

# Hidden Biases of End-to-End Driving Models

Bernhard Jaeger Kashyap Chitta Andreas Geiger  
University of Tübingen Tübingen AI Center

{bernhard.jaeger, kashyap.chitta, a.geiger}@uni-tuebingen.de

## Abstract

*End-to-end driving systems have recently made rapid progress, in particular on CARLA. Independent of their major contribution, they introduce changes to minor system components. Consequently, the source of improvements is unclear. We identify two biases that recur in nearly all state-of-the-art methods and are critical for the observed progress on CARLA: (1) lateral recovery via a strong inductive bias towards target point following, and (2) longitudinal averaging of multimodal waypoint predictions for slowing down. We investigate the drawbacks of these biases and identify principled alternatives. By incorporating our insights, we develop TF++, a simple end-to-end method that ranks first on the Longest6 and LAV benchmarks, gaining 11 driving score over the best prior work on Longest6.*

## 1. Introduction

End-to-end driving approaches have rapidly improved in performance on the CARLA leaderboard [1], the de-facto standard for fair online evaluation. Driving scores have increased from under 20 [10, 28] to over 70 [35, 41] in just two years. However, why recent systems work so well is not fully understood, as both methods and training sets differ largely between submissions. Rigorous ablations are expensive due to the large design space of driving systems and the need to simulate large amounts of driving for evaluation.

In particular, recent methods trained with Imitation Learning (IL) have shown strong performance [7, 9, 12, 35, 41]. They are trained using offline datasets, yet they can surprisingly recover from the classic compounding error problem of IL [27, 33], as indicated by their high route completions [9, 12, 35, 41]. While they do not utilize HD maps as input, they are provided with map-based GNSS locations in the center of the lane (spaced 30 m apart on average) called target points (TPs) that describe the route the car should follow. TPs were introduced as an alternative form of conditioning signals to convey driver intent [10, 28]. Prior work [14, 15] used discrete navigation commands or NCs (i.e. follow lane, turn right, ...) instead.

In this paper, we show that TP conditioned models re-

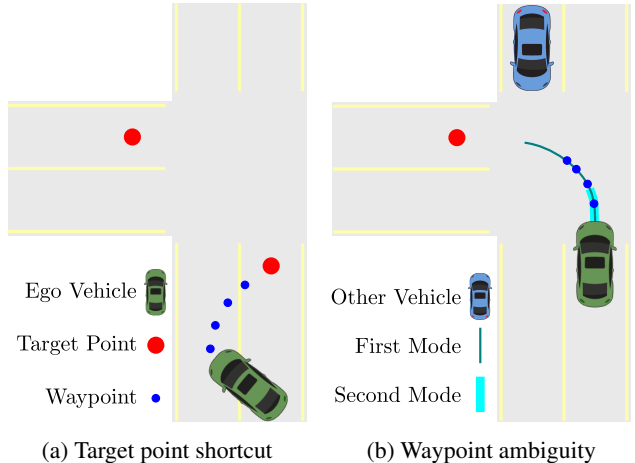


Figure 1: **Hidden biases.** (a) When outside their training distribution, current methods extrapolate **waypoint** predictions to the nearest **target point**, helping them recover. (b) The future velocity is **multi-modal**, but current methods commit to a single plan, which leads to interpolation.

cover from the compounding error problem because they use geometric information contained in the TP to reset steering errors periodically (at every TP). This makes them implicitly rely on accurate map information, even though they are otherwise HD map free. Steering directly towards a TP is a shortcut [17] that these IL methods learn to exploit. When methods accumulate enough steering error to be out of distribution during deployment, we observe that they steer towards the nearest TP. When the TP is close, this has the effect of driving back to the lane center, where it is in distribution again. This is illustrated in Fig. 1a. However, when the TP is far away, this shortcut can lead to catastrophic steering errors (e.g. cutting a turn). We show examples of this behavior for various SotA architectures in Section 3.1. We demonstrate that the shortcut problem is intrinsically related to the decoder architecture and that a transformer decoder [38] can mitigate it.

Another common aspect of the current SotA is that they use waypoints (future positions of an expert driver) as output representations [9, 12, 41]. We point out that this is an ambiguous representation as the future velocity is multi-

modal, yet the model commits to a point estimate. This is illustrated in Fig. 1b. We show that this ambiguity can sometimes be helpful due to the continuous nature of waypoints: the network can continuously interpolate between modes. We propose an alternative that explicitly predicts the uncertainty of the network via target speed classification, and show that interpolating between target speeds weighted by the uncertainty reduces collisions. Our controlled experiments also cover important but sometimes neglected details in the training of end-to-end driving systems, such as augmentation, training schedules, and dataset size. In particular, we revisit the idea of shift and rotation augmentations to aid recovery [2,6]. These were common in early IL methods for CARLA with control outputs [15], but are harder to implement with waypoint outputs and not used by the current SotA. We find that they yield significant improvements.

Using these insights, we develop TransFuser++ (TF++), which sets a new SotA on the Longest6 [12] and LAV [9] benchmarks. Two of the ideas we apply, transformer decoder pooling and path-based outputs instead of waypoints, have been used in Interfuser [35]. However, these are presented as minor details, and their impact is not studied in isolation as in our work. TF++ is significantly simpler than Interfuser, yet outperforms it by a large margin on Longest6, as we show in Section 4. We use  $\sim 4\times$  less data, 1 camera instead of 4, and do not require complex heuristics to extract throttle and brake commands for our path output.

#### Contributions:

- We show that target point conditioned models learn a shortcut that helps them recover from steering errors.
- We point out that the waypoint output representation is ambiguous, but its continuous nature helps models collide less by interpolating to slow down.
- Using the insights gained from controlled experiments, we propose TransFuser++ which places first on the Longest6 and LAV benchmarks.

Code, data, and models are available at <https://github.com/autonomousvision/carla-garage>.

## 2. Related Work

IL for autonomous driving dates back over 30 years [26]. It regained traction with the seminal works of [4, 6, 14] and the release of CARLA [16], a 3D simulator used in numerous recent research advances in autonomous driving [10, 11, 15, 19, 25, 34, 37, 43]. Early IL approaches evaluated in CARLA used a discrete navigation command (NC) [8, 10, 15], but their performance is not competitive to modern approaches [9, 12, 35, 41] which predict waypoints conditioned on TPs. In this work, we revisit an idea used in early systems but neglected in modern TP-conditioned methods: geometric shift augmentations [6, 10, 15, 27].

LAV [9] supervises waypoint outputs with additional data by making predictions for other nearby agents in the

scene. Their waypoints are initially generated with NC conditioning. These are then refined using a GRU [13]. They observe a large (+50) improvement in route completion from this refinement. Our findings suggest that the refinement module improves steering primarily through TP conditioning, which is only provided to the refinement GRU. TCP [41] observes that waypoints are stronger at collision avoidance than directly predicting controls, but sometimes suffer at large turns. They leverage the strength and mitigate the weaknesses of these two representations by creating a situation dependent ensemble. Our study suggests that the waypoints are better at collision avoidance due to an implicit slowdown when the car is uncertain. PlanT [31] investigates planning on CARLA by processing object bounding boxes with a transformer. Its sensorimotor version differs from other SotA models in that it is trained in two stages (not end-to-end). In our setting, we find that end-to-end training is crucial. We show that their observations about dataset scale also hold for end-to-end models.

**TransFuser** [12, 28] is a simple, well-known and widely used baseline for CARLA. We provide an explanation for the large differences in route completion between the target point conditioned TransFuser architecture and its NC conditioned baselines [8, 10, 15]. Further, we propose modifications to its architecture, output representation and training strategy which lead to significant improvements.

**Interfuser** [35] regresses a path for steering, predicts object density maps, and classifies traffic rule flags. This representation is converted by a forecasting mechanism and hand designed heuristics into control. Like our proposed method, it uses a transformer decoder for pooling features, and disentangles future velocities from the path in the output. While these ideas were already present, they were not studied or discussed as significant. Our work adds to the literature by showing that these design choices are indeed critical to performance and providing explanations as to why. Furthermore, our final system is simpler and significantly outperforms Interfuser on Longest6.

**ReasonNet** [36], **ThinkTwice** [22] and **CaT** [42] are concurrent end-to-end driving approaches that all use TP conditioning and waypoints as output representation.

## 3. Hidden Biases of End-to-End Driving

We consider the task of urban navigation from point A to B [12]. The goal is to complete routes with dense traffic, multiple lanes and complex geometries (e.g. roundabouts) without incurring infractions. Along the way, the agent encounters manually designed pre-crash traffic scenarios. Each route is a list of GNSS coordinates called target points (TPs) which can be up to 50 m apart ( $\sim 30$ m on average).

**Metrics:** We use the CARLA online metrics. Route Completion (RC) is the percentage of the route completed. Infraction Score (IS) is a penalty factor starting at 1.0 that

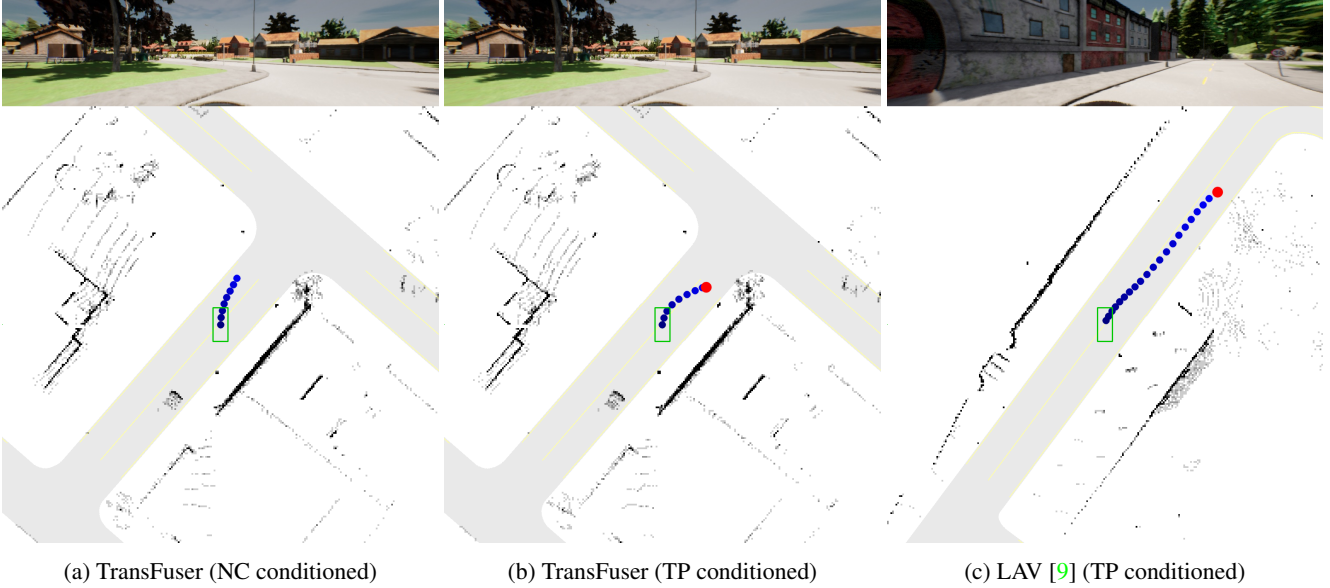


Figure 2: **Extrapolation to target point.** In unknown situations, TP conditioned methods extrapolate their **waypoints** towards **target points**. This periodically resets steering errors and is a form of implicit map based recovery. However, relying on extrapolation is a shortcut that can lead to catastrophic errors in certain situations (see Fig. 3a and Fig. 3b).

gets reduced multiplicatively for every infraction. Our main metric is the Driving Score (DS) which multiplies RC with the IS. Where insightful, we report infraction per kilometer metrics. For a comprehensive description, refer to [12].

