Multi-Fidelity Recursive Behavior Prediction

Mihir Jain
Qualcomm AI Research

Kyle Brown*
Stanford University

Ahmed K. Sadek
Qualcomm AI Research

Abstract

Predicting the behavior of surrounding vehicles is a critical problem in automated driving. We present a novel game theoretic behavior prediction model that achieves state of the art prediction accuracy by explicitly reasoning about possible future interaction between agents. We evaluate our approach on the NGSIM vehicle trajectory data set and demonstrate lower root mean square error than state-of-the-art methods.

1 Introduction

Predicting the future motion of surrounding vehicles is a critical problem in autonomous driving research and development. One of the key difficulties associated with behavior prediction is *interaction* between traffic participants. Existing models vary in the way they reason about interactive driver behavior. Some models ignore interaction completely, predicting the future behavior of a target vehicle based solely on that vehicle's previous motion [2; 3; 11; 13; 15; 18]. Other models *implicitly* reason about interaction by conditioning motion prediction on the local traffic scene (including the current state or motion history of other nearby vehicles) [1; 8; 9; 14; 16; 17]. Still other models reason *explicitly* about interaction, addressing the prediction task from a game-theoretic perspective [10; 12; 19; 20; 21].

We present a novel game theoretic behavior prediction model that we call Multi-Fidelity Recursive Behavior Prediction (MFRBP). MFRBP achieves better prediction accuracy (as measured by root mean squared error) than previous state-of-the-art models by explicitly reasoning about possible future interaction between agents. The proposed algorithm employs a recursive trajectory prediction scheme inspired by the Level-k [7] and Cognitive Hierarchy [4] recursive reasoning paradigms.

This paper gives a condensed overview of the general Multi-Fidelity Recursive Behavior Prediction algorithm. We also discuss several specific implementations of our model, all of which incorporate specific elements from Convolutional Social Pooling for Vehicle Trajectory Prediction as proposed by Deo and Trivedi [8] in 2018. All experiments are conducted with the publicly available NGSIM data set.

2 Methods

Consider a traffic scene consisting of n agents. The motion history of the traffic scene from the initial time t_0 to the current time t can be compactly represented by the set of time histories $\mathcal{X}_{\text{hist}} = \{\mathbf{x}_1^{t_0:t}, \ldots, \mathbf{x}_n^{t_0:t}\}$, where $\mathbf{x}_i^{(t)}$ is the state of agent i at time t. The corresponding set of future trajectories (from time $t+1$ to time t_f) is $\mathcal{X}_{\text{future}} = \{\mathbf{x}_1^{t+1:t_f}, \ldots, \mathbf{x}_n^{t+1:t_f}\}$. The input to our model is $\mathcal{X}_{\text{hist}}$, and the output is $\hat{P}(\mathcal{X}_{\text{future}})$, a probability distribution over the future trajectories of all agents. In all experiments presented here, we choose to model $\hat{P}(\mathcal{X}_{\text{future}})$ as a set of Gaussian trajectory

* Work done during Kyle Brown's internship at Qualcomm AI Research. Kyle's reseach is also supported by the National Science Foundation.

predictions $\{(\hat{\mathbf{x}}_1^{(t+1:t_f)}, \Sigma_1), \ldots, (\hat{\mathbf{x}}_n^{(t+1:t_f)}, \Sigma_n)\}$, where $\hat{\mathbf{x}}_i$ and Σ_i represent the mean vector and covariance matrix of a Gaussian distribution over the future trajectory of the ith agent.

The general Multi-Fidelity Recursive Behavior Prediction algorithm is outlined in Algorithm 1. In the remainder of this section, we discuss two important characteristics of our model: *recursive prediction* and *multi-fidelity modeling*.

Recursive Prediction Our model employs a game-theoretic recursive prediction scheme inspired by Level-k Reasoning [7] and Cognitive Hierarchy [4]. Within this paradigm, a "level k" agent model assumes that all other agents act according to a level $k-1$ (or lower) agent model.

First, the algorithm assigns a reasoning level k_i and a corresponding sequence of policy models $(\pi_{i,0}, \ldots, \pi_{i,k_i})$ to each agent $i \in \{1, \ldots, n\}$ in the scene. At each level k (starting from $k = 0$ and proceeding upward) a predicted trajectory is generated for each agent $i \in \{1, \ldots, n\}$ if that agent's assigned reasoning level k_i is greater than or equal to k. The crucial detail is that, for $k > 0$, the level k trajectory prediction for a given agent is *explicitly conditioned* on the highest level (up to $k-1$) previously computed trajectory predictions associated with each other agent.

