Act Natural! Projecting Autonomous System Trajectories Into Naturalistic Behavior Sets ⋆

Hamzah I. Khan Adam J. Thorpe David Fridovich-Keil

University of Texas at Austin, Austin, TX 78712 USA

Abstract: Autonomous agents operating around human actors must consider how their behaviors might affect those humans, even when not directly interacting with them. To this end, it is often beneficial to be predictable and appear naturalistic. Existing methods to address this problem use human actor intent modeling or imitation learning techniques, but these approaches rarely capture all possible motivations for human behavior or require significant amounts of data. In contrast, we propose a technique for modeling naturalistic behavior as a set of convex hulls computed over a relatively small dataset of human behavior. Given this set, we design an optimization-based filter which projects arbitrary trajectories into it to make them more naturalistic for autonomous agents to execute while also satisfying dynamics constraints. We demonstrate our methods on real-world human driving data from the inD intersection dataset (Bock et al., 2020).

Keywords: Advanced Control Design, Intelligent Road Transportation, Cognitive Control, Human Behavior Modeling, Optimization, Dynamical Systems

1. INTRODUCTION

In interactive settings, human actors require a degree of predictability from autonomous agents to ensure the safe and comfortable operation of all interacting actors. Failing to do so can cause problems: for example, autonomous cars can often behave more cautiously than humans expect, leading nearby human drivers to react in unexpected ways and, potentially, cause collisions (Teoh and Kidd, 2017). The requirement of acting like other actors is even encoded in some traffic laws [1] in which drivers must follow the "flow of traffic" regarding their speeds on highways. As acting in ways that stand out unnecessarily can lead to safety and comfort concerns, autonomous vehicles must be able to understand how humans naturalistically behave. However, naturalistic behavior tends to be opaque and thus difficult to model mathematically, and there is a need for methods that can highlight what naturalistic behavior looks like from observations of human behavior. Existing solutions (Bajcsy et al., 2021; Sadigh et al., 2016a) tend to either model a limited subset of possible influences in human decision-making or do not produce naturalistic behaviors without significant data (Kuefler et al., 2017). These deficiencies present a need for improved naturalistic behavior generation in autonomous prediction and downstream planning tasks.

We present a data-driven method for identifying the set of naturalistic human behaviors using convex hulls and an optimization-based method for projecting autonomously planned trajectories into it. Specifically, given a set of observed human driving trajectories, we compute a sequence of convex hulls around the states at each time instant. With this representation in hand, we project trajectories

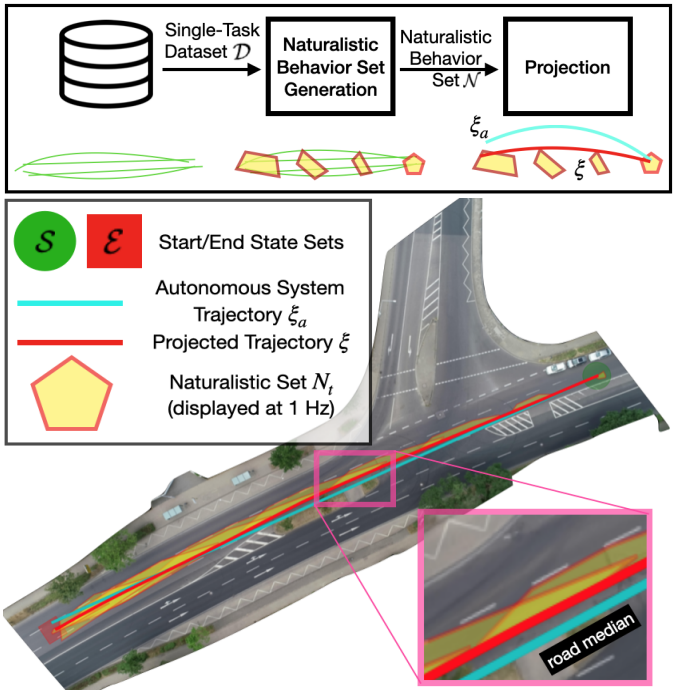

Fig. 1. (top). Given a single-task dataset $\mathcal{D}$, our method first generates a naturalistic behavior set by computing time-indexed convex hulls. Then, we project arbitrary trajectories into this set to make the behaviors more naturalistic. (bottom). We compute the naturalistic behavior set $\mathcal{N}$ using trajectories in the second lane of the upper road beginning in the green circle and ending in the red square. We then project a trajectory into $\mathcal{N}$.

