

TIP: Task-Informed Motion Prediction for Intelligent Vehicles

Xin Huang¹, Guy Rosman², Ashkan Jasour¹, Stephen G. McGill², John J. Leonard^{1,2}, Brian C. Williams¹

Abstract—When predicting trajectories of road agents, motion predictors usually approximate the future distribution by a limited number of samples. This constraint requires the predictors to generate samples that best support the task given task specifications. However, existing predictors are often optimized and evaluated via task-agnostic measures without accounting for the use of predictions in downstream tasks, and thus could result in sub-optimal task performance.

In this paper, we propose a task-informed motion prediction model that better supports the tasks through its predictions, by jointly reasoning about prediction accuracy and the utility of the downstream tasks, which is commonly used to evaluate the task performance. The task utility function does not require the full task information, but rather a specification of the utility of the task, resulting in predictors that serve a wide range of downstream tasks. We demonstrate our approach on two use cases of common decision making tasks and their utility functions, in the context of autonomous driving and parallel autonomy. Experiment results show that our predictor produces accurate predictions that improve the task performance by a large margin in both tasks when compared to task-agnostic baselines on the Waymo Open Motion dataset.

I. INTRODUCTION

Motion prediction is crucial for intelligent systems. It captures the distribution of future behavior of nearby road agents, and allows intelligent systems to plan and act. The output distribution is often approximated via a set of *weighted samples* [1]–[4]. The samples allow the downstream tasks to evaluate the predictions for risk assessment more efficiently, compared to complicated spatial distributions [5]; the weights are necessary to provide an accurate estimate of the risk. In many cases, a predictor only affords a small set of trajectory samples, due to the time complexity of processing the predictions. For instance, evaluating one prediction sample for risk assessment may take up to a few milliseconds [5], making it impractical to evaluate a large number of samples for real-time decision making.

Traditionally, distance-based measures, such as displacement errors, are used to optimize and evaluate the prediction samples, and have proved tremendously useful for reducing and gauging prediction errors. However, they do not account for the relevant downstream task, such as a planner that selects safe and efficient plans by reasoning about the predictions, and a driver assistance system that warns the driver when detecting risky behaviors.

Consider a motivating example in Fig. 1, where the *object agent*, defined as a nearby road agent of interest, may follow

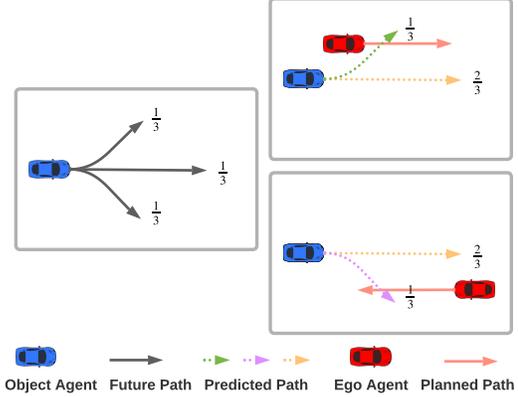


Fig. 1: A motivating example for task-informed prediction, where three equally likely future trajectories for the object agent are illustrated on the left. Given a sample limit of two, a task-informed predictor supports the task by outputting different sets of prediction samples weighted by their importance to the task, depending on the task information, such as the future plan of the ego agent, as illustrated in the two scenarios on the right. In contrast, a task-agnostic predictor tells no difference between these two sample sets and may generate predictions leading to unsafe events.

three possible future trajectories with an equal chance, as shown on the left. Due to the sample constraint, a predictor in this simplified example is only allowed to generate two trajectory samples¹. We present two possible prediction outcomes on the right, denoted by the colored dashed lines. There are no differences between these two outcomes when evaluated by a task-agnostic distance-based metric, such as displacement error. However, when the predictions are used in a planning system for the *ego agent* that aims to navigate safely with the object agent, one outcome may be favored over the other, depending on the ego plan. In the top right scenario, the green and the orange samples are more informative as they help the ego agent identify a potential near collision, whereas generating the purple and the orange predictions in this scenario may lead to an unsafe ego plan. The purple and the orange samples are, however, favored in the bottom right scenario, given a different ego plan. Therefore, it is important to reason about the downstream task when providing a limited set of prediction samples. The predictor should also provide an accurate estimate of the importance (or weight) of each sample with respect to the task. As we show in both scenarios, the predictor estimates

¹Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139, USA
xhuang@csail.mit.edu

²Toyota Research Institute, Cambridge, MA 02139, USA

¹In practice, a predictor affords more than two samples, but from a much larger prediction space.

the weight of each sample to help the planner accurately compute the probability of collision with the object agent.

In short, a predictor should provide a faithful estimate of the future trajectories of nearby road agents through a limited number of weighted samples. These predictions should assist decision making by characterizing nearby road agents and their future actions [6], providing an *approximate sufficient statistics* [7] for the task. Furthermore, its design should accommodate a variety of decision making tasks as opposed to a specific task, for flexibility in broader applications.

With this rationale in mind, we propose *TIP*, a *Task-Informed motion Predictor*, which is learned by jointly optimizing prediction accuracy and the performance of downstream tasks. The training loss leverages a specification of the task, such as its utility function, instead of ignoring the task or requiring the task itself to be co-trained (e.g. optimizing a specific planner with the predictor). Compared to existing prediction methods, *TIP* generates predictions that improve the downstream task performance by providing relevant information to the task given limited samples. In addition, it accommodates a variety of decision making tasks within intelligent systems, through different utility functions that characterize the tasks, compared to models that are constrained to a specific task and require that task to be differentiable [8]. While *TIP* is designed to support an arbitrary decision making task given their utility functions, we present two common decision making tasks, in the context of autonomous driving and parallel autonomy, to demonstrate the effectiveness of our proposed model.

Our contributions are as follows: i) We present a task-informed motion prediction model for intelligent systems that can be used to improve the performance of downstream tasks. Our approach covers a wide range of tasks, through utility function surrogates that characterize the tasks. ii) We show two case studies of our predictor being used in a planning task and a warning task, which are commonly used by state-of-the-art intelligent systems such as autonomous driving and parallel autonomy. iii) We demonstrate that our system helps achieve better task performance across both tasks through detailed quantitative and qualitative comparisons to task-agnostic baselines in a large-scale motion prediction dataset.

