Interpretable Decision-Making for End-to-End Autonomous Driving

Mona Mirzaie Bodo Rosenhahn
Institute for Information Processing, Leibniz University Hannover
mirzaie@tnt.uni-hannover.de

Abstract

Trustworthy AI is mandatory for the broad deployment of autonomous vehicles. Although end-to-end approaches derive control commands directly from raw data, interpreting these decisions remains challenging, especially in complex urban scenarios. This is mainly attributed to very deep neural networks with non-linear decision boundaries, making it challenging to grasp the logic behind AI-driven decisions. This paper presents a method to enhance interpretability while optimizing control commands in autonomous driving. To address this, we propose loss functions that promote the interpretability of our model by generating sparse and localized feature maps. The feature activations allow us to explain which image regions contribute to the predicted control command. We conduct comprehensive ablation studies on the feature extraction step and validate our method on the CARLA benchmarks. We also demonstrate that our approach improves interpretability, which correlates with reducing infractions, yielding a safer, high-performance driving model. Notably, our monocular, non-ensemble model surpasses the top-performing approaches from the CARLA Leaderboard by achieving lower infraction scores and the highest route completion rate, all while ensuring interpretability. Code is available here[1].

1. Introduction

Autonomous driving systems in use today are categorized as either highly modularized hand-engineered, or end-to-end approaches. The modular pipeline is the conventional method widely adopted in the industry, which consists of several subsystems like perception, localization, prediction, planning, and control [30]. The main advantage of the modular pipeline lies in its interpretability, which allows the pinpointing of malfunctioning modules [44, 49]. Nevertheless, a closer look at the perception module shows that even with the integration of certain learning methods, its design still heavily relies on human-defined features, such as lane markings and vehicles [28]. The reliance on human heuristics (e.g. rule-based planners) can result in overly conservative driving policies and difficulties in generalization, as each new driving scenario might require reconfigured heuristics [10, 49]. Furthermore, the development and maintenance of this system are expensive and overly complex [11, 53]. Despite years of sustained efforts, these approaches still fall short of full autonomy [10, 44].

Recent advancements in the realm of autonomous driving have shown great potential for end-to-end approaches. Human drivers often rely on reflex-like actions, requiring minimal high-level reasoning and conscious attention—a concept aligned with end-to-end learning [44]. Moreover, the limitations of the conventional methods can be mitigated through end-to-end autonomous driving techniques, where a driving policy is learned and generalized to new tasks with minimal manual engineering [10, 49]. While ensuring trustworthiness in autonomous driving necessitates interpretability to address safety concerns and support commercialization [29], achieving this remains a major challenge in end-to-end systems [12, 28, 53, 55]. The absence of intermediate outputs results in a black-box learning scheme [33, 52, 54], which makes it difficult to identify the origins of errors and to clarify the rationale behind certain driving decisions [28, 51, 55]. To alleviate this limitation, we have incorporated a technique into our model that enables the acquisition of diverse feature maps. This enhancement significantly improves interpretability within our model, leading to a notable reduction in infractions during driving sessions in the CARLA driving simulator [17]. Several top-ranked approaches on the CARLA public Leaderboard [45] fuse multiple modalities, such as LiDAR and multi-view camera inputs, and leverage ensemble and knowledge distillation techniques, as well as incorporate many auxiliary tasks. Conversely, our work attains state-of-the-art performance by employing solely a monocular camera, without the aid of ensemble or distillation techniques and traffic-rule-specific sub-tasks. It also offers substantial reductions in execution time and computational complexity compared to other competitors on the CARLA Leaderboard. Our main contributions are as follows:

[1] https://github.com/MandM-VisionLab/DTCP

- We present a mechanism that improves interpretability by enforcing diversity in learned features.
- We evaluate our model's interpretability and show significant improvements over the baseline, TCP [50].
- We conduct an in-depth empirical analysis on the LAV benchmark and validate our approach on the official CARLA Leaderboard.

2. Related Work

Visual Explanation in Autonomous Driving: Interpreting a model's decisions is crucial for diagnosing failures and enhancing system performance. Methods to improve interpretability generally fall into five categories: attention-based mechanisms, utilization of semantic inputs, auxiliary tasks, cost learning, and the application of natural language processing (NLP) [12, 54].

Attention-based models reveal internal reasoning by highlighting salient regions [28], or leveraging object-level features [48]. In parallel, auxiliary-task strategies leverage multi-task optimization to decode latent features into other relevant outputs—such as classification [6, 42, 46], semantic segmentation [6, 8, 13, 22, 51], depth estimation [20, 23, 36], affordance predictions [31, 46], or motion prediction [8, 42]—thus promoting interpretable and robust representations. Semantic inputs offer structured high-level context by feeding e. g. bird's-eye-view (BEV) representations [1, 27] or rasterized scenes [16] directly to the network and hence enhancing interpretability through inherently understandable representations. Finally, cost-volume visualization (cost maps) provides transparent insights into the model's decision-making by ranking sampled trajectories based on the predicted cost for reaching specified locations over time [18, 19, 55]. By displaying these maps for multiple future steps, the system's trajectory choices and underlying rationales become more interpretable.

Interpretability in Other Applications: Recent advancements in fine-grained image classification have led to significant improvements in interpretability. B-Cos Networks [4] enhance interpretability by aligning input features with model weights, while ProtoPNet [7] and BotCL [47] achieve this through prototype-based reasoning and refined concept-based learning. Norrenbrock *et al*. [37] and Chang *et al*. [5] develop lightweight training techniques that effectively discern both distinctive and exclusive features, enabling fine-grained class differentiation without complex architectures. Collectively, these works underscore the role of representation learning in enhancing interpretability by ensuring models learn distinct, meaningful features. Motivated by [5, 37], we design a novel feature diversity loss and integrate it into our training process alongside other losses, leading to more refined and localized feature maps. Our approach results in a more interpretable model, as demonstrated in Sec. 4.2. Our autonomous driving framework, which directly enhances interpretability in its model design, is a combination of cost learning and attention-like mechanisms. As the feature diversity loss acts as a regularization term, it penalizes redundant or overlapping feature activations while enforcing diversity in learned representations. Although our framework does not explicitly implement an attention mechanism like those in transformers or self-attention layers, the sparse feature selection process mimics the behavior of attention. It filters out irrelevant information and attends to the most relevant features leading to localized and interpretable feature selection, which is characteristic of attention models that focus on salient regions.

Evaluating Interpretability: Recent models integrate motion prediction and planning within a single network, leveraging these components to enhance interpretability. For instance, PlanT [41] proposes a Relative Filtered Driving Score (RFDS) to quantify how well the planner focuses on the most critical objects for safe driving decisions, while UniAD [21] employs cost-map visualizations to highlight the safety risks leading to potential collisions across different trajectories. Beyond these approaches, models incorporating transformers and attention mechanisms, such as TransFuser [13], ReasonNet [43], and InterFuser [42], visualize attention weights directly to highlight the most influential tokens in decision-making. Comparing saliency maps with chosen ground-truth references, such as semantic concepts or bounding boxes, is another widely used technique, as demonstrated in [57, 60]. Following Boggust *et al*. [3], we report Intersection over Union (IoU), Ground Truth Coverage (GTC), and Saliency Coverage (SC) metrics to quantify our model's interpretability improvements, complemented by visualization results. To gauge how tightly our model's visual attention aligns with its driving decisions, we compute the Pearson correlation between saliency mass and the control signals. Our results on the CARLA benchmarks show a substantial decrease in the infraction scores, further validating the effectiveness of our approach and highlighting the crucial role of interpretability in driving performance.

3. Method

In this section, we introduce our proposed diversity losses and how agents can benefit from using them during training to improve interpretability.

3.1. Problem Setting

We optimize a driving policy π that generates raw control signals $a = \{\text{brake}, \text{throttle}, \text{steer}\}$ in the ranges $[0, 1]$, $[0, 1]$, and $[-1, 1]$ respectively. The inputs to π include the sensor signal s, the vehicle speed v, and high-level navigation information g (i.e., a discrete navigation command and the destination coordinates). To tackle this problem, we

have focused on the Behavior Cloning (BC) approach in imitation learning (IL), which is a supervised learning method that pursues a policy π to mimic the behavior of an expert π^* [38, 41, 50]. It can be formulated as:

$$\arg\min_{\pi} \mathbb{E}_{(\mathcal{X},\mathrm{a}^*,\mathcal{W})\sim \mathrm{D}} \left[\mathcal{L}_{\text{act}}\left(\mathrm{a}^*, \pi_{\text{act}}(\mathcal{X})\right) + \mathcal{L}_{\text{traj}}\left(\mathcal{W}, \pi_{\text{traj}}(\mathcal{X})\right)\right] \tag{1}$$

where the dataset D is comprised of a sequence of state-action-trajectory tuples $\{(\mathcal{X}, a^*, \mathcal{W})\}_{t=0}^{T}$ acquired from the expert by controlling the ego-vehicle and interacting with the environment in T steps. The $\mathcal{L}_{\text{act}}$ loss function measures the deviation between expert action a^* and the model's predicted action $\pi_{\text{act}}(\mathcal{X})$, whereas $\mathcal{L}_{\text{traj}}$ aims to minimize the distance between the ground truth trajectory $\mathcal{W}$ and the model's predicted trajectory $\pi_{\text{traj}}(\mathcal{X})$.

