

Driving Like Yourself: A Benchmark for Closed-Loop Personalized End-to-End Autonomous Driving

Xiaoru Dong^{1,2}[†], Ruiqin Li²[†], Xiao Han²[†], Zhenxuan Wu², Jiamin Wang², Jian Chen¹, Qi Jiang², Siu Ming Yiu¹^{*}, Xinge Zhu³, and
Yuexin Ma²^{*}

¹ The University of Hong Kong

² ShanghaiTech University

³ The Chinese University of Hong Kong

xrdong@cs.hku.hk, mayuexin@shanghaitech.edu.cn

Abstract. Human driving behavior is inherently diverse, yet most end-to-end autonomous driving (E2E-AD) systems learn a single average driving style, neglecting individual differences. Achieving personalized E2E-AD faces challenges across three levels: limited real-world datasets with individual-level annotations, a lack of quantitative metrics for evaluating personal driving styles, and the absence of algorithms that can learn stylized representations from users’ trajectories. To address these gaps, we propose Person2Drive, a comprehensive personalized E2E-AD platform and benchmark. It includes an open-source, flexible data collection system that simulates realistic scenarios to generate scalable, diverse personalized driving datasets; style vector-based evaluation metrics with Maximum Mean Discrepancy and KL divergence to comprehensively quantify individual driving behaviors; and a personalized E2E-AD framework with a style reward model that efficiently adapts E2E models for safe and individualized driving. Crucially, our framework enables plug-and-play personalization by fine-tuning only the trajectory prediction head, preserving the pretrained base model and ensuring safety. Extensive experiments demonstrate that Person2Drive enables fine-grained analysis and effective personalization, while preserving driving performance and success rate even in challenging scenarios.

1 Introduction

Traditional autonomous driving systems rely on modular pipelines that separate perception, prediction, and control, often leading to suboptimal coordination, whereas recent end-to-end (E2E-AD) approaches use unified deep learning models to map sensor inputs directly to control outputs, achieving notable improvements in behavioral consistency and safety [1, 2, 7]. However, most E2E-AD

^{*} Corresponding authors.

[†] Equal contribution.

approaches employ a single model to fit all drivers, effectively learning an average driving style while neglecting individual differences. In reality, human driving behavior is inherently diverse—shaped by cognitive patterns, habits, and personality traits [24]. Overlooking such individuality limits our understanding of human decision-making and impedes the development of human-centered, trustworthy driving systems [13]. Personalized E2E-AD, by contrast, aims to adapt driving behavior to individual preferences and traits, thereby improving comfort, trust, and user acceptance [8, 11, 12].

To address personalization in autonomous driving, existing methods can be roughly divided into three categories. Early approaches introduce rule-based adjustments or trajectory-level tuning on modular pipelines [10], which are inherently incompatible with end-to-end learning. With the rise of E2E-AD, some studies condition models on discrete driving-style labels (e.g., aggressive, normal, conservative) [9]. However, human driving behaviors vary continuously across scenarios, making such coarse categorization insufficient for capturing individual dynamics. More recently, language-conditioned control has been explored for interpretable human–vehicle interaction [4, 27, 32], yet language signals alone cannot represent fine-grained, continuous control characteristics. This calls for a new personalized E2E-AD paradigm capable of adapting driving styles directly from users’ trajectory data.

Against this background, significant challenges remain in achieving truly individual-level personalized E2E-AD. These challenges arise throughout the entire process of this task, including data, evaluation, and algorithms. The challenges can be summarized as follows:

- **Data level:** Existing driving datasets face several limitations. Many are collected in simulated or non-human environments, while real-world datasets lack individual-level style annotations. Furthermore, they rarely support data augmentation or closed-loop evaluation, limiting their applicability to personalized E2E-AD research.
- **Evaluation level:** Current research lacks standardized and quantitative metrics for assessing individual driving styles, leaving a gap in measuring and comparing behavioral differences across drivers.
- **Algorithm level:** Existing approaches struggle to learn stylized representations directly from users’ trajectory data and lack a unified end-to-end framework for personalized autonomous driving.

Therefore, to address these three levels of challenges, we built a comprehensive personalized E2E-AD platform and benchmark, which is called **Person2Drive**. Our framework establishes a closed-loop research pipeline that enables data collection, model training, and quantitative evaluation of individualized driving behaviors. This platform not only provides the first unified foundation for studying personalized E2E-AD but also facilitates systematic comparison and reproducible research in this emerging area. Specifically, to tackle the aforementioned challenges, our work makes the following key contributions:

- **Personal driving data collection.** We develop an open-source and highly extensible data collection platform that simulates realistic driving scenarios at low cost, enabling the construction of large-scale and easily expandable

