Extracting Kinodynamic Constraints from Expert Driver Data for High-Speed, Mobile Robot Navigation in Off-Road Environments*

Eric R. Damm[1], Steven M. Willits[2], Ryan Arciero[3], Thomas M. Howard[1]

***Abstract*— High-speed, off-road driving demands rapid decision-making and adaptability to diverse environments, making it a domain where humans outperform autonomous robots. To enable robots to meet or exceed human driving performance in such environments, motion planners must be able to reason about the kinodynamic limits of their platforms in the context of the observed environment. In this paper, we present a statistics-based method for generating kinodynamic constraint lookup tables from human-expert demonstration. With this method, we aim to expedite in-field development of autonomous robot systems by eliminating lengthy model training and parameter tuning. We present a case where our method was used to investigate the impact of two types of kinodynamic constraints with both physics-based models and human-expert data on planned motions. To analyze the differences, we compare plans generated by a recombinant motion planning search space using the different constraint models. Data for these experiments originates from physical tests on a field robot built from a Polaris RZR: Side-by-Side platform outfitted with sensors and computing for high-speed, off-road navigation. Comparisons of solutions to 4,316 planning problems extracted from a separate set of logs indicate that an average of 28.8% of states in each solution generated by the physics-derived baseline exceeded the data-driven kinodynamic limits from the expert-driven model. Additionally, the velocity limits imposed by the expert-driven model were less conservative in some regions of the curvature-roll-pitch space, leading to 36.6% of solutions exhibiting higher average speeds than the physics-derived baseline.**

Fig. 1: Polaris RZR: Side by Side outfitted with sensors and computers in a mixed-feature environment. Kinodynamic constraints were determined via human-expert driving to enable the vehicle to autonomously navigate off-road environments at high speeds.

I. Introduction

For humans to drive safely at high speeds in off-road environments, a vehicle operator must be able to plan paths that prevent the vehicle from tipping over. For autonomous ground robots, this requires established system models that reason about the influence of roll, pitch, velocity, and steering curvature on the magnitude and direction of centripetal forces. In these scenarios, it is challenging to define physics-based models that enforce kinodynamic constraints compared to operations in more structured environments [7]. In practice, simplifying assumptions must be made about these models, and the inaccuracies of the predictions are accounted for by added factors of safety. Hand-tuned alternatives may also be useful for planning kinodynamic constraint-aware trajectories, however these models may be difficult to parameterize and require added personnel time to tune [10], [4], [9]. This paper presents a statistical method for generating kinodynamic constraint lookup tables without the need for extensive data collection, model training, or parameter tuning. The work comes as a response to recent research leading to the development of Kinodynamic Efficiently Adaptive State Lattices (KEASL), a motion planning algorithm that applies kinodynamic constraints in a recombinant motion planning search space [2]. The work herein was developed during in-field KEASL testing to enable rapid analysis of the impact on system behavior by different kinodynamic constraint parameters. We present a case in which our approach was used to compare planned motions from KEASL when a data-driven driver profile and a physics-based profile were encoded in the search space. In this case, roll, pitch, and steering curvature were chosen as the velocity constraint parameters. However, other parameters, such as terrain semantics, could be utilized to accommodate various analysis criteria. The results from this analysis show how the data-driven model is not simply the physics-based model with a specific safety factor applied. In some conditions, the physics-derived baseline generated plans that exceeded the

*Research was sponsored by the Defense Advanced Research Projects Agency (DARPA), and was accomplished under Contract Number HR001121C0189. Any opinions, findings, conclusions or recommendations expressed herein are those of the author(s) and do not reflect the views of the Defense Advanced Research Projects Agency or Carnegie Mellon University.

[1]Eric R. Damm and Thomas M. Howard are with the Robotics and Artificial Intelligence Laboratory, Hajim School of Engineering and Applied Sciences, University of Rochester, Rochester, NY, USA edamm@ur.rochester.edu

[2]Steven M. Willits is with Carnegie Mellon University, Pittsburgh, PA, USA

[3]Ryan Arciero is with Arciero Racing, Foothill Ranch, CA, USA

limits imposed by the data-driven model. In other cases, the experiments show conditions where the physics-based model is overly conservative and led to longer expected traversal times compared to the data-driven model.

