

Interactively Inferring Intent against Driver Models Constrained by Metric Temporal Logic

Edward Schmerling, Marco Pavone

Abstract—This work addresses the problem of inferring the high-level behavior of a decision-making agent when the specific model governing its actions is unknown; only constraints on candidate models, in the form of metric temporal logic formulae, are known. We develop a worst-case tree search solution framework and apply it to the case of validating driver intent in multi-vehicle pursuit scenarios.

I. INTRODUCTION

Understanding the behavior of other drivers, using data gathered while on the road, is an essential task in driving — both human and autonomous. Much work has been devoted to developing models for driver behavior and then inferring parameters for these models based on observations from the field. Probabilistic action models have been developed which capture the inherent uncertainty in drivers’ route-choice decision making [1]; these models have traditionally been hand-constructed [2], [3] but some recent work has aimed to learn probabilistic driving models directly from data [4], [5]. A fundamental challenge in implementing these models within a decision making framework, however, is that anticipating the actions of any particular driver on the road may require fitting a unique set of parameters, or even inferring a model structure tailored to the driver’s “type.”

Although a driver’s underlying decision-making model may be unknown, we may still make higher level assumptions about how their actions might be constrained. For example, when cruising at highway speed, we may expect that a neighboring vehicle will: (a) maintain a speed bounded by ± 5 m/s with respect to ambient traffic, (b) not cut across multiple lanes of traffic in quick succession, and (c) maintain a buffer zone of at least a car length behind the vehicle ahead of it. These so-called “rules of the road” may differ according to the setting, and may even be violated under extreme circumstances, but under nominal conditions represent assumptions less restrictive than starting with a particular driver model. In this work we codify rules of the road according to metric temporal logic (MTL) formulae [6]. Loosely speaking, logical formulae admit only a subset of all possible paths that a driver may embark upon. The conjunction of successive logical formulae may be regarded as a way to restrict the space of allowable driver decision models. Indeed, temporal logic constraints have traditionally been used in the context of model checking and formal verification to affirm that a particular model, and any path it may generate, satisfies a designer’s specifications. Recently temporal logic has also been employed in the context of robotic motion planning to specify the high-level mission

requirements that a synthesized controller solution must satisfy [7], [8], [9].

In this work we address the problem of inferring intent within an environment of drivers, whose action models are a priori unknown, through control of an embedded ego vehicle. We assume only that these models satisfy a MTL set of road rules. One or more “target” vehicles may be constrained to satisfy additional high-level requirements, e.g., bounded pursuit of the ego vehicle (surveillance) or line-of-sight avoidance of its monitors (surveillance evasion). These target vehicles are assumed to value their individual objectives above all else, but otherwise satisfy the road rules if there is no direct conflict. We allow the ego vehicle to be unconstrained, but with an objective preference towards satisfying the rules, so that it may have sufficient control authority to identify the intentions of the targets. This setup evokes the authority granted to law enforcement officials; we will also refer to the ego vehicle as the “police car.” The motivation of this work is to not only recognize, but to actively seek out aberrant behavior on the road. Especially as cyber-physical systems gain prevalence in society, the possibility of software subterfuge must be kept in check. Routine failures, where a vehicle violates the rules of the road unprompted, might be passively observable, but detecting a hacked agent within an autonomous vehicle network will generally require a more interactive approach. The primary contribution of this work is to provide a *worst-case* control policy for the police car, executed in an iterative fashion, for verifying target vehicle intent. The notion of worst-case is motivated by the lack of a concrete driver model, even probabilistic, in this work. The worst-case assumption allows us to resolve the infinite family of candidate policies for every other vehicle into a single adversarial policy for which we may apply game theoretical computation tools. In this preliminary work we investigate a discretized minimax rule augmented by Monte Carlo tree search over a fixed time horizon.

II. TECHNICAL APPROACH

Before outlining our algorithmic approach, we briefly review metric temporal logic (MTL) which we use in this work to describe invariants over timed sequences of vehicle states. For the purposes of the next subsection, a state \mathbf{x} is a vector encoding the state of the entire road environment, containing all vehicle positions and velocities.

A. Metric Temporal Logic

MTL is built up inductively from a finite set of *atomic propositions* AP , the logical operators *negation* (\neg) and *conjunction* (\wedge), and the temporal operator *until* (\mathcal{U}_I) through the context-free grammar:

$$\phi ::= \text{True} \mid p \mid \neg\phi \mid \phi \wedge \phi \mid \phi \mathcal{U}_I \phi$$

Edward Schmerling is with the Institute for Computational & Mathematical Engineering, Stanford University, Stanford, CA 94305, schmerling@stanford.edu.

