
DiffStack: A Differentiable *and* Modular Control Stack for Autonomous Vehicles

Peter Karkus¹, Boris Ivanovic¹, Shie Mannor^{1,2}, Marco Pavone^{1,3}

¹NVIDIA Research, ²Technion, ³Stanford University

{pkarkus, bivanovic, smannor, mpavone}@nvidia.com

Abstract

Autonomous vehicle (AV) stacks are typically built in a modular fashion, with explicit components performing detection, tracking, prediction, planning, control, etc. While modularity improves reusability, interpretability, and generalizability, it also suffers from compounding errors, information bottlenecks, and integration challenges. To overcome these challenges, a prominent approach is to convert the AV stack into an end-to-end neural network and train it with data. While such approaches have achieved impressive results, they typically lack interpretability and reusability, and they eschew principled analytical components, such as planning and control, in favor of deep neural networks. To enable the joint optimization of AV stacks while retaining modularity, we present DiffStack, a differentiable *and* modular stack for prediction, planning, and control. Crucially, our model-based planning and control algorithms leverage recent advancements in differentiable optimization to produce gradients, enabling optimization of upstream components, such as prediction, via backpropagation through planning and control. Our results on the nuScenes dataset indicate that end-to-end training with DiffStack yields substantial improvements in open-loop and closed-loop planning metrics by, e.g., learning to make fewer prediction errors that would affect planning. Beyond these immediate benefits, DiffStack opens up new opportunities for fully data-driven yet modular and interpretable AV architectures. Project website: <https://sites.google.com/view/diffstack>

1 Introduction

Intelligent robotic systems, such as autonomous vehicles (AVs), are typically architected in a modular fashion and comprised of modules performing detection, tracking, prediction, planning, and control, among others [15; 44; 30; 45; 2; 32; 55; 33]. Modular architectures are generally desirable because of their verifiability, interpretability and generalization performance; however, they also suffer from compounding errors, information bottlenecks, and integration challenges.

A promising line of work tackling these issues focuses on making AV stacks more integrated (by relaxing inter-module interfaces) and data-driven (by optimizing modules jointly with respect to their downstream task). For example, in the context of AV perception, recent work has achieved substantial performance gains by jointly training tracking models with detection [53] and prediction models [49; 48]. To extend such a joint, data-driven approach to decision making, existing approaches replace hand-engineered components, e.g., planning and control algorithms, with deep neural networks [5; 51; 9]. As neural networks are differentiable, they can be optimized end-to-end for a final control objective; however, they offer weaker generalization, little to no interpretability or safety guarantees.

We introduce **DiffStack**, a differentiable AV stack with modules for prediction, planning, and control that combines the benefits of modular and data-driven architectures (Fig. 1). The prediction module

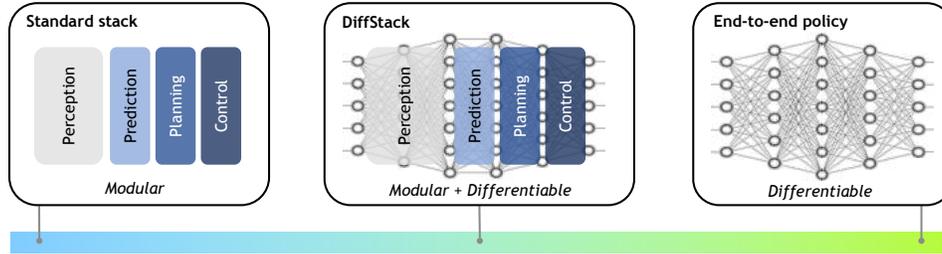


Figure 1: We desire AV architectures that are modular *and* differentiable, combining interpretability, reusability, and generalizability with data-driven end-to-end learning. We take an important step towards this end with DiffStack, a differentiable AV stack for prediction, planning, and control.

in DiffStack is a learned neural network that predicts the future motion of agents; the planning and control modules are principled, hand-engineered algorithms that produce AV actions given the current world state and motion predictions. Importantly, our hand-engineered planning and control algorithms are *differentiable*, enabling the training of the upstream prediction module for a downstream control objective by backpropagating gradients *through* the algorithms. In doing so, DiffStack can jointly optimize the entire stack for the final control objective, as in end-to-end neural networks, however it also maintains the interpretability and formal guarantees of standard modular AV architectures.

To make the planner and controller differentiable, we build upon the growing body of work on differentiable algorithm networks [26]. In particular, we leverage the differentiable Model Predictive Control (MPC) algorithm [1] for our controller. While differentiable algorithms show great promise, most prior works consider relatively simple control problems, e.g., pendulum and cartpole in [1]. To the best of our knowledge, this is the first work to demonstrate that, through careful design choices and integration, differentiable algorithms can be composed into stacks and trained with real-world data for AV prediction, planning, and control.

We evaluate DiffStack in both open-loop and closed-loop simulation settings using the large-scale, real-world nuScenes dataset [6]. Our results show some immediate benefits of differentiable stacks: by training a prediction model with respect to the final control objective, DiffStack increases the effectiveness of predictions for decision making by up to 15% over a large number of diverse scenarios. DiffStack achieves this by, e.g., learning to make fewer prediction errors that would negatively affect planning. Further, DiffStack outperforms alternative handcrafted planning-aware prediction losses; compensates for artificially introduced integration errors; and tunes an interpretable cost function to make resulting plans more similar to human driving. Beyond these immediate benefits, our work demonstrates the feasibility and potential of differentiable AV decision making stacks, making an important step towards a new class of data-driven yet modular AV architectures.

2 Related work

Modular AV architectures. Many state-of-the-art methods for AV perception have achieved substantial gains by relaxing bottlenecks between learned neural network modules and training multiple modules together. Examples include works that more tightly integrate object detection and tracking [53], as well as prediction with object detection [48], classification [19], and tracking [47; 20; 49]. While successful, these works focus on AV modules where neural networks are already commonly employed and thus differentiable, i.e., perception and prediction. Looking downstream, there are many works that focus on more tightly integrating prediction within planning [54; 29; 18; 10; 41], however, these works only *use* predictions in planning rather than *improve* the prediction module with information from planning. Recently, works such as [8; 17; 31] present potential avenues for designing upstream modules in a planning-aware fashion. In particular, Ivanovic and Pavone [17] propose an evaluation metric for motion prediction that considers the impact of predictions on downstream planning. While this metric can be used to train prediction with information from planning, we will show in Section 4 that the resulting improvement is less than with DiffStack.

Data-driven AV architectures. Another line of work eschews modularity entirely, and learns a policy network directly from data [5; 39; 21]. The approach benefits from the removal of information bottlenecks and performance scaling with increasing dataset sizes. However, fully end-to-end



Figure 2: DiffStack is comprised of differentiable modules for prediction, planning, and control. Importantly, this means that gradients can propagate backwards all the way from the final planning objective, allowing upstream predictions to be optimized with respect to downstream decision making.

approaches lack interpretability, important for debugging, verification, and safety guarantees; it is also unclear how well purely data-driven methods can generalize to unseen or rare-but-critical scenarios. Accordingly, recent works have focused on reintroducing structure and interpretability to learned policy networks. Most notably, the Neural Motion Planner [51] and follow-up methods [52; 13; 9] use object detection and agent behavior prediction as interpretable intermediate auxiliary learning objectives while training a joint backbone network from sensors directly to decision making. An important downside of this approach is that all components are neural networks, preventing the incorporation of principled planning and control algorithms, which are otherwise common in practice.

Differentiable algorithms. The idea of making existing model-based algorithms differentiable opens up new design paradigms that blend data-driven deep learning with classic model-based algorithms [26]. When both neural network models and fixed algorithmic computation can be expressed in the same “neural programming” framework, one can perform joint end-to-end learning even in structured modular architectures. Differentiable algorithms have been explored in various robotic domains including localization [23; 25], mapping [22; 27], navigation [24; 16; 26], and robot control [34; 1; 36; 14]. Our work is motivated by the initial successes of these approaches in rather simple robotic domains. To the best of our knowledge, this is the first work to apply differentiable algorithms in the context of autonomous driving stacks including prediction, planning, and control.

