

Tree-structured Policy Planning with Learned Behavior Models

Yuxiao Chen¹, Peter Karkus¹, Boris Ivanovic¹, Xinshuo Weng¹, and Marco Pavone^{1,2}

Abstract—Autonomous vehicles (AVs) need to reason about the multimodal behavior of neighboring agents while planning their own motion. Many existing trajectory planners seek a single trajectory that performs well under *all* plausible futures simultaneously, ignoring bi-directional interactions and thus leading to overly conservative plans. Policy planning, whereby the ego agent plans a policy that reacts to the environment’s multimodal behavior, is a promising direction as it can account for the action-reaction interactions between the AV and the environment. However, most existing policy planners do not scale to the complexity of real autonomous vehicle applications: they are either not compatible with modern deep learning prediction models, not interpretable, or not able to generate high quality trajectories. To fill this gap, we propose Tree Policy Planning (TPP), a policy planner that is compatible with state-of-the-art deep learning prediction models, generates multistage motion plans, and accounts for the influence of ego agent on the environment behavior. The key idea of TPP is to reduce the continuous optimization problem into a tractable discrete Markov Decision Process (MDP) through the construction of two tree structures: an ego trajectory tree for ego trajectory options, and a scenario tree for multi-modal ego-conditioned environment predictions. We demonstrate the efficacy of TPP in closed-loop simulations based on real-world nuScenes dataset and results show that TPP scales to realistic AV scenarios and significantly outperforms non-policy baselines.

I. INTRODUCTION

A key challenge of motion planning for autonomous vehicles (AV)s is reasoning about the interaction between the ego vehicle and neighboring agents. The task is commonly divided into two subproblems: trajectory prediction for other agents, and ego motion planning with prediction. Trajectory prediction has seen substantial progress in recent years, coming from simple kinematic models [1] to powerful deep learning models [2]–[6] capable of generating high-quality, multi-modal predictions. However, motion planning with such high-capacity prediction models remains a challenge.

Typical AV planners simplify the ego behavior planning problem into *trajectory planning*, where a single trajectory is sought that minimizes an expected cost over *all* predicted futures within the planning horizon [7]. Such trajectory planning leads to overly conservative plans because it ignores the new information to be acquired about other agents in subsequent time steps, and the effects of ego actions on the behavior of other agents.

In fact, in decision-making literature, people have long been pursuing a closed-loop policy instead of an open-loop

plan as the former is shown to be optimal for problems such as Markov decision processes (MDP) [8].

To highlight the difference between a *motion policy* and a single trajectory, consider the scenario depicted on the right end of Fig. 1. The ego vehicle (blue) needs to move past the adjacent vehicle (red). Within the planning horizon, the red vehicle either maintains its current lane or cuts in front of the ego vehicle. A traditional planner seeks a single trajectory that performs well in both situations and thus will not choose to pass since lane change is a plausible choice for the red vehicle at current time. Alternatively, a motion policy may choose to nudge forward for the first part of the horizon, observe (by leveraging a forecasting model) the behavior of the adjacent vehicle, and then decide the subsequent motion. If the red vehicle already started a lane change, the ego vehicle brakes to avoid collision; if the red vehicle stays in its own lane, since the ego vehicle is already next to the red vehicle, a sudden lane change is very unlikely, the ego vehicle can then move past the red vehicle.

In the field of robotics, policy planning often takes the form of *contingency planning* [10], [11], with the key idea of preparing multiple plans for different likely futures. To differentiate the two concepts, we view contingency planning as a simplified policy planning variant with only one stage of reasoning. Unfortunately, continuous-space policy planning with high-capacity prediction models is computationally prohibitive, and despite various simplifications [13], [14], [19], practical policy planning for AVs remains an open problem.

In this paper we introduce Tree Policy Planning (TPP), a practical policy planner that meets the requirements of real-world autonomous driving, specifically, TPP is designed with the following desiderata:

- 1) compatible with state-of-the-art deep-learned predictive models for trajectory forecasting
- 2) plans with closed-loop ego-conditioning, i.e., factors in the ego vehicle’s influence on other agents

	DL	EC	Interp	MS	Scalability
Multi-mode planning [9]–[12]	×	×	✓	×	poor
Branch MPC [13]	×	✓	✓	✓	poor
POMDP [14]–[16]	×	✓	✓	✓	poor
Game theoretic [17], [18]	×	✓	✓	×	very poor
Lookout [19]	✓	×	✓	×	good
Cfo [20]	✓	✓	×	×	medium
TPP (ours)	✓	✓	✓	✓	good

TABLE I: Comparison of existing policy planning methods and TPP where “DL” is short for compatibility with deep-learned prediction models, “EC” stands for ego-conditioning, “Interp” is short for interpretable, “MS” stands for multistage, and “Scalability” means the capability to scale to complicated scenes with large number of agents.