3.2. Baseline

3.2.1. End-to-end Driving Model

We adopt TCP (Trajectory-guided Control Prediction) [50] as our base model for its strong performance on the public CARLA Leaderboard 1.0 [45] and minimal reliance on extra modules. Unlike many imitation learning approaches that employ a rule-based expert agent, TCP utilizes Roach [58], an expert agent based on Reinforcement Learning, offering the advantage of being trainable. TCP comprises trajectory and control branches to forecast the planned trajectories and multi time-steps control signals. In the TCP architecture presented in Fig. 1, the front camera image $\mathcal{I}$ is processed by a pre-trained ResNet-34 encoder. Concurrently, the measurement input $\mathcal{C}$, which includes navigational commands and the goal location, along with the velocity is fed into an MLP encoder. The concatenated feature map from both encoders is then utilized by the control and trajectory prediction branches. Lastly, the predicted actions from each branch, i.e., $\mathbf{a}^{traj}$ and $\mathbf{a}^{ctrl}$, will be combined using a situation-based fusion strategy to form the final action for controlling the agent. In the following, each part will be discussed succinctly. For additional details, we refer to [50].

Trajectory Branch: Motivated by the auto-regressive waypoint prediction network in [38], the concatenated feature map from both encoders $\mathbf{F}^{traj}$, along with the current position of the ego agent and the goal location are fed into GRU [14] layer, followed by linear forecasting of T future time-step waypoints $\{\mathbf{w}_t\}_{t=1}^{T}$. This prediction involves computing the distance between each pair of future waypoints $\{\Delta\mathbf{w}_t\}_{t=1}^{T}$ and adding them to the current position $\{\mathbf{w}_{t-1}\}_{t=1}^{T}$ at each time-step, starting from the initial position $(0, 0)$.

To predict control actions $\mathbf{a}^{traj}$ from the forecasted future waypoints, an inverse dynamics model [2] is implemented in the form of a PID controller, as presented in [9]. Particularly, two PID controllers are deployed for lateral and longitudinal control, taking the predicted future waypoints as input. The longitudinal controller adjusts throttle and brake, while the lateral one controls the steering.

Control Branch: For forecasting multiple actions towards the future, a multi-step control strategy is applied. Given the current state $\mathcal{X}_{t-1}$, the control branch predicts future actions $\left\{\mathbf{a}_{t-1}^{ctrl}\right\}_{t=1}^{T}$ for T time-steps based on the policy $\pi_{\text{ctrl}(\mathcal{X}_{t-1})}$ using a temporal module, which delivers information about interaction between the ego-agent and the environment. This module uses a GRU layer that takes the feature at the current time step $\mathbf{F}_{t-1}^{ctrl}$ concatenated with the current predicted action $\mathbf{a}_{t-1}^{ctrl}$ as input and updates the hidden state $\mathbf{h}_{t-1}^{ctrl}$ for the next time-step t. Afterwards, an attention map $\mathbf{w}_t$ is computed using the hidden state from the GRU layer in the trajectory branch $\mathbf{h}_t^{traj}$ and the corresponding one in the temporal module $\mathbf{h}_t^{ctrl}$. The attention weights are multiplied element-wise with the feature map and then combined with the hidden state of the temporal module $\mathbf{h}_t^{ctrl}$. Accordingly, a policy head predicts the control action for the next time-step t using the updated feature map $\mathbf{F}_t^{ctrl}$. Note that, successive GRU layer forward passes receive the hidden state of the previous time-step as the initial hidden state.

Control Actions Fusion: During inference, the final driving action $\mathcal{A}$ is a weighted fusion of $\mathbf{a}^{traj}$ and $\mathbf{a}^{ctrl}$. During turning, $\mathbf{a}^{ctrl}$ receives a higher weight to reduce infractions, as discussed in [50].

3.2.2. Baseline losses

TCP employs common baseline losses for training, including waypoint loss $\mathcal{L}_{\text{traj}}$ [9, 13, 56], control action loss $\mathcal{L}_{\text{ctrl}}$ [25, 26, 58, 59], and sub-task loss $\mathcal{L}_{\text{sub}}$ [15, 56, 58]. Each loss term is weighted by λ to balance its contribution:

$$\mathcal{L}_{\text{baseline}} = \lambda_{\text{traj}} \cdot \mathcal{L}_{\text{traj}} + \lambda_{\text{ctrl}} \cdot \mathcal{L}_{\text{ctrl}} + \lambda_{\text{sub}} \cdot \mathcal{L}_{\text{sub}} \tag{2}$$

In the following, we briefly describe each loss term.

Trajectory Loss: The trajectory loss decreases the L_1 distance between the predicted waypoints $\mathbf{w}$ and the expert (ground truth) waypoints $\tilde{\mathbf{w}}$ for T time-steps:

$$\mathcal{L}_{\text{traj}} = \textstyle\sum_{t=1}^{T} \left\|\mathbf{w}_t - \tilde{\mathbf{w}}_t\right\|_1 \tag{3}$$

Action Loss: Following [58], predicted actions from the control unit $\mathbf{a}^{ctrl}$ and the corresponding expert ones $\tilde{\mathbf{a}}$ are transformed into Beta distributions $\mathcal{B}$. The KL-divergence between these distributions is minimized over current and future time-steps [50]:

$$\mathcal{L}_{\text{ctrl}} = \mathbf{KL}\left(\mathcal{B}\left(\mathbf{a}_0^{ctrl}\right) \| \mathcal{B}\left(\tilde{\mathbf{a}}_0\right)\right) + \tfrac{1}{T}\textstyle\sum_{t=1}^{T} \mathbf{KL}\left(\mathcal{B}\left(\mathbf{a}_t^{ctrl}\right) \| \mathcal{B}\left(\tilde{\mathbf{a}}_t\right)\right) \tag{4}$$

Sub-tasks Losses: As in [58], we include auxiliary tasks—feature matching ($\mathcal{L}_{\text{F}}$), speed regulation ($\mathcal{L}_{\text{S}}$), and

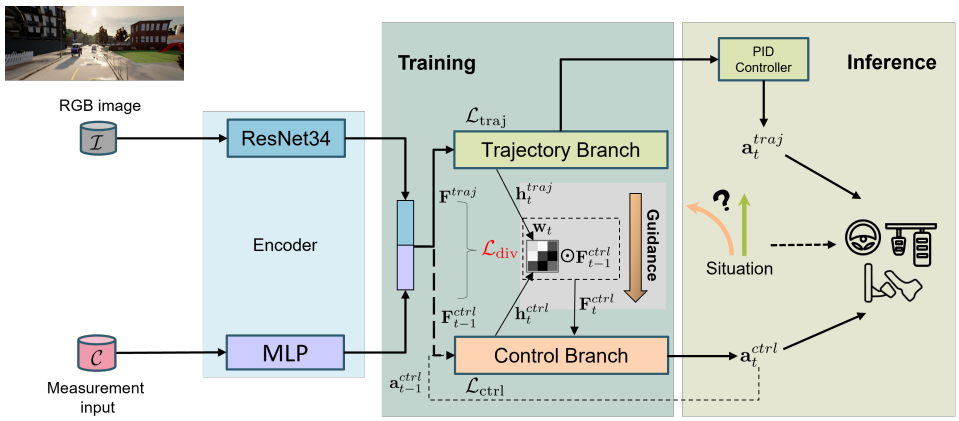

Figure 1. **End-to-End Driving Framework:** The TCP (Trajectory-guided Control Prediction) predicts control signals $\mathbf{a}^{traj,ctrl}$ using a front-camera image $\mathcal{I}$ and a set of measurement data $\mathcal{C}$ (navigational commands, velocity, target point). $\mathbf{F}^{traj}$ denotes concatenated image and measurement encoder features, whereas $\mathbf{F}^{ctrl}_{t-1}$ represents measurement features fused with attention-weighted image features at the current step $t-1$ (dashed arrow). In the trajectory unit, waypoints are predicted using a GRU layer, while the control branch forecasts multi-step control actions leveraging the trajectory branch. During training, two widely adopted loss functions, $\mathcal{L}_{\text{traj}}$ and $\mathcal{L}_{\text{ctrl}}$, together with our proposed $\mathcal{L}_{\text{div}}$ are applied to minimize the difference between the predicted waypoints and actions, and those provided by the expert. During inference, the converted control signals from a PID controller $\mathbf{a}^{traj}$ and the control branch $\mathbf{a}^{ctrl}$ are aggregated using the situational action fusion to form the final control actions.

