Sky-Drive: A Distributed Multi-Agent Simulation Platform for Socially-Aware and Human-AI Collaborative Future Transportation

Zilin Huang†, Zihao Sheng†, Zhengyang Wan†, Yansong Qu, Yuhao Luo, Boyue Wang, Pei Li, Yen-Jung Chen, Jiancong Chen, Keke Long, Jiayi Meng, Yue Leng, Sikai Chen*

***Abstract*—Sky-Drive is the first comprehensive simulation platform that tightly integrates virtual reality (VR), multi-agent interactions, and human-centered AI approaches for multi-agent traffic simulation and human-centered autonomous agent research. Distinct from existing platforms, Sky-Drive introduces several key innovations: (a) a digital twin framework that creates high-fidelity virtual replicas of transportation systems for real-time monitoring and optimization; (b) a distributed multi-agent architecture that enables synchronized simulation across multiple terminals while maintaining precise real-time interactions between autonomous vehicles (AVs), human-driven vehicles (HVs), and pedestrians; (c) a multi-modal human-in-the-loop framework that captures rich human behavioral data through various sensors, including steering wheels, eye-tracking cameras, and smartwatch sensors; (d) integration of fundamental models for enhanced human-machine collaboration and personalized decision-making such as large language models (LLMs) and vision language models (VLMs); (e) a novel human-AI bidirectional mentor mechanism that facilitates effective knowledge exchange between human drivers and AI-enabled autonomous systems through both human feedback and domain knowledge from transportation science; (f) a hardware-in-the-loop module through ROS compatibility that enables direct verification of autonomous driving algorithms on physical platforms. Sky-Drive enables comprehensive research across various applications including VR-enabled vulnerable road user (VRU)-AV interactions, reinforcement learning-enabled autonomous driving policy learning, customized long-tail scenario generation, LLM-enabled personalized driving. Sky-Drive provides a unique environment for accelerating the development of safe and efficient AVs, while laying the groundwork for next-generation human-AI collaborative transportation systems. The demo video and code are available at: https://huang-zilin.com/Sky-Drive-website/.**

***Index Terms*—Autonomous Vehicles, Large Language Models, Human-in-the-Loop, Multi-Agent Simulation, Virtual Reality.**

I. Introduction

AUTONOMOUS driving and related technologies have made significant advancements in recent years, demonstrating increasing maturity in perception, decision-making, and control capabilities [1]–[3]. As these technologies continue to evolve, future transportation systems are expected to comprise complex combinations of intelligent agents, including autonomous vehicles (AVs), human-driven vehicles (HVs), delivery robots, drones, and smart traffic signals. In this emerging ecosystem, each agent must not only ensure its own safe and efficient operation but also continuously align its behavior with human preferences and societal norms through interactions with various road users, including pedestrians and cyclists. Consequently, future research must move beyond validating single-vehicle performance to exploring social awareness and human-AI collaboration in mixed traffic environments.

Validating autonomous driving technologies in real-world environments presents substantial challenges, particularly regarding safety risks and the extensive testing required to demonstrate reliability [4]–[6]. To address these challenges, the autonomous driving community has developed a variety of simulation platforms, such as CARLA [7], AirSim [8], SUMO [9], Highway-Env [10], MetaDrive [11], SMARTS [12], CarSim [13] and IPG CarMaker [14]. These platforms have significantly accelerated development by providing controlled testing environments. However, they face important limitations in addressing the unique needs of future transportation research. First, while existing platforms can simulate multiple agents concurrently on a single computer, they generally lack the capability to synchronize and run multiple agents across distributed systems. This limitation restricts the study of complex interactions where each intelligent agent requires independent control and decision-making capabilities, a key characteristic of future mixed traffic.