¹The authors are with NVIDIA Research, NVIDIA Corporation, 2788 San Tomas Expressway, Santa Clara, CA. {yuxiaoc, pkarkus, bivanovic, xweng}@nvidia.com

²Marco Pavone is with the Department of Aeronautics and Astronautics, Stanford University, and with NVIDIA Research. {pavone@stanford.edu, mpavone@nvidia.com}

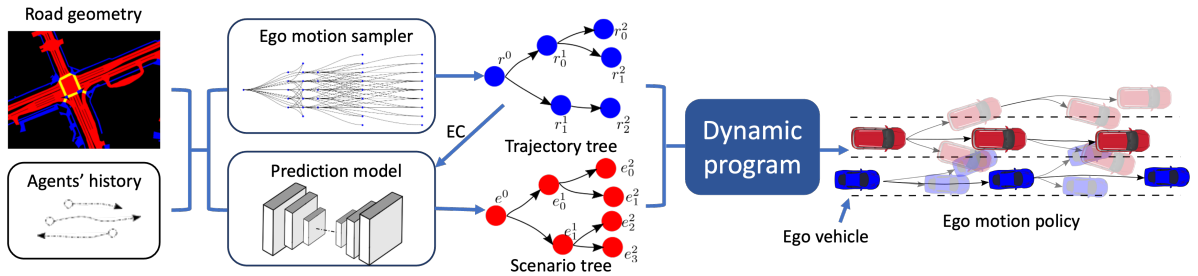


Fig. 1: Structure of TPP: the ego motion sampler generates the trajectory tree, which is fed to the prediction model for ego-conditioned prediction to generate the scenario tree. Then we solve for the ego motion policy via dynamic programming.

- 3) allows for multistage reasoning that leverages the agents' future reactive behavior
- 4) interpretable and easily tunable
- 5) scales to scenarios with a large number of agents

We compare existing policy planning works with TPP in Table I. To the best of our knowledge, TPP is the first algorithm that fulfils our desiderata.

The key idea of TPP is to convert the continuous space motion planning problem into a tractable finite-horizon MDP by constructing two trees: an *ego trajectory tree* carrying ego motion candidates, and a *scenario tree* that contains multiple behavior modes for neighboring agents. The algorithm is illustrated in Fig. 1. It consists of three components: an ego motion sampler that generates the ego trajectory tree, a deep-learned multi-stage prediction model that generates the scenario tree, and a dynamic program module that solves for the optimal policy with respect to the constructed MDP.

We evaluate TPP in closed-loop simulation based on a recently proposed realistic data-driven traffic model [21]. Our results show that TPP achieves good closed-loop performance in real-world scenarios from the nuScenes dataset [22], significantly outperforming non-policy planning baselines. TPP is highly parallelizable, and can run in real time.

II. RELATED WORKS

Various contingency planning algorithms have been proposed and they can be viewed as different approximations of policy planning. A complete policy planning algorithm should be capable of multi-stage reasoning with bi-directional interactions between the agent and the environment. Since the ego is always reacting to the environment, the concept of bi-directional interaction is also referred to as ego-conditioning, emphasizing on the environment reacting to the ego agent. The application of autonomous driving also asks for high quality trajectories, accurate prediction, interpretability, and scalability. All these features combined makes it difficult to realize a non-compromising policy planning algorithm for autonomous vehicles. As a result, different simplification methods have been proposed. Limited to one stage, [10] proposed a trajectory optimization method with one trajectory for each scene evolution mode, similar ideas are also seen in [11], [12] where a predefined set of modes are evaluated online. [19] uses a deep-learned prediction model to generate the modes. [13] performs multistage

reasoning via branch MPC and leveraged ego-conditioning, yet requires a simple differentiable prediction model.

Partially Observed Markov Decision Process (POMDP) [14]–[16] is a classic policy planning method, yet cannot generate high quality trajectories or use powerful deep-learned prediction models. Game theoretic approaches are well-known for reasoning about bi-directional interactions, yet assume some simple fixed behavior model of the environment, e.g. rational [18] or noisy Boltzmann [17]. [23] performs policy planning by training a neural network to react to the environment behavior, yet gives up on multistage reasoning, and is not interpretable and hard to tune.

Table I summarizes the various existing methods and their compromises, and we would like to show in the remainder of the paper that TPP closes the gap to a practical policy planning method for AV.

III. TREE POLICY PLANNING

TPP generates the ego motion policy in three steps as shown in Fig. 1. First, the ego motion sampler generates an *ego trajectory tree*, \mathbf{r} , given the current ego state and a lane map (when available). The ego trajectory tree consists of connected, dynamically feasible trajectory segments that together represent candidate ego-vehicle trajectories the planner can choose from. Second, the prediction module is a deep learned model that generates multi-modal predictions of the environment in the form of a *scenario tree*, \mathbf{e} . Finally, the dynamic programming module converts \mathbf{r} and \mathbf{e} into an MDP and finds an optimal policy via dynamic programming.