value regression ($\mathcal{L}_{\text{V}}$)—to align model predictions closer to expert behavior. Specifically, the feature matching loss $\mathcal{L}_{\text{F}}$ enforces consistency between latent representations of the model and the expert, while the speed regulation loss $\mathcal{L}_{\text{S}}$ minimizes differences between the predicted and expert-desired speeds. The value regression loss $\mathcal{L}_{\text{V}}$ predicts the expected future return, hinting at possible upcoming hazardous incidents. These losses are combined as follows:

$$\mathcal{L}_{\text{sub}} = \lambda_{\text{F}}\mathcal{L}_{\text{F}} + \lambda_{\text{S}}\mathcal{L}_{\text{S}} + \lambda_{\text{V}}\mathcal{L}_{\text{V}} \tag{5}$$

3.3. Our Loss Design

The training loss of our proposed method introduces a novel diversity term $\mathcal{L}_{\text{div}}$, alongside baseline losses (see Eq. (2)), expressed as:

$$\mathcal{L} = \lambda_{\text{div}} \cdot \mathcal{L}_{\text{div}} + \lambda_{\text{baseline}} \cdot \mathcal{L}_{\text{baseline}} \tag{6}$$

Diversity Loss: Inspired by [37], we propose to incorporate a diversity loss during training, which encourages the model to capture more distinct features that are distributed as widely as possible. Our proposed loss is highlighted in red in Fig. 1. For a specific feature map $l \in \{0, 1, \ldots, n_f - 1\}$ within the feature maps $M \in \mathbb{R}^{n_f \times w_M \times h_M}$, the diversity component is formulated as:

$$\hat{s}^l_{ij} = \frac{\exp\left(m^l_{ij}\right)}{\sum_{i'=1}^{h_M}\sum_{j'=1}^{w_M}\exp\left(m^l_{i'j'}\right)} \frac{\bar{f}_l}{\max_{c\in[1,n_f]}\bar{f}_c} \tag{7}$$

For each activation value m_{ij} in the l^{th} feature map of M, we apply a softmax normalization across the spatial dimensions $h_M \times w_M$, followed by scaling based on the ratio of the l^{th} feature map's average activation $\bar{f}_l$ to the maximum average activation $\bar{f}_c$ across all n_f feature maps. Accordingly, the diversity loss $\mathcal{L}_{\text{div}}$ aims to minimize the sum of the maximum activations, meaning that the weighted feature map $\hat{S}^l$ highlights different regions of the input. This yields a rich variety of relevant feature activations, defined as:

$$\mathcal{L}_{\text{div}} = -\sum_{i=1}^{h_M}\sum_{j=1}^{w_M}\max\left(\hat{s}^1_{ij}, \hat{s}^2_{ij}, \ldots, \hat{s}^{n_f}_{ij}\right) \tag{8}$$

This approach prevents redundancy and ensures the network exploits more diverse features across the feature maps. The diversity loss (see Eq. (9)) in our approach is applied to both trajectory and control branches. For the trajectory unit, we consider only the current step feature $\mathbf{F}^{traj}$ containing the concatenated feature map from both trajectory and measurement encoders. In contrast, for the control branch, we average the loss of T future time-steps $\left\{\mathbf{F}^{ctrl}_t\right\}_{t=1}^{T}$. This strategy prioritizes the optimization of the actual control action, thereby underscoring the importance of feature diversity at the current time-step.

$$\mathcal{L}^{branches}_{\text{div}} = \lambda_{\text{div}} \cdot \left(\mathcal{L}^{traj}_{\text{div}}\left(\mathbf{F}^{traj}_0\right) + \mathcal{L}^{ctrl}_{\text{div}}\left(\mathbf{F}^{ctrl}_0\right) + \frac{1}{T}\sum_{t=1}^{T}\mathcal{L}^{ctrl}_{\text{div}}\left(\mathbf{F}^{ctrl}\right)\right) \tag{9}$$

4. Experiments

4.1. Experiment Setup

We conduct our experiments using CARLA 0.9.10.1, comprising 8 publicly available towns [17].

Dataset: Following [8, 50], we use the public TCP dataset containing 420 k frames from 8 cities collected using the Roach driving expert [58] under randomized weather. For the ablation study (Sec. 4.2), we train on 185K frames from 4 towns and test on LAV validation routes [8], including two unseen towns. To further validate our approach, we evaluate the Transfuser model [38] on the Transfuser dataset [40] with our method. For the leaderboard submission (Sec. 4.3), we train on the full TCP dataset.

Metrics: Driving performance is assessed using three metrics [45]: Driving Score (DS), Route Completion (RC), and Infraction Penalty (IP). RC quantifies the percentage of the route successfully driven by the agent. IP penalizes violations of traffic rules starting from an ideal score of 1.0, decreasing according to the severity and frequency of infractions. DS summarizes the agent's performance, balancing route completion and safety, and is derived by weighting RC by infractions incurred. Infractions normalized per kilometer are also reported.

Baselines: To benchmark our approach on the CARLA Leaderboard, we compare against four leading methods alongside TCP [50]. **InterFuser** [42] employs a transformer-based fusion of multi-view cameras and LiDAR data, predicting waypoints, object density maps, and traffic rules. It computes desired velocity using density maps combined with motion forecasts derived from historical vehicle dynamics. **ReasonNet** [43] extends InterFuser by adding temporal reasoning modules for improved motion prediction and uses consistency loss to align occupancy maps with predicted trajectories. **TF++ WP** [24] fuses front-camera images and LiDAR BEV representations, predicting future trajectories with auxiliary segmentation and detection tasks, controlled by a modified PID controller. **LAV** [8] integrates multi-view images and LiDAR into BEV representations, predicting semantic maps and dynamic object bounding boxes. It utilizes a privileged motion planner with a distillation framework, enabling the ego-agent to plan trajectories based solely on its own sensors, controlled by PID controllers coupled with collision prediction.

For our ablation study, in addition to TCP, we also evaluate **TransFuser** [38]. It has a similar architecture to TF++ WP that uses transformers to fuse sensory inputs, including a LiDAR BEV and a front-view camera image, without using any auxiliary tasks, and predicts only the ego-vehicle's future waypoints using the trajectory loss (see Sec. 3.2.2). During inference, two PID controllers process the predicted waypoints for lateral and longitudinal control.

Benchmarks: We evaluate our model on two benchmarks: LAV routes and the CARLA public Leaderboard. The LAV benchmark [8] includes 16 route types in Towns 02 and 05 across 4 weather conditions and random scenarios from the CARLA Leaderboard. The CARLA public Leaderboard 1.0 [45] consists of 100 secret routes in unseen towns under adversarial conditions, following the NHTSA pre-crash typology [35, 45].

Implementation Details: We build our approach, named **DTCP** (Diversity-enhanced TCP), upon the TCP baseline [50] and follow its supplementary material for hyperparameter settings and baseline loss weighting factors λ. Details regarding our method's hyperparameters such as λ_{div} can be found in our supplementary material. We set the future time-steps $T = 4$ in Eq. (9), consistent with the choice in [50] for Eq. (4). Additionally, we reproduce TransFuser [38] as reported in [39] for our ablation study. This earliest version is computationally lighter than later versions, including TransFuser [13] and TF++ WP [24].

Model Validation: Optimal model selection requires validation metrics that comprehensively capture diverse driving demands, a challenge often overlooked in CARLA benchmarks. Methods like Transfuser [38] and TCP [50] prioritize trajectory accuracy but fail to assess robustness in complex scenarios (e.g., traffic avoidance, road sign compliance). To address this, we evaluate our top three checkpoints directly in CARLA on the validation routes and select the best-performing model for benchmarking.

4.2. Ablation Study

Loss Component Analysis: We report our results, which are the mean and standard deviation (std) of five evaluation runs, on LAV benchmark [8] with two different baselines demonstrating the effectiveness of feature diversity in driving performance shown in Tab. 1.

Table 1. **Results on LAV Benchmarks**. We compare the integration of our proposed losses within two baselines, TCP and TransFuser, demonstrating enhanced driving performance and infraction score. Results are averaged over five test runs. The value band *Expert* indicates an upper bound for our IL agent.