⋆ This work was supported by the National Science Foundation under Grant No. 2211548. (e-mail: { hamzah, adam.thorpe, dfk }@utexas. edu).

[1] Texas Transportation Code §545.363

(i.e. from autonomous planners) into the naturalistic behavior set to enforce naturalistic driving constraints.

Existing work establishes that autonomous planners must consider the effect of their behavior on human actors (Sadigh et al., 2016b). Dragan et al. (2013) propose the concept of legible motion from which an observer can easily infer intent or strategy through actor motion. This concept differs from predictability, which involves easily inferred motion without necessarily understanding intent. Sadigh et al. (2016b) notes that the actions of an autonomous vehicle influence human drivers on the road, potentially unintentionally. Such a conclusion implies a need for autonomous vehicles to model human behavior in prediction and planning, as otherwise safe but unexpected actions might cause unsafe reactions by human actors.

Existing methods tend to approach the problem of accounting for human behavior from one of two directions: inferring intent from human behavior models and imitating behavior based on observed data. Human behavior models use approaches ranging from explicitly modeling and inferring specific aspects of human internal states like rationality (Bobu et al., 2018) and target states (Sadigh et al., 2016a) to data-driven predictive human behavior modeling that infers human actors' beliefs of other agents' goals (Bajcsy et al., 2021; Fridovich-Keil et al., 2020), or influencing human actors' (potentially incorrect) internal models of autonomous agents' motion (Tian et al., 2023). However, these approaches tend to fall short because human preference and behavior is naturally opaque and difficult to model mathematically, resulting in multiple unmodeled aspects of the world which may be significant, like how road conditions may influence safety and comfort tolerances. Our approach differs by seeking to frame naturalistic behavior planning as a projection into naturalistic constraints from human behavior data, without otherwise inferring the internal state of human actors.

The second set of approaches to this problem is imitation learning, and these involve motion planning by mimicking provided human behavior data through behavior cloning or inverse reinforcement learning methods. Behavior cloning methods train supervised models on human data but tend to suffer from compounding errors and distribution shift (Bagnell, 2015; Kuefler et al., 2017). Inverse reinforcement learning, which attempts to infer a cost function describing agent objectives, generalizes better than behavior cloning and produces naturalistic human behavior in microscopic settings (i.e., using the local context) but requires significantly more data (Kuefler et al., 2017). Additionally, most inverse reinforcement learning approaches assume very specific reward structures (Song et al., 2018; Huang et al., 2021). For abnormal scenarios in which we have limited data due to dynamic conditions (i.e., bad weather, rare road scenarios), even imitation learning techniques are infeasible for modeling human behavior due to insufficient data.

Our main innovation is a simple and efficient method for identifying convex constraints that capture naturalistic human driving behaviors. We then use our representation to formulate and solve a trajectory projection problem that enforces naturalistic behavior, similarity to an original trajectory, and dynamic feasibility. We make two key contributions: 1) constructing a convex hull around naturalistic driving data, and 2) an optimization-based method for projecting autonomous solutions onto the learned sets. Through numerical results on the inD real-world scenario driving dataset (Bock et al., 2020), we demonstrate the ability of our naturalistic behavior set generation method to capture patterns, which may not be explicitly modeled, in human behavior data. We additionally adjust trajectories generated from other sources to capture such patterns through projection.

2. PROBLEM STATEMENT

Consider a discrete-time dynamical system,

$$\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t), \tag{1}$$

where $\mathbf{x}_t \in \mathcal{X} \subseteq \mathbb{R}^n$ is the state at time t and $\mathbf{u}_t \in \mathcal{U} \subseteq \mathbb{R}^r$ is the control input. The system evolves from an initial condition $\mathbf{x}_0$ over a finite time horizon $t = 0, \dots, T$, where $T \in \mathbb{N}$.

We assume that the dynamics f in (1) are known. We additionally presume access to historical human behavior data consisting of state trajectories of the form

$$\xi^i = [\mathbf{x}_0^{i\intercal}\ \mathbf{x}_1^{i\intercal}\ \cdots\ \mathbf{x}_T^{i\intercal}]^\intercal \in \Xi \subseteq \mathbb{R}^{n(T+1)}. \tag{2}$$

We further expect this data to describe "well-behaved" human behavior. For example, well-behaved human behavior should generally be safe, as humans naturally avoid unsafe behavior.