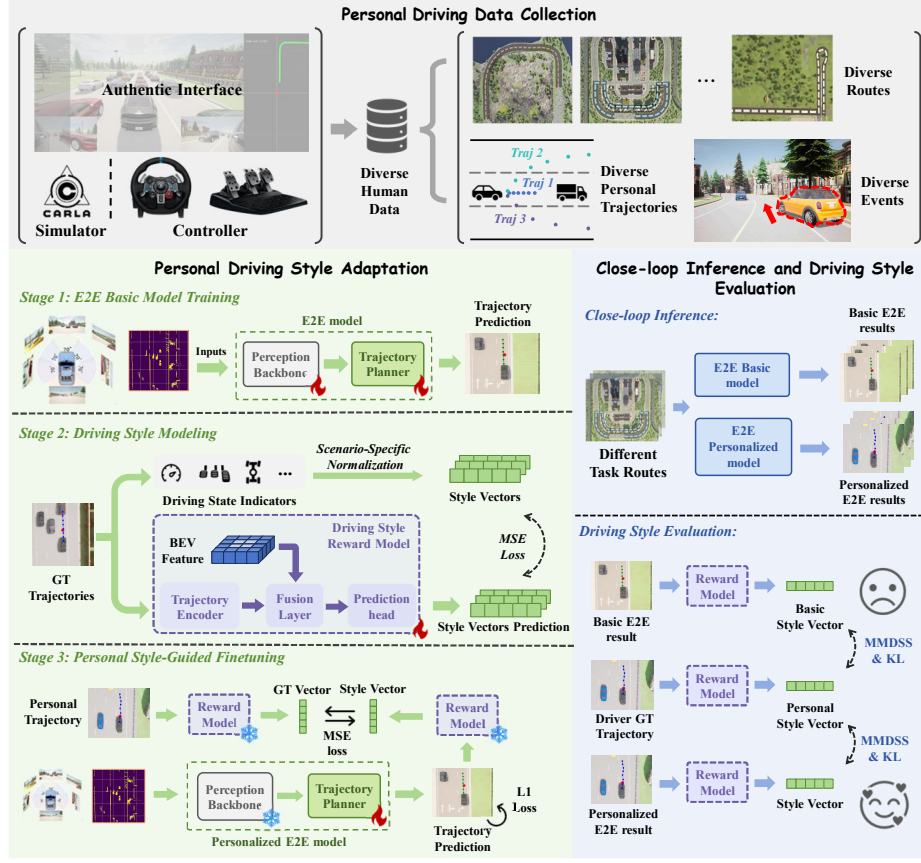


Fig. 1: Person2Drive framework. We develop a low-cost, scalable platform for data collection, featuring a navigation map and multi-camera views that mimic real-world driving habits, enabling authentic data acquisition across diverse scenarios, road segments, and driver styles. To achieve personalized driving style adaptation, we propose a three-stage adaptation method. Finally, the adapted end-to-end model is deployed in a simulated environment for inference, and the style model is used to evaluate the effectiveness of personalization.

personalized driving datasets. The dataset captures individual differences in driving styles across diverse scenes and events, supporting fine-grained analysis at the individual level. Moreover, the platform allows closed-loop validation, which is crucial for evaluating E2E-AD.

- **Driving style evaluation.** We comprehensively consider the dimensions in which individual driving styles manifest, including vehicle control parameters and driver behavior data. Based on these, we design style vector indicators that capture the most significant inter-individual differences and explore the use of Maximum Mean Discrepancy (MMD) and KL divergence to quantify driving style discrepancies. We further propose an MMD similarity score (MMDSS), providing the first quantitative analysis of personalized E2E-AD.

- **Personal driving style adaptation.** Aiming to maintain the performance of the baseline E2E model while realizing personalization, we introduce a dedicated personalized training framework and a style reward model built upon style vector indicators, which extracts stylized representations from users’ driving trajectories. By employing this reward model, we exclusively fine-tune the final trajectory prediction head, enabling efficient and robust personalized E2E-AD.

2 Related Work

2.1 End-to-End Autonomous Driving

E2E-AD directly maps raw sensor inputs to vehicle control commands using unified models, simplifying system design and enabling joint optimization across perception, prediction, and control for more consistent and robust driving. This paradigm not only advances end-to-end autonomous driving systems but also provides a new perspective for modeling and understanding complex driving behaviors. Codevilla et al. [3] demonstrated the practical feasibility of this approach. Recently, transformer-based methods have shown strong capabilities in temporal reasoning and multi-modal fusion, highlighting their potential for E2E-AD [15, 26, 30]. In parallel, the development of generative and world models has enhanced the generalization and interpretability of E2E-AD [17, 21, 34]. Notably, DiffusionDrive introduces diffusion models into E2E-AD, achieving state-of-the-art driving performance with relatively low computational cost [22]. Collectively, these advances not only consolidate the practical foundation of E2E-AD but also provide the technical basis for personalized driving and user-specific adaptation.

2.2 Personalized Autonomous Driving

Personalized autonomous driving (PAD) aims to adapt vehicle behavior to individual driver preferences [6]. Although E2E-AD has advanced significantly, most existing PAD methods are not designed for fully end-to-end deployment. Early studies focused on specific driving behaviors, such as car-following [33], lane-changing [23], and ramp-merging [20], using techniques like Inverse Reinforcement Learning (IRL), digital twins, or context encoding. More recently, large language models (LLMs) have been incorporated to enable interpretable control, for example, D2E translates user intentions into executable instructions [4], and PADriver combines streaming frames with personalized textual prompts [19]. MAVERIC proposes a data-driven approach for personalized driving [29], while StyleDrive leverages a fine-tuned vision-language model to realize end-to-end personalized driving, fusing scene topology with dynamic agent interactions [9]. Despite these advances, existing methods remain limited: style labels are coarse-grained, trajectory-level fine-grained style representations are not learned, and a general end-to-end framework for safe and efficient personalization is still lacking.

2.3 Personalized Driving Datasets

Personalized autonomous driving has attracted considerable attention, leading to the development of several personalized driving datasets. Early datasets explored

individual differences but primarily focused on open-loop evaluation, limiting their applicability to end-to-end systems [18,20,23,31,33]. Bench2Drive provides closed-loop evaluation, yet its data is not collected from human drivers [16]. StyleDrive is the first end-to-end driving dataset with style annotations, but it is semi-closed-loop and restricted to only three style categories: aggressive, normal, and conservative [9]. These limitations highlight the urgent need for a personalized driving dataset that supports end-to-end closed-loop evaluation, enabling true individual-level research in personalized E2E-AD.

3 Personalized Driving Benchmark

We introduce **Person2Drive**, a benchmark for exploring personalized E2E-AD, which emphasizes the fidelity and consistency of **human driving styles**. It comprises an open-source system for collecting multi-modal personalized driving data, a fully annotated dataset with driver identity labels, interpretable driving style evaluation, and a personalized driving style adaptation method for efficient imitation and transfer of individual driving behaviors. The overall framework is illustrated in Figure 1.