Marco Pavone is with the Department of Aeronautics and Astronautics, Stanford University, Stanford, CA 94305, pavone@stanford.edu.

where $p \in AP$ and $I \subset (0, \infty)$ is a time interval. An atomic proposition p is a statement that is assigned a value `True` or `False` for every state \mathbf{x} . In the former case, we say that \mathbf{x} satisfies p and write $\mathbf{x} \models p$. The standard operators of propositional logic may be built up from negation and conjunction in the usual fashion: for two formulae ϕ and ψ we define their *disjunction* $\phi \vee \psi = \neg(\neg\phi \wedge \neg\psi)$, *implication* $\phi \Rightarrow \psi = \neg\phi \vee \psi$, and *equivalency* $\phi \Leftrightarrow \psi = (\phi \Rightarrow \psi) \wedge (\psi \Rightarrow \phi)$. The operator \mathcal{U}_I introduces time dependence into logical formulae; a sequence of states (alternately, path) $\sigma(t)$ satisfies $\phi \mathcal{U}_I \psi$ at time t if there is a time $\tau \in t + I$ such that $\sigma(\tau) \models \psi$ and $\sigma(s) \models \phi$ for all $s \in (t, \tau)$. Then for convenience we define the temporal operators *eventually* $\diamond_I \phi = \text{True} \mathcal{U}_I \phi$, *always* $\square_I \phi = \neg \diamond_I \neg \phi$, and *unless* $\mathcal{W}_I \psi = \neg(\neg \phi \mathcal{U}_I \neg \psi)$. If we write $\sigma \models \phi$ for a path σ without explicitly stating a time t of application, it is assumed that $t = 0$.

B. MTL Formulae for Road Rules

It is unreasonable to assume that drivers are cognizant of the entire road state \mathbf{x} ; rules of the road are typically framed in a much more local perspective. To this end, we define a number of functions that map the environment state \mathbf{x} into the context of a particular car i . Let $s(i; \mathbf{x})$ denote the distance gap between car i and the car immediately ahead of it (we take $s(i; \mathbf{x}) = \infty$ if there is no car ahead of i). Let $v(i; \mathbf{x})$ denote the speed of car i , let $x(i; \mathbf{x})$ denote the longitudinal position of car i (e.g., along the length of a highway), and let $\ell(i; \mathbf{x})$ denote the lane in which car i is driving. We use these functions to construct the atomic propositions which form the basis of our MTL specifications.

For example, to express a rule that cars on the highway must always maintain speed between 25–35 m/s, we write

$$\sigma \models \bigwedge_i \square_{(0, \infty)} (v(i; \sigma) \geq 25 \wedge v(i; \sigma) \leq 35).$$

A rule that cars cannot switch lanes twice within a span of 5 seconds can be written as

$$\sigma \models \bigwedge_i \square_{(0, \infty)} \left(\bigvee_{l_1, l_2} \ell(i; \sigma) = l_1 \mathcal{U}_{(0, 5)} (\square_{(0, 5)} \ell(i; \sigma) = l_2) \right).$$

In this way the expressivity of MTL allows us to specify constraints on driver behavior without having to directly model their action policy. Let RR denote the set of all MTL formulae describing a formulation of the rules of the road; we define the conjunction $\phi_{RR} = \bigwedge \{\phi \mid \phi \in RR\}$. A target vehicle intent corresponds to an additional formula, e.g., bounded pursuit and lane matching by i_v of the police car j :

$$\psi = \square_{(0, \infty)} ((x(i_v; \sigma) < x(j; \sigma) \wedge x(i_v; \sigma) > x(j; \sigma) - 20) \wedge \diamond_{(0, 5)} \ell(i_v; \sigma) < \ell(j; \sigma)). \quad (1)$$

We assume that $\phi_{RR} \mathcal{W}_{(0, \infty)} \psi$ will be satisfied by the road state at all times. We aim to discover, by manipulating the state through ego vehicle action, states satisfying $\neg\psi$ (we have verified that the target vehicle does not display the intent corresponding to ψ) or $\neg\phi_{RR}$ (we have verified that the target displays intent ψ at the expense of following the road rules).