2.1 Estimating the Naturalistic Behavior Set

Our goal is to compute a set-based representation of the state trajectories that captures the naturalistic behavior of the human operator performing a given task, e.g. a driver making a left turn at an intersection or navigating a particular path from point A to point B. In other words, given a dataset $\mathcal{D}$ which consists of m state trajectories (as in (2)) that perform a particular task,

$$\mathcal{D} = \{\xi^i\}_{i=1}^m, \tag{3}$$

we seek to estimate the naturalistic behavior set $\mathcal{N}$. We formulate $\mathcal{N}$ as a collection of sets

$$\mathcal{N} = \{N_0, N_1, \dots, N_T\} \tag{4}$$

indexed by time, such that each N_t is a subset of the state space $\mathcal{X}$ and $\mathcal{N}$ forms a tube over the full time horizon T.

The problem of representing sets of feasible trajectories is fundamentally one of forward reachability (Bansal et al., 2017). Robust approaches that seek a strict overapproximation, while safe, tend to be overly conservative, and rapidly lose value for long-term prediction or estimation (Rober et al., 2023). Bounded interval techniques suffer from the same issues (Ramdani and Nedialkov, 2011). Our approach can be viewed as a purely data-driven approximation of the naturalistic set, in line with some sample-based forward reachability methods (e.g. Lew et al., 2022; Thorpe et al., 2022).

2.2 Projection Into the Naturalistic Behavior Set

Because autonomous system trajectories do not demonstrate naturalistic behavior, the resulting behaviors may be unpredictable, and therefore unsafe, when operating

around humans (Teoh and Kidd, 2017). Thus, we seek to augment an autonomously generated trajectory (i.e., from an arbitrary autonomous planner) to make the trajectory behave more naturally. Specifically, given a trajectory ξ_a computed from an autonomous planner, we seek to project ξ_a into the learned naturalistic behavior set $\mathcal{N}$. Further, we seek a projection that satisfies dynamic feasibility with respect to (1) while remaining as similar to the original trajectory as possible.

3. METHOD

We propose a naturalistic projection technique that identifies a naturalistic behavior set from human behavior data and then projects trajectories into the representation in a dynamically consistent manner.

3.1 Naturalistic Behavior Set Identification

We presume access to a naturalistic dataset $\mathcal{D} = \{\xi^i\}_{i=1}^m$ as in (3) consisting of m trajectories of length T. Using $\mathcal{D}$, we define time-indexed datasets $\{\mathcal{D}_t\}$ at times $t = 0, 1, \ldots, T$ by gathering all states at time t across trajectories in $\mathcal{D}$. For each time $t = 0, 1, \ldots, T$, we then seek to learn $N_t \in \mathcal{N}$ using data

$$\mathcal{D}_t = \{\mathbf{x}_t^0, \mathbf{x}_t^1, \ldots, \mathbf{x}_t^m\}. \tag{5}$$

Representing the Naturalistic Behavior Set. We represent a naturalistic behavior set as a tube of convex hulls. This choice provides a number of practical benefits, including simplicity, flexibility, and data efficiency. First, taking a convex hull produces an equivalent set of linear inequality constraints and ensures that every data point is represented in it. Second, convex hulls do not assume a particular distribution of the underlying data, which is critical to capturing the wide variety of human behavior. Third, using convex hulls avoids requiring a large amount of data. Thus, our method provides the benefit of working on smaller datasets, as compared to more data-intensive learning methods. We discuss additional considerations regarding our choice of convex hulls, including those related to non-convexity, in Section 5.

Efficiently Computing Complex Hulls. In general, the worst-case computational complexity of producing a convex hull from m points in $\mathbb{R}^n$ is $O(m^{\lfloor n/2 \rfloor})$ (Barber et al., 1996). Our method computes a convex hull at each time step, which requires T computations. Thus, the overall computational complexity of naturalistic behavior set generation is $O(Tm^{\lfloor n/2 \rfloor})$.

We note the number of facets on the boundary of a convex hull can rise exponentially with n. Each facet corresponds to a half-space, so introducing additional facets similarly increases the number of constraints required to represent a convex polytope as a half-space intersection (i.e., linear inequality constraints).