3.1 Personal Driving Data Collection

We develop a low-cost, scalable, and open-source data collection system that supports **closed-loop end-to-end operation**, captures authentic human driving styles, and enables diverse data acquisition. With this system, we further construct a personalized driving dataset with multi-modal data.

Data Collection System Existing closed-loop E2E-AD datasets are mostly from expert drivers and lack standardized, extensible collection frameworks, limiting research on individual driving styles. Real-world personalized data is costly and hard to control, confounding style variations with environmental factors. Therefore, there is an urgent need for a low-cost, scalable data collection system. We developed an open-source data collection system based on CARLA [5], enabling human drivers to operate under diverse traffic, scene, and weather conditions while automatically recording multi-modal personalized driving data. Our framework extends the implementation provided by Bench2Drive [16], and can be easily adapted to collect additional data. 1) **Realistic, Low-Cost Environment**. To bridge the gap between simulation and real-world driving, we construct a realistic human-in-the-loop driving environment. A fixed 2D top-down navigation map and rear-view mirrors enhance situational awareness, while a Logitech G29 steering wheel and pedal set provide authentic control dynamics and haptic feedback. Together, these create an immersive and realistic setup that effectively captures human driving behavior. 2) **Diverse Contexts and Multi-modal Data**. Human driving styles vary across environments. To capture diverse behaviors, our system offers multiple CARLA towns, configurable routes, and adjustable traffic densities, alongside events such as intersections, merges, and lane changes to elicit context-dependent driving patterns. Synchronized multi-modal sensor data, including RGB, LiDAR, radar, and semantic segmentation, is recorded under consistent conditions to ensure reproducibility.

This environmental and sensory diversity enables comprehensive modeling and systematic analysis of personalized driving behaviors. 3) **Style-Driven Route Design.** To effectively capture personalized driving behaviors, route design must balance duration and scenario diversity. Unlike existing datasets that focus on short event-centric routes or long evaluation routes [5], we design routes of approximately 1 km, providing sufficient temporal and spatial context for style expression while keeping driver workload manageable. All routes are encoded in a standardized XML format with town and configuration metadata, ensuring consistency and reproducibility across drivers and CARLA towns, forming a unified human-in-the-loop data collection framework. 4) **High Extensibility.** Most importantly, the system is highly extensible, allowing data to be flexibly scaled through this lightweight collection system. Identity labeling, multi-modal synchronization, and style indicator extraction are fully automated, such that drivers only need to operate the vehicle, with no additional manual intervention. This automated pipeline supports efficient scaling to additional drivers and repeated sessions.

Table 1: Datasets comparison. #Subjects is the number of drivers. We give **Personal Multiple Trajectories** on the same route to eliminate randomness of personal driving behavior. “—” in Subjects indicates unavailable driver identities: StyleDrive lacks driver identity annotations and Bench2Drive contains agent-generated trajectories.

Dataset	Evaluation Type	E2E	#Subjects	Human Driving	Individual Annotation	Personal Multiple Trajectories
HITL-RM [20]	Open-Loop	✗	2	✓	✓	✗
HITL-CF [33]	Open-Loop	✗	4	✓	✓	✗
HITL-LC [23]	Open-Loop	✗	2	✓	✓	✗
HITL-Mul [18]	Open-Loop	✗	60	✓	✓	✗
UAH [28]	Open-Loop	✗	6	✓	✓	✗
Brain4Cars [14]	Open-Loop	✗	10	✓	✓	✗
PDB [31]	Open-Loop	✗	12	✓	✓	✗
StyleDrive [9]	Semi-Closed-Loop	✓	-	✓	✗	✗
Bench2Drive [16]	Closed-Loop	✓	-	✗	✗	✗
Person2Drive (ours)	Closed-Loop	✓	50	✓	✓	✓

Personalized Dataset Based on the proposed data collection system, we introduce **Person2Drive**, a large-scale and extensible dataset designed to capture diverse human driving behaviors. The dataset supports closed-loop E2E-AD and includes driver identity annotations, making it suitable for individualized driving studies. Unlike existing datasets that primarily emphasize scene diversity, Person2Drive intentionally records multiple trajectories of the same driver on identical routes. Furthermore, the dataset can be expanded using our collection system to capture additional driver styles.

We collected driving data from 50 drivers, each repeatedly driving four identical one-kilometer routes with multi-modal sensors, resulting in 55 hours of driving data, 2.3 million LiDAR frames and 16.1 million synchronized camera images. The first two routes were repeated eight times and the remaining two routes twice to support both model training and cross-scene evaluation. Moreover, by repeatedly collecting data from the same drivers on identical routes, the dataset enables the analysis of both intra-driver consistency and inter-driver style variations, which are crucial for personalized E2E-AD. The routes span diverse and geometrically complex scenarios, including signalized and unsignal-

ized intersections, roundabouts, lane reductions, curved merging and diverging structures, highway off-ramps, long gentle curves, and varying lane configurations from single-lane to multi-lane roads. This structured design ensures stylistic diversity while maintaining controlled yet realistic traffic conditions.

The dataset offers comprehensive features, including 3D bounding boxes, depth, semantic segmentation, and driver identity labels, all sampled at 10 Hz, serving as a training resource for personalized E2E-AD. To facilitate method reproduction, we adopt a sensor setup similar to Bench2Drive [16], including 1 LiDAR, 6 cameras, 5 radars, 1 IMU & GNSS, 1 BEV camera, and HD-Map. Compared to existing datasets (Table 1), Person2Drive captures finer-grained intra- and inter-driver styles, involves more participants, and includes richer multi-modal data, while emphasizing human-in-the-loop driving and personalized driver labeling. More details about the data collection system and dataset can be found in the supplement.

For style annotation, we move beyond predefined categorical labels and instead model each driver as an **independent style domain**, enabling individual-level personalization. By collecting multiple driving sessions per individual under diverse scenes, our dataset supports fine-grained analysis of intra-driver style changes conditioned on environmental context. Person2Drive explicitly incorporates scene-level semantic annotations that are absent in prior datasets, establishing a comprehensive foundation for benchmarking and advancing research on personalized driving style modeling and fine-grained behavioral understanding.