Second, most platforms offer limited support for human-AI collaboration. Although they can collect human input, they typically treat it as low-level control signals rather than as high-level feedback for improving autonomous driving algorithms. Moreover, they do not enable AI systems to assist human drivers by providing real-time guidance, performance feedback, or personalized training support. In contrast, human-AI collaboration refers to a bidirectional process where human drivers provide feedback not only as commands but as indications of preferences, situational understanding, and normative behaviors, while autonomous systems in turn assist human drivers by offering real-time guidance, performance feedback, and personalized training. This bidirectional exchange enables AI systems to continuously adapt to human needs and expectations while simultaneously enhancing human driving performance through intelligent support. Recent studies highlight the importance of aligning autonomous systems with human expectations [15]–[18], but existing platforms seldom facilitate

This work was funded by the U.S. Department of Transportation, Award #69A3552348305. The corresponding author is Sikai Chen (E-mail: sikai.chen@wisc.edu). † These authors contributed equally to this work.

Zilin Huang, Zihao Sheng, Zhengyang Wan, Yuhao Luo, Boyue Wang, Pei Li, Keke Long, and Sikai Chen are with the Department of Civil and Environmental Engineering, University of Wisconsin-Madison, Madison, WI 53706, USA (E-mails: {zilin.huang, zihao.sheng, zhengyang.wan, yuhao.luo, bwang367, pei.li, klong23, sikai.chen}@wisc.edu). Yansong Qu, and Jiancong Chen are with the Lyles School of Civil and Construction Engineering, Purdue University, West Lafayette, IN 47907, USA (E-mail: {qu120, chen5281}@purdue.edu). Yen-Jung Chen is with the Elmore Family School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47907, USA (E-mail: chen4126@purdue.edu). Jiayi Meng is with the Department of Computer Science and Engineering, The University of Texas at Arlington, Arlington, TX 76019, USA (E-mail: jiayi.meng@uta.edu). Yue Leng is with Google, Sunnyvale, CA 94089, USA (E-mail: lengc@google.com).