II. RELATED WORK

A. Motion Prediction

Motion prediction has been studied extensively in the last few years, from predicting vehicles [9]–[11] to vulnerable road users such as pedestrians and cyclists [4], [12], [13]. Due to the complexity in multi-modal future trajectory distributions, motion predictors often approximate the output distribution through simplified representations, such as a weighted set of samples [3], [4], Gaussian mixture models [1], [9], and occupancy maps [10]. We refer to [14] for a more comprehensive list of literature. In this work, we focus on sampling-based predictors, which offer a good balance between retaining enough continuous information and allowing the downstream tasks to efficiently process the predictions as deterministic positions.

Despite the recent efforts on supplying additional labels to improve prediction accuracy [1], [15], [16], most existing literature quantifies prediction performance through task-agnostic measures, including distance-based metrics such as minimum average displacement error (ADE) [12] and distribution-based metrics such as negative log-likelihood (NLL) [9]. However, such metrics do not account for the use of predictions in the downstream task. In other words, predictions with the same accuracy may lead to different outcomes for tasks such as planning, as illustrated in Fig. 1. Therefore, we propose a prediction model that accounts for not only prediction accuracy, but also the utility of the downstream task given the predictions, to allow better integration between predictions and the task.

A concurrent work [17] is proposing task-aware metrics to evaluate the performance of motion predictors. It presents a proof-of-concept metric that favors the prediction samples based on the sensitivity to the ego agent’s plan. Results show that the proposed metric better provides a measure of planner performance compared to distance-based metrics. Different from [17] that focuses on proposing task-aware metrics to *evaluate* predictor performance, our method goes beyond by leveraging task-specific information to *train* a motion predictor. This allows us to maximize the value of predictions for downstream tasks, as we demonstrate in realistic driving scenarios. More importantly, by training with a task-informed loss, our predictions cover approximate sufficient statistics [7], [18] of the nearby road agents, and leverage even imperfect information of downstream tasks.

B. Prediction for Tasks

Existing works leverage a learned motion predictor to support a variety of tasks, including risk assessment [5], driver safety detection [11], and most commonly, planning for autonomous systems [19]–[24]. They often decouple the optimization of the predictor and the optimization of the task. As a result, the predictor is unaware of its influence on the downstream task and its predictions may not be informative for the downstream task. In this work, we present a more effective predictor that is optimized directly through the utility of the downstream task.

A prior work [11] proposes a multi-task predictor that approximates the utility and its uncertainties of the downstream driver safety detection task, yet the predictor is still learned with a single objective of optimizing the accuracy. In contrast, our model integrates the task utility into the trajectory prediction directly, allowing the prediction results to better support the task.

Our work is closely related to *prediction and planning* (P&P) literature [8], which jointly optimizes the prediction module and the planning module. Compared to end-to-end planning models [25]–[27] that generate planning results from raw sensor inputs, P&P, as a more structured approach, provides better interpretability in its decision making process, and achieves better planning performance [8]. One limitation of P&P is that it requires a fully differentiable pipeline that includes a differentiable planner. In many

practical planning systems that involve black-box modules or off-the-shelf components [28], [29], it is a nontrivial task to differentiate the planner. In contrast, our prediction model only requires the utility function that characterizes the task. The utility function is planner-agnostic, allowing the predictor to support a family of planning algorithms, instead of a specific planner as in P&P. Furthermore, a differentiable utility function is usually easier to acquire than a differentiable planner. This makes our model more broadly applicable to non-differentiable planners and tasks that are more than planners, as we show in our experiments.

Instead of assuming a specific differentiable planner, [30] generates joint predictions for both the ego agent and the object agents, and uses the ego prediction for planning. Such approach offers a great advantage by simply imitating the future behavior of the ego agent without requiring a specific planner, yet the predicted trajectory for the ego agent is prone to input noise and modeling error, making it less reliable to be used for planning in practice.

III. PROBLEM FORMULATION

Motion Prediction The prediction system takes input as i) task-specific input information V and ii) observed agent trajectories $O = \{\mathbf{o}_t\}_{t=-T_p+1}^0$ over a fixed past horizon T_p , where $\mathbf{o}_t = [o_{1,t}, \dots, o_{N,t}]$ includes continuous positions at time step t for up to N agents. The positions are normalized with respect to the center of the last observed positions of all agents. The task input V depends on the specific information from the task, such as an ego planned trajectory, as customary in conditional prediction approaches [31]–[34]. The output is a weighted set of K joint trajectory samples $\mathcal{S} = \{(w^{(k)}, \mathbf{x}^{(k)})\}_{k=1}^K$ for all agents, where $\mathbf{x}^{(k)} = \{\mathbf{x}_t^{(k)}\}_{t=1}^{T_f}$ denotes future trajectory sequences of all agents, i.e. $\mathbf{x}_t^{(k)} = [x_{1,t}^{(k)}, \dots, x_{N,t}^{(k)}]$, up to a fixed future horizon T_f .

Utility-Based Decision Making The task-informed prediction aims to allow accurate estimates of task utility for an arbitrary downstream decision making task. The utility (or reward) serves as a quantitative measurement of task performance and is commonly used in modern decision making tasks. In this work, we define the task specification as a tuple (\mathcal{I}, u) , where \mathcal{I} is a set of candidate decisions for the task, such as plans of the ego agent; u is a task-specific differentiable utility function mapping a decision $I \in \mathcal{I}$ and the task-informed predictions \mathcal{S} to a scalar that quantitatively measures the performance of the decision. For the sake of simplicity, we define $u_I = u(I, \mathcal{S})$ in the rest of the paper. This specification allows us to optimize our model to support different tasks through their task-specific surrogate utility functions, without requiring a specific task pipeline.