We use the following notation. Let s_r denote the state of the ego vehicle and s_e denote the state of the environment, including the state of nearby agents and the location of static obstacles. For the trajectory tree, r_j^i denotes the j -th node of stage i and the root node is simply r^0 . The nodes in the scenario tree are labeled the same way. We enforce that all nodes of the same stage contain trajectories of the same duration with t_0^i and t_f^i being the starting and ending time of the i -th stage. The trajectory of a any node other than the root needs to follow the trajectory of its parent node. For both trees, $\text{Par}(\cdot)$ denotes the parent of a node and $\text{Ch}(\cdot)$ denote the set of children of a node, e.g., $\text{Par}(r_1^1) = r^0$, $\text{Ch}(e_0^1) = \{e_0^2, e_1^2\}$ for the two trees in Fig. 1.

Next, we discuss the three components in detail.

A. Ego-motion sampler

The ego motion sampler generates a set of dynamically feasible trajectory options for planning to consider, and organizes them into a trajectory tree.

Terminal point sampling. We first sample a set of terminal points and then find a feasible trajectory to the terminal point. Terminal points are sampled from a fixed set of acceleration-steering values, as well as along (and near to) lane centerlines when lane information is available.

Spline fitting. We then fit a spline from the current ego state to the terminal point following [24]. With four boundary conditions on each polynomial, two 3rd order (cubic) polynomial can be uniquely determined, which allows us to directly compute the coefficient of $X[t]$ and $Y[t]$ in closed-form and calculate the trajectory from $X[t]$ and $Y[t]$. Finally we drop trajectories that violate dynamic constraints.

Tree construction. In practice, trajectory trees are grown stage by stage, starting from the current state as root r^0 . We set an upper bound on the number of children a node can have, and randomly drop children nodes if the limit is exceeded. In each stage we grow children branches in parallel, thus the time complexity is linear in stage number.

B. Multi-stage prediction model

A critical part of TPP is the prediction model that generates a scenario tree. For policy planning, we need a scene-centric, ego-conditioned, and multi-stage prediction model. We will first review these requirements and then show examples of deep-learning models satisfying our requirements.

Requirements. First, the prediction model needs to predict *multiple modes* of the trajectory distribution for each agent. Second, predictions must be *scene-centric*, i.e., outputting modes of the joint trajectory distribution of all relevant agents in the scene. Branches of our scenario tree corresponds to modes of the joint distribution so that our planner will be able to evaluate ego trajectory candidates against joint futures effectively. Examples for multi-modal scene-centric predictions include [6], [25]. Third, we need *ego-conditioning* (EC), i.e., prediction results should be conditioned on the ego vehicle's hypothetical future motion. Ego conditioning allows the planner to leverage the reactive behaviors of the environment towards the ego vehicle and is a known technique in prediction literature [3], [6]. However, training a EC prediction model remains challenging as the ground truth under counterfactual ego future motion is not known.

Lastly, we need to generate modes of the predicted trajectory distribution in *multiple stages*. That is, the generated scenario tree contains multiple stages with each node branching into multiple children after each stage. Multistage prediction allows the ego agent to depend its future motion on the future multimodal observation of the environment (similar to a decision tree). Details on converting existing models to enable multistage prediction can be found in the full version [26].

Model architectures. To evaluate TPP's flexibility w.r.t. different prediction models, we tested it with three vastly distinct prediction models. The first model is a rasterized

model with CNN backbone and a CVAE; the second model is the PredictionNet [2] model with a Unet backbone, and the third model is the Agentformer model [25]. All models are modified to enable multistage prediction and ego-conditioning. Details can be found in the full version [26].

C. Dynamic programming

The dynamic programming module takes the ego trajectory tree \mathbf{r} and the scenario tree \mathbf{e} , along with a handcrafted cost function l and finds an optimal ego motion policy.

Intuitively, our algorithm constructs a discrete MDP over \mathbf{r} and \mathbf{e} and finds an optimal policy through dynamic programming. More specifically, states of the MDP are nodes of the scenario/trajectory trees that represent the joint ego and environment state; transitions are given by the tree parent/children edges in the two trees; and rewards are defined by the integral of the cost over the time segment. We then solve for the optimal finite-horizon *policy* for the discrete MDP using dynamic programming, i.e., calculating the value function backwards in stage via the Bellman equation. With the optimal policy, the continuous-time trajectory is recovered from the trajectory tree by executing the trajectory segments according to the policy. In practice we use a cost function l with standard terms for collision avoidance, lane keeping, goal reaching, and ride comfort; and a dynamically extended unicycle model [27] for the vehicle dynamics.

Next we provide a formal discussion of the TPP algorithm.

IV. FORMAL DISCUSSION

A. Analysis without ego-conditioning

In our motion planning problem we are interested in minimizing the expected cumulative cost over a horizon:

$$\mathcal{J} = \int_0^T l(s(t), s_e(t)) dt,$$

where l is the running cost function that depends on the ego state s and the environment state s_e , T is the planning horizon. Our planning algorithm operates on the two trees \mathbf{r} and \mathbf{e} with the same number of stages, so the cumulative cost can be rewritten as

$$\mathcal{J} = \sum_{i=0}^N L_i(r^i, e^i), \quad (1)$$

where r^i and e^i are the nodes the ego vehicle and the environment take at stage i , and $L_i(r^i, e^i) = \int_{t_0^i}^{t_1^i} l(s^{r^i}(t), s^{e^i}(t)) dt$. Due to the stochasticity of the environment, we aim to minimize $\mathbb{E}_{e^i} [L_i(r^i, e^i)]$, where the expectation is on e^i , the actual branch the environment takes and the distribution is given as part of the prediction. For simplicity, we shall present the result assuming \mathbf{e} is not ego-conditioned, i.e., there is a single \mathbf{e} , and later show that the result can be easily extended to the case with ego-conditioning.