3.2 Driving Style Evaluation

Evaluating driving style is a critical yet underexplored aspect of personalized E2E-AD. The lack of standardized metrics prevents meaningful evaluation of inter-driver differences and model performance. Therefore, we present a three-step evaluation framework for personalized E2E-AD: 1) constructing style vectors to capture key dimensions of individual driving behavior, 2) quantifying inter-driver differences, and 3) applying scenario-specific normalization for consistent comparison across diverse contexts. Built on this framework, we introduce MMDSS, providing the first reproducible, quantitative analysis of driving styles under the E2E-AD setting.

Style Vector Indicators To comprehensively capture individual driving styles and highlight inter-driver differences, we design style vector indicators based on vehicle control parameters and driver behavior data. To identify the most distinctive trajectory-derived style indicators, we introduce a median normalization and variance-based discriminability method, inspired by classical robust statistical analysis [25]. For each driver, we compute the indicator’s mean and normalize it by the population median, capturing deviations from typical behavior. We then measure the variability of these normalized values across all drivers and select the top ten indicators with the highest standard deviation as representative style features, as Eq. 1.

$$x_{std} = \sqrt{\frac{1}{N} \sum_{i=1}^N (x' - \mu)^2}, \quad (1)$$

where $x' = x - x_{mid}$ and x_{mid} means the median. This process robustly highlights inter-driver differences while mitigating scale and outlier effects. More details can be found in Section 5.1 and Supplementary Material.

Style Differences Quantification To better quantify differences across driving styles, we introduce MMDSS, a distribution-level metric that models driving style as a behavioral distribution rather than isolated trajectory instances. This can capture global style tendencies and provide a more stable measure of inter-style divergence. We also report the KL divergence to further validate the separability between different style distributions.

Specifically, given two sets of trajectory features $X = \{x_i\}_{i=1}^m$ and $Y = \{y_j\}_{j=1}^n$, the empirical MMD is computed as in Eq. 2:

$$\begin{aligned} \text{MMD}^2(X, Y) = & \frac{1}{m^2} \sum_{i=1}^m \sum_{i'=1}^m k(x_i, x_{i'}) \\ & + \frac{1}{n^2} \sum_{j=1}^n \sum_{j'=1}^n k(y_j, y_{j'}) - \frac{2}{mn} \sum_{i=1}^m \sum_{j=1}^n k(x_i, y_j), \end{aligned} \quad (2)$$

where $k(\cdot, \cdot)$ denotes an RBF kernel function here. While MMD quantifies distributional divergence, it is not directly interpretable as a similarity score. Therefore, we introduce MMDSS, a normalized similarity score defined as:

$$\text{MMDSS} = \frac{1}{1 + \text{MMD}}. \quad (3)$$

MMDSS transforms the raw divergence measure into a bounded, interpretable score in $[0, 1]$, where higher values indicate more consistent driving styles. This makes MMDSS particularly suitable for evaluating intra- and inter-individual behavioral similarity in personalized driving.

In addition to MMD, we also employ KL divergence to assess distributional differences in driving styles further. In our implementation, KL is computed using a histogram-based symmetric formulation with 50 bins, independent min-max normalization to $[0, 1]$, and ϵ -smoothing (10^{-8}). Lower KL value indicates a closer match between the two style distributions, providing a complementary measure of style similarity.

By combining both MMD and KL divergence, we can capture a comprehensive understanding of how different drivers' styles diverge not only in terms of distributional distances (MMD) but also in terms of relative entropy (KL divergence). This multi-metric approach offers a robust and interpretable analysis of driving styles.

Scenario-specific Normalization In addition, driving behavior is inherently context-dependent. For example, a speed of 80 km/h may be considered aggressive in an urban road segment but conservative on a highway. To account for such context-dependent variations, we adopt a scenario-specific normalization strategy. Here, each route is treated as a distinct scenario. Specifically, we analyze the feature distributions of all drivers within each route and perform min-max normalization separately. This design mitigates style bias induced by route-dependent feature shifts and enables more consistent comparison of individual driving behaviors across heterogeneous driving environments.

3.3 Personal Driving Style Adaptation

Based on the proposed style evaluation model, we propose a Personal Driving Style Adaptation (**PDSA**) strategy to guide the E2E-AD model to be more similar to the target style while maintaining great performance. Our entire training framework includes three stages: E2E Basic Model training, Driving Style Modeling, and Personal Style-Guided Finetuning, as illustrated in Fig. 1.

E2E Basic Model training. We first employ DiffusionDrive [22] as our basic driving model, leveraging its ability to capture complex temporal dynamics in vehicle motion. The model is trained using historical trajectory sequences alongside the BEV features extracted from the environment on the Bench2Drive dataset. This combination allows the backbone to effectively encode both past motion patterns and spatial context, providing a strong foundation for basic driving performance and even style-aware refinement.

Driving Style Modeling. To bridge raw trajectory data and interpretable driving style features, we introduce a learnable reward model that replaces hand-crafted style formulas. Given a trajectory, the model outputs a 10-dimensional style vector by learning a differentiable mapping between trajectory dynamics and driver style. This continuous supervision enables generalization to unseen behaviors, which fixed analytic functions cannot achieve. The model is implemented as a lightweight MLP with two hidden layers of 256 units each and ReLU activations. To incorporate environmental context, we encode the BEV map into a global scene vector, inject it at each time step of the trajectory encoder, and aggregate temporal features through the MLP to produce the final style vector.