**Baseline:** As a baseline, we reproduce TransFuser [12]. This is a simple method representative of current SotA driving systems on CARLA. The main differences in our reproduction are a 360° field of view (FOV) LiDAR to enable safe lane changes and a dataset where the expert is driving 2× faster to speed up evaluation. A list of other minor changes is provided in the supplementary material.

**Benchmark:** For all experiments in this section, we train on CARLA towns 01, 03, 04, 06, 07 and 10. We use the 16 validation routes from LAV [9] in the withheld towns 02 and 05 for evaluation. In order to reduce the influence of training and evaluation variance which can be large in CARLA [5, 12], every experiment reports the average of 3 training seeds, evaluated 3 times each. We additionally report the training standard deviation for the main metrics.

### 3.1. A shortcut for recovery

While current SotA CARLA methods have fundamentally different architectures [9, 12, 35, 41], they are all trained with conditional IL using fixed pre-recorded datasets. The evaluation task involves challenging routes which are ~1.5km long, so it would be expected that these methods suffer from *compounding errors*, a well-known problem for IL [33]. Geometric shift augmentations [6] are a common approach to teach IL methods how to recover from such compounding steering errors [6, 15]. Surprisingly,

even though such augmentations are not employed by SotA methods on CARLA, they report high RCs, demonstrating that they are not prone to this expected failure mode.

All the aforementioned methods condition their predictions using the next target point (TP) along the route. To understand the importance of the TP, we train 2 versions of our reproduction of TransFuser: one with the original TP conditioning and another with a discrete (NC) condition. This is a one hot vector that indicates whether the car should follow the lane, turn right at the next intersection, etc. It contains no geometric information about the center of the lanes, but still removes the inherent task ambiguity of inner-city driving [14]. For the NC conditioned model, we also implement shift and rotation augmentations by collecting augmented frames during data collection. We deploy a second camera in the simulator, that changes its position and orientation randomly at every time-step. The output labels are transformed accordingly for training the model (implementation details and examples can be found in the supplementary material). The results are shown in Table 1.

Cond.	Aug.	DS ↑	RC ↑	Dev ↓
NC	-	32 ± 8	56 ± 12	0.86
NC	✓	35 ± 3	54 ± 4	0.99
TP	-	<b>39 ± 9</b>	<b>84 ± 7</b>	<b>0.00</b>

Table 1: **Conditioning and Augmentation.**

We observe a significant difference in RC of 28 points when switching between TP and NC conditioning. The TP conditioned model has 0 route deviations per km (“Dev”),

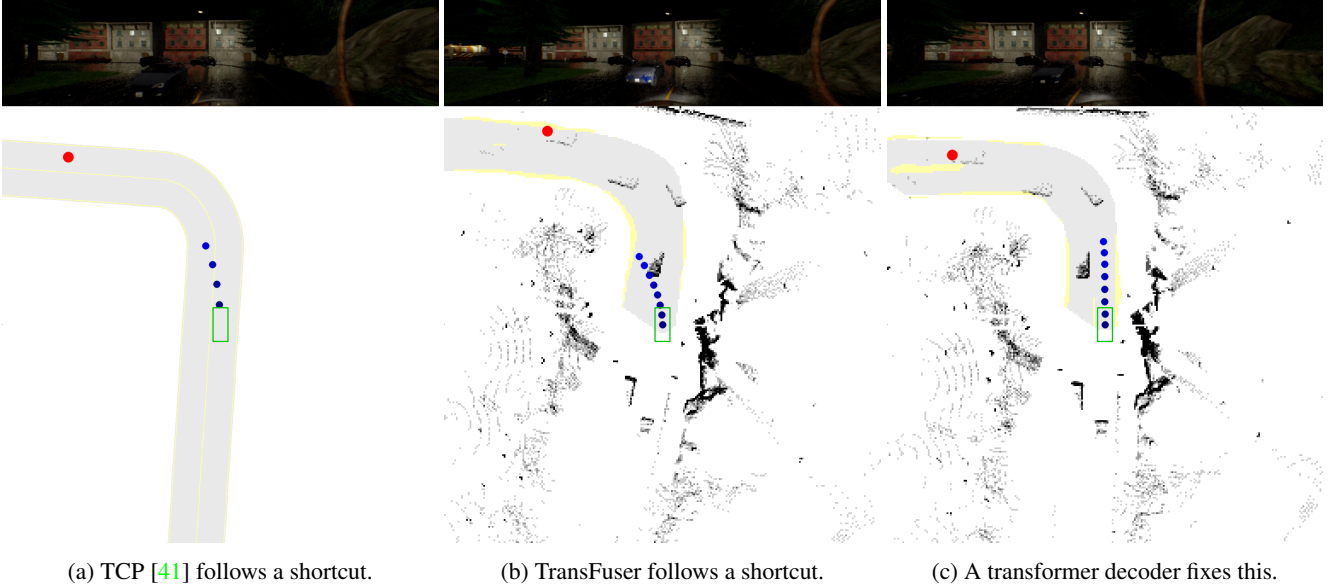


Figure 3: **Target point shortcut.** When TP conditioned methods extrapolate to spatially distant waypoints, they incur large steering errors. Replacing global average pooling in TransFuser with a cross-attention mechanism mitigates the issue.

which indicates that it never takes a wrong turn or drives too far away (more than 30m) from the lane. However, this is not the case for NC conditioned models. While we do see a small improvement in DS with augmentation, its RC is still unsatisfactory. This indicates that the geometric information regarding the lane center available with the TP aids recovery in SotA IL methods, prompting further investigation. We inspect the TransFuser and LAV [9] models by constructing situations where the car is forcefully steered out of lane. Fig. 2 visualizes these scenarios with the camera, LiDAR and ground truth HD maps. For TP conditioned models, the predicted waypoints (shown in blue) extrapolate towards the nearest TP (red) even though the situation is out of distribution. This shows that one reason for their strong route following is that the bird’s eye view (BEV) TP resets steering errors periodically: methods learn to steer towards nearby TPs because the expert trajectory in the dataset always goes through them. This has the effect that TP conditioned methods periodically drive back to the center of the lane, resetting any accumulated errors. Furthermore, it suggests that SotA IL approaches strongly rely on pre-specified geometric information regarding lane centers for recovery.

We identify this as a form of shortcut learning, which can be useful when the car is close to a target point, but can also lead to catastrophic steering errors when the target point is further away. An example would be directly extrapolating to a target point behind a turn, which leads to cutting the turn. In Fig. 3a and Fig. 3b, we show instances where this happens for TCP [41] and TransFuser [12] in a validation town at nighttime. TransFuser also predicts the BEV segmentation as an auxiliary task, which we overlay

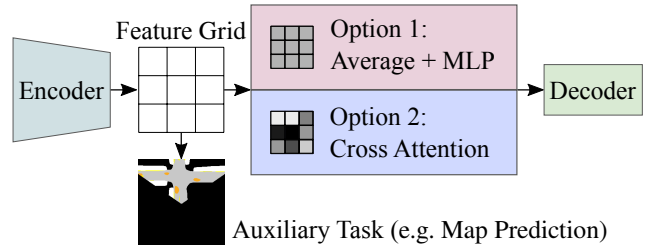


Figure 4: **Pooling.** Existing approaches vectorize feature grids either by global average pooling (top, e.g. [9, 12]) or with attention mechanisms (bottom, e.g. [35, 41]). The latter retains spatial information gained via auxiliary tasks.

in the figure (gray: route, yellow: lane marking). For TCP, we render the ground truth map and omit the LiDAR, since it does not use LiDAR and does not predict BEV segmentation. Both methods predict waypoints that are tilted towards the TP instead of following the street. As a result, they drive into the opposing lane. We refer to this as the *TP shortcut*.

### 3.2. Improved pooling and data augmentation

One design choice which differs across SotA architectures is the pooling applied between the encoder and decoder. In Fig. 4, we summarize these. In particular, TransFuser [12] and LAV [9] employ global average pooling (GAP) followed by an MLP. InterFuser [35] pools features via the cross-attention mechanism of a transformer decoder. Finally, TCP [41], which has two decoders, uses GAP for one decoder and attention-based pooling for the other.

As shown in Fig. 4, the networks are encouraged to learn spatially meaningful feature grids through auxiliary convo-



lutional decoders that predict outputs such as BEV semantic segmentation. However, GAP does not maintain the spatial information in the features. Fig. 3b shows the BEV road segmentation predicted by TransFuser overlaid by its LiDAR input. Unlike the waypoints, the BEV predictions are quite accurate. This indicates strong features coming from the encoder. Nevertheless, the final conditional waypoint predictions focus on the TP signal in this situation. We hypothesize that the GAP operation makes it difficult for the downstream decoder to utilize the strong BEV features. Note that attention-based pooling, such as the transformer decoder of Interfuser, preserves spatial information through the use of positional encodings.

We compare the original TransFuser GAP approach with the spatially preserving design of Interfuser in Table 2. For the latter variant, we remove the GAP operation and MLP in TransFuser. We then process the  $8 \times 8$  BEV feature grid as tokens with a transformer decoder. Implementation details can be found in the supplementary material.

Pooling	Aug.	DS $\uparrow$	RC $\uparrow$	Stat $\downarrow$
GAP + MLP	-	$39 \pm 9$	$84 \pm 7$	1.04
Transformer Decoder	-	$43 \pm 6$	<b><math>93 \pm 3</math></b>	0.55
Transformer Decoder	✓	<b><math>49 \pm 8</math></b>	$90 \pm 4$	<b>0.10</b>

Table 2: Pooling and augmentation.

After replacing GAP with the transformer decoder, the RC increases by 9 points and collisions per km with static objects (“Stat”) reduce by a factor of 2. Static objects (such as poles) only occur outside lanes in CARLA, implying that the network manages to stay in the lane far more consistently. Fig. 3c provides qualitative evidence that the transformer decoder can succeed in situations where GAP failed due to the TP shortcut (additional examples in supplementary). In addition, we investigate the inclusion of shift and rotation augmentations designed to aid lateral recovery (as described in Section 3.1). Collisions with the static environment are reduced further by a factor of 5 indicating that augmentation provides strong benefits.

### 3.3. The ambiguity of waypoints

Waypoints are used in many SotA systems as outputs [9, 12, 41]. They are obtained by recording an expert driver’s GNSS locations at fixed time intervals (e.g. 0.5s) and transforming them into a local coordinate frame. The model is then trained to predict future waypoints, typically with an  $L_1$  regression objective. Waypoints entangle both the path and the future velocities of the vehicle. The velocity after a specific time interval extracted from the waypoints is used by a downstream controller as a target speed.