When the highest reasoning level has been reached (i.e. when $k = \max_{i \in \{1,\ldots,n\}} k_i$), the algorithm returns a set containing the final predicted trajectory for each agent.

Multi-Fidelity Behavior Modeling For a given traffic scene history $\mathcal{X}_{\text{hist}}$, the output of our model depends entirely on the reasoning levels and sequences of policy models assigned to the agents. These assignments are determined by the user-defined methods on lines 3 and 4 of Algorithm 1.

In our experiments, we show that this design flexibility can be used for *multi-fidelity modeling*, meaning higher-fidelity motion prediction for some agents than for others. *Multi-fidelity modeling* can be useful in applications (e.g. automated driving) where we may know and/or care more about some agents than others.

Algorithm 1 Multi-Fidelity Recursive Behavior Prediction

1: **procedure** MULTIFIDELITYRECURSIVEBEHAVIORPREDICTION($\mathcal{X}_{\text{hist}}$)
2:     **for** $i \in 1:n$
3:         $k_i \leftarrow$ ASSIGNREASONINGLEVEL($\mathcal{X}_{\text{hist}}, i$)
4:         $(\pi_{i,0}, \ldots, \pi_{i,k_i}) \leftarrow$ ASSIGNPOLICYMODELS($\mathcal{X}_{\text{hist}}, i, k_i$)
5:         $(\hat{\mathbf{x}}_{i,0}^{(t+1:t_f)}, \Sigma_{i,0}) = \pi_{i,0}\big(\{\mathbf{x}_j^{(t_0:t)} \mid j \in \{1, \ldots, n\}, j \neq i\}\big)$
6:     **for** $k \leftarrow 0, \ldots, \max_{i \in \{1,\ldots,n\}} k_n$
7:         **for** $i \in 1:n$
8:             **if** $k \leq k_i$
9:                 $(\hat{\mathbf{x}}_{i,k}^{(t+1:t_f)}, \Sigma_{i,k}) = \pi_{i,k}\big(\{(\mathbf{x}_j^{(t_0:t)}, \hat{\mathbf{x}}_{j,min(k_j,k-1)}^{(t+1:t_f)}) \mid j \in \{1, \ldots, n\}, j \neq i\}\big)$
10:     **return** $\{(\hat{\mathbf{x}}_{1,k_1}^{(t+1:t_f)}, \Sigma_{1,k_1}), \ldots, (\hat{\mathbf{x}}_{n,k_n}^{(t+1:t_f)}, \Sigma_{n,k_n})\}$

Policy models used in our experiments Our simple experiments incorporate three distinct policy models. The first two policy models condition only on motion *history* $\mathcal{X}_{\text{hist}}$. whereas, the third policy model explicitly conditions on previously computed level 0 trajectory predictions.

The *Constant Velocity (π_{CV}) model* simply predicts that a target vehicle will travel at constant velocity equal to the average velocity vector (both longitudinal and lateral components) over the last second. This can be thought of as "low-fidelity" motion prediction.

The *Convolutional Social Pooling (π_{CSP}) model* was originally proposed by Deo and Trivedi [8]. π_{CSP} combines a Long Short-Term Memory network (LSTM) encoder-decoder architecture with a convolution neural network (CNN) "social pooling" architecture to generate multimodal trajectory predictions. This model is depicted in blue in Figure 1. For each target, π_{CSP} accepts as input the state-histories of both the target vehicle and its neighbors (vehicles that fall in a rectangular region around the target). The final output is a set of six Gaussian distributions, each with an associated likelihood, that represent six possible trajectories. In our experiments, we take only the mode with the highest probability for prediction. In contrast to π_{CV}, π_{CSP} is a "high-fidelity" motion prediction model.