Personal Style-Guided Finetuning. Given a reward model that maps driving trajectories to style vectors, we leverage it to fine-tune the end-to-end driving model, updating only the trajectory prediction head. All predicted trajectories are mapped to style vectors via the reward model and aligned with the corresponding ground-truth style vectors using mean squared error. This alignment loss is incorporated into the fine-tuning process, enhancing the trajectory prediction head’s sensitivity to style while avoiding performance degradation that could result from direct trajectory-level supervision. In this way, PDSA bridges style evaluation and policy optimization, enabling style-consistent driving behaviors without compromising baseline driving competence.

4 Experiment

In this section, we conduct extensive experiments on both our proposed dataset and the StyleDrive dataset to evaluate the effectiveness of the proposed style assessment system and the performance of the proposed fine-tuning framework.

4.1 Driving Style Evaluation

We first validate the proposed style evaluation framework. A reasonable and effective style evaluation system should ensure that intra-style differences are smaller than inter-style differences. Therefore, on the StyleDrive dataset, we compute style vectors from the ground-truth trajectories and calculate MMDSS and KL divergence across different style categories. By comparing intra-style and

inter-style differences, we verify the effectiveness of our driving style evaluation, as shown in Figure 2(a).

It presents MMDSS and KL divergence across intra-style (Intra-A/N/C) and inter-style comparisons. MMDSS behaves as expected—values within the same style remain high (about 0.93), while cross-style values drop notably (0.61–0.80), confirming its reliability in capturing style consistency. In contrast, KL divergence shows partial deviation: although Intra-N and Intra-C exhibit small intra-style KL values, Intra-A is unexpectedly higher. This irregularity may arise from distributional imbalance or estimation instability, highlighting KL’s sensitivity to sample support and noise. Overall, our proposed MMDSS proves to be a robust indicator of intra-style similarity.

To further examine whether the proposed driving style evaluation framework can capture fine-grained individual style differences, we conduct additional experiments on the Person2Drive dataset. Specifically, based on the extracted style vector indicators, we first compute MMDSS and KL divergence across multiple trajectories of each driver to quantify intra-individual style consistency. We then compute the same metrics between different drivers to assess inter-individual style differences, thereby validating the framework’s capability to effectively distinguish personalized driving behaviors.

As shown in Figure 2(b), intra-person MMDSS (0.941) exceeds inter-person MMDSS (0.681), and intra-person KL (0.023) is much lower than inter-person KL (0.470), demonstrating strong within-driver consistency and clear between-driver separability. These results confirm that our evaluation framework reliably captures intra- and inter-person driving style differences and provides an interpretable, quantitative basis for analyzing personalized driving styles.




Fig. 2: Style similarity evaluation results on the GT trajectories of StyleDrive dataset and Person2Drive dataset. (a) Results on the StyleDrive dataset. (b) Results on the Person2Drive dataset.

4.2 Reward Model-based Style Modeling

We evaluate the proposed reward model-based style modeling framework on the StyleDrive dataset by comparing the style discrepancy metrics of the style vectors predicted by our reward model (**Ours**) with those of the ground-truth vectors (**GT**), considering both intra-style and inter-style differences. Scenario-level segmentation is also performed to assess the model’s ability to capture fine-grained style variations. The results, measured by MMDSS and KL divergence

(Table 2), show that the learned style vectors closely match the ground-truth representations. Intra- and inter-style MMDSS indicate that the learned style vectors align well with the real styles, with minor deviations in certain scenarios. Both MMDSS and KL divergence confirm that intra-style differences are smaller than inter-style differences within the same scenario. These results demonstrate that the reward model can effectively learn fine-grained personalized driving styles from user trajectories, with MMDSS providing a robust and interpretable measure of style consistency and separability.

Table 2: Reward model-based style modeling results on the StyleDrive dataset. The table shows the differences (Δ) between the learned styles and the ground-truth across various scenarios, illustrates that the learned representations are highly consistent with the real ones in reflecting both intra-style and inter-style variations.

	Scene	Lane Following			Special Interior Road			Protected Intersections		
		GT	Ours	Δ	GT	Ours	Δ	GT	Ours	Δ
Intra-Style	MMDSS ($\uparrow$)	0.9427	0.9511	0.0084	0.9548	0.9575	0.0027	0.9175	0.9129	-0.0046
	KL ($\downarrow$)	0.0194	0.0171	-0.0023	0.0157	0.0083	-0.0074	0.0323	0.0263	-0.0060
Inter-Style	MMDSS ($\uparrow$)	0.6771	0.6783	0.0012	0.6790	0.6762	-0.0028	0.7052	0.7072	0.0020
	KL ($\downarrow$)	0.0272	0.0233	-0.0039	0.0187	0.0109	-0.0078	0.0419	0.0341	-0.0078

We then also conduct a similar analysis on the Person2Drive dataset to verify the effect of our method. Specifically, we compare the ground-truth and reward-model outputs at both intra- and inter-person levels to validate whether the reward model can effectively capture and represent driving styles across individuals. As illustrated in Figure 3, the reward-based style representations closely align with the ground-truth at both intra- and inter-person levels. MMDSS effectively captures the expected pattern, with smaller intra-person differences and larger inter-person differences, whereas KL values are notably lower, particularly for inter-person comparisons. By directly comparing sample distributions, MMDSS provides a more robust and interpretable measure of style consistency and separability, confirming that the reward model reliably learns fine-grained, individualized driving styles.




Fig. 3: Style Modeling results on the Person2Drive dataset.

4.3 Personal Driving Style Adaptation Results

To verify the performance of the proposed PDSA framework, we first conducted preliminary experiments on the StyleDrive dataset, followed by formal experiments on the Person2Drive dataset. Specifically, for the StyleDrive, we fine-tune the baseline using aggressive driving data, while for the Person2Drive, we adopt each individual’s historical driving trajectory to personalize the model. Then we compare the MMDSS between the style indicators before and after fine-tuning.

We also compare different variants of our proposed PDSA as ablation studies, including: **DFT**, which means directly fine-tuning; **PDSA-WB**, which means using reward models without BEV features; and **PDSA**, which means using reward models with BEV features. Results are shown in Tables 3 and 4.