Method	Driving score %, ↑	Route completion %, ↑	Infraction penalty [0, 1], ↑
Roach Expert [58]	77.96 ± 1.12	98.35 ± 1.73	0.79 ± 0.01
TCP [50]	53.10 ± 2.98	88.39 ± 3.14	0.60 ± 0.02
DTCP (**ours**)	**61.87** ± 1.35	**98.35** ± 1.67	**0.63** ± 0.01
Rule-based Expert	68.52 ± 1.15	97.98 ± 0.31	0.69 ± 0.01
TransFuser [38]	19.71 ± 0.79	48.09 ± 4.98	0.56 ± 0.06
DTCP (**ours**)	**29.60** ± 0.87	**55.84** ± 6.80	**0.58** ± 0.08

Using our diversity loss in TCP significantly improves DS by 8%, RC by 10%, and IP by 0.03. We observe a reduction in nearly all infractions as depicted in Tab. 2. Specifically, it reduces both red light violation and layout collisions by a factor of 4.

Applying our approach to the TransFuser baseline increases the route completion by 7% and driving score by 10%. Our model maintains a better trade-off between safety and driving efficiency, a balance desired in every au-

Table 2. **Infraction Details on LAV Benchmark**. We present a more detailed analysis of infractions across five runs. Our method significantly reduces nearly all types of collisions, matching expert-level performance in some categories.

Benchmark	Method	Collisions pedestrians #km, ↓	Collisions vehicles #km, ↓	Collisions layout #km, ↓	Red light infractions #km, ↓	Stop sign infractions #km, ↓	Off-road infractions #km, ↓	Route deviations #km, ↓	Route timeouts #km, ↓	Agent blocked #km, ↓
LAV Routes [8]	*Roach Expert* [58]	0.00	0.08	0.00	0.00	0.00	0.00	0.00	0.02	0.01
	TCP [50]	0.00	0.09	0.04	0.04	**0.09**	0.05	0.00	0.06	0.06
	DTCP (**ours**)	0.00	**0.07**	**0.01**	**0.01**	0.14	**0.02**	0.00	**0.02**	**0.01**

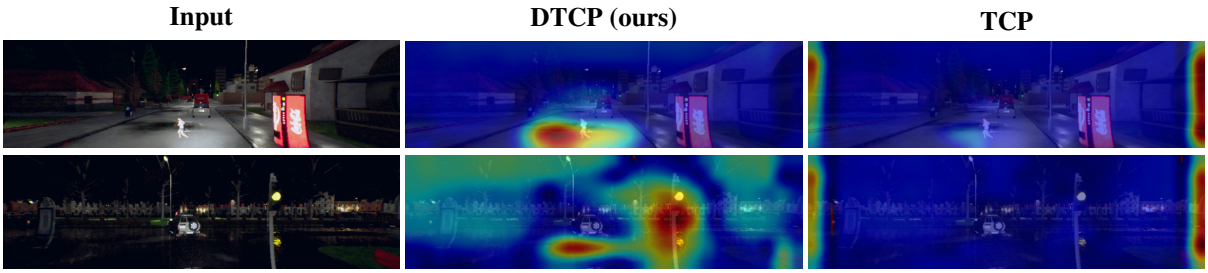

Figure 2. **EigenCam Visualizations in Various Challenging Scenarios**. From left to right: original image, DTCP (ours), and reproduced TCP. TCP shows a uniform feature distribution, primarily focusing on the left and right due to crossing traffic participants. In contrast, our method captures more diverse feature representations, enhancing focus on regions critical for driving decisions.

Table 3. **Quantitative Evaluation of Interpretability for Infraction-Related Categories**. We report the mean for each metrics. Our method achieves the highest GTC, indicating that it allocates greater attention to objects of interest compared to the reproduced TCP. This result demonstrates that incorporating the diversity loss helps to meaningfully localize the feature map.

Routes	Method	Pedestrians			Cyclists			Vehicles			Traffic Lights			Overall		
		IoU	GTC	SC	IoU	GTC	SC	IoU	GTC	SC	IoU	GTC	SC	IoU	GTC	SC
Town01	TCP [50]	0.01	0.07	0.01	0.01	0.12	0.02	0.01	0.06	0.01	**0.08**	0.23	**0.09**	0.02	0.09	0.02
	DTCP (**ours**)	0.01	**0.60**	0.01	**0.02**	**0.44**	0.02	0.01	**0.25**	**0.02**	0.05	**0.35**	0.06	0.02	**0.30**	0.02
Town02	TCP [50]	0.00	0.00	0.00	0.01	0.11	0.01	0.01	0.04	0.01	**0.07**	0.26	**0.08**	0.02	0.08	0.02
	DTCP (**ours**)	**0.01**	**0.45**	**0.01**	**0.03**	**0.55**	**0.03**	**0.02**	**0.33**	**0.02**	0.05	**0.28**	0.06	0.02	**0.33**	0.02
Town03	TCP [50]	0.00	0.00	0.00	0.00	0.02	0.00	0.01	0.09	0.01	0.00	0.00	0.00	0.01	0.06	0.01
	DTCP (**ours**)	**0.01**	**0.87**	**0.01**	**0.01**	**0.40**	0.00	0.01	**0.19**	0.01	0.00	**0.70**	0.00	0.01	**0.28**	0.01
Town05	TCP [50]	0.00	0.02	0.00	0.00	0.00	0.00	0.01	0.07	0.01	0.00	0.00	0.00	0.01	0.06	0.01
	DTCP (**ours**)	**0.01**	**0.59**	**0.01**	0.00	**0.07**	0.00	0.01	**0.27**	0.01	0.00	**0.19**	0.00	0.01	**0.27**	0.01

tonomous driving model.

Interpretability: We visualize the last layer of the image encoder for both the reproduced TCP and our model in Fig. 2 using EigenCam [34], which projects latent features onto eigenvectors to highlight critical regions. Due to incorporating the feature diversity loss ($\mathcal{L}_{\text{div}}$), interpretability improves compared to the reproduced version. Specifically, our model focuses on the crossing pedestrian in the first row, while also considering the risk of yellow traffic lights (last row) [50].

We further validate the interpretability improvement using *Shared Interest* metrics (IoU, GTC, SC) [3], comparing model-generated saliency maps to ground-truth bounding boxes of high-risk infraction categories (details in Supplementary). IoU measures the intersection of the ground truth set ($\mathcal{G}$) and the saliency set ($\mathcal{S}$) relative to their union ($\frac{I}{U}$). GTC indicates the proportion of ground truth features covered by the saliency set ($\frac{I}{\mathcal{G}}$), and SC denotes the proportion of saliency features matching the ground truth ($\frac{I}{\mathcal{S}}$). As shown in Tab. 3, our model achieves substantially higher GTC across all categories than the reproduced TCP, indicating a stronger focus on ground-truth features relevant for decision-making [3, 32]. Comparable IoU and SC scores between models are due to our model's reliance on diverse features beyond ground truth, and targets in each frame constitute a minor portion of the activation map. Additionally, not all bounding boxes directly impact decision-making, and partial overlaps can reduce metric scores despite appropriate saliency allocation (see Supplementary Material for visual examples). Our analysis shows that improved interpretability correlates with lower infraction rate, as prioritizing infraction-related objects enhances driving safety and performance, as corroborated by Tab. 1. In our Supplementary Material, we additionally compute driving-related semantic IoU to demonstrate that DTCP assigns distinct attentions to semantically coherent regions, thereby supporting human-level interpretability.

Tab. 4 reports the Pearson correlation between saliency magnitude and steering signals derived from control actions $\mathbf{a}^{traj,ctrl}$. DTCP exhibits positive correlations, whereas the

Table 4. **Saliency–Steering Pearson Correlation on LAV Routes.** ρ denotes Pearson's correlation between saliency mass and steering; $R^2{=}\rho^2_{\text{DTCP}}/\rho^2_{\text{TCP}}$ is the relative explained variance; z is Fisher's test statistic comparing correlations. DTCP correlations are *positive*, whereas TCP correlations are *negative*, indicating our model's attention is directed toward driving-relevant cues.

Control Signal	ρ_{DTCP}	ρ_{TCP}	R^2	z
steer-ctrl	$+0.122$	-0.223	0.30	7.84
steer-traj	$+0.186$	-0.060	9.76	5.57

reproduced TCP shows negative correlations, resulting in substantially higher explained variance (R^2: $\sim$10$\times$ *traj*, 3$\times$ *ctrl*) and decisive Fisher z-scores (z=5.6, *traj*; 7.8, *ctrl*). This further demonstrates that DTCP's enhanced focus on driving-critical cues improves driving performance (cf. Tab. 1, Tab. 5).