In Fig. 5, we plot a histogram comparing the target speeds of the expert algorithm (i.e. training dataset) to those

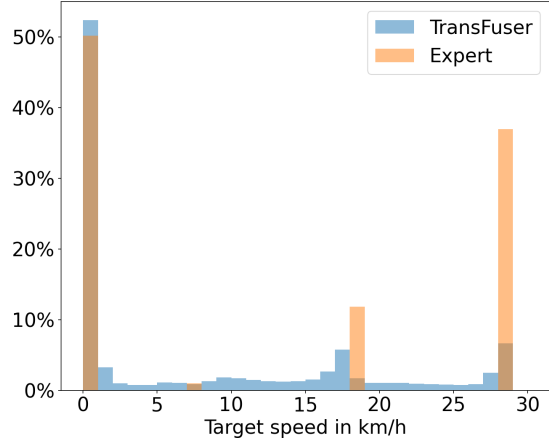


Figure 5: Transfuser interpolates between modes.

extracted from TransFuser (best model in Table 2). Interestingly, the distributions are quite different. By design, the expert chooses one of four target speeds: 29, 18, 7, or 0 km/h. These values cover behaviors needed for city driving, corresponding to four situations: regular driving, slowing down in intersections, slowing down near pedestrians, and stopping. In contrast, for TransFuser, the predicted target speeds cover the entire 0-29 km/h range.

While the future path that our expert follows is deterministic and unambiguous (center of the lane, lane change at pre-defined locations), future velocities are multi-modal. As both are jointly represented by waypoints, this leads to an ambiguous entangled representation. However, existing methods which utilize this entangled waypoint representation predict only point estimates (i.e., a single set of waypoints) as output, hence only a single mode is modeled. From Fig. 5, we observe that the waypoint based TransFuser model indeed interpolates between modes, a behavior that is expected when modeling multi-modal outputs deterministically. We illustrate this behavior in detail using the examples of (1) approaching an intersection with a green light (Fig. 6a), and (2) a cyclist cutting into the vehicle’s path (Fig. 6b). TransFuser slows down in these situations since it is uncertain, e.g., the light may turn red in Fig. 6a. The car stops in time in Fig. 6b because of the decreased speed.

While this averaging is indeed beneficial in some situations, an entangled representation is undesirable as it is less interpretable and does not explicitly expose uncertainty. Moreover, when stopped (all waypoints collapsed to one location), the steering signal is undefined which requires additional heuristics in the controller. To resolve this, we disentangle the future velocities from the path by sampling the expert’s position at fixed distances instead of fixed time intervals for training a (deterministic) *path predictor*. However, when predicting the path instead of time-dependent trajectories, one requires an additional method to determine a target speed for the car. We propose to predict the tar-

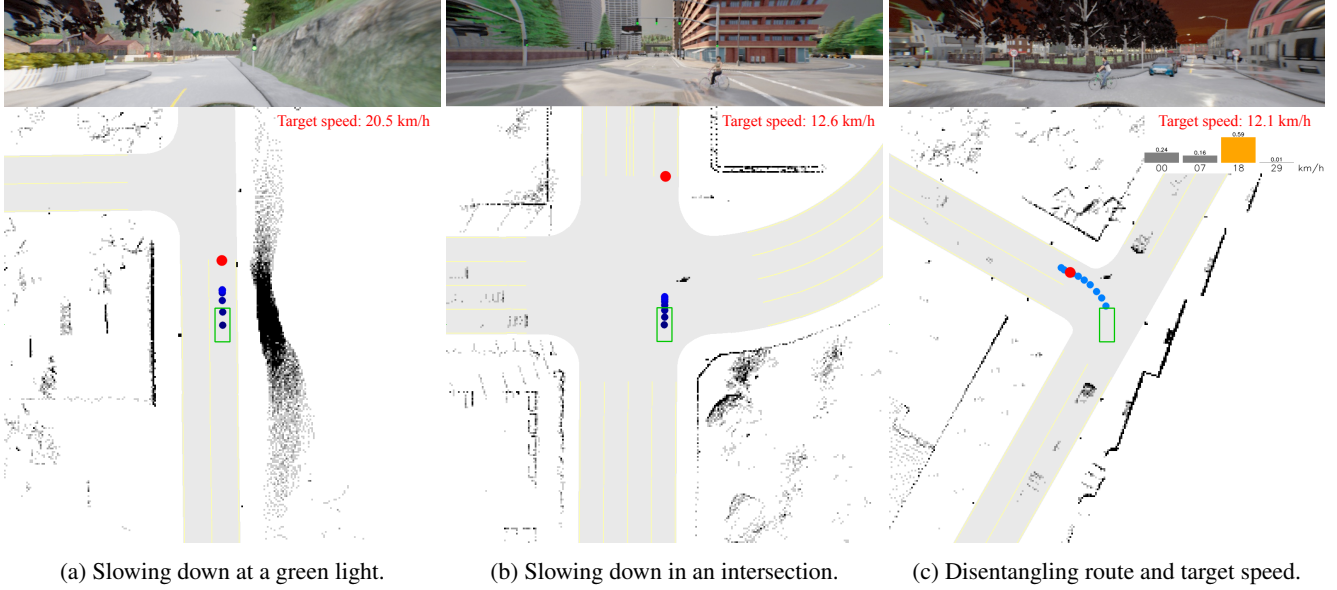


Figure 6: **Waypoints are ambiguous.** The model’s output representation forces it to predict a single mode for future velocities. There is a possibility that the traffic light might turn red in the future, or that the cyclist either cuts in or yields to the agent. The particular mode (or interpolation thereof) that the model converges to, depends on the training run and dataset.

get speed by simple classification using an additional MLP head. We incorporate the prediction uncertainties into the final output by using a confidence weighted average of the predicted target speed as input to the controller. Note that this is just one example of how the resulting confidence estimates may be leveraged by a vehicle controller.

Output	DS $\uparrow$	RC $\uparrow$	Veh $\downarrow$	Stat $\downarrow$
Waypoints	$49 \pm 8$	$90 \pm 4$	<b>0.70</b>	0.10
Path + Argmax	$40 \pm 1$	$88 \pm 2$	1.25	0.03
Path + Weighted	<b><math>50 \pm 3</math></b>	$88 \pm 1$	0.83	<b>0.02</b>

Table 3: **Output representation.**

Compared to the waypoint representation, naively converting the predicted classes to a target speed via the most likely class (Argmax) increases vehicle collisions (“Veh”), as seen in Table 3. Employing the proposed weighting (Weighted) achieves lower collisions. Moreover, the path-based model steers better as indicated by the lower static obstacle collisions (“Stat”). The disentangled representation achieves the same driving score as the entangled waypoint representation, providing an alternative with a simpler and more interpretable controller: (1) it eases design since it has identical parameters to the controller in the expert, (2) it has access to an unambiguous path representation even when driving slowly, and (3) it explicitly exposes and makes use of uncertainties with regard to target speeds.

Fig. 6c shows a qualitative example. The car is slowing down due to its uncertainty. The lower target speed allows it successfully merge behind the cyclist.

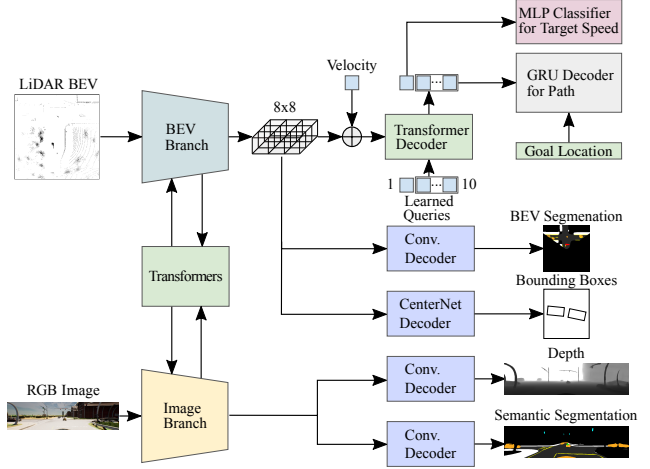


Figure 7: **TransFuser++ architecture.**

### 3.4. Scaling up to TransFuser++

By incorporating these insights, we obtain a significantly improved version of TransFuser [12] that we call TransFuser++ (Fig. 7). We now describe two important implementation changes compared to [12] that involve scaling.

Two stage	Frozen	DS $\uparrow$	RC $\uparrow$	Veh $\downarrow$
-	-	$50 \pm 3$	$88 \pm 1$	0.83
✓	-	<b><math>54 \pm 1</math></b>	$92 \pm 5$	<b>0.79</b>
✓	✓	$44 \pm 6$	<b><math>97 \pm 1</math></b>	1.21

Table 4: **Two stage training.**

**Training schedule:** We observe no benefits from simply doubling the number of training epochs for a single training stage. However, in Table 4, we double the training time with a two-stage approach. First, we pre-train the encoder with only the perception losses (2D depth, 2D and BEV semantic segmentation, and vehicle bounding boxes, as shown in Fig. 7) for the usual number of epochs. We then fine-tune the resulting checkpoint with all losses, after including the GRU decoder and MLP classifier. Initializing the backbone with features that are pre-trained on the auxiliary tasks leads to a 4 DS improvement, consistent with the claims of LAV [9]. We also experiment with freezing the pre-trained backbone and only training the transformer decoder and its heads in the second stage. This leads to a drop of 10 DS, indicating that end-to-end optimization is important.

**Dataset scale:** Recent work on CARLA [31] shows large improvements by scaling the dataset size. However, [31] considers planning models with privileged inputs on training towns. We investigate the impact of scaling when evaluating on held-out validation towns for end-to-end models. We start with 185k training samples (as in [12]) and scale it up by re-running the training routes 3 times with different traffic (555k frames). Table 5 shows a clear improvement of 6 DS via scaling. This shows that current IL approaches can still benefit from larger datasets. We train for the same amount of epochs, which implies that this improvement comes at the cost of  $3\times$  longer training.

Dataset size	DS $\uparrow$	RC $\uparrow$	Veh $\downarrow$
185k	54 $\pm$ 1	92 $\pm$ 5	0.79
555k	<b>60</b> $\pm$ 6	<b>98</b> $\pm$ 1	<b>0.73</b>

Table 5: Effects of scale.