3 DiffStack: A Differentiable and Modular Autonomous Vehicle Stack

DiffStack is an autonomous driving stack with modules for prediction, planning, and control (Fig. 2). At a high level, DiffStack produces a trajectory for the AV to follow given a goal state, the tracked states of other agents, and a lane graph. To do so, it predicts where non-AV agents could move in the future; uses the predictions to select a safe, low-cost, and dynamically-feasible motion plan from a set of candidates; and finally optimizes the plan to yield a continuous and smooth ego-trajectory.

The key feature of DiffStack is that it is *differentiable*, i.e., we can compute $\partial \mathcal{L}_j / \partial \theta_i$, the gradient of the training objective \mathcal{L}_j of downstream module j with respect to the parameters θ_i of upstream module i . Differentiability enables training all parameters jointly for the overall objective of the stack with a gradient-based optimizer. We can also interpret DiffStack as a neural network with interpretable internal representations and fixed computational blocks for planning and control. Section 3.1 will describe how these components are made differentiable.

The main use case considered in this paper is to train the prediction module with respect to a final control objective. Different agents’ predictions have different importance to the ego-vehicle (see examples in Appendix B). Standard metrics used to train prediction models are agnostic to the consequences of downstream planning, treating harmless and dangerous errors equally [8; 17]. By training directly for the end-objective of ego control, we expect DiffStack’s prediction module to not only be accurate, but also synergize with planning and yield better overall ego controls.

We consider two open-loop training settings. In reinforcement learning (RL) experiments we optimize a hindsight cost that measures the quality of an ego trajectory in *hindsight*, i.e., after observing the future trajectory of other agents in the scene. The hindsight cost can be viewed as a negative reward for RL. In imitation learning (IL) experiments we compare the planned ego trajectory to the ego-vehicle’s enacted motion in the data through a mean-squared-error (MSE) loss. We additionally explore using DiffStack to tune the hand-specified planning and control cost functions. In this section, we describe each module in DiffStack and their respective training setups, with more details in Appendix C.

Nomenclature. *Ego* refers to the AV and *agent* denotes a non-AV vehicle or pedestrian. States s are 2D positions and velocities for pedestrians; 2D positions, heading, and longitudinal velocity for vehicles. Control variables u are heading rate and longitudinal acceleration. A trajectory is a sequence of states, or states and controls, depending on the context. Distances between states are measured using the 2D Euclidean distance, whereas root-mean-squared state distance over time is used for trajectories. We perform two types of optimization: *training* (sometimes called *learning*) optimizes parameters of the stack θ over a dataset for a training objective or loss \mathcal{L} ; *planning* and *control* optimizes the ego trajectory s_e to minimize a cost C with θ fixed.

3.1 DiffStack modules

Prediction. We employ Trajectron++ [40], a state-of-the-art CVAE that takes T_H seconds of state history for all agents $a \in A$ as input and outputs multimodal trajectory predictions,

$$\hat{s}_{a,k}^{(1:T)} = \text{CVAE}_\theta \left(s_{a' \in A}^{\text{gt}, (-T_H:0)} \right) \quad \text{for } k \in K, \quad (1)$$

where $k \in K$ is a mode of the output distribution. The encoder of the CVAE processes agent state histories with recurrent LSTM networks and models inter-agent interactions using graph-based attention. The decoder is a GRU that outputs a Gaussian Mixture Model (GMM) for each future timestep. The GMM modes correspond to the CVAE’s $K=25$ discrete latent states. To ensure predictions are dynamically-feasible, GMMs are defined over controls and then integrated through a known (differentiable) dynamics function to produce a trajectory. We use the default model configuration without map and ego conditioning. We augment the input states with an ego-indicator variable to allow for ego-agent relation reasoning. The raw prediction training objective is the InfoVAE loss, $\mathcal{L}_{\text{pred}} = \mathcal{L}_{\text{InfoVAE}}(\hat{s}_{a,k}, s_a^{\text{gt}})$, the same as for the original Trajectron++.

Planning. The planner is a sampling-based algorithm that generates a set of N dynamically-feasible ego trajectory candidates, $\mathcal{P} = \{s_n, u_n\}_{n \in N}$, and selects the candidate with the lowest cost. Namely,

$$s_{\text{plan}}, u_{\text{plan}} = \arg \min_{s_n, u_n \in \mathcal{P}} C(s_n, u_n; \hat{s}_{a \in A}, g, m), \quad (2)$$

where C is the cost function, $\hat{s}_{a \in A}$ are multimodal predictions from the prediction module, g is a given goal, and m is a lane graph. To generate trajectory candidates we sample a set of lane-centric terminal states, fit a cubic spline from the current state to the terminal state, and reject dynamically-infeasible trajectories. The cost function is a weighted sum of standard handcrafted terms,

$$C = w_1 C_{\text{coll}}(s, \hat{s}_{a \in A}) + w_2 C_g(s, g) + w_3 C_{\ell\perp}(s, m) + w_4 C_{\ell\angle}(s, m) + w_5 C_u(u), \quad (3)$$

penalizing collisions, distance to the goal, lateral lane deviation, lane heading deviation, and control effort, respectively. Most notably, the collision term incorporates predictions into planning, $C_{\text{coll}}(s, \hat{s}_{a \in A}) = \sum_{a \in A} \sum_{t \in 1:T} \varphi \left(\sum_{k \in K} \pi_k \|s^{(t)} - \hat{s}_{a,k}^{(t)}\|^2 \right)$, where $\hat{s}_{a,k}$ is the k -th mode of the predicted trajectory distribution for agent a , π_k is the probability of the k -th mode, φ is a Gaussian radial basis function, and $\|\cdot\|$ is the Euclidean norm. For simplicity, we only perform prediction for the vehicle closest to the ego and use GT futures for other agents. The remaining terms of (3) are defined in Appendix C.

We desire the planner to be differentiable, i.e., we wish to compute $\frac{\delta \mathcal{L}_{\text{plan}}(s_{\text{plan}}, u_{\text{plan}})}{\delta \theta}$ where $\mathcal{L}_{\text{plan}}$ is a measure of plan quality given some form of training data, and θ are parameters of the prediction model and/or the cost function. Using the chain rule, $\frac{\delta \mathcal{L}_{\text{plan}}}{\delta \theta} = \frac{\delta \mathcal{L}_{\text{plan}}}{\delta C(\hat{s}_\theta)} \frac{\delta C(\hat{s}_\theta)}{\delta \hat{s}_\theta} \frac{\delta \hat{s}_\theta}{\delta \theta}$, where the last term is differentiable because C is composed of differentiable functions. We make the middle term differentiable by relaxing the min operator of (2) into sampling from a categorical distribution, effectively treating the planner as a classifier during training, where the probability p_n of choosing candidate n is proportional to $\exp(-C(s_n, u_n; \cdot))$. We can then use cross entropy as the training loss, $\mathcal{L}_{\text{plan}} = \mathcal{L}_{\text{CE}} = -\sum_{n \in N} \mathbb{1}(s_n = s^*) \log p_n$, where the target trajectory s^* is chosen depending on the training setup (see Section 3.2).

Control. The control module performs MPC over a finite horizon using an iterative box-constrained linear quadratic regulator (LQR) algorithm [43]. Formally, we aim to solve

$$s_{\text{ctrl}}, u_{\text{ctrl}} = \arg \min_{s, u} C(s, u; \hat{s}_{a \in A}, g, m) \quad \text{s.t.} \quad s^{(0)} = s^{\text{init}}, \quad s^{(t+1)} = f_d(s^{(t)}, u^{(t)}), \quad \underline{u} \leq u \leq \bar{u}, \quad (4)$$

Table 1: Open-loop evaluation results. DiffStack outperforms all baselines in planning and control metrics because it trains its prediction model for the downstream task.