To address this problem, we consider a transformation (e.g. using a selector matrix or principal component analysis) of the state $\mathbf{x}_t$. We call the transformed state a "hull state" in $\mathbb{R}^{n_c}$, where $n_c < n$, produced via $y_t = y(\mathbf{x}_t)$. In other words, for all $t = 0, 1, \ldots, T$, we seek to learn $N_t \in \mathcal{N}$ using data

$$\mathcal{D}_t = \{y(\mathbf{x}_t^0), y(\mathbf{x}_t^1), \ldots, y(\mathbf{x}_t^m)\}. \tag{6}$$

In general, having few hull states (i.e., small n_c) results in a computationally efficient algorithm. However, if more hull states are needed, polynomial-time approximation algorithms for computing convex hulls exist (Sartipizadeh and Vincent, 2016; Balestriero et al., 2022).

We represent N_t via the convex hull of hull states y_t^i,

$$N_t = H(\mathcal{D}_t) = \text{ConvexHull}(\{y_t^i\}_{i=1}^m). \tag{7}$$

The naturalistic behavior set $\mathcal{N}$ can then be constructed as a collection of convex hulls indexed in time, as in (4).

Forming Convex Constraints. Convex sets can be well-approximated as convex polytopes, which can be equivalently represented via half-space intersections. Thus, the condition $y_t \in N_t$ can be equivalently written as the (linear) half-space intersection inequality constraint

$$G_t y_t \leq h_t. \tag{8}$$

Expressing the naturalistic behavior set in this form allows us to utilize (8) as a constraint within an optimization problem.

3.2 Projection Into the Naturalistic Behavior Set

Next, we seek to make a given autonomous system trajectory ξ_a behave more naturally while retaining dynamic feasibility. This can be interpreted as projecting the trajectory into the naturalistic behavior set. We thus seek to identify a set of controls $\mathbf{u}_0, \ldots, \mathbf{u}_{T-1}$ that generates a naturalistic trajectory ξ similar to ξ_a subject to the dynamic constraints f. Let $\mathbf{x}_{init}$ be the initial condition at time $t = 0$ of ξ_a. We define the projection

$$\begin{aligned}
\min_{\mathbf{u}_0,\ldots,\mathbf{u}_{T-1}} \quad & \mathrm{d}(\xi_a, \xi) && \text{(9a)}\\
\text{s.t.} \quad & \xi = [\mathbf{x}_0^\intercal\ \mathbf{x}_1^\intercal\ \cdots\ \mathbf{x}_T^\intercal]^\intercal && \text{(9b)}\\
& \mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t), \quad t = 0, 1, \ldots, T && \text{(9c)}\\
& \mathbf{x}_0 = \mathbf{x}_{init} && \text{(9d)}\\
& y(\mathbf{x}_t) \in N_t, \quad t = 0, 1, \ldots, T, && \text{(9e)}
\end{aligned}$$

where $d : \Xi \times \Xi \to \mathbb{R}$ in (9a) is a distance metric for trajectories, e.g. the Euclidean distance or a trajectory similarity metric like the one proposed by Chen et al. (2011). The constraints in (9c) and (9e) enforce dynamic feasibility and the learned naturalistic behavior constraints, which can be modeled as linear inequalities as in (8). In addition, we note that we can easily augment (9) to enforce additional constraints such as control limits or safety restrictions.

The projection in (9) may be non-convex if the dynamics f are nonlinear. Nevertheless, a variety of well-studied techniques exist to identify local minimizers of non-convex problems like (9). We refer the reader to Nocedal and Wright (1999) for further details.

4. EXPERIMENTS

We demonstrate our naturalistic projection technique on real-world human driving data from the inD dataset (Bock et al., 2020).

4.1 The inD Dataset

The inD datset (Bock et al., 2020) records and labels naturalistic traffic data for vehicles, bicyclists, and pedestrians at four German intersections using a drone camera positioned above each intersection.

Each actor i's trajectory is annotated at time t with state

$$\mathbf{x}_t^{\mathrm{GT},i} = [\boldsymbol{p}_t^{i\intercal}\ \boldsymbol{v}_t^{i\intercal}\ \boldsymbol{a}_t^{i\intercal}\ \theta_t]^\intercal, \tag{10}$$

containing planar position $\boldsymbol{p}_t^i$ tracking the center of the actor, planar velocity $\boldsymbol{v}_t^i$, planar acceleration $\boldsymbol{a}_t^i$, and heading θ_t^i. Trajectory ξ^i for actor i is constructed as in (2) and states are sampled at 25 frames per second, with actor i being visible and recorded from the first frame in which actor i is visible, at $t = 0$, until the last frame in which it is visible, at $t = H^i$.