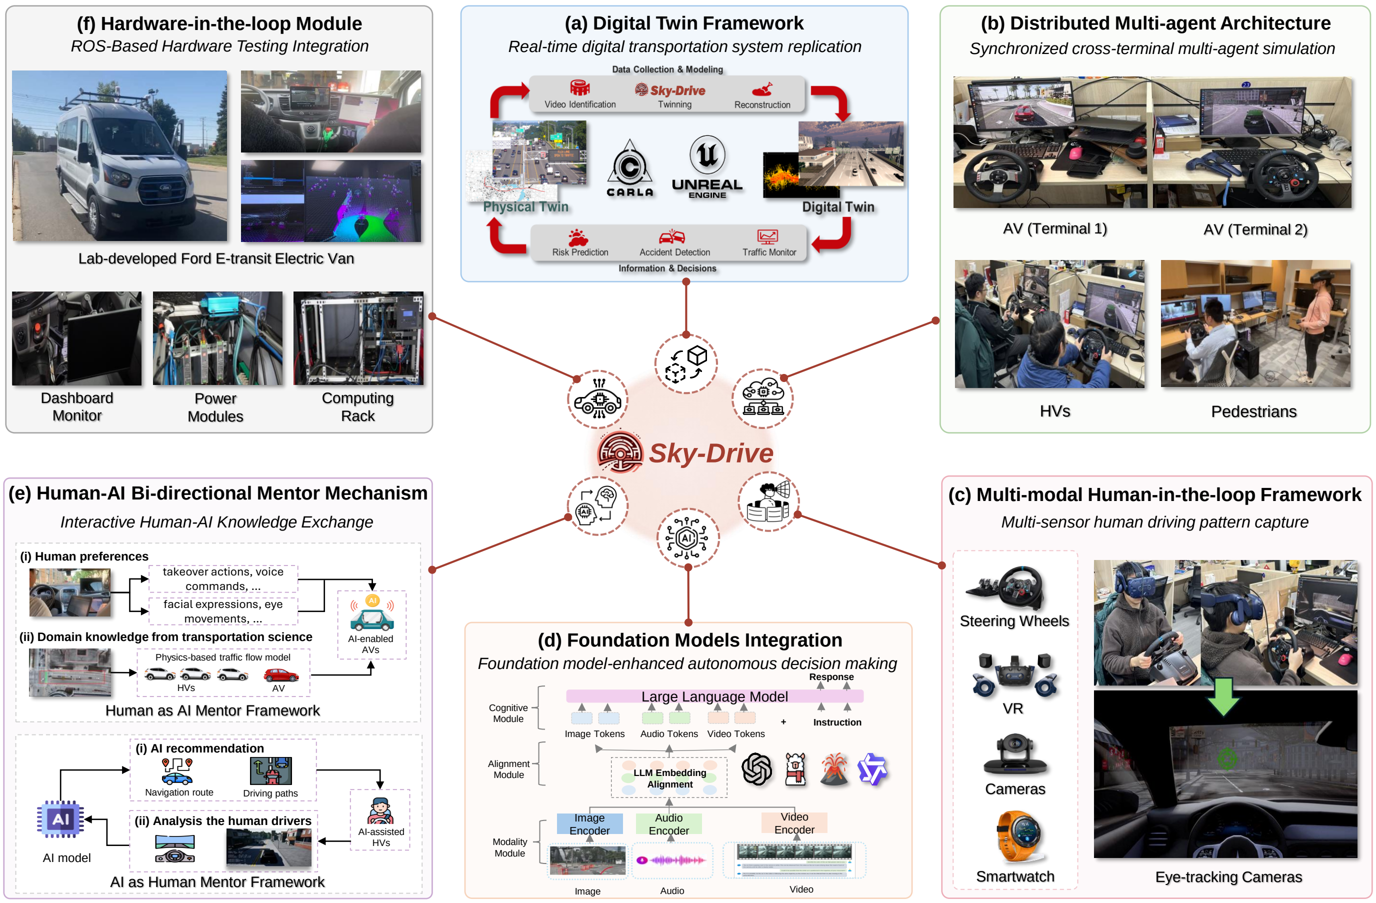

Fig. 1. Overview of Sky-Drive's key components and functionalities. (a) digital twin framework creating high-fidelity virtual replicas of transportation systems through multi-source data integration; (b) distributed multi-agent architecture enabling synchronized simulation across multiple terminals for complex traffic interactions; (c) multi-modal human-in-the-loop framework capturing comprehensive behavioral data through integrated sensor systems; (d) foundation models integration leveraging LLMs and VLMs for enhanced human-machine collaboration; (e) human-AI bi-directional mentor mechanism facilitating knowledge exchange between human drivers and autonomous systems; (f) hardware-in-the-loop module enabling direct validation through ROS compatibility.

real-time bidirectional knowledge exchange between humans and AI. Additionally, the emergence of foundation models, trained on diverse datasets and equipped with broad world knowledge, offers new opportunities for capturing and utilizing human knowledge. Nonetheless, current simulation platforms primarily use foundation models for scenario generation rather than for enabling active, bidirectional human-AI knowledge exchange in decision-making processes.

While some simulation platforms have incorporated reinforcement learning (RL) capabilities to improve autonomous driving policies, they remain primarily focused on optimizing individual vehicle safety and efficiency metrics. However, it is still necessary to incorporate social awareness into decision-making processes to support the development of socially aware autonomous systems. Social awareness refers to an autonomous system's ability to account for the effects of its actions on surrounding road users and the overall transportation system, aiming to promote traffic flow stability, enhance the comfort and safety of other participants, and enable harmonious coexistence between autonomous vehicles and human drivers in mixed traffic environments. In this context, established knowledge from transportation science represents a valuable resource. Specifically, validated traffic flow theories and human behavior models, developed through decades of research, could provide essential insights for designing autonomous systems that achieve socially aware behaviors and facilitate harmonious interactions with human drivers in mixed traffic environments.

