, P. Trautman, and S. Iba. Dynamic channel: A planning framework for crowd navigation. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 5551–5557, 2019.

[8] M. Chang, J. Lambert, P. Sangkloy, J. Singh, S. Bak, A. Hartnett, D. Wang, P. Carr, S. Lucey, D. Ramanan, and J. Hays. Argoverse: 3d tracking and forecasting with rich maps. *CoRR*, abs/1911.02620, 2019.

[9] J. DeCastro, K. Leung, N. Aréchiga, and M. Pavone. Interpretable policies from formally-specified temporal properties. In *Proceedings of the IEEE International Conference on Intelligent Transportation Systems (ITSC)*, pages 1–7, 2020.

[10] Y. Diaz-Mercado and M. Egerstedt. Multirobot mixing via braid groups. *IEEE Transactions on Robotics*, 33 (6):1375–1385, 2017.

[11] I. Dynnikov and B. Wiest. On the complexity of braids. *Journal of the European Mathematical Society*, 009(4): 801–840, 2007.

[12] S. Ettinger, S. Cheng, B. Caine, C. Liu, H. Zhao, S. Pradhan, Y. Chai, B. Sapp, C. Qi, Y. Zhou, Z. Yang, A. Chouard, P. Sun, J. Ngiam, V. Vasudevan, A. McCauley, J. Shlens, and D. Anguelov. Large scale interactive motion forecasting for autonomous driving : The waymo open motion dataset. *CoRR*, abs/2104.10133, 2021.

[13] V. Gadepally, A. Krishnamurthy, and Ü. Özgüner. A Framework for Estimating Long Term Driver Behavior. *Journal of Advanced Transportation*, 2017:1–11, 2017.

[14] J. Grannen, P. Sundaresan, B. Thananjeyan, J. Ichnowski, A. Balakrishna, M. Hwang, V. Viswanath, M. Laskey, J. E. Gonzalez, and K. Goldberg. Untangling dense knots by learning task-relevant keypoints. In *Proceedings of the Conference on Robot Learning (CoRL)*, 2021.

[15] Y.-C. Hsu, S. Gopalswamy, S. Saripalli, and D. 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.

[16] J. Hu, M. Prandini, and S. Sastry. Optimal maneuver for multiple aircraft conflict resolution: a braid point of view. In *Proceedings of the IEEE Conference on Decision and Control (CDC)*, volume 4, pages 4164–4169 vol.4, 2000.

[17] G. Konidaris, L. P. Kaelbling, and T. Lozano-Perez. From Skills to Symbols: Learning Symbolic Representations for Abstract High-Level Planning. *Journal of Artificial Intelligence Research*, 61:215–289, 2018.

[18] R. Krajewski, T. Moers, J. Bock, L. Vater, and L. Eckstein. The round dataset: A drone dataset of road user trajectories at roundabouts in germany. In *Proceedings of the IEEE International Conference on Intelligent Transportation Systems (ITSC)*, pages 1–6, 2020.

[19] X. Li, G. Rosman, I. Gilitschenski, C. I. Vasile, J. A. DeCastro, S. Karaman, and D. Rus. Vehicle trajectory prediction using generative adversarial network with temporal logic syntax tree features. *IEEE Robotics and Automation Letters*, 6(2):3459–3466, 2021.

[20] L. Liebenwein, W. Schwarting, C.-I. Vasile, J. DeCastro, J. Alonso-Mora, S. Karaman, and D. Rus. Compositional and contract-based verification for autonomous driving on road networks. In *Proceedings of the International Symposium on Robotics Research (ISRR)*, pages 163–181, 2020.

[21] J. Lin and J. McCann. An artin braid group representation of knitting machine state with applications to validation and optimization of fabrication plans. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 1147–1153, 2021.

[22] A. Mavrogiannis, R. Chandra, and D. Manocha. B-gap: Behavior-rich simulation and navigation for autonomous driving. *IEEE Robotics and Automation Letters*, 7(2):4718–4725, 2022.

[23] C. Mavrogiannis and R. A. Knepper. Hamiltonian coordination primitives for decentralized multiagent navigation. *The International Journal of Robotics Research*, 40(10-11):1234–1254, 2021.