## 4. Comparison to State of the Art

**Methods:** This section demonstrates the advantages of TF++ over several state-of-the-art models, which we list below, starting with the most recent. (1) **Interfuser** [35] processes multiple cameras and a BEV LiDAR by encoding each of them individually with a CNN. The resulting feature grids are fed into a transformer and decoded into perception outputs and the path to follow. The network does not predict longitudinal controls. Instead, it uses the BEV bounding boxes and a simple motion forecast (extrapolating historical dynamics) and linearly optimizes for the target speed using heuristically chosen objectives. (2) **Perception PlanT** [31] is a transformer based method trained in two independent stages that uses BEV bounding boxes as an intermediate representation. Its perception is based on TransFuser, and its planner outputs waypoints. (3) **TCP** [41] is a camera-only model with two output representations: waypoints and controls. During test time, it ensembles the two

outputs together with a weighted average that changes based on whether the vehicle is turning. (4) **TransFuser** [12] fuses perspective cameras with a BEV LiDAR by processing them with individual CNNs and exchanging features using transformers. It uses global average pooling and waypoint outputs, as described in Section 3. (5) **LAV** [9] processes cameras and LiDAR point clouds into an intermediate BEV representation. It trains a planner on this representation that predicts waypoints for the ego vehicle with a NC conditioned GRU followed by a TP conditioned refinement GRU. During training, the planner is also tasked to predict the trajectory of other surrounding vehicles to increase the number of labels. The authors release two versions of this method, which we call LAV v1 and LAV v2. (6) **WOR** [8] is an IL method for which the labels are enriched using a reward function to provide dense supervision for all possible actions. Unlike the other baselines, it is conditioned with the NC and does not utilize the TP. WOR is the best NC conditioned baseline that is publicly available.

**Benchmarks:** We use two benchmarks to evaluate on seen and unseen towns, with scenarios taken from [12] (type 1,3,4,7,8,9 and 10). Table 6 compares the performance of systems on the Longest6 benchmark [12] which consists of 36 routes in training towns 01 to 06. On Longest6, models trained on any data can be evaluated, so we compare against the author-provided models or directly report numbers from the respective papers. The mean and std of three evaluations is reported. For TF++ we train 3 models and report the average result. Longest6 has the advantage that it evaluates in dense traffic and has diverse towns and routes. The drawbacks are that it does not penalize stop sign infractions or measure generalization to new towns. Therefore, we additionally compare performance on validation towns using the LAV [9] routes in Table 7. These are 4 routes with 4 weathers (16 combined) in Town 02 and 05, which are withheld during training. We re-train methods with the dataset from the corresponding paper 3 times and evaluate each seed 3 times. Reported results are the mean of all runs and the std between the training seeds. We compare against the reported SotA TCP and our reproduced TransFuser baseline on this benchmark and provide additional baselines in the supplementary.

**Dataset:** For TF++ we collect data on the same training routes as in [12] with an improved expert labeling algorithm (described in the supplementary material). We repeat this 3 times on the same routes with different traffic, as in [31]. For validation, we withhold Town 02 and 05 during training, else we train on all towns. In total, we train with 750k frames (550k when withholding Town 02 and 05).

**Results:** We start with the results in training towns (Table 6). The best NC conditioned method, WOR, has significantly lower RC than all TP conditioned systems (22 lower than LAV v1). In particular, the route deviations (Dev) are

Method	DS $\uparrow$	RC $\uparrow$	IS $\uparrow$	Ped $\downarrow$	Veh $\downarrow$	Stat $\downarrow$	Red $\downarrow$	Dev $\downarrow$	TO $\downarrow$	Block $\downarrow$
WOR [8]	21 $\pm$ 3	48 $\pm$ 4	0.56 $\pm$ 0.03	0.18	1.05	0.37	1.28	0.88	0.08	0.20
LAV v1 [9]	33 $\pm$ 1	70 $\pm$ 3	0.51 $\pm$ 0.02	0.16	0.83	0.15	0.96	0.06	0.12	0.45
Interfuser [35]	47 $\pm$ 6	74 $\pm$ 1	0.63 $\pm$ 0.07	0.06	1.14	0.11	0.24	<b>0.00</b>	0.52	<b>0.06</b>
TransFuser [12]	47 $\pm$ 6	93 $\pm$ 1	0.50 $\pm$ 0.06	0.03	2.45	0.07	0.16	<b>0.00</b>	<b>0.06</b>	0.10
TCP [41]	48 $\pm$ 3	72 $\pm$ 3	0.65 $\pm$ 0.04	0.04	1.08	0.23	0.14	0.02	0.18	0.35
LAV v2 [9]	58 $\pm$ 1	83 $\pm$ 1	0.68 $\pm$ 0.02	<b>0.00</b>	<b>0.69</b>	0.15	0.23	0.08	0.32	0.11
Perception PlanT [31]	58 $\pm$ 5	88 $\pm$ 1	0.65 $\pm$ 0.06	0.07	0.97	0.11	0.09	<b>0.00</b>	0.13	0.13
TF++ (ours)	<b>69</b> $\pm$ 0	<b>94</b> $\pm$ 2	<b>0.72</b> $\pm$ 0.01	<b>0.00</b>	0.83	<b>0.01</b>	<b>0.05</b>	<b>0.00</b>	0.07	<b>0.06</b>
Expert	81 $\pm$ 3	90 $\pm$ 1	0.91 $\pm$ 0.04	0.01	0.21	0.00	0.01	0.00	0.07	0.09

Table 6: **Performance on training towns (Longest6).** Released models, std over 3 evaluations.

Method	DS $\uparrow$	RC $\uparrow$	IS $\uparrow$	Ped $\downarrow$	Veh $\downarrow$	Stat $\downarrow$	Red $\downarrow$	Stop $\downarrow$	Dev $\downarrow$	TO $\downarrow$	Block $\downarrow$
TransFuser (ours)	39 $\pm$ 9	84 $\pm$ 7	0.46 $\pm$ 0.06	<b>0.00</b>	0.74	1.04	0.20	1.07	0.00	0.23	0.21
TCP [41]	58 $\pm$ 5	85 $\pm$ 3	0.67 $\pm$ 0.06	<b>0.00</b>	<b>0.35</b>	0.16	<b>0.01</b>	1.05	0.00	0.19	0.19
TF++ (ours)	<b>70</b> $\pm$ 6	<b>99</b> $\pm$ 0	<b>0.70</b> $\pm$ 0.06	0.01	0.63	<b>0.01</b>	0.04	<b>0.26</b>	<b>0.00</b>	<b>0.05</b>	<b>0.00</b>
Expert	94	95	0.99	0.00	0.02	0.00	0.02	0.00	0.00	0.00	0.08

Table 7: **Performance on validation towns (LAV).** Reproduced models, std over 3 trainings (for 3 evaluations each).

10 times higher. Current TP conditioned SotA methods achieve (close to) 0 route deviations with driving scores in the range of 47 DS to 58 DS. TF++ outperforms all baselines on Longest6, showing a 19% relative improvement over the previous SotA Perception PlanT. TF++ achieves close to expert-level performance on all infractions except for vehicle collisions (Veh). On the validation towns (Table 7) TF++ outperforms our reproduced TransFuser by 31 DS improving all metrics, particularly environmental collisions (Stat). It improves the prior SotA result TCP by 21%.

**Runtime:** Table 8 compares the runtime of TF++ and TransFuser which have the same sensor setup. TF++ yields a large DS improvement while being only 14% slower.

Method $\uparrow$	DS $\uparrow$	Time (ms) $\downarrow$
TransFuser (ours)	39	44
TF++ (ours)	70	50

Table 8: **Runtime.** We show the runtime per frame in ms averaged over 300 time steps on a single route on a RTX 3090. TF++ outperforms TransFuser by a wide margin despite using a similar compute budget during inference.

## 5. Conclusion

In this work, we show that recent SotA driving models have exceptional route following abilities because they learn a strong bias towards following nearby TPs. Shortcut learning like this is a general phenomenon in deep neural networks [17] and has been observed in the context of autonomous driving for inputs such as velocity [15] or temporal frames [4, 40]. We add to this literature by observ-

ing shortcut learning with respect to the conditioning signal. While shortcut learning usually has a negative impact on performance, we observe positive (improved recovery) and negative (cutting turns) effects. We show that the negative effects can be mitigated by avoiding global pooling and incorporating data augmentation.

A second commonality in SotA approaches is the use of waypoints as an output representation. We observe that this representation is ambiguous because it predicts a point estimate for multi-modal future velocities. We disambiguate the representation by disentangling future velocities from the deterministic path predictions and classifying target speeds instead. We then weigh target speeds according to their predicted confidence in our controller. Surprisingly, we find that interpolation is helpful for reducing collisions.

We propose TransFuser++ by improving the popular baseline TransFuser with a series of controlled experiments. TF++ is a simple end-to-end method that sets a new state of the art on the LAV and Longest6 benchmarks.

**Limitations:** This study investigates urban driving in CARLA, where all investigated methods drive at relatively low speed ( $< 35\text{km/h}$ ). Therefore, problems specific to high-speed driving are not considered. In addition, lanes are free of static obstacles, hence scenarios requiring navigation around them are not included. We point out the strong reliance of current methods on TPs for recovery. This implies a reliance on accurate localization and mapping to obtain these TPs. They are accurately mapped in CARLA, however, this assumption might not hold in real environments.

**Broader Impact:** We aim to make progress towards autonomous driving. This technology, if realized, could have massive societal impact, reducing road accidents, trans-



portation costs and improving the mobility of elderly people. Potential negative implications include a reduction in jobs for human drivers and possible military applications.

**Acknowledgements:** Andreas Geiger and Bernhard Jaeger were supported by the ERC Starting Grant LEGO-3D (850533), the BMWi in the project KI Delta Learning (project number 19A19013O) and the DFG EXC number 2064/1 - project number 390727645. Kashyap Chitta was supported by the German Federal Ministry of Education and Research (BMBF): Tübingen AI Center, FKZ: 01IS18039A. We thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for supporting Bernhard Jaeger and Kashyap Chitta as well as Niklas Hanselmann and Katrin Renz for proofreading.