Limitations: The results in Tab. 2 indicate that our method still struggles with stop signs. The reproduced TCP experiences route timeouts 3$\times$ and agent blocks 6$\times$ more often than DTCP, reflecting slower and hesitant driving. Stop sign infractions depend on slowing down to 2 m/s and stopping within the simulator-defined trigger area [24]. Hence, the slower reproduced TCP meets this criterion more frequently than the faster DTCP, which achieves higher route completion. This discrepancy arises because the LAV validation routes predominantly feature road-painted 'STOP' markings rather than upright signs. While upright signs can be detected earlier, painted signs become visible for only a brief moment, making braking success primarily speed-dependent rather than perception-driven. Thus, reproduced TCP's better performance here is incidental rather than a result of superior perception. For distant, upright signs, DTCP's diversified features provide better detection, explaining the reversed stop-sign infraction results observed between Tab. 2 and Tab. 6. Stop-sign infractions in our model can be substantially reduced by incorporating bounding-box detection, ensuring the controller initiates braking as soon as the vehicle enters the boundary of a painted stop sign, as suggested in [24]. We further discuss failure cases in our Supplementary Material, illustrating particular limitations of using the *Shared Interest* metrics (see Tab. 3).

4.3. State-of-the-Art Comparison

We evaluate our approach on the public CARLA Leaderboard 1.0 [45], based on CARLA version 0.9.10.1. Although the new version of leaderboard, Leaderboard 2.0, which is based on the CARLA version 0.9.14, is already available, we have decided to use Leaderboard 1.0 as it is more suitable for benchmarking due to its 31 submissions versus only 3 submissions in new version. We will update our results on the new leaderboard in future works.

CARLA Leaderboard 1.0: Table 5 presents the results of our approach, denoted as DTCP, compared to the five top-ranked entries on the public CARLA Leaderboard 1.0 [45]. For a fair comparison, we also submitted a reproduced version of the TCP model. Conversely, the TCP model submitted by Wu *et al*. [50] comprises an ensemble model combination of two models TCP and TCP-SB. TCP-SB incorporates two distinct encoders for both branches instead of utilizing shared ones with the trajectory unit supervision (cf. Fig. 1). Our model secured 4th place among all 31 leaderboard submissions with the highest route completion. It outperforms reproduced TCP by 12% on both DS and RC with an improvement of 0.06 in IP. This demonstrates that despite utilizing only a monocular camera, our model exhibits enhanced object detection ability owing to its increased interpretability, although all other models using multi camera-views and a LiDAR (see Tab. 5) are presumed superior in this task, as explained in [50]. Although TCP reports a DS of 69 in [50], reproductions vary notably: Jaeger *et al*. [24] achieve 48 DS for their reproduced TCP, whereas our reproduction attains 58 DS. The CARLA leaderboard ensures repeatability via docker-based evaluations but does not guarantee reproducibility, as publicly released code and models can differ from the officially submitted models. These discrepancies align with experimental findings by Jaeger *et al*. (cf. Table 17 in [24]), indicating that implementation details not explicitly highlighted in the original papers significantly influence leaderboard outcomes.

Despite not employing any ensemble model like TCP-Ens and TF++ WP or specific tasks for red light and stop sign detection like those found in ReasonNet and InterFuser, we recorded the second-lowest red light infraction rate as shown in Tab. 6. When compared to LAV [8], our approach yields almost the same RC with a 0.11 higher IP, without using any multi-stage training with many pre-trained modular components and the distillation framework utilized by LAV. Our agent experiences the lowest blocking rate without adopting any creeping strategy employed by TF++ WP [24] to address the inertia problem observed in the IL agent [13, 15], which tends to cause a stilled vehicle to remain inactive. Interestingly, in contrast to Tab. 2, DTCP achieves nearly a 3$\times$ lower stop-sign infraction rate compared to the reproduced version. This highlights visual detection as the key bottleneck for distant, upright signs, where DTCP benefits significantly from its diversified features.

We also study the difference between the top six leaderboard submissions regarding the number of used sensors, model art, the inclusion of traffic rule auxiliary tasks, and the runtime duration. Both our model and TCP Ensemble employ only a front-view camera, whereas all other leading state-of-the-arts in Tab. 5 use a LiDAR together with either multi-view cameras or a high-resolution monocular camera

Table 5. **Performance Comparison on the CARLA Public Leaderboard 1.0** [45] (accessed in Jun. 2025). We compare our approach (DTCP) with 6 top-ranked leaderboard submissions. Our model surpasses its reproduced baseline (TCP Reproduced) across all metrics and achieves the highest route completion among all competitors. Despite using only a monocular camera, without ensemble techniques or traffic rules sub-tasks, it has attained state-of-the-art performance while requiring the lowest execution time.

Rank	Method	Sensor Inputs		Ensemble	Auxiliary tasks traffic rules	#Execution time	Driving score %, ↑	Route completion %, ↑	Infraction penalty [0, 1], ↑
		#Cameras	LiDAR						
1	ReasonNet [43]	5	✓	✗	✓	∼ 117 h	**79.95**	89.89	**0.89**
2	InterFuser [42]	4	✓	✗	✓	∼ 80 h	76.18	88.23	0.84
3	TCP Ensemble [50]	1	✗	✓	✗	∼ 80 h	75.14	85.63	0.87
4	DTCP (**ours**)	1	✗	✗	✗	∼ 50 h	70.74	**95.60**	0.75
5	TF++ WP Ensemble [24]	1	✓	✓	✓	∼ 187 h	66.32	78.66	0.84
6	LAV [8]	4	✓	✗	✓	∼ 394 h	61.85	94.46	0.64
-	TCP Reproduced (**our results**)	1	✗	✗	✗	∼ 50 h	58.56	83.14	0.69

Table 6. **Infraction Analysis on CARLA Public Leaderboard 1.0** (accessed in Jun. 2025). We compare our approach (DTCP) against the top 6 baselines regarding infractions incurred. Our model records the second-lowest red light violation rate, achieving this without leveraging traffic-rule-specific auxiliary tasks. With minimal occurrences of agent blockage, off-road incidents, and deviations from the route, our agent demonstrates superior route completion rates, as reported in Tab. 5.

Rank	Method	Collisions pedestrians #km, ↓	Collisions vehicles #km, ↓	Collisions layout #km, ↓	Red light infractions #km, ↓	Stop sign infractions #km, ↓	Off-road infractions #km, ↓	Route deviations #km, ↓	Route timeouts #km, ↓	Agent blocked #km, ↓
1	ReasonNet [43]	0.02	0.13	0.01	0.08	0.00	0.04	0.00	0.01	0.33
2	InterFuser [42]	0.04	0.37	0.14	0.22	0.00	0.13	0.00	0.01	0.43
3	TCP Ensemble [50]	0.00	0.32	0.00	0.09	0.00	0.04	0.00	0.00	0.54
4	DTCP (**ours**)	0.00	0.14	0.06	0.05	0.17	0.04	0.00	0.08	0.10
5	TF++ WP Ensemble [24]	0.00	0.50	0.00	0.01	0.00	0.12	0.00	0.00	0.71
6	LAV [8]	0.04	0.70	0.02	0.17	0.00	0.25	0.09	0.04	0.10
-	TCP Reproduced (**our results**)	0.00	4.64	0.17	0.13	0.48	0.00	0.00	0.11	0.30

with a wide field of view (FOV), as in TF++ WP [24]. In contrast to our model and TCP Ensemble [50], both ReasonNet [43] and Interfuser [42] utilize a traffic sign classifier to predict the traffic light state, determine if the ego agent is at an intersection, and identify an upcoming stop sign. TF ++ WP adopts many auxiliary tasks such as 2D depth, 2D and BEV semantic segmentation, and bounding box detection. For BEV semantic segmentation, TF++ WP predicts 11 classes, while for bounding box predictions, it distinguishes between 4 classes, including traffic light states and stop signs in both tasks. In addition to predicting a BEV semantic map, LAV [8] also leverages a brake predictor to handle traffic sign and hazardous stoppages using 3 camera-view images and an focusing-view (telephoto lens) image. Similar to LAV, ReasonNet and Interfuser also utilize a focus-view image as a strategy to capture distant traffic lights, as the small objects like traffic lights comprise only a few pixels in the whole input image, and could easily be overlooked in the prediction [12]. Among all top models, our approach is the only one that does not rely on any ensemble or distillation framework, yet obtains competitive driving performance. Furthermore, our model requires 2.3× less runtime compared to ReasonNet and thus achieves the lowest execution time by a significant margin from other competitors. Based on the results reported in [50], we believe that using an ensemble model could potentially boost the driving performance of our model further.