4.2 Identifying a Single-Task Dataset $\mathcal{D}$

The inD dataset contains naturalistic trajectories of actors performing a variety of tasks. We generate a single-task dataset $\mathcal{D}$ by defining a heuristic indicator function. Let $\mathcal{V}$ contain the indices of all moving vehicles. We formalize a filtering heuristic

$$h(\xi^i; \mathcal{S}, \mathcal{E}) = i \in \mathcal{V}\ \wedge\ \mathbf{x}_0^i \in \mathcal{S}\ \wedge\ \mathbf{x}_{H^i}^i \in \mathcal{E}. \tag{11}$$

To ensure nontrivial behavior, the first term of (11) considers only moving vehicles. The second and third terms further filter the naturalistic trajectories under the assumption that every actor moving from a start polygon $\mathcal{S}$ to an end polygon $\mathcal{E}$ performs the same task. We specify different polygons $\mathcal{S}, \mathcal{E}$ for each subsequent experiment.

4.3 Generating the Naturalistic Behavior Set

Our method requires naturalistic data over which we can compute convex hulls, so we define the hull state using the information available in (10). We first select dynamics f by modeling each moving vehicle actor as a point with mass M evolving according to planar double-integrator dynamics

$$\mathbf{x}_{t+1} = \begin{bmatrix} p_{x,t+1} \\ v_{x,t+1} \\ p_{y,t+1} \\ v_{y,t+1} \end{bmatrix} = \begin{bmatrix} p_{x,t} + \Delta t v_{x,t} \\ v_{x,t} + \Delta t F_{x,t}/M \\ p_{y,t} + \Delta t v_{y,t} \\ v_{y,t} + \Delta t F_{y,t}/M \end{bmatrix}, \tag{12}$$

where $\mathbf{u}_t = [F_{x,t}\ F_{y,t}]^\intercal$ are forces applied to the point mass. As (12) constitutes a linear equation in $\mathbf{x}_t$ and $\mathbf{u}_t$, we denote the dynamics as $\mathbf{x}_{t+1} = A\mathbf{x}_t + B\mathbf{u}_t$ for brevity. In practice, many systems of interest are differentially flat and admit a representation of state and control in which dynamics are linear (Sastry, 2013, Ch. 9).

Next, we define the hull state

$$y_t = y(\mathbf{x}_t) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \mathbf{x}_t \tag{13}$$

by extracting the two-dimensional position from state $\mathbf{x}_t$ of the dynamics in (12). As the full naturalistic behavior set $\mathcal{N}$ captures sets of positions over time, we neglect higher order kinematics, although they can be included in principle at the expense of additional computation. We note that including higher order kinematics is only possible given appropriate naturalistic data. Since (10) includes velocity and acceleration, we could use planar quadruple-integrator dynamics with jerk controls for experiments on this dataset. Using (13), we generate datasets $\{\mathcal{D}_t\}_{t=1}^T$ as described by (6). As a convex hull can only be generated from a dataset with at least $n_c + 1$ points, we specify T to be the maximum time satisfying $|\mathcal{D}_t| \geq n_c + 1$. Next, we generate the naturalistic set at each time, N_t, by computing the convex hull with the Quickhull algorithm (Barber et al., 1996) as described by (7). Finally, we build the full naturalistic behavior set $\mathcal{N}$ as in (4).

4.4 Framing the Projection Problem

For the subsequent experiments, we frame the projection of ξ_a into $\mathcal{N}$ as an alteration of (9),

$$\min_{\mathbf{u}_0,\ldots,\mathbf{u}_{H_a-1}} \quad \|\xi_a - \xi\|_2^2 \tag{14a}$$

$$\text{s.t. } \xi = [\mathbf{x}_0^\intercal\ \mathbf{x}_1^\intercal\ \cdots\ \mathbf{x}_{H_a}^\intercal]^\intercal \tag{14b}$$

$$\mathbf{x}_{t+1} = A\mathbf{x}_t + B\mathbf{u}_t\ \forall\ t \in \{0, 1, \ldots, H_a\} \tag{14c}$$

$$G_t y_t \leq h_t\ \forall\ t \in \{0, 1, \ldots, T\}, \tag{14d}$$

where ξ_a has horizon H_a and $T = |\mathcal{N}|$. Equation (14a) defines the similarity objective $d(\cdot, \cdot)$ as a Euclidean distance. Equation (14c) enforces (linear) planar double-integrator dynamics over the entire trajectory horizon H_a as described by (12). Note that we adjust the projection to account for the case where $T \neq H_a$. Equation (14d) describes naturalistic behavior constraints as linear inequalities, as described by (8).