## References

- [1] Carla autonomous driving leaderboard. <https://leaderboard.carla.org/>, 2020. 1, 17
- [2] Alexander Amini, Tsun-Hsuan Wang, Igor Gilitschenski, Wilko Schwarting, Zhijian Liu, Song Han, Sertac Karaman, and Daniela Rus. Vista 2.0: An open, data-driven simulator for multimodal sensing and policy learning for autonomous vehicles. In *Proc. IEEE International Conf. on Robotics and Automation (ICRA)*, 2022. 2, 15
- [3] Lei Jimmy Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. Layer normalization. *arXiv.org*, 1607.06450, 2016. 14
- [4] Mayank Bansal, Alex Krizhevsky, and Abhijit S. Ogale. Chauffeurnet: Learning to drive by imitating the best and synthesizing the worst. In *Proc. Robotics: Science and Systems (RSS)*, 2019. 2, 8
- [5] Aseem Behl, Kashyap Chitta, Aditya Prakash, Eshed Ohn-Bar, and Andreas Geiger. Label efficient visual abstractions for autonomous driving. In *Proc. IEEE International Conf. on Intelligent Robots and Systems (IROS)*, 2020. 3
- [6] Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Prasoon Goyal, Lawrence D. Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, Xin Zhang, Jake Zhao, and Karol Zieba. End to end learning for self-driving cars. *arXiv.org*, 1604.07316, 2016. 2, 3
- [7] Sergio Casas, Abbas Sadat, and Raquel Urtasun. Mp3: A unified model to map, perceive, predict and plan. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1
- [8] Dian Chen, Vladlen Koltun, and Philipp Krähenbühl. Learning to drive from a world on rails. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, 2021. 2, 7, 8, 11, 13, 18
- [9] Dian Chen and Philipp Krähenbühl. Learning from all vehicles. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 2, 3, 4, 5, 7, 8, 13, 15, 16, 17
- [10] Dian Chen, Brady Zhou, Vladlen Koltun, and Philipp Krähenbühl. Learning by cheating. In *Proc. Conf. on Robot Learning (CoRL)*, 2019. 1, 2
- [11] Kashyap Chitta, Aditya Prakash, and Andreas Geiger. Neat: Neural attention fields for end-to-end autonomous driving. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, 2021. 2, 17
- [12] Kashyap Chitta, Aditya Prakash, Bernhard Jaeger, Zehao Yu, Katrin Renz, and Andreas Geiger. Transfuser: Imitation with transformer-based sensor fusion for autonomous driving. *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, 2022. 1, 2, 3, 4, 5, 6, 7, 8, 11, 12, 13, 15, 17
- [13] Kyunghyun Cho, Bart van Merriënboer, Çağlar Gülçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoder-decoder for statistical machine translation. In *Proc. of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2014. 2
- [14] Felipe Codevilla, Matthias Müller, Antonio López, Vladlen Koltun, and Alexey Dosovitskiy. End-to-end driving via conditional imitation learning. In *Proc. IEEE International Conf. on Robotics and Automation (ICRA)*, 2018. 1, 2, 3
- [15] Felipe Codevilla, Eder Santana, Antonio M. López, and Adrien Gaidon. Exploring the limitations of behavior cloning for autonomous driving. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, 2019. 1, 2, 3, 8
- [16] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. CARLA: An open urban driving simulator. In *Proc. Conf. on Robot Learning (CoRL)*, 2017. 2
- [17] Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard S. Zemel, Wieland Brendel, Matthias Bethge, and Felix A. Wichmann. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2020. 1, 8
- [18] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv.org*, 1606.08415, 2016. 14
- [19] Anthony Hu, Gianluca Corrado, Nicolas Griffiths, Zak Murez, Corina Gurau, Hudson Yeo, Alex Kendall, Roberto Cipolla, and Jamie Shotton. Model-based imitation learning for urban driving. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 2, 17
- [20] Martin Isenburg. Laszip: lossless compression of lidar data. *Photogrammetric engineering and remote sensing*, 2013. 11
- [21] Bernhard Jaeger. Expert drivers for autonomous driving. Master’s thesis, University of Tübingen, 2021. 18
- [22] Xiaosong Jia, Penghao Wu, Li Chen, Jiangwei Xie, Conghui He, Junchi Yan, and Hongyang Li. Think twice before driving: Towards scalable decoders for end-to-end autonomous driving. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2, 17
- [23] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *Proc. of the International Conf. on Learning Representations (ICLR)*, 2019. 12
- [24] Rudolph Van Der Merwe. *Sigma-point Kalman filters for probabilistic inference in dynamic state-space models*. PhD thesis, Oregon Health & Science University, 2004. 13
- [25] Eshed Ohn-Bar, Aditya Prakash, Aseem Behl, Kashyap Chitta, and Andreas Geiger. Learning situational driving. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2

- [26] Dean Pomerleau. ALVINN: an autonomous land vehicle in a neural network. In *Advances in Neural Information Processing Systems (NIPS)*, 1988. 2
- [27] Aditya Prakash, Aseem Behl, Eshed Ohn-Bar, Kashyap Chitta, and Andreas Geiger. Exploring data aggregation in policy learning for vision-based urban autonomous driving. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1, 2
- [28] Aditya Prakash, Kashyap Chitta, and Andreas Geiger. Multi-modal fusion transformer for end-to-end autonomous driving. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1, 2, 16, 17
- [29] Ilija Radosavovic, Raj Prateek Kosaraju, Ross B. Girshick, Kaiming He, and Piotr Dollár. Designing network design spaces. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2020. 13
- [30] Sashank J. Reddi, Satyen Kale, and Sanjiv Kumar. On the convergence of adam and beyond. In *Proc. of the International Conf. on Learning Representations (ICLR)*, 2018. 12
- [31] Katrin Renz, Kashyap Chitta, Otniel-Bogdan Mercea, Almut Sophia Koepke, Zeynep Akata, and Andreas Geiger. Plant: Explainable planning transformers via object-level representations. In *Proc. Conf. on Robot Learning (CoRL)*, 2022. 2, 7, 8, 15, 16, 17
- [32] Labbe Roger. Kalman and bayesian filters in python. 2020. 13
- [33] Stéphane Ross, Geoffrey J. Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Conference on Artificial Intelligence and Statistics (AISTATS)*, 2011. 1, 3
- [34] Axel Sauer, Nikolay Savinov, and Andreas Geiger. Conditional affordance learning for driving in urban environments. In *Proc. Conf. on Robot Learning (CoRL)*, 2018. 2
- [35] Hao Shao, Letian Wang, RuoBing Chen, Hongsheng Li, and Yu Liu. Safety-enhanced autonomous driving using interpretable sensor fusion transformer. In *Proc. Conf. on Robot Learning (CoRL)*, 2022. 1, 2, 3, 4, 7, 8, 14, 17
- [36] Hao Shao, Letian Wang, Ruobing Chen, Steven L. Waslander, Hongsheng Li, and Yu Liu. Reasonnet: End-to-end driving with temporal and global reasoning. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2
- [37] Marin Toromanoff, Emilie Wirbel, and Fabien Moutarde. End-to-end model-free reinforcement learning for urban driving using implicit affordances. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [38] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems (NIPS)*, pages 5998–6008, 2017. 1, 14
- [39] Eric A Wan and Rudolph Van Der Merwe. The unscented kalman filter for nonlinear estimation. In *Proc. IEEE Adaptive Systems for Signal Processing, Communications, and Control Symposium*, 2000. 13
- [40] Chuan Wen, Jierui Lin, Trevor Darrell, Dinesh Jayaraman, and Yang Gao. Fighting copycat agents in behavioral cloning from observation histories. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 8
- [41] Peng Wu, Xiaosong Jia, Li Chen, Junchi Yan, Hongyang Li, and Yu Qiao. Trajectory-guided control prediction for end-to-end autonomous driving: A simple yet strong baseline. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 1, 2, 3, 4, 5, 7, 8, 12, 13, 16, 17
- [42] Jimuyang Zhang, Zanming Huang, and Eshed Ohn-Bar. Coaching a teachable student. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2, 17
- [43] Zhejun Zhang, Alexander Liniger, Dengxin Dai, Fisher Yu, and Luc Van Gool. End-to-end urban driving by imitating a reinforcement learning coach. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, 2021. 2, 11, 12, 17, 18

# Appendix

## A. Changes to TransFuser

Our reproduced TransFuser largely follows [12] but has some minor differences in implementation details described in this section.

### A.1. Expert

For data collection, we follow the common practice of using an automatic labeling algorithm (expert driver) to generate the imitation labels. The auxiliary perception labels are provided directly by the CARLA simulator except for the bird’s eye view (BEV) segmentation, for which we follow [43] and render the relevant objects into an HD map. To generate the imitation labels, the expert driver needs to solve planning and control, but can bypass perception using privileged access to the simulator. We build our expert upon the model predictive control (MPC) approach of [12]. Lateral control is done by following the next point (at least 3.5 meters away) in a path, created by an A\* algorithm, with a PID controller. Longitudinal control is done via MPC that differentiates between 4 target speeds. For regular driving, we use 8 m/s (double the target speed compared to [12]). Inside intersections, we slow down to 5 m/s.

**Collision avoidance:** The target speed is set to 0 m/s when the MPC algorithm predicts a collision. We predict collisions similarly to [12]. They approximate all agents’ future positions by iteratively unrolling a kinematic bicycle model [8]. Actions for other cars are set to be the same as the current time step while the ego agent’s own action is approximated by using a PID controller to follow the A\* path. In case the ego agent’s bounding box overlaps with another agent’s bounding box at future time step  $t$ , a collision is predicted. To keep a safety distance to the leading vehicle, the MPC expert in [8] additionally has a static bounding box in front of the ego agent that predicts a collision if it intersects with any other agent. Because our agent has double the maximum driving speed, we need to keep a larger safety distance. Using a static area would not suffice here, as the static bounding box will become so large that it sticks into the opposing lane during turns, causing unnecessary braking. Instead, we approximate the area where the expert would end up if it performed a full brake (after 1 meter of driving) at his current speed  $s$ . The distance to stop  $d$  is approximated with:

$$d = 0.5 * \left(\frac{s * 3.6}{10.0}\right)^2 + 2.5 \quad (1)$$

We again use the kinematic bicycle model to compute where the expert will be after this distance. If the bounding box of the expert at that time step intersects with any bounding boxes for the current time step, we set the target speed

to 0. This approach has the advantage that the safety area follows the road and increases with increasing agent speed. A second problem coming from the higher driving speed is scenario 3 where a pedestrian runs in front of the vehicle. The expert will stop as soon as the pedestrian attempts to move forward, but when driving at 8 m/s this is already too late to prevent the collision. To solve this problem, we preemptively slow down to 2 m/s whenever a pedestrian is within a 30-meter radius in front of the vehicle.

**Traffic rules:** Stop signs infractions are addressed by slowing down to 2 m/s when the safety area intersects with the stop sign trigger area (provided by the simulator) and stopping once the car is on the stop sign trigger area. Red lights are resolved by setting the target speed to 0 when the safety area or vehicle intersects with the entrance of an intersection and the corresponding traffic light is active.

Method	DS $\uparrow$	RC $\uparrow$	Veh $\downarrow$	Block $\downarrow$
MPC Expert [12]	77 $\pm$ 2	89 $\pm$ 1	0.28	0.13
MPC Expert (ours)	<b>81 <math>\pm</math> 3</b>	<b>90 <math>\pm</math> 1</b>	<b>0.21</b>	<b>0.09</b>

Table 9: **Expert comparison on Longest6**

**Performance:** We show in Table 9 that our expert outperforms the baseline by 4 DS. This represents a higher upper bound for our imitation learning models. However, since no method can currently reach that upper bound, we see no direct improvement from the better expert (see Table 10).

### A.2. Dataset

We generate our training dataset by executing the expert described in Section A.1 on the training routes from [12] and storing every frame (20 FPS). Weathers are randomized per route instead of per frame to avoid exposure problems. During training, we train on every fifth frame, leading to an effective FPS of 4. For the scaling experiment, we re-run data collection on each route 3 times. This randomizes weathers and traffic (the traffic manager in CARLA 0.9.10 is not deterministic) but the environment is the same.