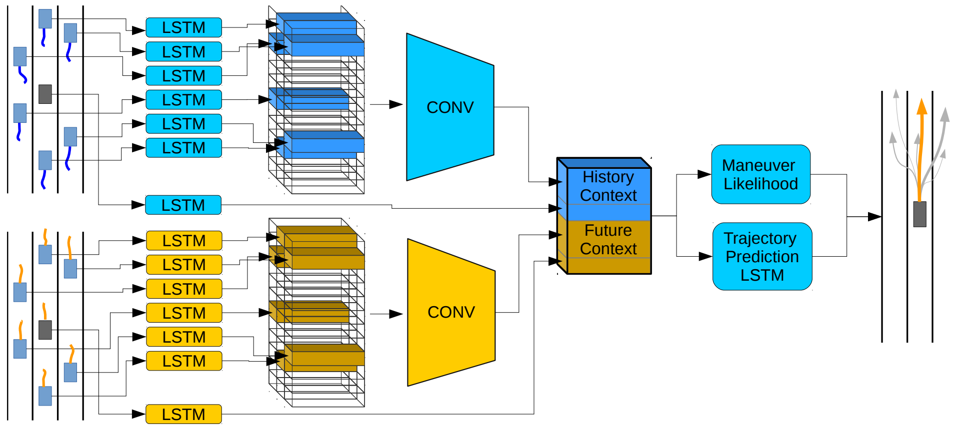

Figure 1: **Future-Conditional Convolutional Social Pooling ($\pi_{\text{FC-CSP}}$) Architecture.** All components shown in blue are part of the original π_{CSP} architecture as proposed by Deo and Trivedi [8]. The yellow elements represent the portion of the model that processes the predicted future trajectories.

The *Future-Conditional Convolutional Social Pooling ($\pi_{FC\text{-}CSP}$) model*, shown in Figure 1, is a novel extension of π_{CSP}, allowing the model to explicitly condition on predicted motion in addition to the motion history. This is accomplished by adding a "future" social pooling block (lower block), which is architecturally identical to the history social pooling block but receives as input the predicted future trajectories of the vehicles surrounding the target vehicle. The decoder LSTM layer in $\pi_{\text{FC-CSP}}$ receives the concatenation of both past and future "social context" tensors, and outputs a multimodal Gaussian trajectory distribution of the same form as the output of π_{CSP}. As with π_{CSP}, our experiments use only the highest probability mode.

3 Experiments

We perform three experiments to evaluate the performance of Multi-Fidelity Recursive Behavior Prediction. The first experiment is designed to compare our model against the existing state-of-the-art (the π_{CSP} baseline), while the other two experiments are intended to evaluate our algorithm in a setting that is more representative of an autonomous driving scenario.

All three experiments are conducted on the publicly available NGSIM I-80 [6] and US101 [5] datasets. There are 3 subsets of both US-101 and I-80, consisting of vehicle trajectories recorded by overhead camera at a frequency of 10 Hz. The test set consists of one quarter of the vehicle trajectories from each of the subsets of the US-101 and I-80 datasets. We split the trajectories into segments of 8 s (sampled at 10 Hz), where we use 3 s of track history and a 5 s prediction horizon. We use the common metric of *Root Mean Square Error* (RMSE) for evaluating the performance of our models.

3.1 Experiment 1: Level 1 Recursive Behavior Prediction

Our first experiment compares a specific implementation of Multi-Fidelity Recursive Behavior Prediction against the performance of the π_{CSP} baseline. This implementation is called Level 1 Recursive Behavior Prediction (L1-RBP). In L1-RBP, all agents are assigned a reasoning level of 1. The level 0 policy model for each agent is *Convolutional Social Pooling* (π_{CSP}), and the level 1 policy for each agent is *Future-Conditional Convolutional Social Pooling* ($\pi_{\text{FC-CSP}}$). Formally: $(k_i, \pi_{i,0}, \pi_{i,1}) := (1, \pi_{\text{CSP}}, \pi_{\text{FC-CSP}})\ \forall i \in \{1, \ldots, n\}$.

We train both the π_{CSP} (level 0 policy) model and $\pi_{\text{FC-CSP}}$ (level 1 policy) models jointly from scratch. As in the original π_{CSP} implementation, we use the leaky-ReLU activation with $\alpha = 0.1$ for all layers and use Adam for optimization.

Results for Experiment 1 We compare the output of (L1-RBP) against the π_{CSP} baseline. The left part of Table 1 shows RMSE values obtained at varying time horizons for π_{CSP} and L1-RBP. Note that ${\pi_{\text{CSP}}}^{+}$ corresponds to the original values reported by Deo and Trivedi [8] and ${\pi_{\text{CSP}}}^{*}$ is our own implementation of the baseline π_{CSP} model.