Method	Pred. Obj.	End Obj.	ADE (m) ↓ ± Std. Err.	NLL (nats) ↓ ± Std. Err.	Planning Loss ↓ $\times 10^{-2}$, ± Std. Err.	Hindsight Cost ↓ $\times 10^{-2}$, ± Std. Err.
No prediction			–	–	0.00 ± 0.00	0.00 ± 0.00
Standard	✓		1.32 ± 0.06	0.42 ± 0.04	-3.76 ± 0.23	-1.61 ± 0.05
Distance weighted	✓		1.34 ± 0.06	1.19 ± 0.10	-4.04 ± 0.22	-1.68 ± 0.04
∇ Cost weighted [17]	✓		1.43 ± 0.10	-1.06 ± 1.52	-4.29 ± 0.20	-1.77 ± 0.07
DiffStack (no $\mathcal{L}_{\text{pred}}$)		✓	1.37 ± 0.06	8.01 ± 0.75	-5.13 ± 0.22	-1.87 ± 0.04
DiffStack	✓	✓	1.27 ± 0.07	0.38 ± 0.05	-5.13 ± 0.18	-1.86 ± 0.04
GT prediction			–	–	-7.56 ± 0.00	-2.25 ± 0.00

where C denotes the cost function, f_d the dynamics, s^{init} the current ego state, and \underline{u}, \bar{u} the control limits. We use the cost defined in (3) for C and the dynamically-extended unicycle [28] for f_d . We initialize the trajectory with $s_{\text{plan}}, u_{\text{plan}}$ from the planner. The algorithm then iteratively forms and solves a quadratic LQR approximation of (4) around the current solution $s^{(i)}, u^{(i)}$ for iteration i , using first- and second-order Taylor approximations of f_d and C , respectively. The trajectory is updated to be close to the LQR optimal control while also decreasing the original non-quadratic cost. We stop iterations upon convergence or a fixed limit.

To make the control algorithm differentiable, we leverage Amos et al. [1]: The iLQR optimal trajectory can be differentiated with respect to the control and dynamics functions by implicitly differentiating the underlying KKT conditions of the last LQR approximation. The gradients can be analytically computed by one additional backward pass of a modified iterative LQR solver. If iLQR fails to converge, we do not backpropagate gradients. In our setting, f_d is fixed and we backpropagate gradients to C and further to the prediction module.

3.2 End-to-end training

An important question for data-driven AV stacks is the training objective and data. Learning in the real world is prohibitively expensive, and building a simulator with realistic traffic agent behavior is an open challenge. Accordingly, standard practice is to perform open-loop training with human driving data [3; 51; 52; 18; 10; 12; 11]. We consider two common types of open-loop training settings: reinforcement learning (RL) and imitation learning (IL).

In the RL setting, we aim to minimize the (hindsight) cost of DiffStack’s output trajectories over training examples in the dataset, $\mathcal{L}_{\text{HC}} = C_h(s_{\text{ctrl}}, u_{\text{ctrl}}; s_{a \in A}^{\text{gt}}, g, m)$. The hindsight cost C_h captures the quality of a trajectory in *hindsight*, i.e., after knowing the future trajectory of non-ego agents $s_{a \in A}^{\text{gt}}$, similar to the concept of rewards in RL. We choose C_h identical to the control cost C , but with GT future trajectories instead of predictions for other agents. Note that our open-loop training setup does not account for the effect of ego actions on other agents; nevertheless, we treat recorded trajectories as GT futures, as in the prediction literature [38]. Given GT futures, \mathcal{L}_{HC} is differentiable wrt. parameters in DiffStack, so $\mathcal{L}_{\text{ctrl}} = \mathcal{L}_{\text{HC}}$ can be directly optimized as the control module’s loss. In the IL setting, we aim to minimize the MSE between DiffStack’s output ego trajectory and the GT in the dataset, $\mathcal{L}_{\text{IL}} = \sum_{t=1:T} \|s_{\text{ctrl}}^{(t)} - s^{\text{gt},(t)}\|^2$.

The total loss for training the stack is a linear combination of module-wise objectives, $\mathcal{L} = \alpha_1 \mathcal{L}_{\text{pred}} + \alpha_2 \mathcal{L}_{\text{plan}} + \alpha_3 \mathcal{L}_{\text{ctrl}}$. We experiment with different combinations of losses, including setting each α_i to zero. In the rest of the paper, we omit α_i for brevity.

4 Experiments

Our experiments investigate the following questions: 1) can DiffStack learn predictions that lead to better plans? 2) how does DiffStack compare to alternative planning-aware training techniques? 3) can DiffStack correct systematic integration errors? 4) can DiffStack learn the control cost from imitation? 5) do our open-loop results translate to closed-loop evaluation?

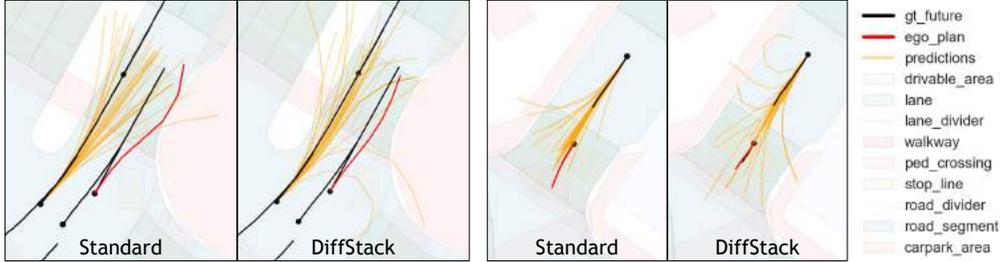


Figure 3: **Left:** The standard stack predicts an unlikely lane intrusion with high confidence, causing the planner to swerve. In contrast, DiffStack’s prediction assigns most of the probability mass to reasonable trajectories that consider the presence of the ego vehicle, and the resulting motion plan stays well within lane boundaries. **Right:** The standard stack incorrectly predicts that a vehicle would collide into the ego-vehicle from behind with high probability, causing the ego-vehicle to move into the intersection. DiffStack assigns more probability to the agent slowing down, as a result the ego-vehicle slows down before the intersection.

We begin with open-loop experiments (1–4), and then present closed-loop results (5) in Section 4.3.

4.1 Experimental setup

Dataset. We use the nuScenes dataset [6], comprised of state annotations for vehicles and pedestrians in 1000 scenes across Boston and Singapore. For training and open-loop evaluation we sample suitable planning scenarios from the dataset with $T_H = 4s$ history and $T = 3s$ future data. For each scenario, we choose one vehicle to act as the ego. Details including dataset splits are in Appendix D.

Metrics. We evaluate DiffStack’s prediction module via standard prediction metrics: average displacement error (ADE) for the most-likely prediction and negative log-likelihood (NLL) for its full distributional output. To evaluate planning and control, we use the cross-entropy planning loss $\mathcal{L}_{\text{plan}} = \mathcal{L}_{\text{CE}}$, and control loss $\mathcal{L}_{\text{ctrl}}$. We report the hindsight cost for reinforcement learning ($\mathcal{L}_{\text{ctrl}} = \mathcal{L}_{\text{HC}}$) and MSE for imitation learning experiments ($\mathcal{L}_{\text{ctrl}} = \mathcal{L}_{\text{MSE}}$). Since control is an AV stack’s end-goal, $\mathcal{L}_{\text{ctrl}}$ reflects the overall performance of the stack. To make these metric values more interpretable, we report them relative to a **No prediction** baseline and a **GT prediction**-based oracle. The baseline ego plans without predictions (ignoring the predicted agent), providing a performance lower bound. The GT oracle plans with GT futures in place of predictions, providing an (imperfect) notion of a performance upper bound. All results are averaged over our validation set, and we report standard errors of the mean over 5 training seeds.