To address these challenges, we propose **Sky-Drive**, an innovative open-source simulation platform designed to advance research in socially aware autonomous driving and human-AI collaboration. Sky-Drive unifies scenario generation, system simulation, data collection, algorithm training, and hardware integration into a comprehensive platform, supporting distributed multi-agent operation and multi-modal human-in-the-loop interaction. As illustrated in Fig. 1, Sky-Drive introduces several key innovations:

- Sky-Drive introduces a distributed multi-agent architecture that enables synchronized simulation across multiple devices through a remote procedure call (RPC) networking model. This design allows independent control of agents on separate terminals while maintaining shared environmental states, better reflecting future mixed traffic.
- Sky-Drive provides a multi-modal human-in-the-loop framework that integrates diverse sensors, including steering wheels, virtual reality (VR) systems, cameras, and

smartwatches, to capture rich human behavioral data. A synchronized data processing pipeline correlates these multi-modal streams, enabling detailed analysis of human driving patterns and responses to complex scenarios.
- Sky-Drive implements an innovative human-AI collaboration mechanism comprising a Human as AI Mentor (HAIM) module that incorporates human feedback and domain knowledge to guide AI learning, and an AI as Human Mentor (AIHM) module that provides real-time guidance and personalized training to human drivers.
- To bridge the gap between simulation and reality, Sky-Drive includes a digital twin framework that builds high-fidelity virtual replicas of transportation systems by integrating data collected from lab-developed AVs, roadside sensors, traffic cameras, and historical records.

To further enhance Sky-Drive's capabilities, two major functionalities are planned:

- Sky-Drive will integrate large foundation models at both the system and agent levels. At the system level, foundation models will provide global observation and feedback to optimize simulation dynamics. At the agent level, they will enhance situational understanding and enable safer, more socially aware, and personalized decision-making.
- Sky-Drive will incorporate a hardware-in-the-loop framework via Robot Operating System (ROS) integration, enabling direct validation of autonomous driving algorithms on physical vehicles and safe evaluation of human-AI collaboration strategies without exposing users to real-world risks.

The remainder of this paper is organized as follows: Section II reviews related work in driving simulators. Section III introduces Sky-Drive's workflow. Section IV details Sky-Drive's features and technical implementation. Section V demonstrates application examples. Section VI discusses planned future enhancements. Finally, Section VII concludes the paper and outlines future research directions.

II. Related Works

A. Driving Simulators

Driving simulation platforms have evolved significantly to address the growing needs of autonomous vehicle research. According to Li et al. [19], these simulators can be categorized based on their primary functions and capabilities.

Comprehensive simulators provide end-to-end virtual environments with complete road networks, diverse traffic agents, pedestrians, and detailed sensor models. CARLA [7] and LGSVL [20] represent prominent open-source examples in this category, offering rich environments for testing autonomous driving systems. Commercial solutions such as Nvidia Drive Sim [21] and rFpro [22], alongside academic developments including DeepDrive [23] and GarchingSim [24], provide similar comprehensive capabilities. Another important category is traffic flow simulators, which focus on modeling network-level vehicle movements, traffic congestion, and large-scale traffic scenarios. Notable examples include SUMO [9], Vissim [25], Flow [26], and CityFlow [27]. Recent developments combine SUMO's traffic modeling with 3D simulators such as CARLA to merge scalability with realism.

Sensory data simulators, such as AirSim [8] and Sim4CV [28], are designed to generate high-fidelity sensor outputs for perception systems. These functionalities are increasingly being integrated into comprehensive simulators while maintaining their critical role in AV perception testing. Driving policy simulators provide configurable environments for evaluating decision-making algorithms. Examples include Highway-Env [10], TORCS [29], SUMMIT [30], MACAD [31], SMARTS [12], and MetaDrive [32]. Additionally, recent data-driven simulators such as Waymax [33], ScenarioNet [34], and Nocturne [35] leverage real-world datasets to generate socially relevant traffic scenarios. Vehicle dynamics simulators, including CarSim [13], IPG CarMaker [14], and Gazebo [36], specialize in accurately modeling vehicle physics, such as suspension responses and tire-road interactions, which are essential for validating control algorithms under realistic conditions.