[24] C. Mavrogiannis, P. Alves-Oliveira, W. Thomason, and R. A. Knepper. Social momentum: Design and evaluation of a framework for socially competent robot navigation. *ACM Transactions on Human-Robot Interaction*, 11(2), 2022.

[25] C. I. Mavrogiannis and R. 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.

[26] C. I. Mavrogiannis, V. Blukis, and R. A. Knepper. Socially competent navigation planning by deep learn-

ing of multi-agent path topologies. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6817–6824, 2017.

[27] K. Murasugi and B. I. Kurpita. *A Study of Braids*. Mathematics and Its Applications. Springer Netherlands, 1999.

[28] A. Orthey, S. Akbar, and M. Toussaint. Multilevel motion planning: A fiber bundle formulation. *arXiv preprint arXiv:2007.09435*, 2020.

[29] S. Pellegrini, A. Ess, K. Schindler, and L. J. V. Gool. You'll never walk alone: Modeling social behavior for multi-target tracking. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 261–268, 2009.

[30] F. T. Pokorny, M. Hawasly, and S. Ramamoorthy. Topological trajectory classification with filtrations of simplicial complexes and persistent homology. *The International Journal of Robotics Research*, 35(1-3): 204–223, 2016.

[31] A. G. Puranic, J. V. Deshmukh, and S. Nikolaidis. Learning from demonstrations using signal temporal logic in stochastic and continuous domains. *IEEE Robotics and Automation Letters*, 6(4):6250–6257, 2021.

[32] J. Roh, C. Mavrogiannis, R. Madan, D. Fox, and S. S. Srinivasa. Multimodal trajectory prediction via topological invariance for navigation at uncontrolled intersections. In *Proceedings of the Conference on Robot Learning*, pages 2216–2227, 2020.

[33] D. Sadigh, S. Sastry, S. A. Seshia, and A. D. Dragan. Planning for autonomous cars that leverage effects on human actions. In *Proceedings of Robotics: Science and Systems (RSS)*, 2016.

[34] T. Salzmann, B. Ivanovic, P. Chakravarty, and M. Pavone. Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 683–700, 2020.

[35] J. Schulz, K. Hirsenkorn, J. Löchner, M. Werling, and D. Burschka. Estimation of collective maneuvers through cooperative multi-agent planning. In *IEEE Intelligent Vehicles Symposium, IV 2017, Los Angeles, CA, USA, June 11-14, 2017*, pages 624–631, 2017.

[36] S. Shalev-Shwartz, S. Shammah, and A. Shashua. Safe, Multi-Agent, Reinforcement Learning for Autonomous Driving. *arXiv:1610.03295*, 2016.

[37] T. Shu, Y. Peng, L. Fan, H. Lu, and S.-C. Zhu. Perception of human interaction based on motion trajectories: From aerial videos to decontextualized animations. *Topics in Cognitive Science*, 10(1):225–241, 2018.

[38] SportVU. Sportvu basketball player tracking. https://www.stats.com/sportvu-basketball/. Accessed: 2021-05-16.

[39] J.-L. Thiffeault. Braids of entangled particle trajectories. *Chaos*, 20(1), 2010.

[40] J.-L. Thiffeault and M. Budišić. Braidlab: A software package for braids and loops, 2013–2021. Version 3.2.4.

[41] R. Tian, N. Li, I. Kolmanovsky, Y. Yildiz, and A. Girard. Game-theoretic Modeling of Traffic in Unsignalized Intersection Network for Autonomous Vehicle Control Verification and Validation. *arXiv e-prints*, art. arXiv:1910.07141, 2019.

[42] E. Tolstaya, R. Mahjourian, C. Downey, B. Vadarajan, B. Sapp, and D. Anguelov. Identifying driver interactions via conditional behavior prediction, 2021.

[43] W. Wang, W. Zhang, J. Zhu, and D. Zhao. Understanding v2v driving scenarios through traffic primitives. *IEEE Transactions on Intelligent Transportation Systems*, 23(1):610–619, 2022.

[44] M. Yan, G. Li, Y. Zhu, and J. Bohg. Learning topological motion primitives for knot planning. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 9457–9464, 2020.