Table 3 reports the MMDSS results for the generated trajectories after fine-tuning with the aggressive driving subset. All three fine-tuning strategies progressively shift the model toward the aggressive style, with the MMDSS between the aggressive style (A) and the ground truth increasing from 0.8453 to 0.8605. Among the methods, PDSA achieves the largest improvement (+0.0152), indicating its superior capability in capturing the desired stylistic characteristics.

Table 3: Evaluation of reward model training and fine-tuning with Aggressive driving dataset based on the driving style labels (**A**ggressive - **N**ormal - **C**onservative) of StyleDrive.

Method	A to GT	Δ	N to GT	Δ	C to GT	Δ
baseline	0.8453	0	0.9439	0	0.7158	0
DFT	0.8551	+0.0098	0.9327	-0.0112	0.7103	-0.0055
PDSA-WB	0.8601	+0.0148	0.9290	-0.0149	0.7080	-0.0078
PDSA	0.8605	+0.0152	0.9308	-0.0131	0.7088	-0.0070

Table 4 reports the performance gains across both scenarios after applying the PDSA framework. DFT yields modest gains over the baseline, indicating personal driving histories provide useful style cues. Incorporating the reward model without BEV features further amplifies these improvements, demonstrating the benefits of style-guided supervision. Our full model, PDSA, achieves the best performance in both scenarios, showing that adding BEV context enables more accurate style alignment. In addition, per-driver gains Δ (method–baseline) indicate that most drivers benefit, with all improvements statistically significant according to **Wilcoxon signed-rank tests** ($p < 0.001$). Overall, these results validate that PDSA effectively injects personalized style into the trajectory prediction head while maintaining stable gains across diverse driving scenarios.

Table 4: Personalization results on Person2Drive. The table presents MMDSS and per-driver robustness statistics (Δ = method–baseline, multiplied by 10^{-3}).

Method	Route 1			Route 2			MMDSS Avg $\uparrow$
	MMDSS $\uparrow$	Δ Mean $\pm$ Var	%($\Delta > 0$) $\uparrow$	MMDSS $\uparrow$	Δ Mean $\pm$ Var	%($\Delta > 0$) $\uparrow$	
Baseline	0.6410	–	–	0.6894	–	–	0.6652
DFT	0.6425	1.597 $\pm$ 1.009	82%	0.7095	20.163 $\pm$ 0.845	70%	0.6760
PDSA-WB	0.6439	2.945 $\pm$ 0.467	76%	0.7131	23.720 $\pm$ 1.094	58%	0.6785
PDSA	0.6465	5.590 $\pm$ 0.225	84%	0.7135	24.151 $\pm$ 0.948	76%	0.6800

4.4 Comparison with coarse style categories (ANC).

To further compare with the prior coarse-grained style supervision (e.g., StyleDrive), we build three category-level fine-tuning baselines by training the base model on trajectories from each ANC group (Aggres-

Table 5: ANC coarse-style fine-tuning vs. personalized fine-tuning on Person2Drive.

Method	A	N	C	Personal (ours)
Average	0.6156	0.6125	0.6174	0.6800

sive/Normal/Conservative), resulting in three ANC-specific models. For each driver, we select the ANC-specific model whose driving style is most similar to the driver and compute their style similarity. We then compare this category-level baseline with fine-tuning on the driver’s own trajectories. As shown in Table 5, ANC-specific model yield weaker style alignment than fine-tuning on individual trajectories, while our personalized approach enables the model to better align its driving style with that of the driver. Overall, this comparison highlights the limitations of coarse category-level supervision and demonstrates the effectiveness of personalized fine-tuning.

4.5 Additional Validation

Alternative Formulations. We compare PDSA with several alternative style-learning formulations, including direct style-

vector supervision(DSV), behavior cloning (BC), inverse reinforcement learning (IRL), and preference learning(PL). As shown in Table 6, PDSA achieves the highest MMDSS among all compared methods, indicating more accurate style alignment than alternative formulations.

Human Preference Study. We conduct a human preference study involving 50 participants. As shown in Table 7, PDSA is consistently preferred over the baseline in both scenes, suggesting that the improvements measured by MMDSS align well with human perception of driving style.

Driver-Specific Behavior Analysis. To distinguish driver-specific behavior from scenario-dependent effects, we compare trajectories generated by different personalized models under identical scenes. The inter-driver MMDSS scores are 0.6383 and 0.6716 in Route 1 and Route 2, respectively, which are lower than the corresponding baseline-personal trajectory similarities (0.6410 and 0.6894), indicating that the learned policies capture driver-specific behavioral characteristics rather than converging to a common scenario-dependent policy.

Table 6: Alternative Formulations

	DSV	BC	IRL	PL	PDSA
MMDSS	0.6690	0.6811	0.6591	0.6796	0.7055

Table 7: Human Preference Study

Method	Route 1 (%) ↑	Route 2 (%) ↑	Participants
Baseline	24.0	28.0	50
PDSA (Ours)	76.0	72.0	50

5 Analysis and Discussion

5.1 Driving Style Indicator Analysis

In this section, we analyze the effectiveness of the proposed style indices. Specifically, during the collection of the Person2Drive dataset, we record the corresponding potential style indicators. The measured differences across drivers for the top 20 most discriminative indicators are presented in Figure 4, demonstrating the ability of these indices to capture inter-driver variability. The figure reveals that the scores capture the divergence of driving behaviors across different scenarios. Metrics such as TTC and Lane Change Frequency (per km)

achieve higher scores on highways, indicating greater variability among individual driving trajectories in these metrics. Urban scenarios show elevated scores for metrics like Longitudinal Acceleration, reflecting pronounced differences in stop-and-go behavior across drivers. The total scores highlight the most influential metrics, with the tenth-ranked metric annotated to delineate high-impact indicators from secondary ones. Overall, the grouped bar plot conveys each metric’s relative contribution to driving style variability, providing quantitative guidance for driving style analysis and E2E-AD personalization.