While existing platforms offer valuable simulation capabilities, certain challenges remain in supporting future transportation research. As shown in Tab. I, most simulators are designed to run on single devices, limiting their ability to model distributed multi-agent scenarios requiring independent processing. Moreover, current platforms offer limited support for socially-aware algorithms that must consider complex interactions with diverse road users. Sky-Drive is designed to address these challenges by introducing a distributed architecture that enables synchronized simulation across multiple terminals, combined with a comprehensive digital twin framework for high-fidelity environmental replication.

B. Human–AI Collaboration Environments

Several simulation platforms have made notable contributions to human-AI collaboration in the context of autonomous driving. NVIDIA's DRIVE Sim and Omniverse platform [38] facilitates collaboration through physics-based synthetic data generation for training autonomous systems. Nevertheless, their approach primarily supports one-way knowledge transfer, where simulated scenarios inform AI models, rather than enabling bidirectional knowledge exchange. Applied Intuition offers human-in-the-loop testing capabilities that allow human operators to validate the decisions made by autonomous systems, but its framework is primarily designed for validation purposes, rather than fostering continuous collaborative learning [39]. MORAI provides digital twin environments that visualize AI decision-making processes for human operators, yet its interaction mechanisms are limited to basic feedback collection, lacking the integration of knowledge for mutual learning [40].

More specialized platforms have attempted to advance human-AI collaboration. The VISTA simulation system developed by MIT enables domain adaptation between virtual and real environments, but it mainly addresses perception tasks rather than knowledge exchange mechanisms [41]. The GAMMA framework generates mixed-reality traffic data that incorporates aspects of human driving behavior, though it lacks explicit mechanisms for integrating human expertise into the

TABLE I
COMPARISON OF REPRESENTATIVE SIMULATORS WITH SKY-DRIVE

	Distributed Multi-agent Simulation	**Digital Twin Environment**	**Hardware-in-the-Loop**	**Traffic Flow Modeling**	**AI Framework Integration**	**Human-in-the-loop Interface**
Closed Source						
Nvidia Drive Sim [21]	-	✓	✓	-	✓	✓
rFpro [22]	-	✓	✓	-	-	✓
CarSim [13]	-	-	✓	-	-	✓
Matlab [37]	-	✓	✓	-	✓	✓
Open Source						
DeepDrive 2.0 [23]	-	-	-	-	✓	-
GarchingSim [24]	✓	-	✓	-	✓	✓
CARLA [7]	-	✓	✓	-	✓	✓
SUMO [9]	-	✓	✓	✓	-	-
Flow [26]	-	-	-	✓	✓	-
CityFlow [27]	-	-	-	✓	✓	-
TORCS [29]	-	-	-	-	✓	-
SUMMIT [30]	-	✓	-	-	✓	-
MACAD [31]	-	-	-	-	✓	✓
MetaDrive [32]	-	✓	-	-	✓	✓
SMARTS [12]	-	-	-	-	✓	-
Nocturne [35]	-	✓	-	✓	✓	-
Waymax [33]	-	✓	-	✓	✓	-
Gazebo [36]	-	✓	-	-	-	✓
Sky-Drive (Ours)	✓	✓	✓	✓	✓	✓

Note: The "Distributed Multi-agent Simulation" functionality in this table refers to the capability of simulators to synchronize and run multiple agents (e.g., AVs, HVs, and pedestrians) across different computers in real-time simulations. This is distinct from simply running multiple agents concurrently on a single computer, which most simulators can accomplish.