Finally, we define the task objective to obtain the optimal decision that maximizes the utility among all candidate decisions:

$$I_O = \arg \max_{I \in \mathcal{I}} u_I. \quad (1)$$

At training time, we can obtain the utilities of all candidates and find the optimal decision using Eq. (1) given ground truth future trajectories. We optimize the model by

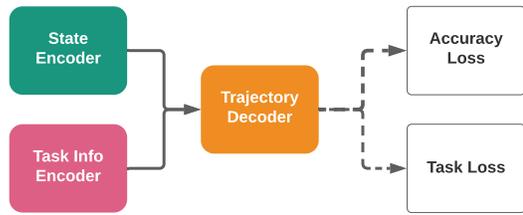


Fig. 2: Diagram of the proposed task-informed prediction model, which includes a state encoder that encodes observed past agent states, a task information encoder that encodes additional task input, and a trajectory decoder that decodes future trajectory predictions. The model is trained through an accuracy loss term that optimizes prediction accuracy and a task loss term that guides the model to favor predictions supporting the downstream task.

maximizing the following *softmax* over the decision utilities:

$$R_{task} = \frac{\exp(u_{I_O})}{\sum_{I' \in \mathcal{I}} \exp(u_{I'})}, \quad (2)$$

which is widely used in reinforcement learning [35] to encourage the optimal decision to have higher utility compared to the other decisions.

IV. APPROACH

Our predictor leverages an encoder-decoder model, as depicted in Fig. 2, following a standard architecture in [36]. The state encoder encodes the observed trajectory for up to N agents. For each agent i , an MLP is used to encode the position at each time step, and an LSTM is used to encode the position encodings from time step $-T_p+1$ to 0 into a hidden state h_{S_i} . The hidden states of each agent are concatenated into a joint hidden state h_S . If there are fewer than N agents in a scene or there exist invalid positions for any agent, we zero-mask the encodings of missing positions. The task information encoder encodes task-specific inputs V , such as the future plan of the ego agent, through a separate model into a separate hidden state h_V . The structure of the task information encoder depends on the input representation, and we defer a detailed description in our experiments. The trajectory decoder model takes the concatenated encoded states from both encoders, i.e. $h = h_S \oplus h_V$, and predicts a set of K joint trajectory samples \mathcal{S} and their weights.

We train the model by jointly optimizing prediction accuracy and task performance,

$$\mathcal{L} = \mathcal{L}_{acc} + \alpha \mathcal{L}_{task}, \quad (3)$$

where α determines the relative weight between two terms. In the experiments, we present the trade-off between prediction accuracy and task performance, by varying α values.

The accuracy loss \mathcal{L}_{acc} measures the accuracy of the prediction samples compared to the ground truth future trajectory \hat{S} . We follow the standard motion prediction literature, i.e. [1], [37], based on a variant of variety loss

proposed in [12]:

$$\mathcal{L}_{acc} = \sum_{k=1}^K \mathbb{1}(k = \hat{k}) (-\log w^{(k)} + \|\mathbf{x}^{(k)} - \hat{S}\|_2), \quad (4)$$

where \hat{k} is the index of the best prediction sample, in terms of L2 distance to the ground truth trajectory.

The task term \mathcal{L}_{task} minimizes the negation of Eq. (2) so that the optimal decision would have higher utility compared to other candidate decisions given the predictions.

$$\mathcal{L}_{task} = -R_{task}. \quad (5)$$

While TIP can be used for an arbitrary decision making task that requires behavior prediction given its task specifications, in the rest of this section, we present two use cases of tasks and their utility functions to demonstrate the flexibility of our approach over different tasks. We defer implementation details on both tasks in Section V.

A. Example Task 1: Planning for Autonomous Systems

In the planning task, we assume that the ego agent is equipped with an arbitrary planner that generates a set of M motion plan candidates $\mathcal{I}_P = \{\tau_1, \dots, \tau_M\}$. The planning utility function is defined to capture safety and efficiency:

$$u_P(\tau, \mathcal{S}) = u_{efficiency}(\tau) + \beta u_{safety}(\tau, \mathcal{S}), \quad (6)$$

where τ is an ego plan candidate, and $\mathcal{S} = \{(w^{(k)}, x_{object}^{(k)})\}_{k=1}^K$ is a weighted set of object prediction samples generated from our predictor. In this task, the predictor takes the input² of the observed agent states and the ego plan τ as the task input V to produce the prediction samples \mathcal{S} . The efficiency term $u_{efficiency}$ measures the travelled distance of the ego plan. The safety term u_{safety} measures the expected closest distance between the ego plan and the object predictions, computed as follows:

$$u_{safety}(\tau, \mathcal{S}) = \sum_{k=1}^K w^{(k)} \min_{t=1 \dots T_f} \|\tau_t - x_{object,t}^{(k)}\|_2. \quad (7)$$

In practice, the improvement of the safety utility diminishes if the agents are far away from each other. Therefore, we upper bound the utility by a safety threshold d_{safe} :

$$u_{safety}(\tau, \mathcal{S}) = \min(d_{safe}, \sum_{k=1}^K w^{(k)} \min_{t=1 \dots T_f} \|\tau_t - x_{object,t}^{(k)}\|_2). \quad (8)$$

B. Example Task 2: Warning for Parallel Autonomy

A pre-collision warning system is widely adopted in parallel autonomy (or shared autonomy), as a vehicle shared-control framework [11], [40], [41] that monitors driver actions and warns before an unsafe event could happen. The warning system differs from the planning system in a few ways. First, it requires a *joint* predictor for both ego agent

²We follow existing conditional behavior prediction literature [31], [33] to generate predictions conditioning on the future ego plan. Such predictors have demonstrated to predict reactive agent behaviors and improve accuracy, and are useful for a wide range of planners that generate initial candidate plans using simplified prediction models, such as [38], [39].

and object agent, as the ego agent is controlled by a driver and the future path is unknown to the predictor. This requires predicting the joint behavior in the future to determine if a near collision is likely. Second, it provides *no task-specific input* to the predictor, as it only sends a warning to the driver and does not induce any actual interactions with the world. As a result, the predictor produces prediction samples \mathcal{S} by taking only the observed agent states as inputs.

The warning system is a binary decision making system that chooses an action from $\mathcal{I}_W = \{\text{warn}, \neg\text{warn}\}$. The utility of a warning action is equivalent to the likelihood of near collision between the object agent and the ego agent. To compute the near collision likelihood, we follow the two-step procedure given the joint prediction samples $\mathcal{S} = \{(w^{(k)}, x_{ego}^{(k)}, x_{object}^{(k)})\}_{k=1}^K$ for the ego agent and the object agent. First, the system computes the collision score $r^{(k)} \in \{0, 1\}$ as a Boolean value for each trajectory sample:

$$r^{(k)} = \left(\min_{t=1 \dots T_f} \|x_{ego,t}^{(k)} - x_{object,t}^{(k)}\|_2 < d_{warn} \right), \quad (9)$$

where d_{warn} is the minimum safety distance threshold allowed. The collision score is 1 if the closest distance between two agents is smaller than this threshold, and 0 otherwise.