Fig. 4: Average style indices differences among all individuals.

5.2 Does style-guided fine-tuning affect driving ability?

The style-guided fine-tuning enables the model to align the generated trajectory with the real driving style. However, it remains unclear whether such stylized fine-tuning affects driving ability. Here, we compare the basic driving ability before and after PDSA in terms of driving score and success rate. We further extend this comparison to diverse and difficult cases in Bench2Drive to evaluate performance under **challenging scenarios**. As shown in Table 8, style-guided fine-tuning (PDSA) improves driving performance across datasets. On the Person2Drive dataset, the driving score increases by 11.91 points, with only a negligible decrease of 0.5% in success rate. On the Bench2Drive dataset, the driving score increases by 0.78 points and success rate improved by 2.18%. These results suggest that aligning trajectories with personalized styles does not degrade driving performance, demonstrating that our fine-tuning approach effectively balances style adaptation with fundamental driving ability.

Table 8: Driving safety evaluation.

Method	Person2Drive		Bench2Drive	
	Driving Score↑	Success Rate↑	Driving Score↑	Success Rate↑
Baseline	58.22	98.9%	54.91	13.27%
PDSA-WB	67.45	98.6%	54.16	15.22%
PDSA	70.13	98.4%	55.69	15.45%

5.3 Fine-tuning Trajectory Visualization

To further validate the effectiveness of our PDSA method, we randomly sample a subset of instances and provide qualitative visualizations across scenarios. The selected lane-changing and post-stop following examples are shown in Figure 5 and Figure 6. Figure 5 shows that, in contrast to the other model variants—which primarily exhibit a tendency to drive faster—PDSA is the only method that successfully reproduces the human driver’s lane-change behavior. It captures both

the timing and execution of the maneuver, producing trajectories that align most closely with the ground-truth driving style. Figure 6 shows that baseline and DFT show clear discrepancies from human ground truth, often stopping either too close to or too far from the leading vehicle. In comparison, PDSA produces stopping behaviors that align most closely with human drivers, demonstrating its superior capability to capture and reproduce individualized post-stop following style. Overall, the visualization confirms the quantitative results: PDSA effectively adapts the model toward each driver’s unique style, achieving noticeable convergence in both motion dynamics and trajectory geometry.



Fig. 5: Comparison of driving behaviors across different models, visualized from the vehicle’s first-person view.



Fig. 6: Comparison of post-stop following gaps across different models, visualized from the vehicle’s first-person view.

6 Conclusion

In this work, we address the gap in achieving truly individual-level personalized E2E-AD. To this end, we develop an open-source and extensible data collection platform and construct Person2Drive. This large-scale personalized driving dataset enables fine-grained style analysis across multiple drivers and scenarios. Building upon this foundation, we introduce an explainable evaluation system for quantitative driving style assessment, establishing the first unified standard for evaluating individual driving styles. Furthermore, we propose a style-guided lightweight fine-tuning framework that allows autonomous models to generate personalized yet robust trajectories. Comprehensive experiments and analyses on both Person2Drive and public datasets validate the effectiveness of our framework. We believe this work provides a unified benchmark and methodological foundation for future research in user-aware and trustworthy autonomous driving, paving the way toward truly human-centered E2E-AD.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (Project Number 62595774), MoE Key Laboratory of Intelligent Perception and Human-Machine Collaboration (KLIP-HuMaCo), HPC Platform of ShanghaiTech University.

References

1. Chen, L., Wu, P., Chitta, K., Jaeger, B., Geiger, A., Li, H.: End-to-end autonomous driving: Challenges and frontiers. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2024)
2. Chib, P.S., Singh, P.: Recent advancements in end-to-end autonomous driving using deep learning: A survey. *IEEE Transactions on Intelligent Vehicles* **9**(1), 103–118 (2023)
3. Codevilla, F., Müller, M., López, A., Koltun, V., Dosovitskiy, A.: End-to-end driving via conditional imitation learning. In: 2018 IEEE international conference on robotics and automation (ICRA). pp. 4693–4700. IEEE (2018)
4. Cui, C., Yang, Z., Zhou, Y., Ma, Y., Lu, J., Li, L., Chen, Y., Panchal, J., Wang, Z.: Personalized autonomous driving with large language models: Field experiments. In: 2024 IEEE 27th International Conference on Intelligent Transportation Systems (ITSC). pp. 20–27. IEEE (2024)
5. Dosovitskiy, A., et al.: Carla leaderboard 2.0: Evaluating closed-loop autonomous driving. CARLA Challenge Technical Report (2023)
6. Eboli, L., Mazzulla, G., Pungillo, G.: How drivers’ characteristics can affect driving style. *Transportation research procedia* **27**, 945–952 (2017)
7. Gruyer, D., Magnier, V., Hamdi, K., Claussmann, L., Orfila, O., Rakotonirainy, A.: Perception, information processing and modeling: Critical stages for autonomous driving applications. *Annual Reviews in Control* **44**, 323–341 (2017)
8. Hang, P., Lv, C., Xing, Y., Huang, C., Hu, Z.: Human-like decision making for autonomous driving: A noncooperative game theoretic approach. *IEEE Transactions on Intelligent Transportation Systems* **22**(4), 2076–2087 (2020)
9. Hao, R., Jing, B., Yu, H., Nie, Z.: Styledrive: Towards driving-style aware benchmarking of end-to-end autonomous driving. *arXiv preprint arXiv:2506.23982* (2025)
10. Hasenjaeger, M., Heckmann, M., Wersing, H.: A survey of personalization for advanced driver assistance systems. *IEEE Transactions on Intelligent Vehicles* **5**(2), 335–344 (2019)
11. Hasenjaeger, M., Wersing, H.: Personalization in advanced driver assistance systems and autonomous vehicles: A review. In: 2017 IEEE 20th international conference on intelligent transportation systems (itsc). pp. 1–7. IEEE (2017)
12. Häuslschmid, R., Von Bülow, M., Pfleging, B., Butz, A.: Supporting trust in autonomous driving. In: Proceedings of the 22nd international conference on intelligent user interfaces. pp. 319–329 (2017)
13. Holzinger, A., Saranti, A., Angerschmid, A., Retzlaff, C.O., Gronauer, A., Pejakovic, V., Medel-Jimenez, F., Krexner, T., Gollob, C., Stampfer, K.: Digital transformation in smart farm and forest operations needs human-centered ai: challenges and future directions. *Sensors* **22**(8), 3043 (2022)
14. Jain, A., Koppula, H.S., Soh, S., Raghavan, B., Singh, A., Saxena, A.: Brain4cars: Car that knows before you do via sensory-fusion deep learning architecture. *arXiv preprint arXiv:1601.00740* (2016)
15. Jia, X., Wu, P., Chen, L., Xie, J., He, C., Yan, J., Li, H.: 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. pp. 21983–21994 (2023)
16. Jia, X., Yang, Z., Li, Q., Zhang, Z., Yan, J.: Bench2drive: Towards multi-ability benchmarking of closed-loop end-to-end autonomous driving. *Advances in Neural Information Processing Systems* **37**, 819–844 (2024)