Baselines. We compare DiffStack with standard AV stacks. For a fair comparison, we use the same set of modules as in DiffStack, but without making the planner and controller differentiable. We only train the prediction model, using increasingly planning-aware training objectives. **Standard** trains the prediction model with only the prediction loss $\mathcal{L}_{\text{pred}}$, unaware of downstream planning. Next, we re-weight prediction losses (in each batch) based on a handcrafted measure of relevance for planning. **Distance weighted** weights losses with the inverse distance between the ego and agent GT futures. **∇ Cost weighted** uses the magnitude of the control cost gradient with respect to the GT future trajectory of the agent, proposed as a planning-aware prediction metric in [17]. Finally, DiffStack backpropagates gradients of the final loss to the prediction module. **DiffStack (no $\mathcal{L}_{\text{pred}}$)** uses $\mathcal{L} = \mathcal{L}_{\text{plan}} + \mathcal{L}_{\text{ctrl}}$; **DiffStack** uses $\mathcal{L} = \mathcal{L}_{\text{pred}} + \mathcal{L}_{\text{plan}} + \mathcal{L}_{\text{ctrl}}$.

Implementation details. We implement DiffStack in PyTorch [35] and build on the open-source code of Trajectron++ [40] and Differentiable MPC [1]. We use $T_H = 4s$, $T = 3s$, and $\Delta t = 0.5s$. We train all models for 20 epochs using 4 NVIDIA Tesla V100 GPUs, taking 10–20 hours. Additional details are in the Appendix. Code is available at <https://sites.google.com/view/diffstack>.

4.2 Open-loop results

End-to-end training is useful. Our main results are summarized in Table 1. The key observation is that DiffStack improves the effectiveness of predictions on ego planning by up to 15.5% compared to standard training ($\mathcal{L}_{\text{ctrl}} = -1.86$ vs. -1.61); and reduces the gap to the cost attainable with a

Table 2: System integration results. DiffStack compensates for errors in module integration. Even if GT prediction targets are erroneously offset by 1m, training with downstream losses corrects the prediction model’s bias and improves overall system performance.

Method	Pred. Obj.	End Obj.	ADE (m) ↓ ± Std. Err.	NLL (nats) ↓ ± Std. Err.	Planning Loss ↓ ×10 ⁻² , ± Std. Err.	Hindsight Cost ↓ ×10 ⁻² , ± Std. Err.
No prediction			–	–	0.00 ± 0.00	0.00 ± 0.00
Standard-biased	✓-biased		3.64 ± 0.34	8.43 ± 0.72	16.22 ± 2.73	3.13 ± 0.40
DiffStack-biased	✓-biased	✓	2.01 ± 0.08	4.20 ± 0.09	-3.58 ± 0.96	-1.61 ± 0.12
GT prediction			–	–	-7.56 ± 0.00	-2.25 ± 0.00

GT-informed oracle by 39.1% ($-1.86 + 2.25$ vs. $-1.68 + 2.25$). DiffStack also improves planning without significantly impacting raw prediction accuracy (ADE=1.27 vs. 1.32, within standard error). Alternative methods that re-weight the prediction loss in a planning-aware manner (rows 2 and 3) also help, but less than DiffStack. Surprisingly, we can even recover comparable prediction performance in terms of ADE when training solely for planning and control objectives (row 4). The poor distribution fit (high NLL) is due to our control cost being agnostic to the variance of the predicted GMMs.

DiffStack makes fewer prediction errors that affect planning. Qualitatively analysing predictions and plans shows that the main source of improvement from end-to-end training is the reduction of spurious and/or unrealistic predictions that lead to plans with large (unnecessary) deviations from the lane center or the goal position. Fig. 3 shows two particular examples. In both cases, DiffStack’s predictions are more realistic and accurate, yielding much more reasonable downstream plans.

DiffStack can compensate for system integration errors. Integrating independently developed modules into an AV stack can be challenging due to, e.g., misaligned module interfaces. Table 2 shows results for an experiment that explores this issue. We introduce an artificial interface mismatch between prediction and planning by adding a fixed 1m offset to all GT prediction targets. As a result, the standard model yields very poor plans (note the positive value in row 1, indicating worse performance than baseline planning). However, DiffStack can compensate for interface mismatch and substantially improves the overall plans, suggesting that it can learn to correct erroneous predictions that affect planning. More broadly, these results demonstrate the potential for differentiable AV stacks to reduce various development/engineering costs associated with developing AVs, e.g., by replacing tedious manual parameter tuning with end-to-end data-driven optimization. We further explore this topic in the cost tuning experiments below.

DiffStack can imitate humans better. Imitation learning results are summarized in Table 4. The overall performance of the stack is now measured by $\mathcal{L}_{\text{ctrl}} = \mathcal{L}_{\text{MSE}}$, the MSE between planned and human expert trajectories relative to the MSE for the baseline planner (without predictions). We observe a similar trend as before; however, the results also reveal some limitations. The standard stack performs worse in terms of MSE than the baseline stack which neglects predictions entirely (note the positive sign for relative \mathcal{L}_{MSE} in row 1). While DiffStack improves MSE significantly, the absolute difference is small. We hypothesize this is caused by our cost function not being rich enough to fully capture human driving behavior; and because agent-agent interactions that affect planning are rare in nominal driving data (comprised mainly of lane- and speed-keeping). These limitations could be addressed by a richer cost function and alternative goal definitions, which we leave to future work.

DiffStack can learn a better control cost. While our main use case in this paper is to learn planning-aware prediction, DiffStack opens up various other opportunities for data-driven optimization of various components of the AV stack. For example, an important challenge in practice is to design a cost function

for planning and control. In this experiment, we explore DiffStack’s potential to learn interpretable control costs that optimize the quality of output plans from data. Table 3 reports results in the imitation learning setting, where we first train a prediction module (as before), then we train for an additional 20 epochs allowing DiffStack to update the weights w_i of the control cost (3) by backpropagating gradients from the final control loss $\mathcal{L}_{\text{ctrl}}$. Compared to the default hand-tuned weights (row 1), DiffStack significantly decreases plan MSEs (row 2). The resulting learned weights w_i are lower for

Table 3: DiffStack learns a better control cost.

Method	Planning Loss ↓	MSE (m ²) ↓
Hand-tuned	2.67 ± 0.00	0.38 ± 0.00
Learned	2.41 ± 0.01	0.32 ± 0.00

Table 4: Imitation learning results. DiffStack’s plans better match the GT ego trajectories from data.

Method	Pred. Obj.	End Obj.	ADE (m) ↓ ± Std. Err.	NLL (nats) ↓ ± Std. Err.	Planning Loss ↓ ×10 ⁻³ , ± Std. Err.	MSE (m ²) ↓ ×10 ⁻² , ± Std. Err.
No prediction			–	–	0.00 ± 0.00	0.00 ± 0.00
Standard	✓		1.25 ± 0.02	0.53 ± 0.23	–0.50 ± 0.05	0.60 ± 0.15
DiffStack	✓	✓	1.43 ± 0.14	0.30 ± 0.23	–0.80 ± 0.04	–1.80 ± 0.16
GT prediction			–	–	–1.10 ± 0.00	–2.80 ± 0.00

the control effort term, and higher for the goal and lane keeping terms (see Appendix A). To ensure safety, we fix the weights for collision avoidance. We leave comparison with alternative techniques for learning control costs to future work. We expect DiffStack to be more effective compared to, e.g., Bayesian optimization, where the number of learned parameters is large.

Ablations. Results of an ablation study can be found in Appendix A. In short, improvements from DiffStack are consistent with our observations when training for $\mathcal{L} = \mathcal{L}_{\text{pred}} + \mathcal{L}_{\text{plan}}$; $\mathcal{L} = \mathcal{L}_{\text{pred}} + \mathcal{L}_{\text{ctrl}}$; when changing the relative scale of loss components; with higher time resolution $\Delta t = 0.1$ for planning and control; and when only using one of the planning or control modules.