Next, we compute the overall collision likelihood by taking the expected collision score r as the weighted sum of individual warning scores: $u_W(\text{warn}) = r = \sum_{k=1}^K w^{(k)} r^{(k)}$. Intuitively, the utility of $\neg\text{warn}$ is the likelihood of no near collision, i.e. $u_W(\neg\text{warn}) = 1 - u_W(\text{warn})$.

To compute the ground truth optimal decision, we compute the likelihood of near collision from the observed future trajectories following the same procedure in Eq. (9). Since the observed future trajectories are deterministic, the resulting likelihood is either 0 or 1.

V. EXPERIMENTS

In this section, we show experimental results in two different tasks to demonstrate the advantage of our proposed task-informed predictor on a naturalistic driving dataset.

A. Dataset

We train and validate our model in the Waymo Open Motion dataset [36]. It is one of the largest motion prediction datasets in terms of the number of scenes, total time, and prediction horizon, i.e. 8 seconds with 80 time steps. More specifically, we focus on the interactive dataset that is mined to cover interesting interactions. This allows us to demonstrate the effectiveness of our model in complicated long-term interacting scenarios. We follow the standard train/validation split from the dataset.

B. Example Task 1: Planning for Autonomous Systems

1) *Model Details:* The MLP in the state encoder has 32 neurons, followed by ReLU and dropout layers with a rate of 0.1. The LSTM has a hidden size of 32 and an output dimension of 32. The task information encoder encodes the planned trajectory of the ego agent, as the task-specific input V , through an MLP with 32 neurons, followed by ReLU and dropout layers with a rate of 0.1. The trajectory decoder uses

a two-layer MLP with 32 neurons to output \mathcal{S} that includes the predicted trajectory samples and their weights. We choose $\alpha = 20$ to keep the two loss magnitudes on the same scale, and $\beta = 5$ to prioritize safe driving. The model is trained for 20 epochs and is optimized using Adam [42], with a batch size of 32 and a learning rate of 10^{-3} . It takes approximately 2 milliseconds to generate all samples in each example.

2) *Task Details*: To simulate the planning task, we select the ego agent and the object agent randomly from the interactive pair in the Waymo data. The ego planner simulates three planned trajectories based on the observed future trajectory of the ego agent. It interpolates and scales the trajectory coordinates by 0.8x, 1.0x, and 1.2x at each time step to simulate conservative driving, normal driving, and aggressive driving, while imposing limits on acceleration and speed following [43]. This simple approach provides multiple driving options while ensuring the plans are realistic and closely follow the agent intention from data (see examples in Fig. 3), yet one can use an arbitrary planner to supply the plan candidates in practice.

In order to find the ground truth optimal plan, we also have to simulate the reactive behavior of the object agent in the future, which depends on the interaction type between two agents. When the object agent is yielding to the ego agent in the data, we modify its future trajectory with an equal chance to either speed up to pass or slow down to keep yielding, in response to the *conservative* ego plan. We keep its future trajectory unchanged in response to the other two ego plans. In contrast, when the object agent is being yielded by the ego agent in the data, we modify its future trajectory with an equal chance to either slow down to yield or speed up to maintain the lead, in response to the *aggressive* ego plan. Simulating realistic object agent behavior for simulation purposes is a topic of ongoing research [44]. In our experiments, we follow a simple heuristic based on how humans normally react to others and find it to be realistic and effective. In addition, to validate the representativeness of our simulation model, we have experimented with an additional simulation model based on IDM, as in [45], and observed consistent improvements in task performance using our proposed approach. Full results are available in Table IV.

We choose the threshold d_{safe} to be 3.64 meters as the 10% percentile of the pairwise closest distances in the Waymo dataset. This value is smaller than the radius of a regular car with a length of approximately 4.5 meters and a width of approximately 2.0 meters [46].

3) *Quantitative Results*: We compare our proposed model TIP_P , with a Task-Agnostic Predictor (we refer to it as TAP in the rest of the experiments) that uses the same model as ours, but is trained with only the *accuracy loss* in Eq. (4). This baseline is equivalent to the baseline model proposed in [36] and represents a broad prediction literature that ignores predictions in downstream tasks or decouples prediction and the tasks. In addition, we demonstrate that our proposed approach can be applied to a different utility function, TIP_{Pa} , that is trained for the same planning task

Model	minADE↓	minFDE↓	AUC-ROC $_P$ ↑	AUC-ROC $_{Pa}$ ↑
TAP	2.80	6.44	0.594	0.616
TIP_P	2.89	6.54	0.667	0.586
TIP_{Pa}	2.93	6.51	0.555	0.697

TABLE I: Comparison between our TIP and the baseline model, in terms of prediction accuracy and task performance, on two tasks P and Pa . The task-informed predictors trade off little accuracy to much better task performance. Our approach supports multiple utility functions, as suggested by the relevant metrics highlighted in the colored cells.

α	minADE↓	minFDE↓	wADE↓	wFDE↓	AUC-ROC↑
0	2.80	6.44	5.76	14.32	0.594
1	2.83	6.47	5.82	14.46	0.613
5	2.85	6.52	5.89	15.07	0.623
20	2.89	6.54	6.66	16.69	0.667
100	4.01	8.67	9.72	23.14	0.676

TABLE II: Performance of TIP_P as a function of the task loss coefficient α .

but represents a drastically different altruistic planner (P_a):

$$u_{Pa}(\tau, \mathcal{S}) = u_{\text{efficiency}}(\tau_{\text{object}}) + \beta u_{\text{safety}}(\tau, \mathcal{S}), \quad (10)$$

where τ_{object} is the simulated trajectory of the object agent that reacts to the ego plan τ . This utility function models an altruistic planner that favors the object agent as opposed to the ego agent. Such a planner is commonly seen in robotics social navigation that minimally interferes with humans [21].