Table 1: **Comparison on RMSE.**

Prediction Horizon (s)	**Baselines** $\pi_{\text{CSP}}{}^{\dagger}$	$\pi_{\text{CSP}}{}^{*}$	**Experiment 1** L1-RBP	**Experiment 2** L1-MFRBP	**Experiment 3** L1-MFRBP (planning)
1	0.62	0.54	**0.53**	0.54	*0.54*
2	1.29	1.20	**1.19**	1.20	*1.19*
3	2.13	2.03	**1.95**	1.99	*1.95*
4	3.20	3.09	**2.87**	2.97	*2.88*
5	4.52	4.39	**3.97**	4.16	*4.01*

3.2 Experiment 2: Level 1 Multi-Fidelity Recursive Behavior Prediction

For experiment 2, we introduce *Level 1 Multi-Fidelity Recursive Behavior Prediction (L1-MFRBP)* (L1-MFRBP). L1-MFRBP targets the "ego-centric" prediction task, which is more representative of the autonomous driving use case: We randomly select a vehicle to treat as an "ego" agent, limiting the set of other agents in the scene to those within a plausible "sensor range" of this agent. This is repeated for many different "ego" agents during both training and testing. L1-MFRBP is identical to L1-RBP, except that it incorporates a multi-fidelity scheme wherein agents at the periphery of the designated ego agent's sensor range are assigned to a lower reasoning level ($k_i = 0$) and a lower fidelity ($\pi_{i,0} = \pi_{\text{CV}}$) policy model. The L1-MFRBP π_{CSP} and $\pi_{\text{FC-CSP}}$ policies are jointly trained. The training process includes generating and using lower-fidelity constant velocity trajectory predictions for the peripheral agents.

Results for Experiment 2 It would take too long to exhaustively evaluate L1-MFRBP on the full test set (i.e. by treating every single vehicle in turn as the ego agent). Instead, we sample enough ego agents to ensure that a single level 1 prediction can be computed for each vehicle in the test set. We report the average results of the level 1 predictions from 10 full iterations through the test set in this manner. Our results are presented in Table 1 alongside the results from Experiment 1. L1-MFRBP exhibits slight improvement over π_{CSP}, indicating that a multi-fidelity recursive prediction scheme can enhance performance even if the "low-fidelity" models are very naive. Improvement is more pronounced over longer prediction horizons.

3.3 Experiment 3: L1-MFRBP conditioned on Ego Future

Experiment 3 seeks to quantify the performance improvement that results from conditioning motion prediction on a candidate future ego trajectory. To explore this question, we use the ground truth future trajectory as a surrogate for the ego agent's "planned" trajectory. This takes the place of the ego agent's level 0 trajectory in a "planning-aware" version of L1-MFRBP. In other words, we "cheat" by allowing the model to observe the ground truth future trajectory for each designated ego agent during training and testing. We wish to make it clear, therefore, that experiment 3 is not meant to compete with the other models.

Results for Experiment 3 RMSEs for planning-aware L1-MFRBP are shown in italics in the last column of Table 1. The numbers are better than for L1-MFRBP, which suggests that conditioning on a planned trajectory (as in a real automated driving scenario) can improve motion prediction.

4 Conclusion

We have demonstrated that motion prediction in traffic scenes can be improved by recursively reasoning about future interaction between agents. We have also shown that multi-fidelity modeling can be effectively incorporated in the recursive prediction process.

Immediate directions for future work include extending our method to reason about multiple possible future scenarios (i.e. multi-modal scene motion prediction), incorporating a more flexible and diverse set of policy models, reasoning about input state uncertainty, and devising a more comprehensive set of experiments and performance metrics to evaluate our models. We aim to eventually implement a refined version of our algorithm on a real automated vehicle.

Acknowledgment We appreciate the support of Prof. Mykel J. Kochenderfer for his thorough reviews and insightful suggestions.

References

[1] F. Altché and A. de La Fortelle. An LSTM network for highway trajectory prediction. In *IEEE International Conference on Intelligent Transportation Systems (ITSC)*, pages 353–359, 2017.