For the final model in the main paper we recollect the data at 4 FPS and train on every frame. The model sees equivalent data, but this makes the second dataset is 5x smaller and hence easier to release. To reduce storage requirements further, we compress the camera images with JPG (we add the compression during inference as well to avoid distribution shifts), perception labels with PNG (depth maps are stored at 8 bit resolution), text files with zip, and LiDAR point clouds with LASzip [20] (which compresses point clouds much better than standard zip). We store full 360° LiDAR sweeps and realign all points into the coordinate system of the current frame. Since the expert is not a perfect driver, we log the driving score of each training route and only train on routes with 100 DS. The labels

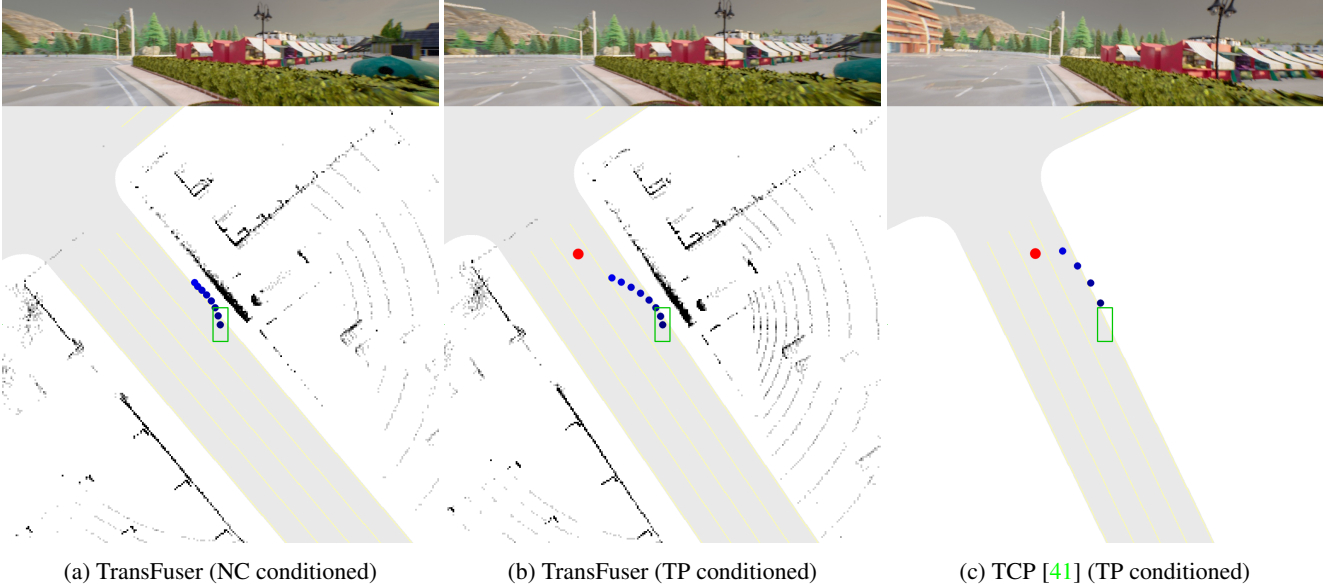


Figure 8: **Extrapolation to target point.** Additional example of an unknown situation where TP conditioned methods extrapolate their waypoints towards target points. The TP conditioned models succeed to drive back to the lane, while the discrete conditioned model gets stuck on the sidewalk.

for the disentangled path are generated by storing the next 10 2D points (spaced 1 meter apart) from the A\* path the expert is following.

### A.3. Training and Architecture

**Training:** We use the same loss functions as in [12] for all 5 outputs. For the new target speed classification task we use a class frequency weighted cross-entropy loss with label smoothing 0.1. As classes we use the 4 different target speeds of the expert (see Section A.1). We observe some failure cases where the target speed predictions are overconfident, and the car starts braking a few frames too late when stopping behind another vehicle, leading to close rear collisions. Label smoothing addresses this problem because it makes the car drive a bit slower (due to the probability weighting). This gives the car a few more frames time to brake and avoids these collisions. To make this more explicit, we do not use label smoothing in the final model in the main paper, and instead reduce the target speeds by 2 m/s during inference, achieving the same effect.

During training, each individual loss is added together with the same weight. The loss weights are normalized to sum to 1. We use AdamW [23] with amsgrad [30] and a slightly higher learning rate of 0.0003 which is reduced by a factor of 10 after epoch 30. We train for 31 epochs and always use the last epoch as our model. All models in Section 3 of the main paper are trained on 8 2080ti GPUs with distributed data parallel and total batch size of 48. For the dense dataset, we subsample by a factor of 5 but shift the first frame by the GPU index, so that all the frames are seen

during training (a subtle form of data augmentation, close by frames are mostly redundant).

The final model is trained with 4 A100 GPUs and batch size 128 to speed up training. We also add standard color augmentations to the camera and a discrete conditioning concatenated with the velocity input, so that the target speed branch also has a conditioning signal. This is conceptually more sound, but we observed no significant impact on driving performance. Our path labels have a subtle ambiguity due to conversion from a path in global coordinates to the local ego coordinate system (points can be shifted by 1 meter depending where the car is on the path). To resolve this, we linearly interpolate between the stored points at 1 meter distances in ego coordinates during training. This makes the predictions look qualitatively more consistent, but again has no significant impact on driving score.

**Architecture:** For the BEV semantic segmentation we predict 11 classes (unlabeled, road, sidewalk, non cross-able lane marker, cross-able lane marker, stop signs, traffic light stop line green / yellow / red, vehicle, pedestrians) [43]. We only predict pixels that are visible in the camera, since some classes need RGB features. For the perspective segmentation we predict 7 classes (unlabeled, road, sidewalk, lane markings, traffic light, vehicle, pedestrian). For the bounding boxes we predict 4 classes (red and yellow traffic light stop line, stop sign, vehicles and pedestrians). Traffic lights and stop sign boxes are only predicted if they affect the ego vehicle. Vehicles and pedestrians only when they are in the LiDAR range. Architecture wise we add 1x1 convolution layers in the LiDAR branch before each transformer that



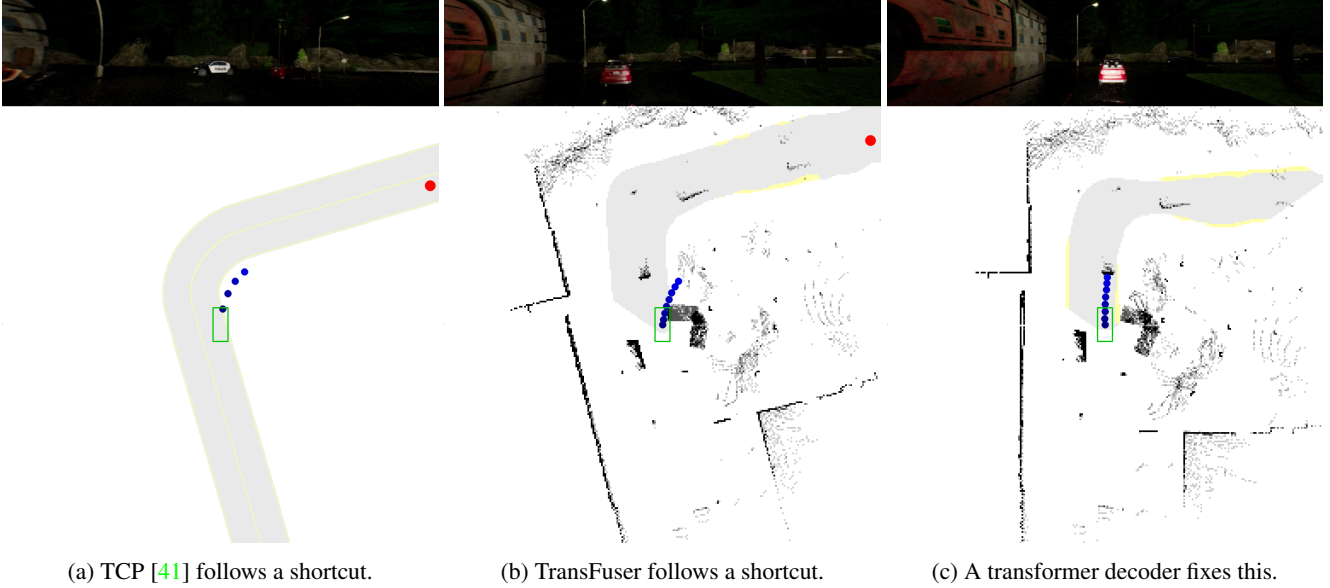


Figure 9: **Target point shortcut.** When TP conditioned methods extrapolate to spatially distant waypoints, they incur large steering errors. Replacing global average pooling in TransFuser with a cross-attention mechanism mitigates the issue.

match the channel dimension of the LiDAR to the one in the image branch (allows using different backbones for the two branches, though we keep using RegNetY-3.2GF [29] for both). Our reproduced TransFuser predicts 8 waypoints, each placed 250 ms apart (up to 2 seconds into the future).

**Sensors:** We follow [41] and use a single high resolution camera. It has a  $110^\circ$  horizontal field of view (FOV), is mounted at  $(-1.5, 0.0, 2.0)$  and has a resolution of  $256 \times 1024$ . Our LiDAR is mounted at  $(0.0, 0.0, 2.5)$  and has a full  $360^\circ$  FOV. LiDAR points are realigned into the current vehicle coordinate system by using the filter described in Section A.4. The LiDAR points are voxelized into a  $256 \times 256$  grid representing a  $64 \times 64$  meter area with the vehicle at its center. Each pixel covers a  $0.25 \times 0.25$  meter area. We remove the LiDAR ground plane by removing all points with a height of less than 0.2 meters.

**Performance:** The reproduced TransFuser has a slightly lower DS to the original on the Longest6 benchmark as shown in Table 10. The difference in driving score likely is a result of the driving score weighting the dissimilar failure modes differently. Note, that our reproduced TransFuser is a single model, whereas the original result was from an ensemble of 3 models.

Method	DS $\uparrow$	RC $\uparrow$	Veh $\downarrow$	Stat $\downarrow$
TransFuser (ours)	40	82	<b>1.17</b>	0.57
TransFuser [12]	<b>47</b>	<b>93</b>	2.45	<b>0.07</b>

Table 10: **Reproduced TF vs original on Longest6**

#### A.4. Localization

We localize our vehicle with a GNSS sensor. The GNSS signal has Gaussian noise applied to it, leading to average localization errors of  $\sim 0.7$  meters. Like prior work [9, 12] we use a filtering algorithm to reduce the noise. In particular, we use an Unscented Kalman Filter (UKF) [32, 39] with Van der Merwe’s scaled sigma point algorithm [24]. As its model of the car, our UKF uses the same kinematic bicycle model as the expert [8]. The filter tracks the position, orientation and speed of the ego vehicle. The parameters of the filter are tuned manually by reducing localization error on a small dataset consisting of ground truth localization paired with GNSS signals (similar to tuning hyperparameters on a validation set in supervised learning). The filter reduces the average localization error to below  $\sim 0.1$  meters.