5. Conclusion

In this research, we propose an interpretable end-to-end autonomous driving method by diversifying the feature representations. Our approach matches or outperforms state-of-the-art results on various benchmarks, particularly in reducing traffic infractions and achieving high route completion, while eschewing specialized auxiliary tasks. We demonstrate our model's enhanced adeptness at prioritizing critical safety considerations, thereby effectively reducing traffic violations and enhancing overall driving safety. Our findings illustrate the crucial role of feature diversity in bolstering the decision-making capabilities of autonomous driving agents, leading to a more streamlined and interpretable model. We show that improving interpretability is closely related to reducing infractions and, consequently, boosting driving performance. Notably, our proposed method has achieved state-of-the-art performance with minimized computational complexity on the CARLA public Leaderboard without the aid of ensemble techniques or extensive sensory inputs, which highlights its efficiency and scalability. As future work, we plan to integrate an ensemble model and evaluate its performance on the new CARLA Leaderboard.

References

[1] Mayank Bansal, Alex Krizhevsky, and Abhijit Ogale. Chauffeurnet: Learning to drive by imitating the best and synthesizing the worst. *arXiv preprint arXiv:1812.03079*, 2018. 2

[2] Richard Bellman and Robert Kalaba. On adaptive control processes. *IRE Transactions on Automatic Control*, 4(2):1–9, 1959. 3

[3] Angie Boggust, Benjamin Hoover, Arvind Satyanarayan, and Hendrik Strobelt. Shared interest: Measuring human-ai alignment to identify recurring patterns in model behavior. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pages 1–17, 2022. 2, 6

[4] Moritz Böhle, Mario Fritz, and Bernt Schiele. B-cos networks: Alignment is all we need for interpretability. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10329–10338, 2022. 2

[5] Dongliang Chang, Yifeng Ding, Jiyang Xie, Ayan Kumar Bhunia, Xiaoxu Li, Zhanyu Ma, Ming Wu, Jun Guo, and Yi-Zhe Song. The devil is in the channels: Mutual-channel loss for fine-grained image classification. *IEEE Transactions on Image Processing*, 29:4683–4695, 2020. 2

[6] Raphael Chekroun, Marin Toromanoff, Sascha Hornauer, and Fabien Moutarde. Gri: General reinforced imitation and its application to vision-based autonomous driving. *Robotics*, 12(5):127, 2023. 2

[7] Chaofan Chen, Oscar Li, Daniel Tao, Alina Barnett, Cynthia Rudin, and Jonathan K Su. This looks like that: deep learning for interpretable image recognition. *Advances in neural information processing systems*, 32, 2019. 2

[8] Dian Chen and Philipp Krähenbühl. Learning from all vehicles. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17222–17231, 2022. 2, 5, 6, 7, 8, 3

[9] Dian Chen, Brady Zhou, Vladlen Koltun, and Philipp Krähenbühl. Learning by cheating. In *Conference on Robot Learning*, pages 66–75. PMLR, 2020. 3

[10] Jianyu Chen, Shengbo Eben Li, and Masayoshi Tomizuka. Interpretable end-to-end urban autonomous driving with latent deep reinforcement learning. *IEEE Transactions on Intelligent Transportation Systems*, 23(6):5068–5078, 2021. 1

[11] Long Chen, Yuchen Li, Chao Huang, Bai Li, Yang Xing, Daxin Tian, Li Li, Zhongxu Hu, Xiaoxiang Na, Zixuan Li, et al. Milestones in autonomous driving and intelligent vehicles: Survey of surveys. *IEEE Transactions on Intelligent Vehicles*, 8(2):1046–1056, 2022. 1

[12] Li Chen, Penghao Wu, Kashyap Chitta, Bernhard Jaeger, Andreas Geiger, and Hongyang Li. End-to-end autonomous driving: Challenges and frontiers. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024. 1, 2, 8

[13] Kashyap Chitta, Aditya Prakash, Bernhard Jaeger, Zehao Yu, Katrin Renz, and Andreas Geiger. Transfuser: Imitation with transformer-based sensor fusion for autonomous driving. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 2, 3, 5, 7

[14] Kyunghyun Cho, B van Merrienboer, Caglar Gulcehre, F Bougares, H Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. In *Conference on Empirical Methods in Natural Language Processing (EMNLP 2014)*, 2014. 3

[15] Felipe Codevilla, Eder Santana, Antonio M López, and Adrien Gaidon. Exploring the limitations of behavior cloning for autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9329–9338, 2019. 3, 7

[16] Nemanja Djuric, Vladan Radosavljevic, Henggang Cui, Thi Nguyen, Fang-Chieh Chou, Tsung-Han Lin, Nitin Singh, and Jeff Schneider. Uncertainty-aware short-term motion prediction of traffic actors for autonomous driving. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2095–2104, 2020. 2

[17] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. Carla: An open urban driving simulator. In *Conference on robot learning*, pages 1–16. PMLR, 2017. 1, 4, 3

[18] Paul Drews, Grady Williams, Brian Goldfain, Evangelos A Theodorou, and James M Rehg. Aggressive deep driving: Model predictive control with a cnn cost model. *arXiv preprint arXiv:1707.05303*, 2017. 2

[19] Paolo Falcone, Francesco Borrelli, Jahan Asgari, Hongtei Eric Tseng, and Davor Hrovat. Predictive active steering control for autonomous vehicle systems. *IEEE Transactions on control systems technology*, 15(3):566–580, 2007. 2

[20] Jeffrey Hawke, Richard Shen, Corina Gurau, Siddharth Sharma, Daniele Reda, Nikolay Nikolov, Przemysław Mazur, Sean Micklethwaite, Nicolas Griffiths, Amar Shah, et al. Urban driving with conditional imitation learning. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 251–257. IEEE, 2020. 2

[21] Yihan Hu, Jiazhi Yang, Li Chen, Keyu Li, Chonghao Sima, Xizhou Zhu, Siqi Chai, Senyao Du, Tianwei Lin, Wenhai Wang, et al. Planning-oriented autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17853–17862, 2023. 2

[22] Zhiyu Huang, Chen Lv, Yang Xing, and Jingda Wu. Multi-modal sensor fusion-based deep neural network for end-to-end autonomous driving with scene understanding. *IEEE Sensors Journal*, 21(10):11781–11790, 2020. 2

[23] Keishi Ishihara, Anssi Kanervisto, Jun Miura, and Ville Hautamaki. Multi-task learning with attention for end-to-end autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2902–2911, 2021. 2

[24] Bernhard Jaeger, Kashyap Chitta, and Andreas Geiger. Hidden biases of end-to-end driving models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8240–8249, 2023. 5, 7, 8

[25] Xiaosong Jia, Yulu Gao, Li Chen, Junchi Yan, Patrick Langechuan Liu, and Hongyang Li. Driveadapter: Breaking the coupling barrier of perception and planning in end-to-end autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7953–7963, 2023. 3

[26] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21983–21994, 2023. 3

[27] Jinkyu Kim and Mayank Bansal. Attentional bottleneck: Towards an interpretable deep driving network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 322–323, 2020. 2

[28] Jinkyu Kim and John Canny. Interpretable learning for self-driving cars by visualizing causal attention. In *Proceedings of the IEEE international conference on computer vision*, pages 2942–2950, 2017. 1, 2

[29] Anton Kuznietsov, Balint Gyevnar, Cheng Wang, Steven Peters, and Stefano V Albrecht. Explainable ai for safe and trustworthy autonomous driving: A systematic review. *IEEE Transactions on Intelligent Transportation Systems*, 2024. 1

[30] Jesse Levinson, Jake Askeland, Jan Becker, Jennifer Dolson, David Held, Soeren Kammel, J Zico Kolter, Dirk Langer, Oliver Pink, Vaughan Pratt, et al. Towards fully autonomous driving: Systems and algorithms. In *2011 IEEE intelligent vehicles symposium (IV)*, pages 163–168. IEEE, 2011. 1

[31] Ashish Mehta, Adithya Subramanian, and Anbumani Subramanian. Learning end-to-end autonomous driving using guided auxiliary supervision. In *Proceedings of the 11th Indian Conference on Computer Vision, Graphics and Image Processing*, pages 1–8, 2018. 2

[32] Katelyn Morrison, Ankita Mehra, and Adam Perer. Shared interest... sometimes: Understanding the alignment between human perception, vision architectures, and saliency map techniques. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3776–3781, 2023. 6

[33] Matthias Mueller, Alexey Dosovitskiy, Bernard Ghanem, and Vladlen Koltun. Driving policy transfer via modularity and abstraction. In *Conference on Robot Learning*, pages 1–15. PMLR, 2018. 1

[34] Mohammed Bany Muhammad and Mohammed Yeasin. Eigen-cam: Class activation map using principal components. In *2020 international joint conference on neural networks (IJCNN)*, pages 1–7. IEEE, 2020. 6, 1

[35] Wassim G Najm, John D Smith, Mikio Yanagisawa, et al. Pre-crash scenario typology for crash avoidance research. Technical report, United States. National Highway Traffic Safety Administration, 2007. 5