The end result is a policy that chooses which node to execute for the next stage given the current ego and environment node, i.e., $r^{i+1} = \pi^*(r^i, e^i)$, where π^* is the optimal policy. The form of the policy is determined by the assumption about

the information flow, i.e., during stage i , the ego vehicle can observe the environment's response e^i , and choose the ego motion among the children nodes of the current node r^i . To obtain the optimal policy, a value function is needed:

$$V(r^i, e^i) = \mathbb{E}_{e^{i:N}} \left[\sum_{k=i}^N L_k(r^k, e^k) \right], \quad (2)$$

which is the expected cost-to-go. The dynamic program goes backwards in stage and is trivial for the last stage:

$$V(r^N, e^N) = L_N(r^N, e^N).$$

For stage i where $i < N$, we have

$$V(r^i, e^i) = L_i(r^i, e^i) + \min_{r^{i+1} \in \text{Ch}(r^i)} \mathbb{E}_{e^{i+1}} [V(r^{i+1}, e^{i+1}) | e^i].$$

Since e^{i+1} is enumerable, the expectation is calculated as

$$\mathbb{E}_{e^{i+1}} [V(r^{i+1}, e^{i+1}) | e^i] = \sum_{e_j^{i+1} \in \text{Ch}(e^i)} \mathbb{P}[e_j^{i+1} | e^i] V(r^{i+1}, e_j^{i+1}),$$

where the branching probability of the scenario tree is given by the prediction model. Further define

$$\begin{aligned} Q(r^{i+1}, e^i) &:= L_i(r^i, e^i) + \mathbb{E}_{e^{i+1}} [V(r^{i+1}, e^{i+1}) | e^i] \\ &= L_i(r^i, e^i) + \sum_{e_j^{i+1} \in \text{Ch}(e^i)} \mathbb{P}[e_j^{i+1} | e^i] V(r^{i+1}, e_j^{i+1}). \end{aligned}$$

where $r^i = \text{Par}(r^{i+1})$. We name this function Q function as it is analogous to the Q function in reinforcement learning where here b^{k+1} is the action we need to choose.

The value function is then obtained backwards in stage:

$$V(r^i, e^i) = \begin{cases} L_i(r^i, e^i), & i = N \\ \min_{r^{i+1} \in \text{Ch}(r^i)} Q(r^{i+1}, e^i), & i < N \end{cases} \quad (3)$$

The optimal policy is then

$$\pi^*(r^i, e^i) = \arg \min_{r^{i+1} \in \text{Ch}(r^i)} Q(r^{i+1}, e^i). \quad (4)$$

Theorem 1 (Optimality): The policy in (4) is the optimal policy with respect to the cost in (1).

Proof: The proof is by induction. For stage N , the cost is a constant, π^* is trivially optimal, and $V(r^N, e^N)$ is the cost-to-go. For $0 \leq i < N$, assume $V(r^{i+1}, e^{i+1})$ is the expected cost-to-go function at stage $i+1$, then by the principle of optimality, the optimal node to take at stage $i+1$ is

$$r^{i+1,*} = \arg \min_{r^{i+1} \in \text{Ch}(r^i)} L_i(r^i, e^i) + \mathbb{E}_{e^{i+1}} [V(r^{i+1}, e^{i+1}) | e^i],$$

which is exactly $\pi^*(r^i, e^i)$. This in turn shows that

$$\begin{aligned} V(r^i, e^i) &= L_i(r^i, e^i) + \mathbb{E}_{e^{i+1}} [V(r^{i+1,*}, e^{i+1}) | e^i] \\ &= \min_{r^{i+1} \in \text{Ch}(r^i)} Q(r^{i+1}, e^i). \end{aligned}$$

This is true for all nodes in \mathbf{r} and \mathbf{e} , then by induction, $\pi^*(s^i, e^i)$ is the optimal policy for all $i \in \{0, 1, \dots, N\}$. ■

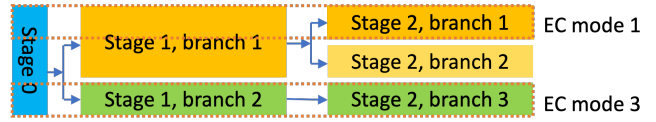


Fig. 2: Flattened trajectory tree as ego-conditioning modes

B. Extension to ego-conditioned prediction.

In the case where the scenario tree is ego-conditioned, the algorithm structure stays roughly the same. To generate the ego-conditioned trajectory prediction, we first run the ego motion sampler to generate the trajectory tree, then the trajectory tree is "flattened" to create the ego-conditioning (EC) modes. E.g., for trees in Fig. 1, the flattened EC modes are shown in Fig. 2. Under each EC mode, the prediction model generates a distinct scenario tree conditioned on the ego vehicle's motion. To keep the planning problem under ego-conditioning well-defined, the ensemble of EC scenario trees needs to satisfy a condition called causal consistency.