We evaluate each model on prediction accuracy and task performance. The prediction accuracy is measured by minADE/minFDE metrics [12], as standard in motion prediction benchmarks [36], [47]. In addition, we present weighted ADE (wADE) and weighted FDE (wFDE) metrics that measure the expected errors given the predicted weights. The unit of all accuracy metrics is in meters. We measure the task performance through recall and fall-out. Recall measures the percentage of optimal plans that are successfully recognized. Fall-out measures the percentage of false alarms, i.e. sub-optimal plans that are wrongly recognized as optimal. We plot the recall and fall-out at various thresholds, as receiver operating characteristic (ROC) curve [48], and compute its area under the curve (AUC) score to determine the task performance. As our planner is dealing with a multi-class decision making problem, we compute the AUC-ROC score using the one-vs-one methodology, which is insensitive to data balance [48].

We report two separate AUC-ROC metrics, AUC-ROC $_P$ and AUC-ROC $_{Pa}$, depending on which utility function is used to determine the ground truth optimal plan. The results are summarized in Table I, where we color the cell to highlight the results measured by their relevant metrics.

Trade off accuracy for better task performance While the task-agnostic model (TAP) achieves the best prediction accuracy, our model achieves much better task performance at the cost of little accuracy. More specifically, compared to TAP, TIP_P improves the task performance by 12.29% at the cost of 1.55% accuracy loss in terms of minFDE, and TIP_{Pa} improves the task performance by 13.15% at

the cost of 1.09% accuracy loss. In most cases, the task performance matters more than the prediction accuracy, as the planner interacts directly with the world and a small error may lead to undesirable outcomes. We present the trade-off between prediction accuracy and task performance by varying α values, as in Table II.

Adapt to multiple utility functions We show that our model can be adapted to multiple utility functions within the same planner task, through results highlighted by the yellow cell. For instance, when training with a different utility function u_{Pa} that favors altruistic behavior, our model achieves good task performance on the relevant metric (e.g. TIP_{Pa} achieves a high $AUC-ROC_{Pa}$ score), although it tends to perform poorly on the drastically different task, as these two tasks adopt competing objectives on being “selfish” versus being “altruistic”. In practice, given a specific task, the user can specify the utility function that best describes the task objective to serve their own need.

4) *Qualitative Results:* In Fig. 3, we present a representative scenario to demonstrate the advantage of our model compared to the task-agnostic baseline. In this scenario, the planner proposes two candidate plans³ (in magenta) for the ego agent, whose observed trajectory is in red: one *normal* plan that yields to the object agent, and one *aggressive* plan that speeds up. For each plan, we visualize the simulated reactive trajectories of the object agent in cyan and its observed trajectory in blue. In this example, the normal plan is favored over the aggressive plan, as the latter leads to a near collision. The predictions (in olive) of our predictor and the baseline are visualized in Fig. 3(a) and (b), respectively. In Fig. 3(b), the task-agnostic predictor TAP generates predictions that indicate near collisions for both plans, which lead to a higher utility for the aggressive plan as it travels further. In contrast, in Fig. 3(a), our predictor TIP_P generates predictions that help better approximate the utility for each plan – the normal plan comes with a higher utility as no near collision is detected, and the aggressive plan results in a lower utility due to a near collision indicated by the predictions. As a result, TIP_P helps find the correct decision to choose the normal plan.

C. Example Task 2: Warning for Parallel Autonomy

1) *Model Details:* The task-informed predictor model for the warning task leverages the same structure as described in Sec. V-B.1, except that it does not include a task information encoder (e.g. the predictions are conditioned only on past observations). In addition, the utility defined in Eq. (9) is not differentiable due to the Boolean comparison operation. So we utilize a soft warning score using the sigmoid function:

$$r^{(k)} = \text{sigmoid}(d_{\text{warn}} - \min_{t=1 \dots T_f} \|x_{\text{ego},t}^{(k)} - x_{\text{object},t}^{(k)}\|_2). \quad (11)$$

The soft score is close to 1 when the closest distance is smaller than the safety distance threshold, and close to 0 otherwise. We use the same distance threshold of 3.64 meters as in the planning task.

³We omit the conservative plan in the discussion for the sake of simplicity.

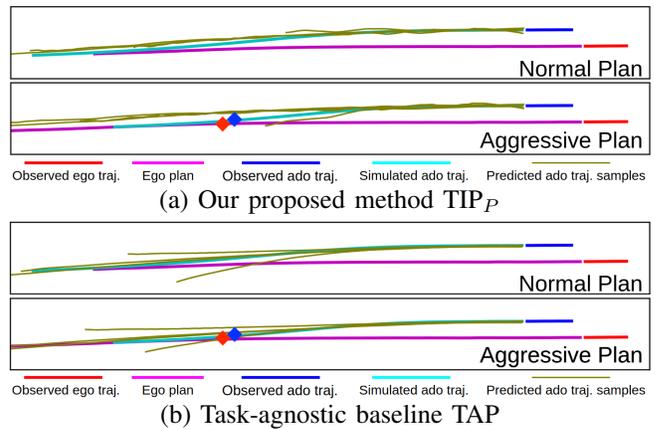


Fig. 3: Predictions from TIP_P (a) and TAP (b) in a representative example, where the aggressive plan (colored in magenta on the bottom of each subfigure) is less favorable than the normal plan (colored in magenta on the top of each sub-figure), as it leads to a near collision with the object agent, as indicated by the diamond markers. TIP_P generates predictions (in olive), which indicate a near collision for the aggressive plan and no collisions for the normal plan, to help the planner choose the correct decision. TAP generates predictions that indicate higher utility for the aggressive plan, leading to the wrong decision.

2) *Quantitative Results:* We evaluate the performance of our model using the same accuracy and task metrics as in the planning task. In the following, we present a series of experiments to validate the advantage of our model.

Trade off accuracy for better task performance We perform a study on the trade-off between prediction accuracy and task performance, by tuning the coefficient of the task loss α in Eq. (3). The results are reported in Table III. When α is 0, the model is equivalent to the Waymo baseline [36] as a task-agnostic predictor (TAP) that is optimized only for accuracy. This model achieves the best accuracy metrics. As α increases, the task performance improves at the cost of prediction accuracy. We use $\alpha = 20$ in the rest of the experiments, as it achieves a good balance between accuracy and task performance. In practice, the choice of α depends on the specific task requirement.