We note that the Euclidean distance $d(\cdot, \cdot)$ is convex. Moreover, both constraints (14c) and (14d) are linear, indicating that (14) is a convex optimization problem. For this reason, we solve (14) with an efficient convex optimization library, CVXPY (Diamond and Boyd, 2016). If $\xi^i \in \mathcal{D}$ is dynamically feasible, then projection is guaranteed to find a dynamic feasibile trajectory $\xi \in \mathcal{N}$.

4.5 Curved Road

Fig. 1 depicts two curved roads separated by a median running through a T-intersection. We define $\mathcal{D}$ as including all moving vehicles beginning and ending in the second lane of the upper road, where $\mathcal{S}$ is given by the green circle and $\mathcal{E}$ by the red square. Filtering based on these criteria results in 39 trajectories. We expect the naturalistic behavior set to be influenced by the curve of the road. The set generated from $\mathcal{D}$ is shown in Fig. 1.

Analysis. The naturalistic sets $\{N_t\}$ begin compact but lengthen along the lane over time, indicating that vehicles drive at different speeds along this lane. At the most curved point in the lane, the naturalistic behavior set covers the outside portion of the lane but not the inside. This observation suggests that drivers naturally hug the outside of a curved lane.

Projection. Fig. 1 also depicts ξ_a, a constant-velocity trajectory moving straight through the second lane. We call this trajectory non-naturalistic because it is not within the naturalistic behavior set. In particular, we note that it gets abnormally close to the median, which could negatively impact safety or comfort. As expected, applying a naturalistic projection to ξ_a results in a trajectory that curves along the outside of the road, replicating the behavior we see from human drivers. Our method reproduces naturalistic behavior within this trajectory without explicitly modeling factors like comfort.

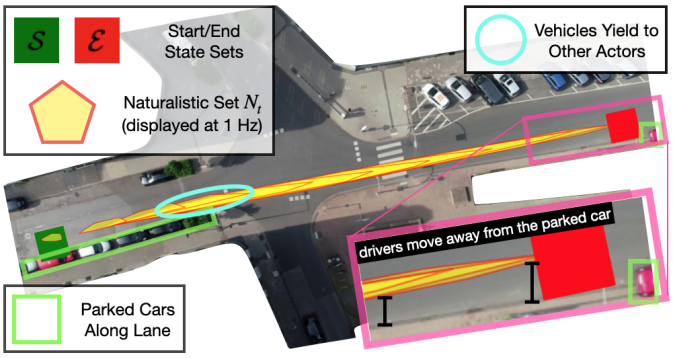

Fig. 2. We compute the naturalistic behavior set over two-dimensional position using all trajectories of moving vehicles that begin in the green square and end in the red. We generate the naturalistic behavior set based on these trajectories and plot it. Parked cars along the driving lane are boxed in light green. We additionally circle (in light blue) the area of the lane in which vehicles tend to yield to crossing road users.

4.6 Busy Intersection with Parked Cars

Fig. 2 captures a main road running through a four-way intersection and a pedestrian crossing, and the road is lined with parked cars. We define $\mathcal{D}$ as including all moving vehicles beginning and ending in the eastbound lane, where $\mathcal{S}$ is given by the green square and $\mathcal{E}$ by the red square. Filtering based on these criteria results in 49 trajectories. We expect the naturalistic behavior set to be influenced by the delays caused by naturally maintaining distance from parked cars (highlighted in light green) and by waiting for crossing pedestrians. The naturalistic behavior set generated from $\mathcal{D}$ is shown in Fig. 2.

Analysis. The naturalistic sets $\{N_t\}$ begin compact but lengthen along the lane over time, indicating a variety of speeds driving along this lane. We make two observations about the vehicles driving along the specified path. First, the naturalistic behavior set reflects a larger distance from parked vehicles. Near the beginning of the task, the naturalistic behavior set polygons are in the left portion of the lane, not centered. More dramatically, as shown in the inset of Fig. 2 near the end of the task, the naturalistic behavior set is longer along the center of the lane than along the right side. This observation indicates that vehicles move faster towards the center to avoid the parked red car. These behaviors indicate a naturalistic preference to maintain a larger distance from parked vehicles, without explicitly modeling it.