4.3 Closed-loop evaluation

Simulation. We perform closed-loop simulation using a simple log-replay setup, where ego states are unrolled based on the planned control outputs and known dynamics, while non-ego agents follow their fixed trajectories recorded in the dataset. The evaluation scenarios are $T_{\text{sim}} = 10s$ long. The goal and lane inputs for planning are updated in each simulation step based on the logged ego trajectory. We compute the following metrics. **Trajectory Cost** captures closed-loop performance: it evaluates the cost function C on unrolled simulation trajectories, $1/T_{\text{sim}} \sum_{t=1:T_{\text{sim}}} C(s_{\text{sim}}^{(t)}, u_{\text{sim}}^{(t)}; s_{a \in A}^{\text{gt}, t}, m^{(t)})$, where s_{sim} and u_{sim} are the unrolled ego state and control. We exclude the goal cost term because the goal is updated in each simulation step. **Open-loop Cost** is the average of hindsight costs calculated open-loop at each simulation step, $\mathcal{L}_{\text{HC}} = C(s_{\text{ctrl}}, u_{\text{ctrl}}; s_{a \in A}^{\text{gt}}, g, m)$. **Collision Cost** is the collision term in the trajectory cost, $w_1 C_{\text{coll}}(s_{\text{sim}}^{(t)}, s_{a \in A}^{\text{gt}, (t)})$. **Lane Cost** is the sum of lane keeping terms, $w_3 C_{\ell \perp} + w_4 C_{\ell \angle}$. **Control Effort** is the control effort term, $w_5 C_u$. **Deviation** is the average distance between unrolled and logged ego trajectories from the data, $\|s_{\text{sim}}^{(t)} - s^{\text{gt}, (t)}\|$.

Results. Simulation results are in Table 5. Most importantly, we observe similar relative improvement for DiffStack in terms of closed loop trajectory cost as in terms of open-loop cost. DiffStack performs substantially better than the *no prediction* baseline, but somewhat surprisingly the standard stack performs worse, both in terms of closed-loop trajectory cost and open-loop cost. Analysing the cost components sheds light on possible reasons. As expected, incorporating predictions into planning results in lower collision costs, but higher lane keeping costs. The contribution of these two terms to the average cost depends on the data distribution, e.g., the frequency of close interactions where predictions are useful. As we saw earlier in Fig. 3, poor predictions in the standard stack frequently lead to unnecessary lane deviations in the planner. This is reflected in the substantially higher lane cost and control effort cost for the standard stack compared to DiffStack in Table 5.

5 Limitations & Conclusions

Limitations. One limitation of our work is the open-loop training and log-replay based simulation setup. In lieu of a real-world AV or a simulator with strong behavioral realism, this is standard practice; however, recent efforts on accurate behavior simulation could be leveraged in the future [4; 42; 7; 37; 50]. Our implementation of DiffStack also has limitations. First, it is not differentiable wrt. *all* possible parameters, e.g., no gradients flow from the control loss to the planner’s trajectory candidate generator. Future work may develop more sophisticated differentiable algorithms and explore ideas for gradient approximation for non-differentiable components [46]. Second, individual modules could be improved, e.g., by adding a more sophisticated prediction model, improving candidate sampling in the planner, and adding trust-region constraints to the controller. Finally, even though hand-engineered components, modularity, and intermediate training objectives in differentiable stacks

Table 5: Closed-loop results. DiffStack outperforms the standard stack similarly to open-loop results.

Method	Trajectory Cost $\times 10^{-2} \downarrow$	Open-loop Cost $\times 10^{-2} \downarrow$	Collision Cost $\times 10^{-2} \downarrow$	Lane Keep Cost $\times 10^{-2} \downarrow$	Control Effort $\times 10^{-2} \downarrow$	Deviation (m) $\times 10^{-2} \downarrow$
No prediction	0.00	0.00	0.00	0.00	0.00	0.00
Standard	0.87	0.33	-8.21	8.91	0.18	1.05
DiffStack	-4.37	-2.40	-8.20	4.84	-0.99	0.98
GT prediction	-6.30	-2.58	-8.54	3.34	-1.08	0.70

remedy challenges of learning AV policies end-to-end, other challenges naturally remain. For instance, designing an overall performance metric, or reducing the scarcity of and cost to obtain (interesting) driving data remain open problems, each of which are impactful areas of future work.

Conclusions. In this paper we take a step towards fully-differentiable and modular AV stacks by introducing key differentiable components for prediction, planning, and control. Our experimental results show the potential benefits of jointly training modules of AV stacks for downstream performance. For motion prediction in particular, our results indicate that there is value in moving from purely prediction-oriented evaluation metrics towards downstream task-oriented metrics, in line with arguments in recent work [8; 17; 31]. While our experiments focused on learning planning-aware predictions, DiffStack opens up various exciting opportunities for task-oriented learning in modular stacks. For example, we may learn a rich neural network as part of the control cost to capture hard-to-engineer concepts, e.g., reasoning with occlusions or accounting for uncertainties. Eventually, we envision having differentiable modules for the entire AV stack, allowing any subset of modules to be learned and optimized for a downstream task. Overall, this would relax information bottlenecks and enable uncertainty to more easily propagate through the stack without needing to forgo interpretability, modularity, and verifiability of the various components.

Acknowledgements. We thank Yuxiao Chen and Edward Schmerling for generous help with our planner implementation and closed-loop evaluation experiment; Christopher Maes for fruitful discussions on differentiable MPC and cost learning experiments; and Nikolai Smolyanskiy for valuable high-level feedback.

References

- [1] Brandon Amos, Ivan Jimenez, Jacob Sacks, Byron Boots, and J Zico Kolter. Differentiable MPC for end-to-end planning and control. In *Advances in Neural Information Processing Systems*, pages 8299–8310, 2018.
- [2] Argo AI. Developing a self-driving system you can trust, April 2021. Available at <https://www.argo.ai/wp-content/uploads/2021/04/ArgoSafetyReport.pdf>.
- [3] M. Bansal, A. Krizhevsky, and A. Ogale. ChauffeurNet: Learning to drive by imitating the best and synthesizing the worst. In *Robotics: Science and Systems*, 2019.
- [4] Luca Bergamini, Yawei Ye, Oliver Scheel, Long Chen, Chih Hu, Luca Del Pero, Błażej Osiński, Hugo Grimmer, and Peter Ondruska. SimNet: Learning reactive self-driving simulations from real-world observations. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 5119–5125. IEEE, 2021.
- [5] Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Praseon Goyal, Lawrence D Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, et al. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*, 2016.
- [6] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuScenes: A multimodal dataset for autonomous driving. In *Conference on Computer Vision and Pattern Recognition*, 2020.
- [7] Holger Caesar, Juraj Kabzan, Kok Seang Tan, Whye Kit Fong, Eric M. Wolff, Alex H. Lang, Luke Fletcher, Oscar Beijbom, and Sammy Omari. NuPlan: A closed-loop ml-based planning benchmark for autonomous vehicles. In *Conference on Computer Vision and Pattern Recognition (CVPR) ADP3 Workshop*, 2021.