II. Technical Approach

This research presents a statistics-based method for deriving kinodynamic constraints from human-expert drivers to enable rapid system performance analysis and development. Previous research found success in applying similar lookup-table methods to dynamic robot systems [8], [1]. This method was implemented in the field during KEASL testing, where different constraint parameter combinations were applied to the search space. To provide a basis to compare to when using expert-driven constraints, physics-derived constraints were tested by using functions derived from physical interpretations of the system. The equations and free-body diagrams defining the physics-derived constraints can be found in Appendix I [3]. Equation 1 and Figure 6a show the resulting velocity constraints from the force balance between gravity and the centripital force on the robot. Equations 2 and 3, and Figure 6b show the resulting velocity constraints from ensuring a braking distance of 2 vehicle lengths.

For both the data-driven and physics-derived profiles, two lookup tables were generated. The first is a 1-dimensional table that applies velocity limits based on the pitch of the vehicle, and the second is a 2-dimensional lookup table that applies velocity constraints based on the roll and steering curvature of the vehicle. Note that steering curvature is the reciprocal of the radius of curvature.

A. Expert Driven Kinodynamic Constraints

The profile generation process involved scraping the vehicle states and dividing the driven velocities into bins according to the respective vehicle roll and pitch values. Subsequently, the top 1% of values within each bin were averaged, leading to the lookup tables containing the upper limit of safety from the expert drivers. For this case, bins without data were given a velocity constraint of 0m/s because the drivers either avoided those maneuvers, or there was insufficient data to safely assign an informed velocity constraint. While not utilized here, interpolation can fill in missing constraint values when employing higher dimension or higher resolution lookup tables, or when dealing with more sparse datasets.

This method allowed us to capture the strongest aspects of each driver's abilities and combine them into a single driver profile, while remaining below the absolute threshold of safety. These resulting expert-driven constraints are shown in Figures 2b and 2e.

The pitch-dependent velocity constraints for both the data-driven and physics-derived driver profiles follow similar trends of consistently higher speeds at negative pitch (uphill) values, and lower constraints at positive (downhill) pitch values. This is expected from the expert drivers, as maintaining momentum uphill and limiting momentum downhill are key parts in successful traversal of steep hills on non-flat, low-friction surfaces like dirt and sand. This is also expected from the physics-derived model because of the aid in braking from the gravitational component when going uphill, and the inverse effect on braking when going downhill.

The roll and curvature dependent velocity constraint plots differ considerably, with the expert drivers maintaining higher speeds across primarily low roll values and all curvature values, and the physics-derived model preferring the opposite. While the physics-based model dictates high allowable velocities for all roll values at zero steering curvature, the human drivers either maintained low speeds through those scenarios or avoided them altogether. This nuance further motivates generating kinodynamic constraints from human-experts instead of predicting system limitations through physics-based models.

III. Experimental Design

Experiments for expert-driver data collection and autonomous testing were performed in multiple locations consisting of trail and off-road driving with distinct hazards and features including high grass, ditches, sloping hills, bushes, and boulders. Figure 3 shows images of some locations from a front-facing camera on the robot during autonomous traversal. These locations were chosen because of the diverse sets of environmental features.

Fig. 3: Images from the robot's perspective of four different environments during autonomous traversal. Top Left: An environment in a valley containing trees, and a dried riverbed with steep edges. Top Right: An open, sandy area, littered with boulders and small rocks. Bottom Left: A non-flat, sandy area with dense clusters of bushes. Bottom Right: An obstacle-cluttered area with steep hills and drop-offs.

The first part of experimentation involved data collection for the expert-driver kinodynamic profile. To do this, the experts were instructed to drive quickly, while still maximizing safety, through different locations while remaining below a velocity threshold of 12 m/s. The set of expert driven data contained samples from driving in a wide range of driving scenarios. Both kinodynamic constraint profiles were generated at the same resolution, and with the same