We further compare the performance of our model to an interactive prediction model M2I [13] that achieves the state-of-the-art performance in the Waymo benchmark, which yields a minADE of 3.79 meters, a minFDE of 8.40 meters, and an AUC-ROC score of 0.362. The comparison shows that our model is able to achieve much better task performance, i.e. at $\alpha = 20$, despite using a simple prediction backbone.

Support task with limited samples We examine the performance of our model by varying the number of prediction samples. As depicted in Fig. 4, our proposed model TIP_W achieves much better task performance, by sacrificing little accuracy at different numbers of samples. For instance, at 4 samples, it improves the AUC-ROC score by more than 3 times at the cost of 8.13% minFDE (and 5.00% wFDE) lost compared to TAP, which produces predictions without

α	minADE↓	minFDE↓	wADE↓	wFDE↓	AUC-ROC↑
0	4.00	10.57	7.06	20.24	0.165
1	4.05	10.68	7.13	20.46	0.299
5	4.19	10.98	7.32	20.92	0.449
20	4.65	11.43	7.87	21.25	0.655
100	5.29	12.01	12.42	26.46	0.776

TABLE III: Performance of TIP_W as a function of the task loss coefficient α . Task performance significantly improves at the cost of prediction accuracy.

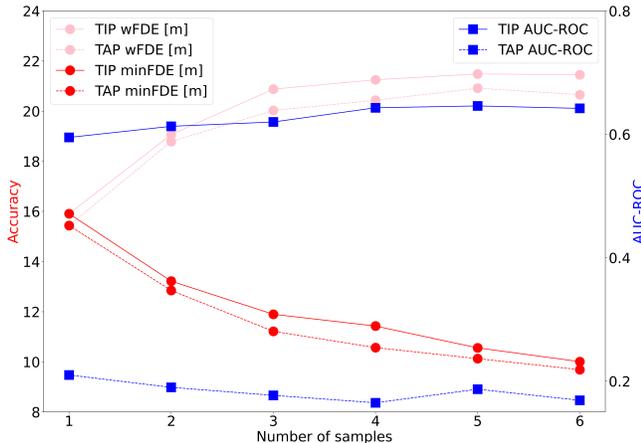


Fig. 4: Prediction accuracy (minFDE and wFDE) and warning task performance (AUC-ROC) as a function of the number of prediction samples K from TIP_W (solid lines) and the baseline TAP (dashed lines). Compared to TAP, TIP_W achieves much better task performance across all samples at the cost of little accuracy. The task performance for TIP_W converges at $K = 4$, indicating its predictions serve as approximate sufficient statistics for the warning task with only a few samples.

considering the effects on downstream tasks. Besides, the task performance for TIP_W stabilizes at 4 samples. This demonstrates that TIP_W provides predictions that summarize the influence by nearby agents as approximate sufficient statistics for the warning task using only a few samples.

Robust to noisy utility estimates We further investigate how robust our predictor is when trained with noisy utility estimates, to probe our approach under imperfect information about the task. At training time, we add to the estimated utilities a random Gaussian noise with zero mean and increasing variance levels and observe that the task performance decreases slightly by 3.08%, within a noise level of 25% of the magnitude of the utility, compared to a standard TIP_W model that is trained without noise. The observation demonstrates that our approach is robust to perturbations in the utility estimates during training, as an example of imperfect task information, or an example of adapting to a task with a slightly different utility function.

Handle multiple object agents We performed additional experiments by training TIP_W to predict joint future trajectories of up to 4 agents, including 3 object agents. Compared to the TAP baseline that has a minFDE of 10.03 meters and an AUC-ROC score of 0.226, TIP_W achieves a much higher

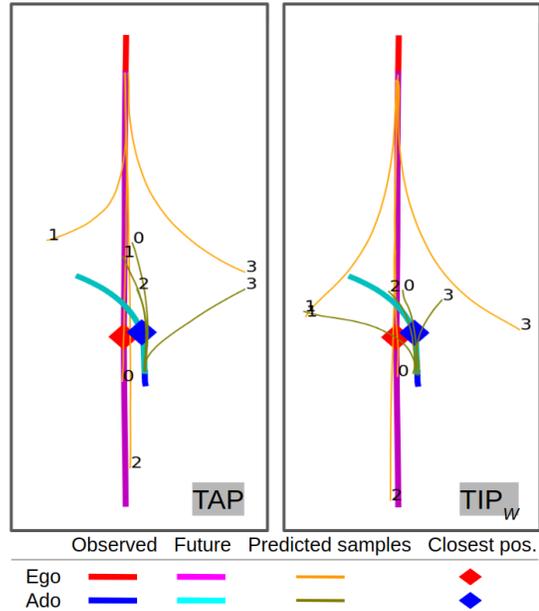


Fig. 5: Comparison between TAP (left) and TIP_W (right) in a warning scenario, where the closest distance is indicated by the diamond markers. The sample indices are annotated to help associate joint predictions. Our model TIP_W generates predictions (in olive and orange) that help identify multiple instances of near collisions, especially through joint samples #2 that match with the observed future trajectories.

task score of 0.314 with a slightly worse minFDE of 10.57. This demonstrates that our predictor works well for more than two agents. Full results are available in Table V.

3) *Qualitative Results:* We present a representative warning example in Fig. 5, where the two agents are getting too close according to their observed future trajectories. The closest distance is indicated by the diamond markers. The sample index k is labelled to help identify joint predictions. From the left plot, we see that the task-agnostic baseline fails to identify a likely near collision. Although it predicts a near collision with joint samples #3, they are predicted with a very low probability, i.e. smaller than 1%. In contrast, our predictor on the right identifies multiple near collision instances, especially through samples #2 that successfully predicts the object agent is going to cross the ego agent’s path, which matches with what happens in the data. As a result, its predictions indicate a high likelihood of collision and lead to the correct warning decision. This example also verifies that in the downstream task, it is usually not required to generate *perfect* predictions, as long as the predictions cover sufficient statistics for the task (e.g. the predictions indicate a collision).

VI. CONCLUSION

We propose a task-informed motion prediction system, in which predictors are trained to both make accurate predictions and support correct decision making in a downstream task. By leveraging a specification of the task, we allow the predicted samples to provide approximate sufficient statistics of the environment for the task and support a variety of tasks,

without requiring a full differentiable task for co-training. We demonstrate our predictor in two tasks on the Waymo dataset, and show its advantage through quantitative and qualitative experiments. Future work includes further improving the task performance through a stronger backbone model and performing experiments in additional benchmarks.