[36] Oskar Natan and Jun Miura. End-to-end autonomous driving with semantic depth cloud mapping and multi-agent. *IEEE Transactions on Intelligent Vehicles*, 8(1):557–571, 2022. 2

[37] Thomas Norrenbrock, Marco Rudolph, and Bodo Rosenhahn. Take 5: Interpretable image classification with a handful of features. In *Workshop on Trustworthy and Socially Responsible Machine Learning, NeurIPS 2022*, 2022. 2, 4

[38] Aditya Prakash, Kashyap Chitta, and Andreas Geiger. Multi-modal fusion transformer for end-to-end autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7077–7087, 2021. 3, 5

[39] Aditya Prakash, Kashyap Chitta, and Andreas Geiger. Supplementary material for multi-modal fusion transformer for end-to-end autonomous driving. 2021. 5

[40] Aditya Prakash, Kashyap Chitta, and Andreas Geiger. Transfuser open dataset. https://github.com/autonomousvision/transfuser/tree/cvpr2021/, 2021. Accessed: 2025-05-12. 5

[41] Katrin Renz, Kashyap Chitta, Otniel-Bogdan Mercea, A Sophia Koepke, Zeynep Akata, and Andreas Geiger. Plant: Explainable planning transformers via object-level representations. In *Conference on Robot Learning*, pages 459–470. PMLR, 2023. 2, 3

[42] Hao Shao, Letian Wang, Ruobing Chen, Hongsheng Li, and Yu Liu. Safety-enhanced autonomous driving using interpretable sensor fusion transformer. In *Conference on Robot Learning*, pages 726–737. PMLR, 2023. 2, 5, 8

[43] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13723–13733, 2023. 2, 5, 8

[44] Ardi Tampuu, Tambet Matiisen, Maksym Semikin, Dmytro Fishman, and Naveed Muhammad. A survey of end-to-end driving: Architectures and training methods. *IEEE Transactions on Neural Networks and Learning Systems*, 33(4): 1364–1384, 2020. 1

[45] CARLA team. Carla autonomous driving leaderboard. https://leaderboard.carla.org/, 2022. Accessed: 2025-05-12. 1, 3, 5, 7, 8

[46] Marin Toromanoff, Emilie Wirbel, and Fabien Moutarde. End-to-end model-free reinforcement learning for urban driving using implicit affordances. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7153–7162, 2020. 2

[47] Bowen Wang, Liangzhi Li, Yuta Nakashima, and Hajime Nagahara. Learning bottleneck concepts in image classification. In *Proceedings of the ieee/cvf conference on computer vision and pattern recognition*, pages 10962–10971, 2023. 2

[48] Dequan Wang, Coline Devin, Qi-Zhi Cai, Fisher Yu, and Trevor Darrell. Deep object-centric policies for autonomous driving. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 8853–8859. IEEE, 2019. 2

[49] Guan Wang, Haoyi Niu, Desheng Zhu, Jianming Hu, Xianyuan Zhan, and Guyue Zhou. A versatile and efficient reinforcement learning framework for autonomous driving. *arXiv preprint arXiv:2110.11573*, 2021. 1

[50] Penghao 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. *Advances in Neural Information Processing Systems*, 35:6119–6132, 2022. 2, 3, 5, 6, 7, 8, 1

[51] Yi Xiao, Felipe Codevilla, Akhil Gurram, Onay Urfalioglu, and Antonio M López. Multimodal end-to-end autonomous driving. *IEEE Transactions on Intelligent Transportation Systems*, 23(1):537–547, 2020. 1, 2

[52] Zhenhua Xu, Yujia Zhang, Enze Xie, Zhen Zhao, Yong Guo, Kwan-Yee K Wong, Zhenguo Li, and Hengshuang Zhao.

Drivegpt4: Interpretable end-to-end autonomous driving via large language model. *IEEE Robotics and Automation Letters*, 2024. 1

[53] Ekim Yurtsever, Jacob Lambert, Alexander Carballo, and Kazuya Takeda. A survey of autonomous driving: Common practices and emerging technologies. *IEEE access*, 8:58443–58469, 2020. 1

[54] Éloi Zablocki, Hédi Ben-Younes, Patrick Pérez, and Matthieu Cord. Explainability of deep vision-based autonomous driving systems: Review and challenges. *International Journal of Computer Vision*, 130(10):2425–2452, 2022. 1, 2

[55] Wenyuan Zeng, Wenjie Luo, Simon Suo, Abbas Sadat, Bin Yang, Sergio Casas, and Raquel Urtasun. End-to-end interpretable neural motion planner. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8660–8669, 2019. 1, 2

[56] Jimuyang Zhang, Zanming Huang, Arijit Ray, and Eshed Ohn-Bar. Feedback-guided autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15000–15011, 2024. 3

[57] Quan-shi Zhang and Song-Chun Zhu. Visual interpretability for deep learning: a survey. *Frontiers of Information Technology & Electronic Engineering*, 19(1):27–39, 2018. 2

[58] Zhejun Zhang, Alexander Liniger, Dengxin Dai, Fisher Yu, and Luc Van Gool. End-to-end urban driving by imitating a reinforcement learning coach. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 15222–15232, 2021. 3, 5, 6

[59] Albert Zhao, Tong He, Yitao Liang, Haibin Huang, Guy Van den Broeck, and Stefano Soatto. Sam: Squeeze-and-mimic networks for conditional visual driving policy learning. In *Conference on Robot Learning*, pages 156–175. PMLR, 2021. 3

[60] Bolei Zhou, David Bau, Aude Oliva, and Antonio Torralba. Comparing the interpretability of deep networks via network dissection. *Explainable AI: Interpreting, explaining and visualizing deep learning*, pages 243–252, 2019. 2

Interpretable Decision-Making for End-to-End Autonomous Driving

Supplementary Material

Mona Mirzaie Bodo Rosenhahn

Institute for Information Processing
Leibniz University Hannover, Germany

mirzaie@tnt.uni-hannover.de

In this supplementary document, we first analyze failure cases in challenging scenarios and discuss the underlying causes (Sec. 6). Next, we present a detailed analysis of our model's behavior through visualizations, including activation maps that highlight key decision-making regions for critical traffic participants (Sec. 7). We then describe how interpretability metrics and mask generation are applied in our evaluations (Sec. 8), and subsequently provide a quantitative analysis of concept-level interpretability (Sec. 9). Finally, we conduct an ablation study to assess the impact of hyperparameter choices on model performance (Sec. 10).

6. Failure Case Analysis

This section presents examples from Fig. 3 demonstrating scenarios that lead to lower Intersection over Union (IoU) and Ground Truth Coverage (GTC) scores. Although our model successfully identifies relevant regions, the activation maps do not always fully overlap with ground-truth bounding boxes, resulting in lower scores. This partial coverage results in lower metric values, making it challenging to fully capture the interpretability advantages of our model based solely on these metrics.

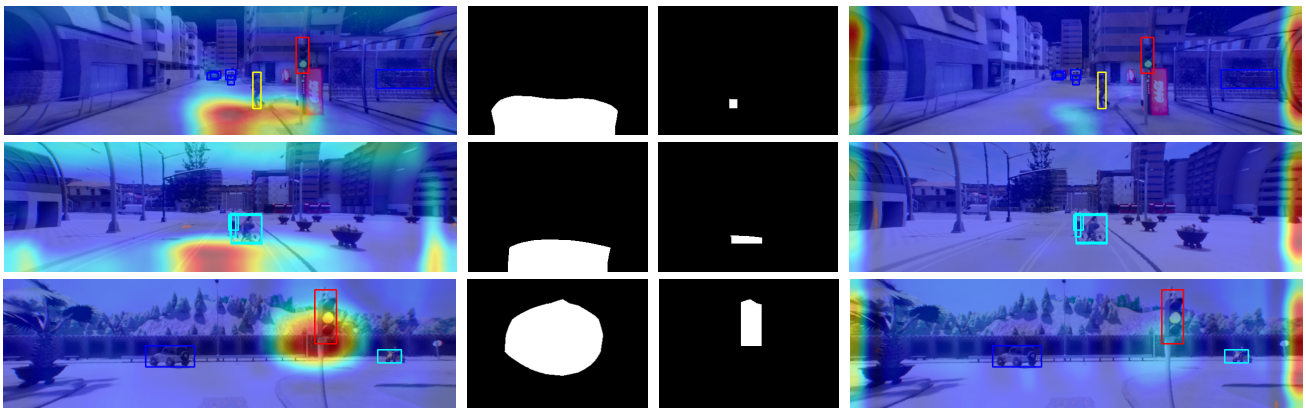

Figure 3. **Example of Failure Cases**. From left to right: heatmap of our model, heatmap binary mask, intersection of the bounding box and heatmap binary mask, and heatmap of the reproduced TCP.