Second, we note that as vehicles approach the intersection, the naturalistic polygons lengthen until they consistently stretch from before the intersection (within the blue circle) until the end of the lane. This shape indicates that some vehicles pass through the intersection quickly, but that others wait to pass the intersection for longer. This effect can also be observed by noting that the polygons become more opaque in this region of the lane, indicating that a vehicle can be anywhere in the lane at that time. We attribute this observation to vehicles needing to wait for other actors to cross ahead of them. Thus, the convex hulls provide an over-approximation of the naturalistic set N_t at each time. In this experiment, actors have two distinct types of behavior: either they stop and yield or they can proceed along the lane, which results in an especially large over-approximation due to the branching nature of the task. This over-approximation motivates further work to represent these situations with unions of polygons at each time and propose a means of projecting onto that union.

5. CURRENT LIMITATIONS & FUTURE WORK

We present our approach as a practical first step towards identifying and using naturalistic behavior sets. Convex hulls enable capturing naturalistic data in a computationally efficient and data efficient representation, especially when compared to learned models. While convex hulls sometimes fail to be resilient to outliers, we note that the simplicity of the convex hull representation is extremely powerful due to the lack of assumptions on the underlying data and that techniques exist to reject outliers when necessary (Fischler and Bolles, 1981).

Non-Convexity. While our representation does not currently handle non-convexity of the naturalistic set, we note that many on-road maneuvers are not affected by this restriction (i.e. adjusting angular velocity thresholds for differently curved roads). To describe one example where assumptions of convexity are insufficient, consider a maneuver in which vehicles drive straight along a road which has a pothole. Naturalistic behavior may dictate that vehicles drive to either side of the pothole, meaning that no convex polygon can capture such a scenario without including behaviors that drive over the pothole. An over-approximated set would thus be inadequate for modeling this situation. Instead, we would need to identify two variants of naturalistic behavior within this non-convex maneuver: driving to the right of the pothole and driving to the left. Introducing multiple naturalistic behavior subsets requires adjustments to ensure our method works effectively. First, we would need to ensure that our method could represent each variant appropriately, and one possible method of doing so is using multiple convex hulls, though further work is needed to explore this open problem. Second, we would need to adjust the optimization problem (9) to include discrete variables for selecting between naturalistic subsets at a given time, and provide an efficient solver to the now mixed-integer problem.

Interaction. We note, as in Section 4.6, that our technique does not currently model interaction. We anticipate that for most cases, this case falls under the broader case of non-convexity. For a maneuver like that in Fig. 2, which involves cars which either proceed straight unimpeded or slow down for a pedestrian crossing, we anticipate that an autonomously generated trajectory would need to select a choice closer to one of these two subsets of naturalistic behavior. This example would lead to two possible convex hulls after a branching point in the maneuver. In Fig. 2, this branching point might in the blue circle, after which we would model two types of nonconvex naturalistic behavior as described previously.

Undesirable Road Behaviors. While we note that our method is able to capture naturalistic behavior effectively,

there remains an open question as to whether all naturalistic behavior should be reproduced. For example, consider the case of an intersection in which actors regularly violate the law by ignoring stop signs. In such a case, we would likely prefer that an autonomous planner prefer the legal requirements over enforcing the naturalistic constraints. As previously noted, we can encode further preferences of this sort, such as enforcing a safety set, by introducing new constraints as in (9) and (14).

6. CONCLUSION

In this work, we propose a method for computing a naturalistic behavior set over observations of human behavior with a set of time-indexed convex hulls. We subsequently describe an optimization problem that, when solved, projects trajectories into the naturalistic behavior set to produce dynamically feasible trajectories which resemble recorded human behavior. We demonstrate our method on real-world naturalistic driving data, and we show that it can capture and reproduce patterns in behaviors without explicitly modeling those patterns.

REFERENCES

Bagnell, J.A. (2015). An invitation to imitation. *Robotics Inst., Carnegie-Mellon Univ., Pittsburgh, PA, USA*.

Bajcsy, A., Siththaranjan, A., Tomlin, C.J., and Dragan, A.D. (2021). Analyzing human models that adapt online. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2754–2760. IEEE.

Balestriero, R., Wang, Z., and Baraniuk, R.G. (2022). Deephull: Fast convex hull approximation in high dimensions. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 3888–3892.