- [8] S. Casas, C. Gulino, S. Suo, and R. Urtasun. The importance of prior knowledge in precise multimodal prediction. In *International Conference on Intelligent Robots and Systems (IROS)*, 2020.
- [9] Sergio Casas, Abbas Sadat, and Raquel Urtasun. MP3: A unified model to map, perceive, predict and plan. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [10] Yuxiao Chen, Ugo Rosolia, Chuchu Fan, Aaron D. Ames, and Richard Murray. Reactive motion planning with probabilistic safety guarantees. In *Conference on Robot Learning*, 2020.
- [11] Yuxiao Chen, Boris Ivanovic, and Marco Pavone. ScePT: Scene-consistent, policy-based trajectory predictions for planning. In *Conference on Computer Vision and Pattern Recognition*, 2022.
- [12] Yuxiao Chen, Ugo Rosolia, Wyatt Ubellacker, Noel Csomay-Shanklin, and Aaron D. Ames. Interactive multi-modal motion planning with branch model predictive control. *Robotics and Automation Letters*, 7(2):5365–5372, 2022.
- [13] Alexander Cui, Sergio Casas, Abbas Sadat, Renjie Liao, and Raquel Urtasun. Lookout: Diverse multi-future prediction and planning for self-driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16107–16116, 2021.
- [14] Sebastian East, Marco Gallieri, Jonathan Masci, Jan Koutnik, and Mark Cannon. Infinite-horizon differentiable model predictive control. In *International Conference on Learning Representations*, 2019.
- [15] General Motors. Self-driving safety report, 2018. Available at <https://www.gm.com/content/dam/company/docs/us/en/gmcom/gmsafetyreport.pdf>.
- [16] Saurabh Gupta, James Davidson, Sergey Levine, Rahul Sukthankar, and Jitendra Malik. Cognitive mapping and planning for visual navigation. *arXiv preprint arXiv:1702.03920*, 2017.
- [17] B. Ivanovic and M. Pavone. Injecting planning-awareness into prediction and detection evaluation. In *IEEE Intelligent Vehicles Symposium*, 2022.
- [18] B. Ivanovic, A. Elhafi, G. Rosman, A. Gaidon, and M. Pavone. MATS: An interpretable trajectory forecasting representation for planning and control. In *Conference on Robot Learning*, 2020.
- [19] B. Ivanovic, K-H. Lee, P. Tokmakov, B. Wulfe, R. McAllister, A. Gaidon, and M. Pavone. Heterogeneous-agent trajectory forecasting incorporating class uncertainty. *arXiv preprint arXiv:2104.12446*, 2021. URL <https://arxiv.org/abs/2104.12446>.
- [20] B. Ivanovic, Y. Lin, S. Shrivastava, P. Chakravarty, and M. Pavone. Propagating state uncertainty through trajectory forecasting. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2022.
- [21] Ashesh Jain, Luca Del Pero, Hugo Grimmer, and Peter Ondruska. Autonomy 2.0: Why is self-driving always 5 years away? *arXiv preprint arXiv:2107.08142*, 2021.
- [22] Krishna Murthy Jatavallabhula, Ganesh Iyer, and Liam Paull. gradSLAM: Dense SLAM meets automatic differentiation. In *International Conference on Robotics and Automation (ICRA)*, pages 2130–2137, 2020.
- [23] Rico Jonschkowski, Divyam Rastogi, and Oliver Brock. Differentiable particle filters: End-to-end learning with algorithmic priors. *Proceedings of Robotics: Science and Systems*, 2018.
- [24] Peter Karkus, David Hsu, and Wee Sun Lee. QMDP-net: Deep learning for planning under partial observability. In *Advances in Neural Information Processing Systems*, pages 4697–4707, 2017.
- [25] Peter Karkus, David Hsu, and Wee Sun Lee. Particle filter networks with application to visual localization. In *Proceedings of the Conference on Robot Learning*, pages 169–178, 2018.

- [26] Peter Karkus, Xiao Ma, David Hsu, Leslie Pack Kaelbling, Wee Sun Lee, and Tomás Lozano-Pérez. Differentiable algorithm networks for composable robot learning. In *Robotics: Science and Systems (RSS)*, 2019.
- [27] Peter Karkus, Shaojun Cai, and David Hsu. Differentiable SLAM-net: Learning particle SLAM for visual navigation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2021.
- [28] S. M. LaValle. Better unicycle models. In *Planning Algorithms*, pages 743–743. Cambridge University Press, 2006.
- [29] C. Liu, S. Lee, S. Varnhagen, and H. E. Tseng. Path planning for autonomous vehicles using model predictive control. In *IEEE Intelligent Vehicles Symposium*, 2017.
- [30] Lyft. Self-driving safety report, 2020. Available at https://2eg1kz1onwffq1dj1lo2xh4bb-wpengine.netdna-ssl.com/wp-content/uploads/2020/06/Safety_Report_2020.pdf.
- [31] Rowan McAllister, Blake Wulfe, Jean Mercat, Logan Ellis, Sergey Levine, and Adrien Gaidon. Control-aware prediction objectives for autonomous driving. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2022.
- [32] Motional. Voluntary safety self-assessment, 2021. Available at https://drive.google.com/file/d/1JjfQByU_hWvSfkWzQ8PK2Z0ZfVCqQGDB/view.
- [33] NVIDIA. Self-driving safety report, 2021. Available at <https://images.nvidia.com/content/self-driving-cars/safety-report/auto-print-self-driving-safety-report-2021-update.pdf>.
- [34] Masashi Okada, Luca Rigazio, and Takenobu Aoshima. Path integral networks: End-to-end differentiable optimal control. *arXiv preprint arXiv:1706.09597*, 2017.
- [35] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, pages 8024–8035, 2019.
- [36] Marcus Pereira, David D Fan, Gabriel Nakajima An, and Evangelos Theodorou. MPC-inspired neural network policies for sequential decision making. *arXiv preprint arXiv:1802.05803*, 2018.
- [37] Davis Rempe, Jonah Philion, Leonidas J. Guibas, Sanja Fidler, and Or Litany. Generating useful accident-prone driving scenarios via a learned traffic prior. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [38] A. Rudenko, L. Palmieri, M. Herman, K. M. Kitani, D. M. Gavrila, and K. O. Arras. Human motion trajectory prediction: A survey. *International Journal of Robotics Research*, 39(8): 895–935, 2020.
- [39] Ahmad EL Sallab, Mohammed Abdou, Etienne Perot, and Senthil Yogamani. Deep reinforcement learning framework for autonomous driving. *Electronic Imaging*, 2017(19):70–76, 2017.
- [40] Tim Salzmann, Boris Ivanovic, Punarjay Chakravarty, and Marco Pavone. Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data. In *European Conference on Computer Vision*, 2020.
- [41] Simon Schaefer, Karen Leung, Boris Ivanovic, and Marco Pavone. Leveraging neural network gradients within trajectory optimization for proactive human-robot interactions. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9673–9679, 2021.
- [42] Simon Suo, Sebastian Regalado, Sergio Casas, and Raquel Urtasun. TrafficSim: Learning to simulate realistic multi-agent behaviors. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10400–10409, 2021.

- [43] Yuval Tassa, Nicolas Mansard, and Emo Todorov. Control-limited differential dynamic programming. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 1168–1175. IEEE, 2014.
- [44] Uber Advanced Technologies Group. A principled approach to safety, 2020. Available at <https://uber.app.box.com/v/UberATGSafetyReport>.
- [45] Waymo. Safety report, 2021. Available at <https://waymo.com/safety/safety-report>. Retrieved on July 4, 2021.
- [46] Théophane Weber, Nicolas Heess, Lars Buesing, and David Silver. Credit assignment techniques in stochastic computation graphs. *arXiv preprint arXiv:1901.01761*, 2019.
- [47] X. Weng, Y. Yuan, and K. Kitani. PTP: Parallelized tracking and prediction with graph neural networks and diversity sampling. *IEEE Robotics and Automation Letters*, 6(3):4640–4647, 2021.
- [48] Xinshuo Weng, Boris Ivanovic, Kris Kitani, and Marco Pavone. Whose track is it anyway? Improving robustness to tracking errors with affinity-based trajectory prediction. In *Conference on Computer Vision and Pattern Recognition*, 2022.
- [49] Xinshuo Weng, Boris Ivanovic, and Marco Pavone. MTP: Multi-hypothesis tracking and prediction for reduced error propagation. In *IEEE Intelligent Vehicles Symposium*, 2022.
- [50] Danfei Xu, Yuxiao Chen, Boris Ivanovic, and Marco Pavone. Bits: Bi-level imitation for traffic simulation. *arXiv preprint arXiv:2208.12403*, 2022.
- [51] Wenyuan Zeng, Wenjie Luo, Simon Suo, Abbas Sadat, Bin Yang, Sergio Casas, and Raquel Urtasun. End-to-end interpretable neural motion planner. In *Conference on Computer Vision and Pattern Recognition*, pages 8660–8669, 2019.
- [52] Wenyuan Zeng, Shenlong Wang, Renjie Liao, Yun Chen, Bin Yang, and Raquel Urtasun. Dsdnet: Deep structured self-driving network. In *European Conference on Computer Vision*, pages 156–172. Springer, 2020.
- [53] Xingyi Zhou, Vladlen Koltun, and Philipp Krähenbühl. Tracking objects as points. In *European Conference on Computer Vision*, 2020.
- [54] J. Ziegler, P. Bender, T. Dang, and C. Stiller. Trajectory planning for BERTHA – A local, continuous method. In *IEEE Intelligent Vehicles Symposium*, 2014.
- [55] Zoox. Safety report volume 2.0, 2021. Available at <https://zoox.com/safety/>.