[2] S. Ammoun and F. Nashashibi. Real time trajectory prediction for collision risk estimation between vehicles. In *IEEE International Conference on Intelligent Computer Communication and Processing*, pages 417–422, 2009.

[3] A. Barth and U. Franke. Where will the oncoming vehicle be the next second? In *IEEE Intelligent Vehicles Symposium*, pages 1068–1073, 2008.

[4] C. F. Camerer, T.-H. Ho, and J.-K. Chong. A cognitive hierarchy model of games. *The Quarterly Journal of Economics*, 119(3):861–898, 2004.

[5] J. Colyar and J. Halkias. US highway 101 dataset. *Federal Highway Administration (FHWA), Tech. Rep.*, (FHWA-HRT07-030), 2007.

[6] J. Colyar and J. Halkias. US highway I-80 dataset. *Federal Highway Administration (FHWA), Tech. Rep.*, (FHWA-HRT07-030), 2007.

[7] M. Costa-Gomes, V. P. Crawford, and B. Broseta. Cognition and Behavior in Normal-form Games: An Experimental Study. *Econometrica*, 69(5):1193–1235, 2001.

[8] N. Deo and M. M. Trivedi. Convolutional Social Pooling for Vehicle Trajectory Prediction. *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 1468–1476, 2018.

[9] N. Deo and M. M. Trivedi. Multi-modal trajectory prediction of surrounding vehicles with maneuver based lstms. *IEEE Intelligent Vehicles Symposium (IV)*, pages 1179–1184, 2018.

[10] E. Galceran, A. G. Cunningham, R. M. Eustice, and E. Olson. Multipolicy decision-making for autonomous driving via changepoint-based behavior prediction: Theory and experiment. *Autonomous Robots*, 41(6):1367–1382, 2017.

[11] J. Hillenbrand, A. M. Spieker, and K. Kroschel. A Multilevel Collision Mitigation Approach–Its Situation Assessment, Decision Making, and Performance Tradeoffs. *IEEE Transactions on Intelligent Transportation Systems*, 7(4):528–540, 2006.

[12] J. Hong Yoo and R. Langari. A Stackelberg Game Theoretic Driver Model for Merging. In *ASME Dynamic Systems and Control Conference*, 2013.

[13] N. Kaempchen, K. Weiß, M. Schaefer, and K. Dietmayer. IMM object tracking for high dynamic driving maneuvers. *IEEE Intelligent Vehicles Symposium*, pages 825–830, 2004.

[14] A. Khosroshahi, E. Ohn-Bar, and M. M. Trivedi. Surround vehicles trajectory analysis with recurrent neural networks. In *IEEE International Conference on Intelligent Transportation Systems (ITSC)*, pages 2267–2272, Nov 2016.

[15] R. Miller and Qingfeng Huang. An adaptive peer-to-peer collision warning system. In *IEEE Vehicular Technology Conference*, volume 1, pages 317–321, 2002.

[16] J. Morton, T. A. Wheeler, and M. J. Kochenderfer. Analysis of Recurrent Neural Networks for Probabilistic Modeling of Driver Behavior. *IEEE Transactions on Intelligent Transportation Systems*, 18(5):1289–1298, 2017.

[17] D. J. Phillips, T. A. Wheeler, and M. J. Kochenderfer. Generalizable intention prediction of human drivers at intersections. In *IEEE Intelligent Vehicles Symposium*, 2017.

[18] A. Polychronopoulos, M. Tsogas, A. Amditis, and L. Andreone. Sensor Fusion for Predicting Vehicles' Path for Collision Avoidance Systems. *IEEE Transactions on Intelligent Transportation Systems*, 8(3):549–562, 2007.

[19] D. Sadigh, S. Sastry, S. A. Seshia, and A. D. Dragan. Planning for Autonomous Cars that Leverage Effects on Human Actions. *Proceedings of Robotics: Science and Systems*, 2016.

[20] E. Schmerling, K. Leung, W. Vollprecht, and M. Pavone. Multimodal Probabilistic Model-Based Planning for Human-Robot Interaction. In *IEEE International Conference on Robotics and Automation*, 2018.

[21] J. H. Yoo and R. Langari. Stackelberg Game Based Model of Highway Driving. In *ASME Annual Dynamic Systems and Control Conference*, pages 499–508. ASME, 2012.