Definition 1 (Causal consistency): The ensemble of EC scenario tree is causally consistent (CC) if for any stage $i \leq N$, any two scenario trees under different EC modes, if the conditioned ego trajectories are identical up to stage i , the two scenario trees are identical up to stage i .

This condition guarantees the causality of the policy planning problem as for any particular stage i within the prediction horizon, no future information beyond stage i affects the prediction up to stage i . While causal consistency seems trivial, it may be violated by a multi-stage prediction model without special care. For example, if the EC trajectory is encoded as a whole for all stages, the scenario trees generated are not causally consistent.

Proposition 1 (Sufficient condition for CC): Assume that a multi-stage ego-conditioned prediction model generates the trajectory prediction stage by stage and the EC prediction is deterministic, then the generated ensemble of EC scenario trees is causally consistent.

Remark 1: For stochastic prediction models, the generated scenario trees are causally consistent if the random seeds are shared among all EC modes with the same EC trajectory.

The proof is omitted here. It is straightforward to show that the 3 prediction models in Section III-B are CC.

Given causally consistent EC predictions, only two terms change in the policy planning algorithm. First, all nodes of the scenario tree now depend on one of the nodes of trajectory tree of the same stage, and are denoted as $e^i(r^i)$ to emphasize the dependence. Note that the dependence is on r^i instead of the whole ego trajectory due to causal consistency. Second, the branching probability now depends on the ego node, and is denoted as $\mathbb{P}[e_j^{i+1} | e^i(r^i), r^i]$.

Formally, we are solving a finite-horizon MDP with a joint state of (r^i, e^i) and the cost defined in (1). The action space is $A(r^i) = \text{Ch}(r^i)$, the next ego maneuver, and the transition probability is

$$\mathbb{P}[(r^{i+1}, e^{i+1}) | (r^i, e^i(r^i)), r_j^{i+1}] = \begin{cases} \mathbb{P}[e^{i+1}(r_j^{i+1}) | e^i(r^i), r^i], & \text{if } r^{i+1} = r_j^{i+1} \\ 0, & \text{otherwise} \end{cases}$$

	Rasterized	PredictionNet	Agentformer
Mean ADE	0.67	0.25	0.33
Mean FDE	1.3	0.85	1.01

TABLE II: Prediction metrics average displacement error (ADE) and final displacement error (FDE) of the three models. The mean is taken over all modes.

Interestingly, the value function can be defined exactly the same way as

$$V(r^i, e^i(r^i)) = \mathbb{E}_{e^{i:N}} \left[\sum_{k=i}^N L_k(r^k, e^k(r^k)) \right]. \quad (5)$$

Similarly, the Q function is now

$$\begin{aligned} Q(r^{i+1}, e^i(r^i)) &:= L_i(r^i, e^i(r^i)) + \mathbb{E}_{e^{i+1}} [V(r^{i+1}, e^{i+1}(r^{i+1})) | e^i(r^i)] \\ &= L_i(r^i, e^i(r^i)) + \sum_{e^{j+1} \in \text{Ch}(e^i(r^i), r^{i+1})} \mathbb{P}[e^{j+1} | e^i(r^i), r^i] V(r^{i+1}, e^{j+1}), \end{aligned}$$

where $\text{Ch}(e^i(r^i), r^{i+1})$ denotes the set of children nodes following $e^i(r^i)$ on the tree associated with r^{i+1} . The optimal policy is then defined as

$$\pi^*(r^i, e^i(r^i)) = \arg \min_{r^{i+1} \in \text{Ch}(r^i)} Q(r^{i+1}, e^i(r^i)). \quad (6)$$

Theorem 2: Given a causally consistent ensemble of EC scenario trees, the policy in (6) is the optimal policy with respect to the cost function $\mathcal{J} = \sum_{i=0}^N L_i(r^i, e^i(r^i))$.

Proof: The proof essentially follows the same steps as the proof of Theorem 1. Starting from stage N , the cost is constant, thus π^* is trivially optimal and $V(r^N, e^N(r^N))$ is the cost-to-go. For all $0 \leq i < N$, assume for all r^{i+1} , $V(r^{i+1}, e^{i+1}(r^{i+1}))$ is the expected cost-to-go function at stage $i+1$. Then for any r^i , the cost-to-go to take any children node $r^{i+1} \in \text{Ch}(r^i)$ is $L_i(r^i, e^i(r^i)) + \mathbb{E}[V(r^{i+1}, e^{i+1}(r^{i+1})) | e^i(r^i)]$, which is exactly $Q(r^{i+1}, e^i(r^i))$. It follows that $\pi^*(r^i, e^i(r^i))$ gives the optimal b^{i+1} within $\text{Ch}(r^i)$ and the cost-to-go function of stage i is computed as

$$V(r^i, e^i(r^i)) = \min_{r^{i+1} \in \text{Ch}(r^i)} Q(r^{i+1}, e^i(r^i)).$$

The optimality of π^* is then proved by induction. ■

V. EXPERIMENTS

A. Training of prediction model

We configure TPP with three different prediction models. Apart from the model architecture, the experimental setup is the same. Prediction models are trained with the nuScenes dataset [28], comprised of 1000 scenes lasting for 20 seconds collected in urban areas of Boston and Singapore. For simplicity, only vehicles are considered. All prediction models use 2 stages and the branching factor is 4.