A Additional results

A.1 Ablation study

Table 6: Ablation study results.

Method	$\mathcal{L}_{\text{pred}} \times 10^0$	$\mathcal{L}_{\text{plan}} \times 10^2$	$\mathcal{L}_{\text{ctrl}} \times 10^3$	ADE (m) ↓ ± Std. Err.	NLL (nats) ↓ ± Std.Err.	Planning Loss ↓ $\times 10^{-2}, \pm \text{Std.Err.}$	Hindsight Cost ↓ $\times 10^{-2}, \pm \text{Std.Err.}$
No prediction				–	–	0.00 ± 0.00	0.00 ± 0.00
Standard	1	0	0	1.32 ± 0.06	0.42 ± 0.04	–3.76 ± 0.23	–1.61 ± 0.05
DiffStack (default)	1	1	1	1.27 ± 0.07	0.38 ± 0.05	–5.13 ± 0.18	–1.86 ± 0.04
DiffStack (no $\mathcal{L}_{\text{plan}}$)	1	0	1	1.33 ± 0.08	0.48 ± 0.16	–4.98 ± 0.11	–1.86 ± 0.04
DiffStack (no $\mathcal{L}_{\text{ctrl}}$)	1	1	0	1.48 ± 0.21	0.55 ± 0.15	–5.24 ± 0.10	–1.85 ± 0.06
DiffStack	1	5	5	1.64 ± 0.17	0.85 ± 0.08	–5.53 ± 0.07	–1.89 ± 0.03
DiffStack	1	0.2	0.2	1.28 ± 0.04	0.66 ± 0.19	–4.79 ± 0.18	–1.81 ± 0.02
GT prediction				–	–	–7.56 ± 0.00	–2.25 ± 0.00
No pred. (no plan.)				–	–	–	0.0 ± 0.00
Standard (no plan.)	1	–	0	1.27 ± 0.06	0.31 ± 0.19	–	–1.8 ± 0.23
DiffStack (no plan.)	1	–	1	1.40 ± 0.15	0.19 ± 0.40	–	–2.0 ± 0.23
GT pred. (no plan.)				–	–	–	–2.6 ± 0.00
No pred. (no contr.)				–	–	0.00 ± 0.00	0.0 ± 0.00
Standard (no contr.)	1	0	–	1.23 ± 0.02	0.56 ± 0.22	–4.08 ± 0.27	–1.1 ± 0.03
DiffStack (no contr.)	1	1	–	1.46 ± 0.18	0.53 ± 0.13	–5.29 ± 0.10	–1.2 ± 0.01
GT pred. (no contr.)				–	–	–7.56 ± 0.00	–1.5 ± 0.00
No pred. ($\Delta t=0.1$)				–	–	0.00 ± 0.00	0.0 ± 0.00
Standard ($\Delta t=0.1$)	1	0	–	1.23 ± 0.02	0.56 ± 0.22	–5.65 ± 1.67	–2.1 ± 0.06
DiffStack ($\Delta t=0.1$)	1	1	–	1.49 ± 0.12	0.88 ± 0.27	–14.01 ± 0.63	–2.4 ± 0.05
GT pred. ($\Delta t=0.1$)				–	–	–25.11 ± 0.00	–2.9 ± 0.00

We present results for an ablation study in Table 6. The table is divided into sections in which decision making metrics are comparable. Decision making metrics are not comparable across sections because of the different configuration of the stack and a different criteria for rejecting samples from the validation set. Each section reports planning loss and hindsight cost relative to a **No prediction** baseline, and we report relative results for a **GT prediction** oracle, same as before. In rows of the table we train stacks with different weighted sum of losses, $\mathcal{L} = \alpha_1 \mathcal{L}_{\text{pred}} + \alpha_2 \mathcal{L}_{\text{plan}} + \alpha_3 \mathcal{L}_{\text{ctrl}}$. The second to fourth columns indicate the α_i parameters.

In the first section of Table 6 we compare DiffStack trained with different weighted combination of losses. We observe similar improvements when using only the control loss or only the planning loss. The relative weight of the prediction loss vs. the downstream decision making losses behave as expected: the higher/lower the weight for the downstream planning and control losses the more/less planning and control metrics improve at the expense of slightly worse/better raw prediction metrics.

In the second section we remove the planner module, and initialize the control algorithm with a zero control sequence. In the third section we remove the control module and use the output of the planner as the final ego trajectory. In the last section we change the planning and control time resolution from $\Delta t=0.5$ s to $\Delta t=0.1$ s. We observe similar trends in terms of improvement from DiffStack in all cases. In terms of absolute metrics we observe a benefit for using both the planning and the control modules (not shown in Table 6).

A.2 Learned control cost experiment

Table 7: Detailed results for the control cost learning experiment.

Cost weights	$w'_1 (C_{\text{coll}})$	$w'_2 (C_{\text{g}})$	$w'_3 (C_{\ell\perp})$	$w'_4 (C_{\ell\angle})$	$w'_5 (C_u)$	α	Planning Loss ↓	MSE (m ²) ↓
Hand-tuned	5.00	0.50	0.30	0.30	1.00	1.00	2.67	0.38
Learned	5.00	0.68	0.36	0.43	0.63	1.10	2.41	0.32

We report detailed results for the control cost learning experiment in Table 7. Values are averages over 5 training seeds. All standard errors are ≤ 0.01 (omitted).

We first train a prediction module (as before), then we train for an additional 20 epochs allowing DiffStack to update the weights w_i of the control cost (3) by backpropagating gradients from the final control loss $\mathcal{L}_{\text{ctrl}} = \mathcal{L}_{\text{MSE}}$. We do not use a planning loss $\mathcal{L}_{\text{plan}}$. For practical reasons we reparameterize the cost function such that the weights always sum to a fixed constant, and the total cost is multiplied with a scaler α . Specifically, the cost weights are given by $w_i = \alpha w'_i$, and $w'_i = c \exp(\theta_i) / \sum_{j \in 1:5} \exp(\theta_j)$. Here α and θ_i for $i \in 2:5$ are trainable parameters, c is a constant. To ensure safety we fix θ_1 , the weight parameter for the collision term C_{coll} . We initialize all parameters, θ_i , α , c such that the default hand-tuned weights (shown in row 1) are recovered. Compared to the default hand-tuned weights (row 1), DiffStack significantly decreases plan MSEs (row 2). The resulting learned weights are lower for the control effort term, and higher for the goal and lane keeping terms. We observed similar results when freezing the prediction module while learning the cost weights.