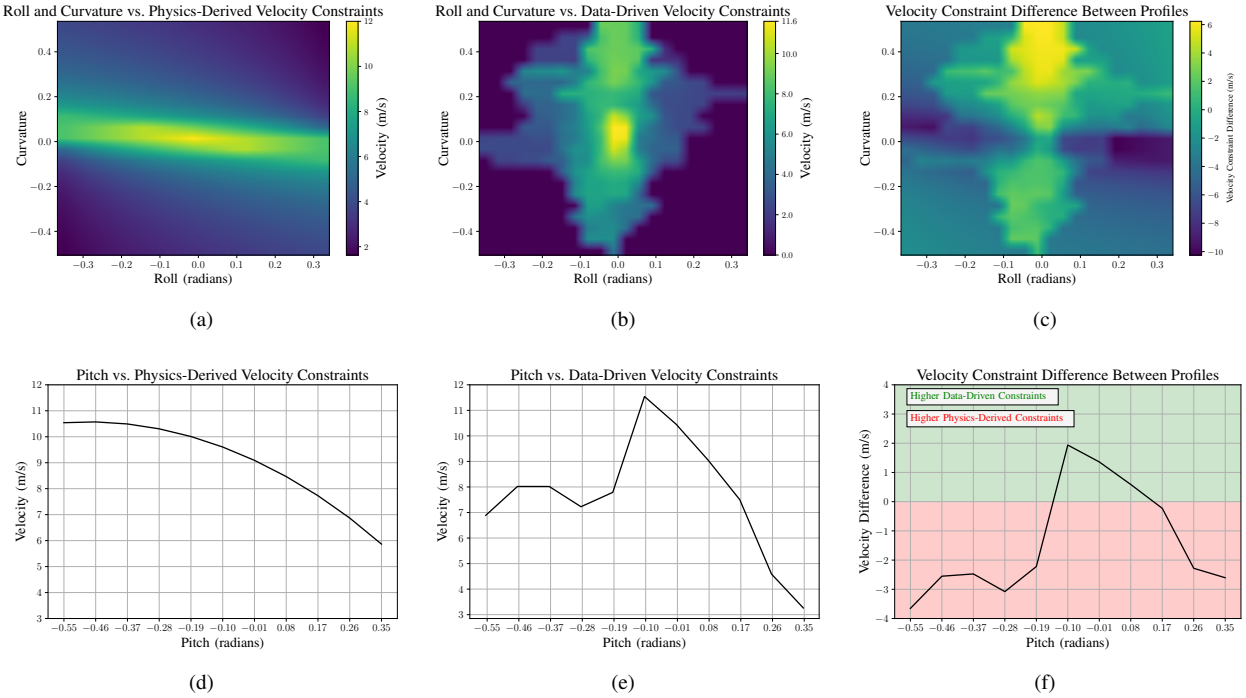

Fig. 2: (a) Surface plot of the physics-derived velocity constraints as a function of roll and steering curvature. (b) Surface plot of the data-driven velocity constraints as a function of roll and steering curvature from the expert-driven demonstration. (c) Surface plot of the differences in velocity constraints between the data-driven and physics-derived roll and curvature dependent velocity constraints. Yellow indicates higher allowable velocities from the data-driven model. (d) Plot of the physics-derived velocity constraints for a required braking distance of two vehicle lengths. (e) Plot of the data-driven velocity constraints as a function of pitch from the expert-driven demonstration. (f) Plot of the differences in velocity constraints between the data-driven and physics-derived pitch dependent velocity constraints. The green and red area indicate higher and lower data-driven constraints, respectively.

bin edges. The physics-derived profile was subject to the same 12 m/s max velocity constraint as the expert-drivers.

Following the lookup table generation, further field-testing was performed to compare the performance of the data-driven and physics-derived driver profiles. Both sets of kinodynamic driver profiles were implemented with KEASL, and 4,316 real-world planned trajectories were collected from autonomous driving through the different locations. To generate these trajectories, a primary goal point was set far beyond the perception horizon. As the robot autonomously traversed toward the primary goal point, a new primary plan would be generated at 1 Hz. Accompanying every primary plan were eight additional shorter plans toward evenly spaced goal points around the vehicle at the perception horizon limit. The horizon was considered a range of 150 meters from the vehicle for these experiments. The 8 additional trajectories generated from each planning cycle were accumulated and used for the analysis discussed in Section IV.

KEASL was given a maximum allowable planning time of 1 second per plan. The map layers considered during search were an obstacle map (a map indicating whether cell contains a lethal obstacle), and a height map (a map with the ground height at each cell). The accumulated plans were compared using the total time of traversal as the performance metric.