Loss function. Our training loss is a weighted sum of the following terms: prediction loss that penalizes the error between the predicted trajectories and the ground truth; EC collision loss that penalizes collisions between the predicted trajectories and the conditioned ego trajectory; collision loss that penalizes collisions between predicted trajectories of different agents in the scene; and additional standard regularization losses specific to the model architecture.

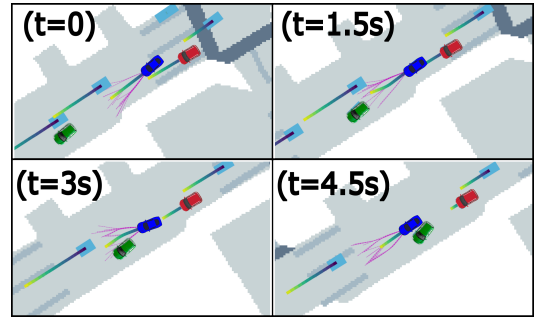


Fig. 3: Example scene from simulation: the ego vehicle (blue) cuts in front of the red vehicle and then swerve around the green static vehicle in front. The ego trajectory tree is shown in magenta, the colored line shows agents' actual trajectory.

Ego conditioning. One natural choice of the ego trajectory for ego-conditioning is the ground truth future trajectory. In addition, to expose the model to more diverse ego trajectories, we perturb the ground truth ego trajectory with noise generated via the *Ornstein–Uhlenbeck* process as alternative ego trajectories. However, since the ground truth for surrounding agents under the counterfactual ego motion is not available, for the lack of a better choice, we still use the original ground truth as the target for the EC trajectory prediction. We penalize collision between the predicted trajectories of surrounding agents and the conditioned ego trajectory so that the model learns to avoid the conditioned ego trajectory.

Training results. We report standard ADE/FDE prediction metrics for the three prediction models in Table II. The results show that all three models learn reasonable predictions.

B. Experimental setup

Closed-loop simulation. We evaluate TPP in BITS, a closed-loop AV simulation environment proposed in [21]. Each simulation scene is initialized with a scene from the nuScenes dataset and then evolves forward with a policy for each agent. One agent is chosen as the ego vehicle and controlled by TPP or its alternatives. All other agents runs the learned BITS policy that generates realistic and diverse behavior of road vehicles. Additionally, to incite more interactions between agents, we "spawn" new agents around the ego vehicle to challenge the planner throughout the simulation duration.

Metrics. We use three key metrics: collision rate, offroad rate, and area coverage. The first two are calculated as the percentage of time steps where the ego vehicle is in collision (offroad). The coverage metric is calculated with a kernel density estimation (KDE) procedure following [21], and it captures the effectiveness of the ego plan in terms of liveness and distance travelled. The evaluation is done on 100 scenes from the evaluation split of nuScenes (not used for training).

Baselines. We consider three baselines, the first two are non-contingent counterpart of TPP. For fair comparison, the two baselines share the same ego-motion sampler and ego-conditioning prediction model as TPP, the first baseline, named non-contingency robust (NCR), considers all possible trajectories predicted by the prediction model (similar to

Planning algorithm	Prediction model	Crash rate ↓	Offroad rate ↓	Coverage ↑
TPP (ours)	Rasterized	0.46%	0.64%	164
Non-Cont. Robust	Rasterized	0.80%	0.61%	163
Non-Cont. Greedy	Rasterized	1.23%	1.00%	168
TPP (ours)	PredictionNet	1.02%	0.10%	174
Non-Cont. Robust	PredictionNet	1.34%	2.07%	166
Non-Cont. Greedy	PredictionNet	1.60%	0.45%	170
TPP (ours)	Agentformer	0.97%	0.20%	174
Non-Cont. Robust	Agentformer	1.08%	0.63%	158
Non-Cont. Greedy	Agentformer	2.52%	0.98%	163
TPP (ours)	Agentformer (Gaussian)	3.92%	5.53%	150
Non-Cont. Robust	Agentformer (Gaussian)	1.40%	0.40%	144
Non-Cont. Greedy	Agentformer (Gaussian)	1.04%	0.25%	149
TPP (ours)	Kinematic model	1.75%	0.51%	160

TABLE III: Closed loop simulation results with TPP and the baselines, the default Agentformer model uses a discrete latent space and Agentformer (Gaussian) is its variant with a Gaussian latent space. TPP significantly outperform the baselines with the same ego trajectory tree and prediction model (bold means best).

[29]); while the second baseline, named non-policy greedy (NCG), only avoids the most likely prediction mode (similar to [7]). The third baseline uses a kinematic prediction model with 2 modes: maintaining speed and braking, and share the same multi-stage tree structure.