7. Comparative Analysis of Activation Maps

In this section, we present additional qualitative analyses of activation maps generated using EigenCam [34] for our proposed model DTCP compared to the baseline TCP [50]. As depicted in Fig. 4 **(a-g)**, DTCP effectively localizes attention on regions critical for driving decisions, significantly enhancing interpretability. Detecting traffic lights, especially those situated at higher and farther positions, is challenging due to their relatively small size in camera images. However, as illustrated in Fig. 4 **(a)**, DTCP effectively addresses this difficulty by accurately focusing its activation map within the bounding box of distant traffic lights. Fig. 4 **(b-c)** illustrates that our model assigns significant attention to yellow traffic lights. Specifically, in **(c)**, upon recognizing the yellow traffic light, the model effectively identifies vehicles within the intersection in our driving path as potential collision risks. As depicted in Fig. 4 **(d)**, our model carefully monitors the surrounding environment during turning maneuvers, proactively identifying potential risks from dynamic objects, such as cyclists (marked by a cyan bounding box) or pedestrians, that could unexpectedly enter our driving path. Fig. 4 **(e-g)** demonstrates that our model remains highly

vigilant in detecting unexpected situations, identifying hazards such as vehicles and pedestrians that violate traffic rules or run red lights, thereby reducing potential collision risks.

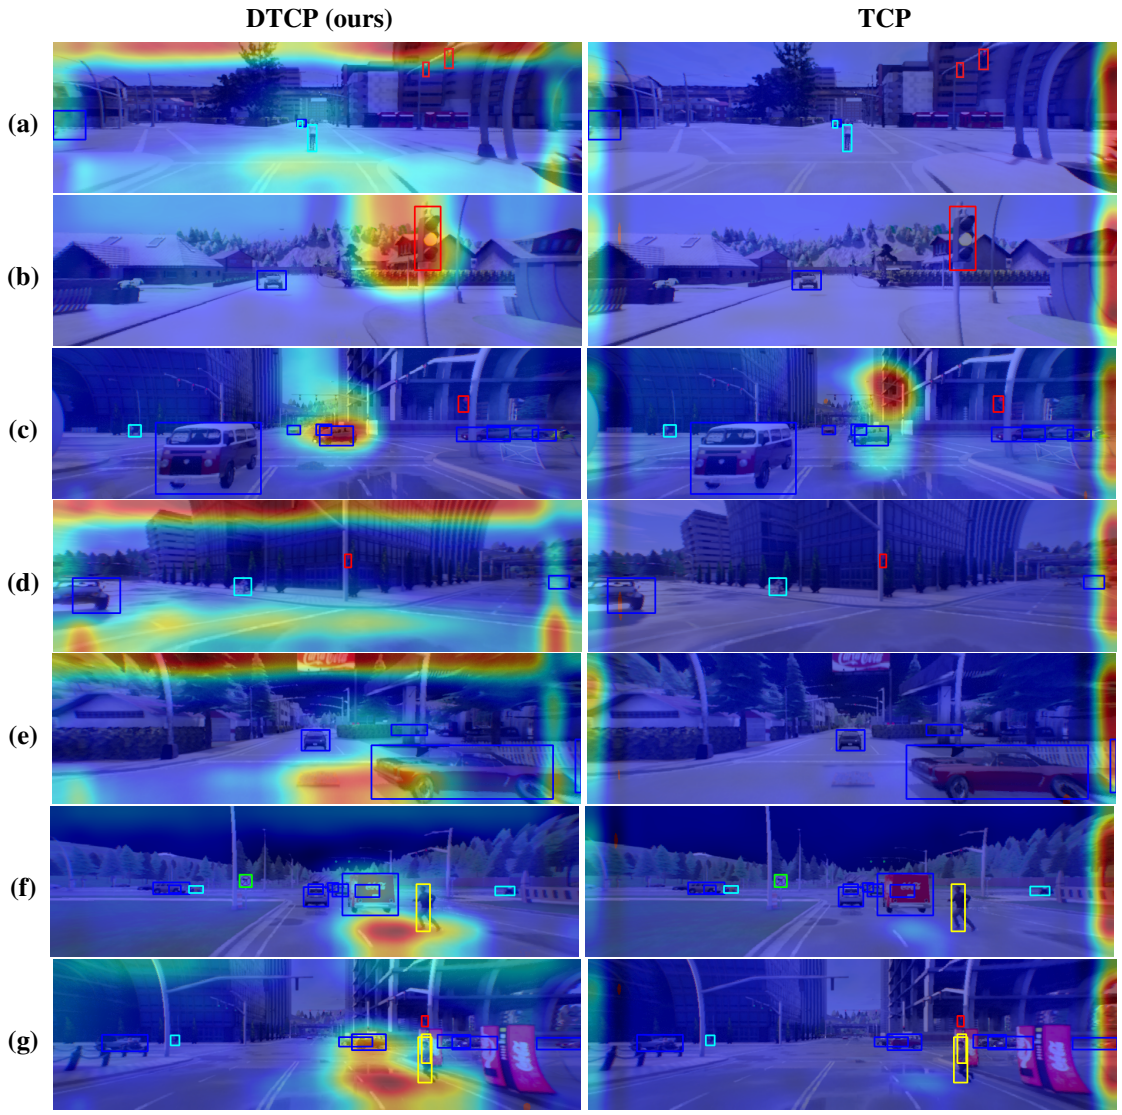

Figure 4. **Comparison of activation maps**. Left: Heatmap of our model DTCP; right: heatmap of the reproduced TCP. Our model achieves superior performance by activating key regions relevant to driving decisions, thus enhancing interpretability.

8. Mask Generation and Interpretability Metrics

For measuring interpretability using *Shared Interest* metrics (IoU, GTC, SC) [3], we generate binary masks from object bounding boxes and our model's heatmap, as illustrated in Fig. 5. Ground-truth bounding boxes are created for categories most likely involved in infractions—pedestrians, cyclists, vehicles, and traffic lights.
For a ground-truth set ($\mathcal{G}$) and a saliency set ($\mathcal{S}$), *Shared Interest* metrics are defined as:

$$\text{IoU} = \frac{|G \cap S|}{|G \cup S|} \tag{10}$$

$$\text{GTC} = \frac{|G \cap S|}{|G|} \tag{11}$$

Figure 5. **Example of Generated Binary Mask**. From left to right: heatmap with overlapping bounding boxes, binary mask derived from the heatmap, and binary mask of the bounding boxes.

$$\mathrm{SC} = \frac{|G \cap S|}{|S|} \tag{12}$$

9. Concept-Level Interpretability

To further demonstrate that DTCP attends more effectively to human-understandable concepts compared to the reproduced model, we compute driving-related semantic IoU, as reported in Tab. 7. Specifically, we measure IoU between model-generated saliency maps and key semantic classes: *road*, *roadline*, and *sidewalk*. The advantage of DTCP in allocating greater attention to dynamic traffic participants and traffic lights is already illustrated in Tab. 3 of our main paper. Results in Tab. 7 further confirm DTCP's superior ability to attend to driving-critical regions, reflecting human-like behavior essential for safe driving.

Table 7. **Interpretability evaluation for driving-related concepts.** We report the average IoU scores. Our method allocates greater attention to decision-critical areas in driving scenarios.

Routes	**Method**	**IoU ↑**		
		Roadline	**Road**	**Sidewalk**
Town01	TCP [50]	0.00	0.01	0.04
	DTCP (**ours**)	**0.01**	**0.30**	**0.21**
Town02	TCP [50]	0.00	0.02	0.03
	DTCP (**ours**)	**0.01**	**0.31**	**0.14**
Town03	TCP [50]	0.00	0.02	0.08
	DTCP (**ours**)	**0.02**	**0.41**	**0.21**
Town05	TCP [50]	0.00	0.02	0.06
	DTCP (**ours**)	**0.02**	**0.40**	**0.11**

10. Diversity Loss Weight Effect

Tab. 8 presents the impact of different loss weights (λ_{div}) for our proposed diversity loss, as introduced in Sec. 3.3 of the main paper. We report the mean and standard deviation of the driving score over three evaluation runs in the CARLA simulator [17] on the LAV benchmark [8]. As shown, setting $\lambda_{\text{div}} = 0.00005$ yields the best driving performance with the lowest standard deviation.

Table 8. **Effects of Diversity Loss Weight**. An empirical analysis to determine the optimal value of λ_{div} in our model on the LAV Routes benchmark [8].

λ_{div}	5×10^{-1}	5×10^{-2}	5×10^{-3}	5×10^{-4}	5×10^{-5}	5×10^{-6}
Driving Score	42.87 ± 3.25	55.58 ± 6.44	48.22 ± 4.01	44.01 ± 1.29	$\mathbf{60.42 \pm 1.27}$	52.65 ± 1.78