AI learning process [42]. Wayve's LINGO architecture makes significant strides by providing natural language explanations for AI decisions, thereby enhancing transparency in human-AI interaction [43]. SafeMod introduces bidirectional planning through large language models, incorporating human-like reasoning patterns into autonomous decision-making [44]. The SurrealDriver framework leverages large language models to generate realistic driving behaviors that align with human expectations in urban contexts [45]. DarwinAI's GenSynth platform demonstrates how human designers and AI can collaborate to accelerate the development of neural networks for autonomous driving applications [46].

Despite these advancements in collaborative interfaces, existing platforms face several limitations in enabling human-AI knowledge exchange. Current platforms generally lack mechanisms for the continuous integration of human feedback, resulting in open-loop, rather than closed-loop, learning systems. Few platforms are capable of comprehensive multimodal data collection from human operators, which is essential for gaining a deeper understanding of driving patterns and improving AI's adaptation to human behaviors. Sky-Drive addresses these limitations through its HAIM and AIHM modules, its multi-modal human-in-the-loop architecture, and its closed-loop learning mechanisms that continuously integrate human expertise into AI development.

III. SKY-DRIVE WORKFLOW

A. Overview

Sky-Drive introduces a modular architecture designed as an integrated workflow for studying socially-aware autonomous driving and human-AI collaboration. As illustrated in Fig. 1, the platform's workflow seamlessly connects four currently implemented functional modules, with two additional modules planned for future development.

The workflow begins with the *digital twin framework*, which feeds high-fidelity virtual replicas of transportation systems into the *distributed multi-agent architecture*. This architecture then enables synchronized simulation across multiple devices, providing the foundation for complex interactions between autonomous agents. The simulation environment created by these two modules serves as the testing ground for the *multi-modal human-in-the-loop framework*, which captures comprehensive behavioral data from human participants. This data is subsequently processed and utilized by the *human-AI collaboration mechanism* to facilitate knowledge exchange between human and autonomous systems. The *foundation models integration* will enhance capabilities at both system and agent levels, providing global observation for system performance feedback while helping individual agents better understand human behavior patterns identified through the human-in-the-loop framework. The *hardware-in-the-loop module* will connect with the digital twin framework to enable direct validation of algorithms on physical platforms, completing the