C. Problem Formulation and Brief Algorithmic Approach

We consider an action space consisting of a finite set of discrete accelerations/decelerations and instantaneous lane changes for all vehicles (including the ego vehicle) driving on a one-way highway. We fix a time horizon T within which the ego vehicle must establish $\neg\psi \vee \neg\phi_{RR}$ and consider discrete timesteps $k\Delta t$ for $k = 0, \dots, N$ (where $T = N\Delta t$). We suppose a worst-case Stackelberg leadership model in which the police car chooses its action before the rest of the vehicles conspire at each timestep to choose actions consistent with $\phi_{RR} \mathcal{W}_{(0, \infty)} \psi$ but delaying realizing a state satisfying $\neg\psi \vee \neg\phi_{RR}$ for as long as possible. This enables us to compute policies on a decision tree, where paths in the decision tree correspond to valid prefixes for ϕ_{RR} . We emphasize that this worst-case assumption is used only to specify a policy for the environment; the target vehicle, and especially the other ambient vehicles, may not actually choose the action sequences implied by the decision tree. When controlling the police car, we account for these deviations from the worst-case in an iterative fashion. Depending on how many vehicles are in the system and how stringent the road rules ϕ_{RR} are in terms of restricting the set of available actions, the branching factor of the decision tree may be very large — to combat this we (a) implement Monte Carlo tree search out to the horizon T after a certain tree depth, and (b) use a scoring function measuring how close a MTL constraint is to violation (e.g., distance of pursuit $x(j; \sigma) - x(i_v; \sigma)$ for example Eq. (1) above), if one exists, to evaluate nodes in the tree.

III. RESULTS

We will present results on a four lane ring road for both pairwise (police car + target) and three vehicle (police car + target + bystander obstacle vehicle) interactions demonstrating identification of the bounded pursuit with lane matching behavior Eq. (1). We consider target and bystander vehicle policies adhering to the worst-case Stackelberg optimum, as well as random policies within the bounds of MTL feasibility.




Fig. 1. We compute an action policy for the ego vehicle, in green, to disambiguate whether the red follower is in determined pursuit, or merely wandering within the bounds of standard road rules.

IV. CONCLUSIONS

Although this work represents only a first step within the problem framework outlined herein, it already suggests future extensions. Incorporating some degree of probabilistic logic would increase the expressivity of the framework, enabling more realistic setups, as well as making elements of decision tree evaluation more tractable by providing an inherent notion of score. Additionally, research along these lines may inform the specification of road rules for future autonomous vehicle networks that are more amenable to automated fault detection.

REFERENCES

- [1] Q. Yang and H. N. Koutsopoulos, "A microscopic traffic simulator for evaluation of dynamic traffic management systems," *Transportation Research Part C: Emerging Technologies*, vol. 4, no. 3, pp. 113–129, 1996.
- [2] K. Ahmed, M. Ben Akiva, H. Koutsopoulos, and R. Mishalani, "Models of freeway lane changing and gap acceptance behavior," *Transportation and traffic theory*, vol. 13, pp. 501–515, 1996.
- [3] T. Toledo, H. N. Koutsopoulos, and M. Ben Akiva, "Integrated driving behavior modeling," *Transportation Research Part C: Emerging Technologies*, vol. 15, no. 2, pp. 96–112, 2007.
- [4] T. Gindele, S. Brechtel, and R. Dillmann, "Learning context sensitive behavior models from observations for predicting traffic situations," in *Proc. IEEE Int. Conf. on Intelligent Transportation Systems*, 2013, pp. 1764–1771.
- [5] T. A. Wheeler, P. Robbel, and M. J. Kochenderfer, "A Probabilistic Framework for Microscopic Traffic Propagation," in *Proc. IEEE Int. Conf. on Intelligent Transportation Systems*, 2015, pp. 262–267.
- [6] R. Koymans, "Specifying real-time properties with metric temporal logic," *Real-Time Systems*, vol. 2, no. 4, pp. 255–299, 1990.
- [7] H. Kress-Gazit, G. E. Fainekos, and G. J. Pappas, "Temporal-logic-based reactive mission and motion planning," *IEEE Transactions on Robotics*, vol. 25, no. 6, pp. 1370–1381, 2009.
- [8] S. L. Smith, J. Tumova, C. Belta, and D. Rus, "Optimal path planning for surveillance with temporal logic constraints," *International Journal of Robotics Research*, 2011.
- [9] S. Karaman and E. Frazzoli, "Vehicle routing problem with metric temporal logic specifications," in *Proc. IEEE Conf. on Decision and Control*. IEEE, 2008, pp. 3953–3958.