REFERENCES

- [1] Y. Chai, B. Sapp, M. Bansal, and D. Anguelov, "Multipath: Multiple probabilistic anchor trajectory hypotheses for behavior prediction," in *CoRL*, 2019.
- [2] M. Liang, B. Yang, R. Hu, Y. Chen, R. Liao, S. Feng, and R. Urta- sun, "Learning lane graph representations for motion forecasting," in *ECCV*. Springer, 2020, pp. 541–556.
- [3] X. Huang, S. G. McGill, J. A. DeCastro, L. Fletcher, J. J. Leonard, B. C. Williams, and G. Rosman, "DiversityGAN: Diversity-aware vehicle motion prediction via latent semantic sampling," *IEEE RA-L*, 2020.
- [4] J. Gu, C. Sun, and H. Zhao, "DenseTNT: End-to-end trajectory prediction from dense goal sets," in *ICCV*, 2021, pp. 15 303–15 312.
- [5] A. Wang, X. Huang, A. Josaur, and B. Williams, "Fast risk assessment for autonomous vehicles using learned models of agent futures," in *RSS*, 2020.
- [6] S. Worrall, G. Agamennoni, J. Nieto, and E. Nebot, "A context-based approach to vehicle behavior prediction," *IEEE Intelligent Transportation Systems Magazine*, vol. 4, no. 3, pp. 32–44, 2012.
- [7] Y. Chen, D. Zhang, M. U. Gutmann, A. Courville, and Z. Zhu, "Neural approximate sufficient statistics for implicit models," in *ICLR*, 2021.
- [8] W. Zeng, W. Luo, S. Suo, A. Sadat, B. Yang, S. Casas, and R. Urta- sun, "End-to-end interpretable neural motion planner," in *CVPR*, 2019.
- [9] J. Wiest, M. Höffken, U. Kreßel, and K. Dietmayer, "Probabilistic trajectory prediction with gaussian mixture models," in *IV*. IEEE, 2012, pp. 141–146.
- [10] B. Kim, C. M. Kang, J. Kim, S. H. Lee, C. C. Chung, and J. W. Choi, "Probabilistic trajectory prediction over occupancy grid map via recurrent neural network," in *ITSC*. IEEE, 2017, pp. 399–404.
- [11] X. Huang, S. G. McGill, J. A. DeCastro, L. Fletcher, J. J. Leonard, B. C. Williams, and G. Rosman, "CARPAL: Confidence-aware intent recognition for parallel autonomy," *IEEE RA-L*, 2021.
- [12] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese, "Social LSTM: Human trajectory prediction in crowded spaces," in *CVPR*, 2016, pp. 961–971.
- [13] Q. Sun, X. Huang, J. Gu, B. Williams, and H. Zhao, "M2I: From factored marginal trajectory prediction to interactive prediction," in *CVPR*, 2022.
- [14] S. Mozaffari, O. Y. Al-Jarrah, M. Dianati, P. Jennings, and A. Mouza- kitis, "Deep learning-based vehicle behavior prediction for au- tonomous driving applications: A review," *T-ITS*, 2020.
- [15] N. Deo and M. M. Trivedi, "Convolutional social pooling for vehicle trajectory prediction," in *CVPR Workshops*, 2018, pp. 1468–1476.
- [16] J. Mercat, T. Gilles, N. El Zoghby, G. Sandou, D. Beauvois, and G. P. Gil, "Multi-head attention for multi-modal joint vehicle motion forecasting," in *ICRA*. IEEE, 2020, pp. 9638–9644.
- [17] B. Ivanovic and M. Pavone, "Injecting planning-awareness into prediction and detection evaluation," *preprint arXiv:2110.03270*, 2021.
- [18] B. Jiang, T.-y. Wu, C. Zheng, and W. H. Wong, "Learning summary statistic for approximate bayesian computation via deep neural network," *Statistica Sinica*, pp. 1595–1618, 2017.
- [19] E. Schmerling, K. Leung, W. Vollprecht, and M. Pavone, "Multimodal probabilistic model-based planning for human-robot interaction," in *ICRA*. IEEE, 2018, pp. 3399–3406.
- [20] A. Sadat, M. Ren, A. Pokrovsky, Y.-C. Lin, E. Yumer, and R. Urta- sun, "Jointly learnable behavior and trajectory planning for self-driving vehicles," in *IROS*. IEEE, 2019, pp. 3949–3956.
- [21] S. Schaefer, K. Leung, B. Ivanovic, and M. Pavone, "Leveraging neural network gradients within trajectory optimization for proactive human-robot interactions," in *ICRA*, 2021.
- [22] H. Nishimura, B. Ivanovic, A. Gaidon, M. Pavone, and M. Schwager, "Risk-sensitive sequential action control with multi-modal human trajectory forecasting for safe crowd-robot interaction," in *IROS*, 2020.
- [23] S. Eiffert, H. Kong, N. Pirmarzdashii, and S. Sukkarieh, "Path planning in dynamic environments using generative rnns and monte carlo tree search," in *ICRA*. IEEE, 2020, pp. 10 263–10 269.
- [24] S. Casas, A. Sadat, and R. Urta- sun, "Mp3: A unified model to map, perceive, predict and plan," in *CVPR*, 2021, pp. 14 403–14 412.
- [25] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang *et al.*, "End to end learning for self-driving cars," *arXiv preprint arXiv:1604.07316*, 2016.
- [26] M. Bansal, A. Krizhevsky, and A. Ogale, "ChauffeurNet: Learning to drive by imitating the best and synthesizing the worst," in *RSS*, 2019.
- [27] F. Codevilla, M. Müller, A. López, V. Koltun, and A. Dosovitskiy, "End-to-end driving via conditional imitation learning," in *ICRA*. IEEE, 2018, pp. 4693–4700.
- [28] A. Bacha, C. Bauman, R. Faruque, M. Fleming, C. Terwelp, C. Rein- holtz, D. Hong, A. Wicks, T. Alberi, D. Anderson *et al.*, "Odin: Team victortango's entry in the darpa urban challenge," *Journal of field Robotics*, 2008.
- [29] Y. Kuwata, G. A. Fiore, J. Teo, E. Frazzoli, and J. P. How, "Motion planning for urban driving using rrt," in *IROS*. IEEE, 2008.
- [30] C. Tang and R. R. Salakhutdinov, "Multiple futures prediction," *Neurips*, vol. 32, 2019.
- [31] E. Tolstaya, R. Mahjourian, C. Downey, B. Vadarajan, B. Sapp, and D. Anguelov, "Identifying driver interactions via conditional behavior prediction," in *ICRA*. IEEE, 2021.
- [32] J. Ngiam, B. Caine, V. Vasudevan, Z. Zhang, H.-T. L. Chiang, J. Ling, R. Roelofs, A. Bewley, C. Liu, A. Venugopal *et al.*, "Scene Transformer: A unified architecture for predicting multiple agent trajectories," in *ICLR*, 2022.
- [33] T. Salzmann, B. Ivanovic, P. Chakravarty, and M. Pavone, "Trajec- tron++: Multi-agent generative trajectory forecasting with heteroge- neous data for control," *ECCV*, 2020.
- [34] H. Song, W. Ding, Y. Chen, S. Shen, M. Y. Wang, and Q. Chen, "Pip: Planning-informed trajectory prediction for autonomous driving," in *ECCV*. Springer, 2020, pp. 598–614.
- [35] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.
- [36] S. Ettinger, S. Cheng, B. Caine, C. Liu, H. Zhao, S. Pradhan, Y. Chai, B. Sapp, C. R. Qi, Y. Zhou *et al.*, "Large scale interactive motion forecasting for autonomous driving: The waymo open motion dataset," in *ICCV*, 2021, pp. 9710–9719.
- [37] H. Cui, V. Radosavljevic, F.-C. Chou, T.-H. Lin, T. Nguyen, T.-K. Huang, J. Schneider, and N. Djuric, "Multimodal trajectory predictions for autonomous driving using deep convolutional networks," in *ICRA*. IEEE, 2019, pp. 2090–2096.
- [38] E. Frazzoli, M. A. Dahleh, and E. Feron, "Real-time motion planning for agile autonomous vehicles," *Journal of guidance, control, and dynamics*, vol. 25, no. 1, pp. 116–129, 2002.
- [39] J. v. d. Berg, S. J. Guy, M. Lin, and D. Manocha, "Reciprocal n-body collision avoidance," in *Robotics research*. Springer, 2011, pp. 3–19.
- [40] L. Saleh, P. Chevrel, F. Claveau, J.-F. Lafay, and F. Mars, "Shared steering control between a driver and an automation: Stability in the presence of driver behavior uncertainty," *T-ITS*, vol. 14, no. 2, pp. 974–983, 2013.
- [41] W. Schwarting, J. Alonso-Mora, L. Paull, S. Karaman, and D. Rus, "Parallel autonomy in automated vehicles: Safe motion generation with minimal intervention," in *ICRA*. IEEE, 2017, pp. 1928–1935.
- [42] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimiza- tion," in *ICLR*, 2015.
- [43] P. S. Bokare and A. K. Maurya, "Acceleration-deceleration behaviour of various vehicle types," *Transportation research procedia*, vol. 25, pp. 4733–4749, 2017.
- [44] H. Caesar, J. Kabzan, K. S. Tan, W. K. Fong, E. Wolff, A. Lang, L. Fletcher, O. Beijbom, and S. Omari, "nuPlan: A closed-loop ml- based planning benchmark for autonomous vehicles," *arXiv preprint arXiv:2106.11810*, 2021.
- [45] J. Bernhard, K. Esterle, P. Hart, and T. Kessler, "BARK: Open behavior benchmarking in multi-agent environments," in *IROS*. IEEE, 2020.
- [46] N. Niroomand, C. Bach, and M. Elser, "Vehicle dimensions based passenger car classification using fuzzy and non-fuzzy clustering methods," *Transportation Research Record*.
- [47] M.-F. Chang, J. Lambert, P. Sangkloy, J. Singh, S. Bak, A. Hartnett, D. Wang, P. Carr, S. Lucey, D. Ramanan *et al.*, "Argoverse: 3D tracking and forecasting with rich maps," in *CVPR*, 2019.
- [48] T. Fawcett, "An introduction to ROC analysis," *Pattern recognition letters*, vol. 27, no. 8, pp. 861–874, 2006.