#### A.5. PID Controller

In [12] the PID controller converts the waypoints into steering by computing the car’s angle towards the average between the first two waypoints. The angle is then input to a PID controller to minimize. Computing the angle based on entangled waypoints makes the steering angle towards the path depend on the predicted speed of the vehicle. This works well at the low driving speeds of [12] but becomes hard to tune when driving at higher (and more diverse) speeds like our expert. Instead, we follow the first waypoint that is at least a certain aim distance  $a$  away from the center of the car (or the last one). We keep  $a$  similar to [12] and use  $a = 2.25$  when driving slower than 5.5 m/s (inside intersections) and  $a = 3.0$  when driving faster. Longitudi-

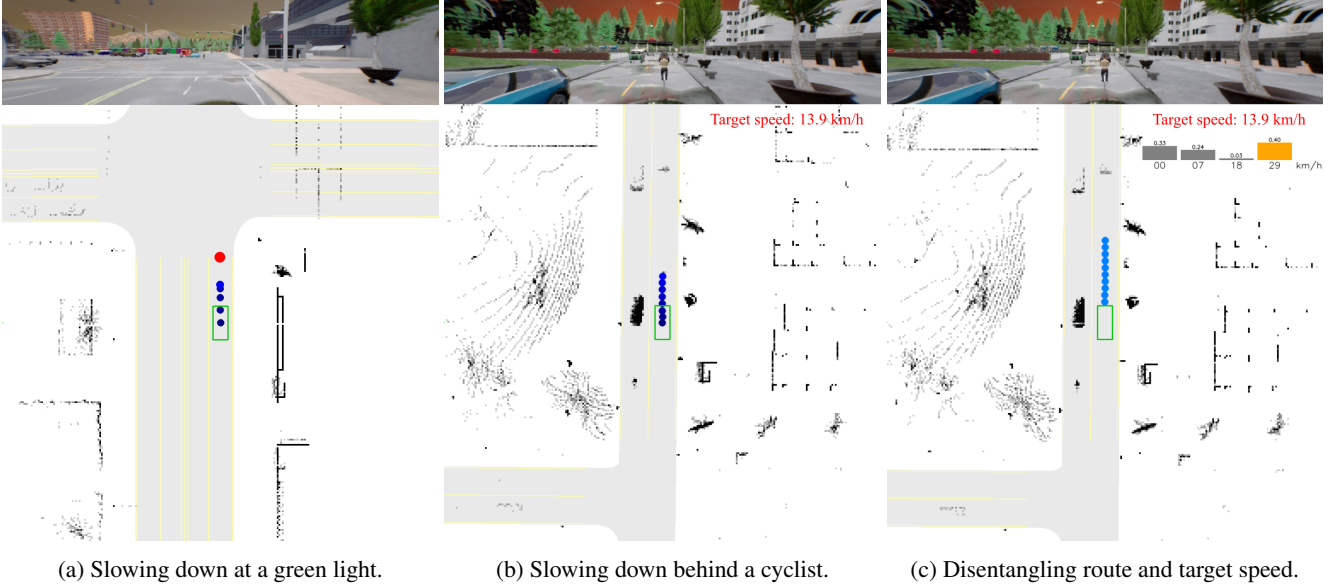


Figure 10: **Waypoints are ambiguous.** The model’s output representation forces it to predict a single mode for future velocities.

nal control is kept similar, we use the velocity between the waypoint 0.5 second into the future and 1 second into the future as target speed. The models using the disentangled waypoint representation do not require tuning an additional controller. Since we directly predict the input to the expert’s PID controller, we can reuse the same PID controllers from the expert (see Section A.1), which we know work well.

The task of yielding to stop signs is defined in CARLA to stop the vehicle on a STOP painting printed on the road. This stop painting is largely occluded by the motor hut once the car is on the sign because our camera is mounted at the back roof of the vehicle. We observe many cases where TF++ detects the stop sign initially and slows down but starts driving again once the painting becomes occluded. We can address this problem in the controller because TF++ detects stop sign bounding boxes in BEV as auxiliary task. For our final model, we keep the last detected stop sign in a buffer and transform it to the current ego coordinate system using the UKF motion estimation. If the car is on the stop sign bounding box we set the action to brake in the controller until the stop sign is cleared. Table 11 shows the impact of the stop

Stop Contr.	DS $\uparrow$	RC $\uparrow$	Stop $\downarrow$
-	60 $\pm$ 5	99 $\pm$ 1	1.35
✓	70 $\pm$ 6	99 $\pm$ 0	0.26

Table 11: **Stop sign controller.**

sign controller change on the validation routes. TF++ detects most stop signs and the controller is able to reduce Stop infractions by 5x, leading to a +10 DS improvement.

## B. TransFuser++ Implementation Details

### B.1. Attention Pooling Implementation

The core implementation of our attention pooling with a transformer decoder is similar to [35]. We use the 8x8 features coming from the BEV branch and reduce the number of channels (1512 with our RegNetY-3.2GF backbone) to 256 with a 1x1 convolution. We add a sinusoidal positional embedding to the features and flatten them afterward. The velocity input, we normalize with a 1D Batchnorm and embed it with a 2 layer MLP to 256 features. Afterward, we add a learnable positional encoding. The velocity token then gets concatenated to the BEV tokens. These tokens are then processed by a standard Transformer decoder [38] that has 8 heads, gelu activation function [18], 6 layers and uses layernorm [3]. We use 1 learned queries for every predicted waypoint and 1 additional one if we predict target speeds. The tokens are then fed as inputs into a GRU, whose hidden state is initialized with an (MLP embedded) target point. A linear layer converts the output to 2 dimensions, and a cumulative sum is applied to the final vector (forcing the network to predict offsets from the first point).

### B.2. Data Augmentation

When collecting data with the expert driver, we mount an additional camera on the vehicle that collects augmented frames and labels at every time step. This means for every frame in the dataset, we have an augmented counterpart. We shift the camera by  $\pm 1$  meter to the left or right (of the vehicle) and rotate it by  $\pm 5^\circ$  around the yaw axis. The particular values are drawn from a uniform distribution. During train-

ing, we load the perturbed camera with 50% probability. We transform the 2D way point labels and other data such that they have the perturbed camera as center. The shifted and rotated waypoint labels do not describe an actual recovery trajectory that an expert would take, and are instead the original waypoints in the center of the lane. However, the PID controller will steer the car back to the center of the lane, when the network predicts these augmented waypoints, achieving a similar effect. The range of errors the network is trained to recover from equals the range of the perturbations. To generate the perturbed camera, we need access to a rendering engine (CARLA in our case). On real data, a similar effect can be achieved using novel view synthesis [2].

## C. Additional Results

### C.1. Longest6 Ablations

As a sanity check, we test our additions to TransFuser by repeating the experiments on the training towns (Longest6) while training with all data. The results are presented in Table 13. We start with our reproduced version of TransFuser and iteratively add 1 change in every row. All results are the average of 3 training seeds, each evaluated 3 times. The reported standard deviation is across the training seeds. The main problems of the reproduced TransFuser are high collisions with other vehicles (Veh) and collisions with the static environments (Stat). Replacing the global average pooling and MLP with a Transformer decoder improves the driving score by 17 points, improving both vehicle collisions and collisions with the environment (indicating improved steering). Adding shift and rotation augmentations has a similar effect, improving the driving score by +9 and reducing collisions. The disentangled representation has a similar effect than on the validation routes, increasing vehicle collisions and reducing collisions with the environment. Since Longest6 has denser traffic, the increased vehicle collisions have a larger impact here, leading to an overall decrease in 2 DS. Two stage training improves DS by +3 and scaling the dataset by 3x leads to a result of 69 DS. Since the waypoint representation is slightly better on longest6, we additionally report the result of a variant called TF++ WP in Table 13. It was trained with the released dataset and uses waypoints as output representation, but is otherwise identical to TF++. TF++ WP again has a slightly higher DS of +1. We also report the result of an ensemble of the 3 training seeds following prior work [12]. The ensemble reduces vehicle collisions by 0.10 and leads to a SotA result of 73 DS on longest6.

### C.2. Brake Threshold

The confidence weighted PID controller from the models using target speed prediction has a hyperparameter that

determines at which confidence threshold the action is set to full brake. In Table 14 we investigate the effects of different thresholds on validation towns.

Lower thresholds reduce vehicle collisions at the cost of lower route completion due to false positive braking. Overall, the choice of threshold is robust, only changing by 3 DS between the best and worst tested threshold. We use the default threshold of 50% for all validation town experiments. On Longest6, the brake threshold has a slightly larger impact because the dense traffic puts a higher focus on collision avoidance. We therefore use the threshold 33% for Longest6 experiments.

### C.3. Additional Baselines

We report 3 additional results on the validation towns in Table 12. TF++ (all towns) is TF++, but its training includes the validation towns. This is not a fair comparison to other methods, and rather serves to illustrate what part of the remaining problems are due to generalization issues. Including the validation towns during training increases the driving score by 20 points, improving both vehicle collisions and stop sign infractions. TF++ is close to expert level performance in this setting, “underfitting” by 4 DS. The expert still has slightly lower vehicle collisions (Veh).

We retrain Perception PlanT [31] on its released dataset. It achieves a DS of 37. Perception PlanT uses a handcrafted intermediate representation (bounding boxes) as visual abstraction. This makes the method more interpretable, but the downside of human designed representations are that they might miss important things. In this case, the bounding boxes do not include stop signs, leading to the method ignoring all stop signs (Stop) and incurring a large penalty for that infraction (1.87). Stop signs were not considered in the benchmark this method was developed on (Longest6).

The reproduced LAV v2 [9] also achieves a surprisingly low driving score of 27. The reason for this are its high stop sign and red light infractions. The red light infractions occur almost exclusively in Town 02 which has European style traffic lights. We evaluate on the LAV routes with additional scenarios (7,8,9,10) compared to the original benchmark. Due to the added scenarios, many traffic lights will turn yellow (and then red) just as the agent approaches an intersection. LAV v2 ignores those situations, incurring many red light infractions, bringing the overall score down. Besides the failure to adhere to traffic rules, LAV v2 achieves SotA results in route completion and vehicle collision avoidance.

### C.4. Additional Examples

Fig. 8 shows another example of the importance of the target point for recovery. This time we forcefully steer the ego car onto the sidewalk, which is also an out-of-distribution situation. The target point conditioned methods (here TransFuser and TCP) extrapolate their waypoints to-

Method	DS $\uparrow$	RC $\uparrow$	IS $\uparrow$	Ped $\downarrow$	Veh $\downarrow$	Stat $\downarrow$	Red $\downarrow$	Stop $\downarrow$	Dev $\downarrow$	TO $\downarrow$	Block $\downarrow$
LAV v2 [9]	27 $\pm$ 1	98 $\pm$ 1	0.27 $\pm$ 0.01	<b>0.00</b>	<b>0.35</b>	0.04	2.89	1.42	<b>0.00</b>	0.09	<b>0.00</b>
Perception PlanT [31]	37 $\pm$ 5	86 $\pm$ 7	0.45 $\pm$ 0.09	<b>0.00</b>	0.92	0.31	0.09	1.87	<b>0.00</b>	0.25	0.14
TransFuser (ours)	39 $\pm$ 9	84 $\pm$ 7	0.46 $\pm$ 0.06	<b>0.00</b>	0.74	1.04	0.20	1.07	<b>0.00</b>	0.23	0.21
TCP [41]	58 $\pm$ 5	85 $\pm$ 3	0.67 $\pm$ 0.06	<b>0.00</b>	<b>0.35</b>	0.16	<b>0.01</b>	1.05	<b>0.00</b>	0.19	0.19
TF++ (ours)	<b>70</b> $\pm$ 6	<b>99</b> $\pm$ 0	<b>0.70</b> $\pm$ 0.06	0.01	0.63	<b>0.01</b>	0.04	<b>0.26</b>	<b>0.00</b>	<b>0.05</b>	<b>0.00</b>
TF++ (all towns)	90 $\pm$ 1	99 $\pm$ 1	0.91 $\pm$ 0.02	0.00	0.18	0.00	0.05	0.00	0.00	0.02	0.00
Expert	94	95	0.99	0.00	0.02	0.00	0.02	0.00	0.00	0.00	0.08

Table 12: **Performance on validation towns (LAV).** Reproduced models, std over 3 trainings and 3 evaluations.