B Motivation for planning-aware prediction models

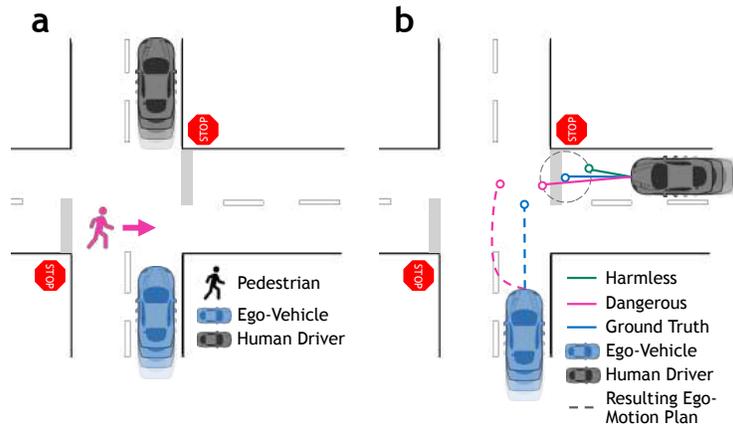


Figure 4: Not all predictions are equally important. (a) The motion of a jaywalker in front of the ego-vehicle is more important than a far-away vehicle. (b) Metrically-equal prediction errors can have severe consequences; incorrectly predicting that a vehicle will not stop at an intersection may trigger a dangerous evasive maneuver in the planner. The visualization was inspired by [17] and the figures include similar graphical elements.

C DiffStack implementation details

C.1 Planning module details

The planner generates trajectory candidates by first sampling a set of lane-centric terminal states. We consider all lanes that are within 4.5m from the ego goal g , and the heading difference between the lane and ego goal state is less than 90° . For all lanes we generate a set of terminal states based on distance traveled along the lane for a set of fixed acceleration values $a \in \{-3, -2, -1, -0.5, 0, 0.5, 1, 2\} \text{m/s}^2$ and lateral lane offsets ℓ_\perp in $\{-0.5, 0, 0.5\} \text{m}$. We then fit a cubic spline from the current state to the terminal state that minimizes mean control effort, and drop the dynamically infeasible trajectories given the control limits of the ego vehicle.

C.2 Control module details

The iLQR algorithm iteratively forms and solves a quadratic LQR approximation of the problem around the current solution $s^{(i)}, u^{(i)}$ for iteration i , using a second-order Taylor approximation of the cost function C , and a first-order Taylor approximation of the dynamics function f_d .

In each iLQR iteration we are solving an LQR problem. The solution of the LQR problem is a control sequence which we use as a candidate for a potential trajectory update. Note that the LQR-optimal control sequence is only optimal according to the approximated quadratic cost, and does not necessarily improve the cost for the original non-quadratic problem. Therefore we perform a line-search for the largest change in control sequence in the direction of the LQR-optimal controls that results in a reduction in the original cost. We start with a down-scaling factor of 1 for the change in controls, unroll controls through the dynamics function, and evaluate the original non-quadratic cost for the resulting trajectory. If this cost is larger than the cost of the unchanged trajectory, we decrease the down-scaling factor by a factor of 5 and repeat the search process up to 5 times. Otherwise we update the trajectory and move on to the next iLQR iteration, which makes a new quadratic approximation around the updated trajectory.

We continue iLQR iterations until convergence or up to 5 iterations. Convergence is defined by a threshold (0.05) on the change in control vector magnitude between iterations.

C.3 Cost function details

The cost function is a weighted sum of a set of handcrafted cost terms for penalizing collisions, distance to the goal, lateral lane deviation, lane heading deviation, and control effort:

$$C = w_1 C_{\text{coll}}(s, \hat{s}_{a \in A}) + w_2 C_g(s, g) + w_3 C_{\ell_\perp}(s, m) + w_4 C_{\ell_\angle}(s, m) + w_5 C_u(u).$$

The cost function terms are defined as followed:

- $C_{\text{coll}}(s, \hat{s}_{a \in A}) = \sum_{a \in A} \sum_{t \in 1:T} \varphi \left(\sum_{k \in K} \pi_k \|s^{(t)} - \hat{s}_{a,k}^{(t)}\|^2 \right)$, the collision term, where $\hat{s}_{a,k}$ is the k -th mode of the predicted trajectory distribution for agent a , π_k is the probability of the k -th mode, φ is a Gaussian radial basis function, and $\|\cdot\|$ is the Euclidean norm. The collision term encourages keeping distance from other agents in expectation.
- $C_g(s, g) = \|s^{(T)} - g\|^2$, the goal term, that captures the distance to goal at the end of the planning horizon. We choose the goal g to match the ‘‘ground-truth’’ ego trajectory in the dataset.
- $C_{\ell_\perp}(s, m) = \sum_{t \in 1:T} \|s^t - s_{\ell_\perp}\|^2$, the lateral lane deviation term, where s_{ℓ_\perp} is the position on the lane closest to the ego position s .
- $C_{\ell_\angle}(s, m) = \sum_{t \in t:T} (h_e^t - h_\ell)^2$, the lane heading deviation term, where h_e and h_ℓ denote ego vehicle heading and lane heading, respectively.
- $C_u(u) = \sum_{t \in t:T} \left((u_{\text{steer}}^{(t)})^2 + (u_{\text{acc}}^{(t)})^2 \right)$, the control effort term, where u_{steer} denotes the steering control and u_{acc} the acceleration control.

The default weight terms are chosen manually as follows: $w_1=5.0$; $w_2=0.5$; $w_3=0.3$; $w_4=0.3$; $w_5=1.0$.

D Experimental setup details

D.1 Dataset details

We use the nuScenes dataset which contains a large amounts of driving data, 1000 scenes with annotated states for vehicles and pedestrians, and a map containing lane information. The data is collected in Boston and Singapore. Scenes are 20s long, and trajectories are recorded at 2Hz ($\Delta t=0.5s$). We process the dataset into train and test scenarios. Each scenario has $T_H = 4s$ of state history and $T = 3s$ of GT future states for all agents. We choose one vehicle in the scene to act as the ego vehicle, and one vehicle to perform prediction for, such that ego and the predicted vehicle are close. For simplicity we use the GT futures for other agents for the purpose of planning. We train models with 75% of the training split, and use the remaining 15% as a validation set. We skip data with insufficient annotations to form a scenario, and where our planner has less than 2 trajectory candidates.

D.2 Training details

We train all models for 20 epochs using the Adam optimizer with gradient clipping, batch size of 256, learning rate of 0.003 by default. We have not performed extensive search over these training parameters. When training for multiple objectives, $\mathcal{L} = \alpha_1 \mathcal{L}_{\text{pred}} + \alpha_2 \mathcal{L}_{\text{plan}} + \alpha_3 \mathcal{L}_{\text{ctrl}}$, we use the following default loss scaling parameters: $\alpha_1=1$, $\alpha_2=100$, $\alpha_3=1000$ for reinforcement learning experiments, and $\alpha_1=1$, $\alpha_2=10$, $\alpha_3=1000$ for imitation learning experiments. We choose the default values such that the average magnitude of the resulting gradients are similar, followed by a simple hyperparameter search on the validation set with a different, fixed random seed. We explore the effect of the α_i loss scaling parameters in an ablation study presented in Table 6.

D.3 Closed-loop simulation details

In our log-replay simulation setup ego states are unrolled based on the planned control outputs, while non-ego agents follow their fixed trajectories recorded in the dataset, and replan every 0.5s. We create $T_{\text{sim}} = 10s$ long evaluation scenarios, one for each scene in our validation set. In each step of the simulation we run the stack to get control commands, update the ego state based on the control command and the known dynamics, and update non-ego agent states according to their logged trajectories in the dataset. The goal and lane inputs for planning are updated in each simulation step based on the logged ego trajectory. For simplicity in the stack we only perform prediction for the most relevant non-ego agent, and ignore all other agents during planning and control. We consider all agents when computing evaluation metrics.