Bansal, S., Chen, M., Herbert, S., and Tomlin, C.J. (2017). Hamilton-jacobi reachability: A brief overview and recent advances. In *2017 56th Annual Conference on Decision and Control (CDC)*, 2242–2253. IEEE Press.

Barber, C.B., Dobkin, D.P., and Huhdanpaa, H. (1996). The quickhull algorithm for convex hulls. *ACM Trans. Math. Softw.*, 22(4), 469–483.

Bobu, A., Bajcsy, A., Fisac, J.F., and Dragan, A.D. (2018). Learning under misspecified objective spaces. In *Conference on Robot Learning*, 796–805. PMLR.

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

Chen, J., Wang, R., Liu, L., and Song, J. (2011). Clustering of trajectories based on hausdorff distance. In *2011 International Conference on Electronics, Communications and Control (ICECC)*, 1940–1944. IEEE.

Diamond, S. and Boyd, S. (2016). Cvxpy: A python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83), 1–5.

Dragan, A.D., Lee, K.C., and Srinivasa, S.S. (2013). Legibility and predictability of robot motion. In *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 301–308. IEEE.

Fischler, M.A. and Bolles, R.C. (1981). Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6), 381–395.

Fridovich-Keil, D., Bajcsy, A., Fisac, J.F., Herbert, S.L., Wang, S., Dragan, A.D., and Tomlin, C.J. (2020). Confidence-aware motion prediction for real-time collision avoidance1. *The International Journal of Robotics Research*, 39(2-3), 250–265.

Huang, Z., Wu, J., and Lv, C. (2021). Driving behavior modeling using naturalistic human driving data with inverse reinforcement learning. *IEEE transactions on intelligent transportation systems*, 23(8), 10239–10251.

Kuefler, A., Morton, J., Wheeler, T., and Kochenderfer, M. (2017). Imitating driver behavior with generative adversarial networks. In *2017 IEEE intelligent vehicles symposium (IV)*, 204–211. IEEE.

Lew, T., Janson, L., Bonalli, R., and Pavone, M. (2022). A simple and efficient sampling-based algorithm for general reachability analysis. In *Learning for Dynamics and Control Conference*, 1086–1099. PMLR.

Nocedal, J. and Wright, S.J. (1999). *Numerical optimization*. Springer.

Ramdani, N. and Nedialkov, N.S. (2011). Computing reachable sets for uncertain nonlinear hybrid systems using interval constraint-propagation techniques. *Nonlinear Analysis: Hybrid Systems*, 5(2), 149–162. Special Issue related to IFAC Conference on Analysis and Design of Hybrid Systems (ADHS'09).

Rober, N., Katz, S.M., Sidrane, C., Yel, E., Everett, M., Kochenderfer, M.J., and How, J.P. (2023). Backward reachability analysis of neural feedback loops: Techniques for linear and nonlinear systems. *IEEE Open Journal of Control Systems*, 2, 108–124.

Sadigh, D., Sastry, S.S., Seshia, S.A., and Dragan, A. (2016a). Information gathering actions over human internal state. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 66–73.

Sadigh, D., Sastry, S., Seshia, S.A., and Dragan, A.D. (2016b). Planning for autonomous cars that leverage effects on human actions. In *Robotics: Science and systems*, volume 2, 1–9. Ann Arbor, MI, USA.

Sartipizadeh, H. and Vincent, T.L. (2016). Computing the approximate convex hull in high dimensions. *ArXiv*, abs/1603.04422.

Sastry, S. (2013). *Nonlinear systems: analysis, stability, and control*, volume 10. Springer.

Song, J., Ren, H., Sadigh, D., and Ermon, S. (2018). Multi-agent generative adversarial imitation learning. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.

Teoh, E.R. and Kidd, D.G. (2017). Rage against the machine? google's self-driving cars versus human drivers. *Journal of Safety Research*, 63, 57–60.

Thorpe, A., Lew, T., Oishi, M., and Pavone, M. (2022). Data-driven chance constrained control using kernel distribution embeddings. In *Learning for Dynamics and Control Conference*, 790–802. PMLR.

Tian, R., Tomizuka, M., Dragan, A.D., and Bajcsy, A. (2023). Towards modeling and influencing the dynamics of human learning. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*, HRI '23, 350–358. Association for Computing Machinery, New York, NY, USA.