Model	AUC-ROC $_P$ ↑	AUC-ROC $_{Pa}$ ↑
TAP	0.732	0.680
TIP $_P$	0.802	0.620
TIP $_{Pa}$	0.623	0.825

TABLE IV: Comparison between our TIP $_P$ and the baseline model TAP, in terms of the task performance, by using an **IDM** model as in BARK to simulate realistic ado agent behaviors. We observe consistent improvements in task performance with a different simulation model.

α	minADE↓	minFDE↓	wADE↓	wFDE↓	AUC-ROC↑
0	6.16	10.03	7.34	13.48	0.226
1	6.21	10.18	7.43	13.60	0.232
5	6.44	10.46	7.57	13.81	0.251
20	6.59	10.57	8.52	15.07	0.314
100	7.32	11.52	14.51	22.48	0.477

TABLE V: Performance of TIP $_W$ as a function of the task loss coefficient α when **predicting for 4 agents**. Task performance significantly improves at the cost of prediction accuracy.

APPENDIX

A. Additional Experiments in Planning Task

In this section, we present additional experiments in the planning task that demonstrate our model works with an additional simulation system.

1) *IDM Simulation Model*: We performed additional experiments by adopting the IDM model used in BARK [45] to simulate ado agent behaviors. The results are summarized in Table IV. We observe that our proposed approach shows improved performance in the task performance under a different simulation model, which is consistent with what we observe in Table I.

B. Additional Experiments in Warning Task

In this section, we present additional experiments in the warning task that demonstrate our model works for more than two agents.

1) *Multiple Object Agents*: We performed additional experiments in the warning task by using our proposed predictor to predict *joint* future trajectories for 4 agents, including the ego agent and 3 closest object agents in each scenario. The warning signal is triggered if any of the object agents is getting too close to the ego vehicle trajectory, using the same criteria as in Eq. (9). The comparison is summarized in Table V. The results show that our proposed predictor is able to improve the task performance by a large margin (39%), at the cost of little accuracy (5.4% for minFDE), with a reasonable α value of 20, compared to the task-agnostic baseline ($\alpha = 0$).