Method	DS $\uparrow$	RC $\uparrow$	Veh $\downarrow$	Stat $\downarrow$
TransFuser (ours)	40 $\pm$ 3	82 $\pm$ 2	1.17	0.57
+ transformer decoder	57 $\pm$ 3	90 $\pm$ 3	0.93	0.19
+ data augmentation	66 $\pm$ 4	94 $\pm$ 2	0.64	0.07
+ disentangled	64 $\pm$ 4	96 $\pm$ 1	0.88	0.02
+ two stage	67 $\pm$ 2	96 $\pm$ 1	0.82	0.01
+ 3x data	69 $\pm$ 1	<b>97</b> $\pm$ 1	0.79	<b>0.00</b>
TF++ WP (ours)	70 $\pm$ 4	94 $\pm$ 2	0.66	0.01
+ ensemble	<b>73</b>	<b>97</b>	<b>0.56</b>	0.01

Table 13: **Longest6 Ablations.** Each row has the previous one as baseline. Std over 3 training and 3 evaluation runs.

Threshold	DS $\uparrow$	RC $\uparrow$	IS $\uparrow$	Veh $\downarrow$
50%	60 $\pm$ 6	<b>98</b> $\pm$ 1	0.61 $\pm$ 0.06	0.73
40%	58 $\pm$ 1	96 $\pm$ 5	0.61 $\pm$ 0.04	0.78
33%	<b>61</b> $\pm$ 3	96 $\pm$ 4	<b>0.64</b> $\pm$ 0.07	<b>0.61</b>
25%	59 $\pm$ 3	95 $\pm$ 6	0.63 $\pm$ 0.08	0.64

Table 14: **Brake threshold.**

wards the nearby target point and drive back to the center of the lane. The discrete conditioned TransFuser that does not have access to the geometric information of the target point gets stuck on the sidewalk instead.

Fig. 9 shows another example of harmful target point extrapolation. TCP and TransFuser both predict waypoints leaving the road in a right turn where the target point is far behind the turn. Changing the global average pooling + MLP approach in TransFuser to a transformer decoder mitigates the problem.

### C.5. Additional Experiments

**Multimodal waypoints:** Instead of disentangling the velocity from the waypoints one could also allow the network to predict multiple sets of waypoints as is sometimes done in trajectory forecasting. To test this approach we train a model with two waypoint GRUs and a selection head to classify the better mode. During training we compute

the loss as the minimum L1-loss from both predicted waypoints. The classification head is supervised with a binary cross entropy loss to classify which of the two waypoint losses has the lower L1 loss. Table 15 shows that the rep-

Output	DS $\uparrow$	RC $\uparrow$	Veh $\downarrow$	Stat $\downarrow$
Multimodal WP	47 $\pm$ 8	<b>96</b> $\pm$ 1	1.13	0.20
Unimodal WP	<b>49</b> $\pm$ 8	90 $\pm$ 4	<b>0.70</b>	<b>0.10</b>

Table 15: **Multimodal waypoint representation.**

resentation performed 2 DS worse than using standard unimodal waypoints. This is not a big difference but the multimodal waypoints are more complex so there is not really a reason to use them.

**NC conditioned AIM:** To test whether recovery from the target point bias depends on the architecture, we reproduce the AIM [28] approach. Compared to TransFuser it uses no LiDAR, no auxiliary losses and no transformers, but has the same target point conditioned waypoint GRU as decoder. Like with TransFuser we run two variants, one with TP conditioning and one with NC. The result presented in Table 16

Cond.	DS $\uparrow$	RC $\uparrow$	Dev $\downarrow$
NC	<b>27</b> $\pm$ 2	50 $\pm$ 3	0.96
TP	26 $\pm$ 2	<b>86</b> $\pm$ 8	<b>0.00</b>

Table 16: **Conditioning effect on AIM [28]**

also show a strong impact of TP conditioning on RC and route deviations (Dev) for the AIM architecture, suggesting that the target point bias does not depend on the particular architecture used. Consistent with [28] we also observe that AIM performs worse than TransFuser overall (-13 DS).

**Target Point statistics:** To show the importance of the TP for the network predictions we train a model where we input the TP as a TF decoder token so that we can analyse the attention weights. We average the attention across all layers, heads and waypoint queries. We run the model on the validation routes and observe an avg. of 25% attention on the TP token which is 16.5x higher than uniform. We also test how steering correlates with the target point. The



sign of the steering angle (when larger than  $1^\circ$ ), of our final model, is the same as that of the target point (TP) 92% of the time (93% for the expert) on the validation routes.

### C.6. CARLA Leaderboard

The CARLA leaderboard [1] (we are considering version 1 in this discussion) is a test server where groups can submit agents enclosed within a docker container. These agents are then evaluated on a set of 10 routes across 2 secret towns. Each route is traversed  $2\times$  with different weather conditions. Additionally, these 20 routes are then repeated 5 times with different random seeds and average metrics across all routes are reported to the user. The CARLA leaderboard serves as a standardized evaluation platform, comparable to a test set, and some works solely rely on it to compare to other work [9, 35, 41]. Significant progress has been achieved on this benchmark. Driving scores have increased from approximately 10 DS to 70-80 DS over a span of three years. The top performing methods on the leaderboard have released code and models alongside the release of their research papers. In this study, we try to reproduce the top 4 reported results by submitting the released code to the leaderboard, either with the released model (\*) file or a retrained one ( $\dagger$ ).

Method	Reproduced			Reported		
	DS $\uparrow$	RC $\uparrow$	IS $\uparrow$	DS $\uparrow$	RC $\uparrow$	IS $\uparrow$
LAV* [9]	25	46	0.74	62	<b>94</b>	0.64
Interfuser* [35]	34	75	0.45	<b>76</b>	88	0.84
TCP $\dagger$ [41]	48	66	0.77	70	83	<b>0.85</b>
TF++ (ours)	53	71	0.76	-	-	-
TransFuser $\dagger$ [12]	55	<b>90</b>	0.63	61	87	0.71
TF++ WP (ours)	62	78	0.81	-	-	-
TF++ WP Ens. (ours)	<b>66</b>	79	<b>0.84</b>	-	-	-

Table 17: CARLA leaderboard reproduced results.

Table 17 shows significant disparities between reported results and the outcomes obtained when utilizing the official codebases. Notably, the released model of the SotA method Interfuser achieves less than half of the reported score. With the exception of TransFuser, these differences are larger than what we would expect from training or evaluation variance. The CARLA leaderboard ensures repeatability (within the bounds of evaluation variance) because one can rerun the submitted docker container. It does however not guarantee reproducibility, since released code and models can differ from what was used to achieve the reported score. Currently, it is publicly not known how to achieve scores above 60 DS, even though some groups have achieved such scores in the past. Given that the official codebases incorporate the main concepts discussed in the respective papers, these results indicate that factors not ex-

plicitly emphasized in the papers significantly influence the outcomes observed on the CARLA leaderboard.

Note, that it has been documented in [12] (Table 5) that some changes can have a significant impact on performance on the leaderboard (+19 DS), even though the same changes do not generalize to the public towns (-2 DS on Longest6). This discrepancy can be attributed, in part, to substantial fluctuations in RC, which we don’t observe in the publicly available towns. The evaluation is secret, so the exact reasons for these fluctuations are unclear. The LAV results are crossed out because the low score can be attributed to CUDA out of memory errors stemming from the released software, that appear after route 53. The CARLA leaderboard fails silently and gives 0 DS on routes where the software crashes. We contacted the organizers in this case because some of the auxiliary metrics looked unusual.

Furthermore, the auxiliary metrics on the CARLA leaderboard are typically not reported because they are known to be incorrectly computed<sup>1</sup>. The results presented in Table 17 were obtained with released code and models around April 2023. It should be noted that these results are subject to change if the authors decide to update their repositories. We submit TF++ and TF++ WP (see Section C.1) to the CARLA leaderboard. Surprisingly, we observe a large difference of 9 DS between the representations. The reason for that difference is unclear because the effect is not reproducible on the publicly available towns. Similar to prior work [12], we submit an ensemble with 3 training seeds of our best model to the leaderboard. The ensemble increases the driving score by 4 points and is the best publicly available model, at the time of writing.

**Discussion:** The CARLA leaderboard aimed at reproducing success in past benchmarks like ImageNet for autonomous driving. It has led to similar fast progress, although the resulting methods are not always reproducible and not very well understood. There are structural changes the community could make that we think should be considered when designing future benchmarks.

The CARLA leaderboard does not have an official validation benchmark. Since validation is fundamental to machine learning development, authors have proposed various validation benchmarks [9, 11, 12, 28, 43] using the publicly available towns. They differ along various axes: the towns, route length, scenarios and weathers used for validation. Based on our experience, none of the available validation benchmarks can reliably predict performance on the CARLA leaderboard. It is hard for authors to make validation routes that are well aligned because the test routes are secret. As a consequence, results on the CARLA leaderboard are often unexpected, requiring authors to run additional investigations [12]. Other authors choose not to use the leaderboard in their publications [19, 22, 31, 42, 43].

<sup>1</sup><https://github.com/carla-simulator/leaderboard/issues/117>

Submissions on the CARLA leaderboard may encounter a pending status if all servers are in use. During this project, we have encountered pending times of up to 3 weeks. Submissions themselves then take more than four days to evaluate, leading to total evaluation times of up to 4 weeks. For comparison, we can run a similar amount of simulation on our cluster in 3–6 hours, by evaluating all routes in parallel. Achieving this efficiency requires a scalable infrastructure but uses the same amount of computational resources as when evaluating routes sequentially. We encourage organizers of future testing benchmarks to release an aligned public validation benchmark that authors can use to do rapid experiments on.

The CARLA leaderboard is a benchmark that is not associated with any dataset. Consequently, authors are required to create their own datasets for training their methods. In the early stages, the quality of these datasets, including sensor configurations, label accuracy, and scale, was often subpar. Substantial progress has been made by enhancing the quality of the data itself [21]. Most papers, including this one, introduce a new dataset alongside their method. While this practice allows for innovation in dataset curation techniques, it also complicates fair comparisons between different methods. Additionally, it can obscure the source of performance improvements if the dataset is not studied as well. Although there have been some efforts in dataset curation [8,43], such efforts could be more fruitful if there was a standardized dataset to compare to.