C. Results

Fig. 3 shows snapshots of the simulation with TPP where the ego vehicle (blue) cuts in front of the red vehicle and then change lane again to avoid the parked green vehicle. The magenta lines show the ego trajectory tree, and the colorful line shows the rollout trajectories. The quantitative results are shown in Tables III and IV, and we make the following observations. Simulation videos can be found [here](#).

TPP outperforms the baselines. As shown in Table III, TPP significantly outperforms the two non-policy baselines on crash and offroad rate under all 3 prediction models while achieving similar coverage. Even under a simple kinematic model, the performance is better than non-policy planners with sophisticated prediction models. TPP with deep-learned prediction models performs better than with a simple kinematic model, indicating the deep-learned prediction model benefits the closed-loop performance.

Better predictions do not mean better plan. Interestingly, from Table III, the "better" prediction model measured by ADE/FDE does not necessarily lead to better closed-loop performance. For example, the rasterized model is the simplest and "worst" model among the three, yet TPP with it outperforms the other two in terms of crash rate.

Ego-conditioning is useful. Table IV compares TPP with or without ego-conditioning. Ego-conditioning improves the closed-loop performance considerably in most cases, especially the coverage metric, hinting that the ego vehicle is more comfortable moving through traffic when it expects reaction from other agents.

Temporal consistency is important. Additionally, we tested the Agentformer model with Gaussian latent space, where the prediction requires sampling of the latent variable at every inference call, compared to using discrete latent space where

	Crash rate ↓	Offroad rate ↓	Coverage ↑
Rasterized	0.46% / 0.53%	0.60% / 1.05%	164 / 163
PredictionNet	1.00% / 0.63%	0.10% / 0.71%	174 / 151
Agentformer	0.97% / 1.30%	0.20% / 0.53%	174 / 163

TABLE IV: Ablation study of ego-conditioning, the first and second number are the metrics with and without ego conditioning. Most metrics see degradation when ego-conditioning is disabled (bold means better).

	Rasterized	PredictionNet	Agentformer
Prediction w/ EC	150ms	5.00s	2.60s
Prediction w/o EC	37ms	75ms	150ms
Ego-motion sampling		30ms	
Dynamic programming		11ms	

TABLE V: Runtime of the three key components

sampling is not needed. While the prediction ADE/FDE is similar, the closed-loop performance with the Gaussian latent is significantly worse. We suspect this is due to the poor temporal consistency between predictions of consecutive time steps. While both TPP and the two baselines saw performance degradation under the Gaussian latent, the policy planner relies on more sophisticated prediction/planning interaction, thus suffers more than the two baselines.

Runtime analysis. Table V shows the runtime of the three components of TPP on an Nvidia 3090 GPU (most computation is done on GPU). It should be noted that all runtime is recorded on unoptimized Python code and there is significant space of improvement. For example, the long runtime for PredictionNet is mainly due to rasterization, and we know its runtime can be as low as 5ms with the proper engineering [2], e.g., efficient CUDA code for rasterization. While the runtime for prediction models is not very informative and is not the focus of this paper, it is quite clear that ego motion sampling and dynamic programming is negligible comparing to the runtime of prediction model. With proper optimization, we believe that TPP can run in real time. We plan to release the code upon publication.

VI. CONCLUSION

We present TPP, a policy planner capable of generating multistage motion policies that react to the environment. It is centered around an ego trajectory tree and a scenario tree that predicts the environment behavior, both containing multiple stages. Ego-conditioning is applied to leverage the reactive behavior of the environment under the ego vehicle's presence. The closed-loop simulation result shows that TPP significantly outperform two non-policy benchmarks and the runtime test suggest that such sophisticated policy planner can be run in real time. Our planned future works include (1) study more flexible topology of the trees and enable event-triggered branching (2) improve the efficiency of computation, especially under ego-conditioning (3) adaptive and personalized policy planning.

ACKNOWLEDGEMENT

We would like to thank Bilal Kartal, Siva Hari, and Edward Schmerling for the insightful discussion.