IV. Results

In 36.6% (1,580) of the generated trajectories, the expert driven profile generated plans with shorter traversal times than those from the physics-derived limits. On average 28.3% of states in each trajectory were in violation of the roll and curvature dependent velocity constraints, 0.7% of states in each trajectory were in violation of the pitch dependent velocity constraints, and 28.8% were in violation of some combination of the two. 100% of trajectories generated by the physics-derived profile violated kinodynamic constraints of the expert-driven profile (1,462 pitch-dependent constraint violations, 4,314 roll and curvature dependent constraint violations). This shows that, although the geometric kinodynamic limits are technically safe, the expert drivers either treat those scenarios with more caution, or avoid them altogether. Figures 4 and 5 depict the constraint violation data.

Although planning time was not considered as a primary metric for kinodynamic profile performance comparison, it was still recorded. The plans were generated from the expert profile in an average time of 0.85 seconds and the physics-derived profile in an average time of 0.88 seconds. Both profiles had median planning times of 1.00 seconds,

because Anytime A* was used as KEASL's underlying search algorithm [6], [5]. In 51.5% (2,221) of the cases, the expert profile generated plans faster than the physics-derived kinodynamic profile. This is due to the expert driven profile restricting the velocities in more cases than the physics-derived profile, causing KEASL to only expand from edges containing low vehicle roll. The ranges of planning times are similar and expected because there was no difference to the algorithmic methods used by KEASL for the search process.

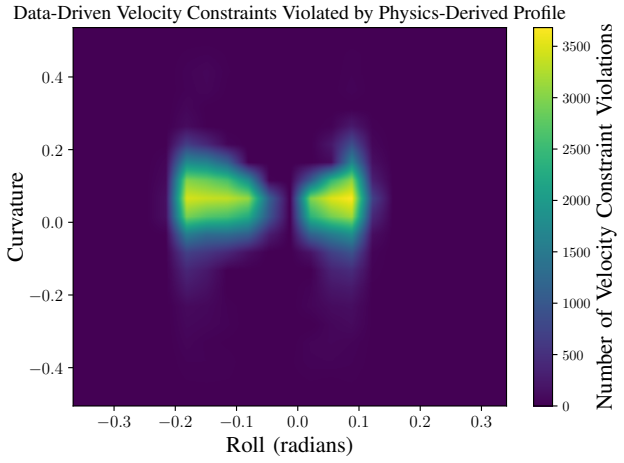

Fig. 4: Surface plot showing the number of physics-derived profile trajectories in violation of each expert profile roll and curvature dependent velocity constraint.

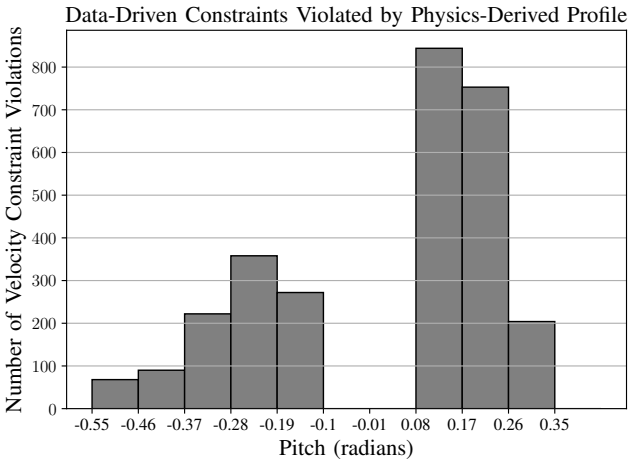

Fig. 5: Plot showing the number of physics-derived profile trajectories in violation of each expert profile pitch dependent velocity constraint.

V. Discussion

The results indicate that the kinodynamic constraints generated from the expert-driven data lead to shorter traversal times than those generated by geometric models in 36.6% of 4,316 trajectories generated from real-world experimental data. Although the physics-derived profile from Section II does not account for every kinematic constraint, the underlying model is meant to be a general representation of the upper limits of vehicle safety for the autonomous vehicle used during field-testing. Additionally, the range of traversal times remains consistent between both the data-driven approach and the physics-derived approach, indicating a sound model that accurately captures the physics of the system. Our results do not indicate a need to replace accurate theoretical models, but they do indicate that generating kinodynamic lookup tables from expert-driven data is an effective method for estimating the system constraints when there are opportunities to do so. Combining physics-derived and data-driven kinodynamic constraints into one profile is another area for exploration, although it is not discussed here. 100% of trajectories from the physics-derived profile violated constraints set by the data-driven profile, and an average of 28.8% of individual vehicle states in each profile violated the data-driven profile constraints. This percentage suggests that there are regularly occurring gaps in what is technically considered safe and what humans consider safe in the scope of off-road driving. The majority of violated constraints were those dependent on roll and curvature (28.3% of each trajectory), with far fewer dependent on pitch (0.7% of each trajectory). This is due to the much steeper drop-off in the roll-curvature dependent constraints than the pitch dependent constraints shown in Figures 2c and 2f. These results show that, when available, human-demonstration should be used as the basis for determining kinodynamic constraints for high-speed, off-road motion planning.