17. Jia, X., You, J., Zhang, Z., Yan, J.: Drivetransformer: Unified transformer for scalable end-to-end autonomous driving. arXiv preprint arXiv:2503.07656 (2025)
18. Ke, Z., Jiang, Y., Wang, Y., Cheng, H., Li, J., Wang, J.: D2e: An autonomous decision-making dataset involving driver states and human evaluation of driving behavior. In: 2024 IEEE 27th International Conference on Intelligent Transportation Systems (ITSC). pp. 2294–2301. IEEE (2024)
19. Kou, G., Jia, F., Mao, W., Liu, Y., Zhao, Y., Zhang, Z., Yoshie, O., Wang, T., Li, Y., Zhang, X.: Padriver: Towards personalized autonomous driving. arXiv preprint arXiv:2505.05240 (2025)
20. Li, S., Wei, C., Wu, G., Barth, M.J., Abdelraouf, A., Gupta, R., Han, K.: Personalized trajectory prediction for driving behavior modeling in ramp-merging scenarios. In: 2023 Seventh IEEE International Conference on Robotic Computing (IRC). pp. 1–4. IEEE (2023)
21. Li, Y., Fan, L., He, J., Wang, Y., Chen, Y., Zhang, Z., Tan, T.: Enhancing end-to-end autonomous driving with latent world model. arXiv preprint arXiv:2406.08481 (2024)
22. Liao, B., Chen, S., Yin, H., Jiang, B., Wang, C., Yan, S., Zhang, X., Li, X., Zhang, Y., Zhang, Q., et al.: Diffusiondrive: Truncated diffusion model for end-to-end autonomous driving. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 12037–12047 (2025)
23. Liao, X., Zhao, X., Wang, Z., Zhao, Z., Han, K., Gupta, R., Barth, M.J., Wu, G.: Driver digital twin for online prediction of personalized lane-change behavior. IEEE Internet of Things Journal **10**(15), 13235–13246 (2023)
24. Liao, X., Zhao, Z., Barth, M.J., Abdelraouf, A., Gupta, R., Han, K., Ma, J., Wu, G.: A review of personalization in driving behavior: Dataset, modeling, and validation. IEEE Transactions on Intelligent Vehicles (2024)
25. Moore, D.S., McCabe, G.P., Craig, B.A.: Introduction to the Practice of Statistics, vol. 4. WH Freeman New York (2009)
26. Prakash, A., Chitta, K., Geiger, A.: Multi-modal fusion transformer for end-to-end autonomous driving. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 7077–7087 (2021)
27. Qin, Z., Yao, X., Wei, C., Ji, A., Wu, G., Sun, Z.: Contextual-personalized adaptive cruise control via fine-tuned large language models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1754–1762 (2025)
28. Romera, E., Bergasa, L.M., Arroyo, R.: Need data for driver behaviour analysis? presenting the public uah-driveset. In: 2016 IEEE 19th international conference on intelligent transportation systems (ITSC). pp. 387–392. IEEE (2016)
29. Schrum, M.L., Sumner, E., Gombolay, M.C., Best, A.: Maveric: A data-driven approach to personalized autonomous driving. IEEE Transactions on Robotics **40**, 1952–1965 (2024)
30. Shao, H., Wang, L., Chen, R., Li, H., Liu, Y.: Safety-enhanced autonomous driving using interpretable sensor fusion transformer. In: Conference on Robot Learning. pp. 726–737. PMLR (2023)
31. Wei, C., Qin, Z., Li, S., Zhang, Z., Zhao, X., Abdelraouf, A., Gupta, R., Han, K., Barth, M.J., Wu, G.: Pdb: Not all drivers are the same—a personalized dataset for understanding driving behavior. arXiv preprint arXiv:2503.06477 (2025)
32. Xu, C., Liu, J., Guo, Y., Zhang, Y., Hang, P., Sun, J.: Towards human-centric autonomous driving: A fast-slow architecture integrating large language model guidance with reinforcement learning. arXiv preprint arXiv:2505.06875 (2025)

- 33. Zhao, Z., Wang, Z., Han, K., Gupta, R., Tiwari, P., Wu, G., Barth, M.J.: Personalized car following for autonomous driving with inverse reinforcement learning. In: 2022 International Conference on Robotics and Automation (ICRA). pp. 2891–2897. IEEE (2022)
- 34. Zheng, W., Song, R., Guo, X., Zhang, C., Chen, L.: Genad: Generative end-to-end autonomous driving. In: European Conference on Computer Vision. pp. 87–104. Springer (2024)