REFERENCES

- [1] J. Kong, M. Pfeiffer, G. Schilbach, and F. Borrelli, "Kinematic and dynamic vehicle models for autonomous driving control design," in *2015 IEEE intelligent vehicles symposium (IV)*. IEEE, 2015, pp. 1094–1099.
- [2] A. Kamenev, L. Wang, O. B. Bohan, I. Kulkarni, B. Kartal, A. Molchanov, S. Birchfield, D. Nistér, and N. Smolyanskiy, "Predictionnet: Real-time joint probabilistic traffic prediction for planning, control, and simulation," *arXiv preprint arXiv:2109.11094*, 2021.
- [3] T. Salzmann, B. Ivanovic, P. Chakravarty, and M. Pavone, "Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data," in *European Conference on Computer Vision (ECCV)*. Springer, 2020, pp. 683–700.
- [4] N. Rhinehart, R. McAllister, K. Kitani, and S. Levine, "Precog: Prediction conditioned on goals in visual multi-agent settings," in *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019, pp. 2821–2830.
- [5] S. Casas, C. Gulino, S. Suo, K. Luo, R. Liao, and R. Urtasun, "Implicit latent variable model for scene-consistent motion forecasting," in *European Conference on Computer Vision (ECCV)*. Springer, 2020, pp. 624–641.
- [6] Y. Chen, B. Ivanovic, and M. Pavone, "Scept: Scene-consistent, policy-based trajectory predictions for planning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 17 103–17 112.
- [7] W. Xu, J. Pan, J. Wei, and J. M. Dolan, "Motion planning under uncertainty for on-road autonomous driving," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 2507–2512.
- [8] M. L. Puterman, "Markov decision processes," *Handbooks in operations research and management science*, vol. 2, pp. 331–434, 1990.
- [9] J. Hardy and M. Campbell, "Contingency planning over probabilistic hybrid obstacle predictions for autonomous road vehicles," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2010, pp. 2237–2242.
- [10] —, "Contingency planning over probabilistic obstacle predictions for autonomous road vehicles," *IEEE Transactions on Robotics*, vol. 29, no. 4, pp. 913–929, 2013.
- [11] W. Zhan, C. Liu, C.-Y. Chan, and M. Tomizuka, "A non-conservatively defensive strategy for urban autonomous driving," in *2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2016, pp. 459–464.
- [12] E. Galceran, A. G. Cunningham, R. M. Eustice, and E. Olson, "Multi-policy decision-making for autonomous driving via changepoint-based behavior prediction: Theory and experiment," *Autonomous Robots*, vol. 41, no. 6, pp. 1367–1382, 2017.
- [13] Y. Chen, U. Rosolia, W. Ubellacker, N. Csomay-Shanklin, and A. D. Ames, "Interactive multi-modal motion planning with branch model predictive control," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 5365–5372, 2022.
- [14] H. Bai, S. Cai, N. Ye, D. Hsu, and W. S. Lee, "Intention-aware online pomdp planning for autonomous driving in a crowd," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 454–460.
- [15] W. Liu, S.-W. Kim, S. Pendleton, and M. H. Ang, "Situation-aware decision making for autonomous driving on urban road using online pomdp," in *2015 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2015, pp. 1126–1133.
- [16] T. Bandyopadhyay, K. S. Won, E. Frazzoli, D. Hsu, W. S. Lee, and D. Rus, "Intention-aware motion planning," in *Algorithmic foundations of robotics X*. Springer, 2013, pp. 475–491.
- [17] J. F. Fisac, E. Bronstein, E. Stefansson, D. Sadigh, S. S. Sastry, and A. D. Dragan, "Hierarchical game-theoretic planning for autonomous vehicles," in *2019 International conference on robotics and automation (ICRA)*. IEEE, 2019, pp. 9590–9596.
- [18] D. Sadigh, S. Sastry, S. A. Seshia, and A. D. Dragan, "Planning for autonomous cars that leverage effects on human actions," in *Robotics: Science and systems (RSS)*, vol. 2. Ann Arbor, MI, USA, 2016, pp. 1–9.
- [19] A. Cui, S. Casas, A. Sadat, R. Liao, and R. Urtasun, "Lookout: Diverse multi-future prediction and planning for self-driving," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021, pp. 16 107–16 116.
- [20] N. Rhinehart, J. He, C. Packer, M. A. Wright, R. McAllister, J. E. Gonzalez, and S. Levine, "Contingencies from observations: Tractable contingency planning with learned behavior models," *arXiv preprint arXiv:2104.10558*, 2021.
- [21] D. Xu, Y. Chen, B. Ivanovic, and M. Pavone, "BITS: Bi-level imitation for traffic simulation," *arXiv preprint arXiv:2208.12403*, 2022.
- [22] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nuscnets: A multimodal dataset for autonomous driving," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 11 621–11 631.
- [23] N. Rhinehart, J. He, C. Packer, M. A. Wright, R. McAllister, J. E. Gonzalez, and S. Levine, "Contingencies from observations: Tractable contingency planning with learned behavior models," *arXiv preprint arXiv:2104.10558*, 2021.
- [24] E. Schmerling, K. Leung, W. Vollprecht, and M. Pavone, "Multimodal probabilistic model-based planning for human-robot interaction," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 3399–3406.
- [25] Y. Yuan, X. Weng, Y. Ou, and K. M. Kitani, "Agentformer: Agent-aware transformers for socio-temporal multi-agent forecasting," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021, pp. 9813–9823.
- [26] Y. Chen, P. Karkus, B. Ivanovic, X. Weng, and M. Pavone, "Tree-structured policy planning with learned behavior models," *arXiv preprint arXiv:2301.11902*, 2023.
- [27] S. LaValle, "Better unicycle models," *Planning Algorithms*, p. 743, 2006.
- [28] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nuscnets: A multimodal dataset for autonomous driving," in *IEEE/CVF conference on computer vision and pattern recognition (CVPR)*, 2020, pp. 11 621–11 631.
- [29] Y. Chen, U. Rosolia, C. Fan, A. D. Ames, and R. Murray, "Reactive motion planning with probabilistic safety guarantees," *Conference on Robot Learning (CoRL)*, 2020.