APPENDIX I: PHYSICS-DERIVED KINODYNAMIC CONSTRAINT EQUATIONS AND FREE-BODY DIAGRAMS

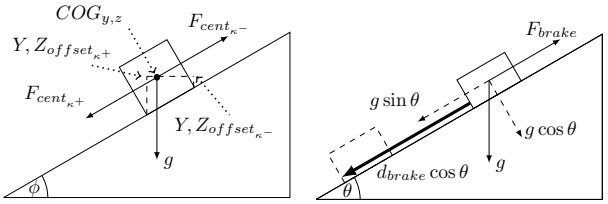

(a) (b)

Fig. 6: (a) A free-body diagram showing the front of a robot sitting at a roll of ϕ. (b) A free-body diagram showing the left side of a robot performing a braking maneuver at a pitch of θ.

$$V_{max} = \sqrt{\left| \frac{g * Y_{offset}}{\kappa * COG_z} \right|} \tag{1}$$

$$F_{brake} = g * (C_{friction} * cos(\theta) - sin(\theta)) \tag{2}$$

$$V_{max} = \sqrt{F_{brake} * d_{brake} * cos(\theta) * 2} \tag{3}$$

REFERENCES

[1] Christopher G. Atkeson, Andrew W. Moore, and Stefan Schaal. Locally weighted learning. *Artificial Intelligence Review*, 11, 1997.

[2] Eric R. Damm, Jason M. Gregory, Eli S. Lancaster, Felix A. Sanchez, Daniel M. Sahu, and Thomas M. Howard. Terrain-aware kinodynamic planning with efficiently adaptive state lattices for mobile robot navigation in off-road environments. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023.

[3] Antonio Diaz-Calderon and Alonzo Kelly. On-line stability margin and attitude estimation for dynamic articulating mobile robots. *The International Journal of Robotics Research*, 24(10):845–866, 2005.

[4] Markus Kuderer, Shilpa Gulati, and Wolfram Burgard. Learning driving styles for autonomous vehicles from demonstration. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2641–2646, 2015.

[5] Maxim Likhachev, David I Ferguson, Geoffrey J Gordon, Anthony Stentz, and Sebastian Thrun. Anytime dynamic a*: An anytime, replanning algorithm. In *ICAPS*, volume 5, pages 262–271, 2005.

[6] Maxim Likhachev, Geoffrey J Gordon, and Sebastian Thrun. Ara*: Anytime a* with provable bounds on sub-optimality. *Advances in neural information processing systems*, 16, 2003.

[7] Neal Seegmiller, Forrest Rogers-Marcovitz, Greg Miller, and Alonzo Kelly. Vehicle model identification by integrated prediction error minimization. *The International Journal of Robotics Research*, 32(8):912–931, 2013.

[8] Kazutoshi Tanaka, Satoshi Nishikawa, Ryuma Niiyama, and Yasuo Kuniyoshi. Immediate generation of jump-and-hit motions by a pneumatic humanoid robot using a lookup table of learned dynamics. *IEEE Robotics and Automation Letters*, 6(3):5557–5564, 2021.

[9] Xuesu Xiao, Zizhao Wang, Zifan Xu, Bo Liu, Garrett Warnell, Gauraang Dhamankar, Anirudh Nair, and Peter Stone. Appl: Adaptive planner parameter learning. *Robotics and Autonomous Systems*, 154:104132, 2022.

[10] Zlatko Zografski. Geometric and neuromorphic learning for nonlinear modeling, control and forecasting. In *Proceedings of the 1992 IEEE International Symposium on Intelligent Control*, pages 158–163, 1992.