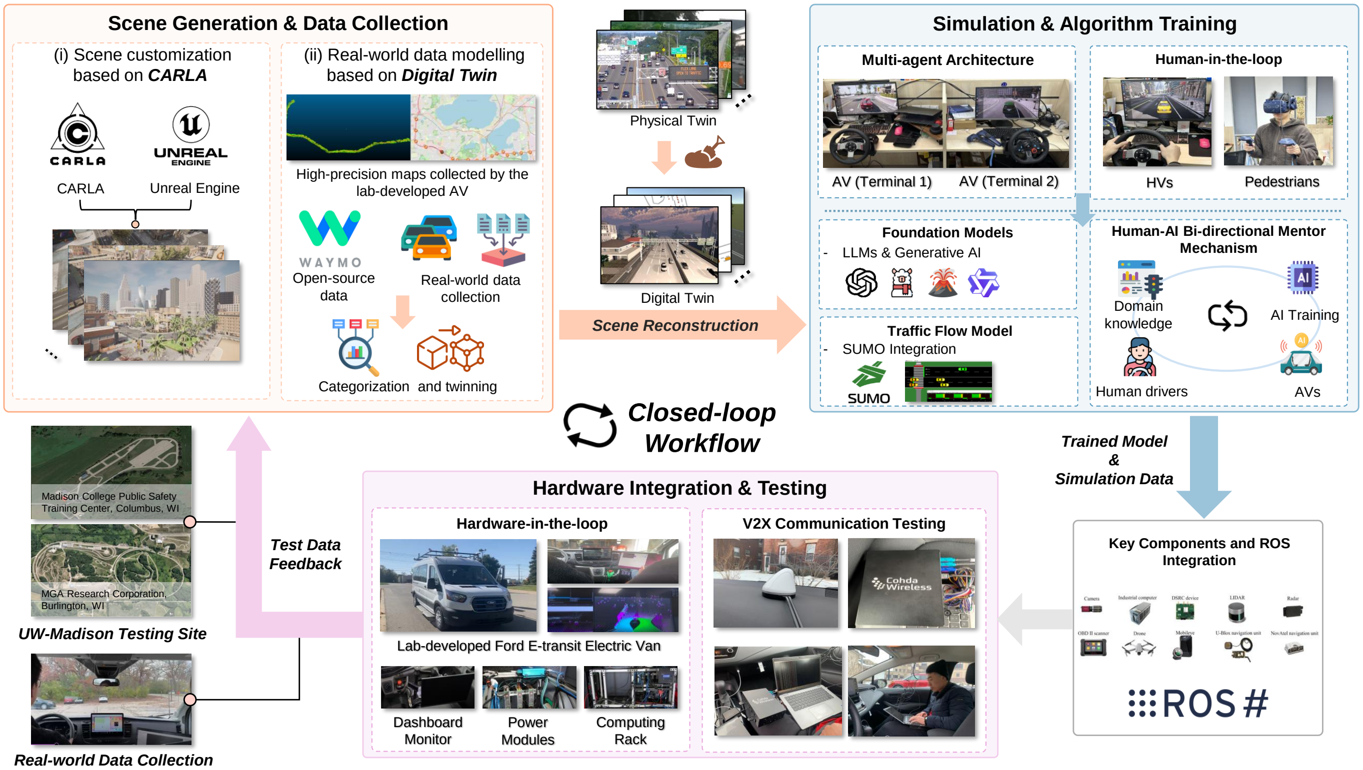

Fig. 2. Workflow of Sky-Drive. (a) scenario generation & data collection through CARLA-based synthetic environments and digital twin integration of real-world traffic data; (b) simulation & algorithm training enabled by distributed multi-agent architecture and human-AI bi-directional mentor mechanism; (c) hardware integration & testing utilizing ROS compatibility for direct validation of autonomous driving algorithms on physical platforms.

cycle by feeding real-world performance data back into the simulation environment.

B. Workflow

The workflow of Sky-Drive, shown in Fig. 2, consists of three primary stages that form a continuous feedback loop:

1) Scenario Generation & Data Collection: As depicted in Fig. 2(a), this stage employs two complementary approaches to ensure comprehensive scenario coverage: (i) Sky-Drive leverages CARLA and Unreal Engine to generate customizable urban environments with detailed road networks, traffic rules, and environmental conditions, enabling controlled testing of specific driving scenarios. (ii) The digital twin framework imports real-world data through multi-source integration. This includes high-precision maps collected by lab-developed autonomous vehicles, open-source data, and real-world traffic data collection. The collected data undergoes sophisticated categorization and twinning processes to create digital replicas of physical environments.

Through this dual approach, Sky-Drive achieves both precise controllability in synthetic scenarios and high fidelity in recreating real-world driving conditions, providing researchers with flexible testbeds for developing and validating autonomous driving algorithms.

2) Simulation & Algorithm Training: As shown in Fig. 2(b), this stage processes the generated scenarios through an integrated learning pipeline with four interconnected components: (i) The distributed multi-agent architecture enables the concurrent operation of multiple agents across different terminals. This allows independent control of agents on separate devices while maintaining synchronized simulation, facilitating complex traffic interactions in a shared environment. (ii) The Human-in-the-Loop component integrates multiple human participants into a simulation that directly captures human behavior through an immersive interface. This allows researchers to study human responses in a variety of traffic settings. (iii) Sky-drive will integrate LLMs/VLMs to enhance simulation capabilities. These models will facilitate natural communication between human participants and autonomous systems for more intuitive interaction and knowledge transfer. (iv) This human-AI collaboration mechanism integrates human feedback and domain knowledge with AI training processes, creating a continuous learning loop where humans inform AI systems provide feedback to human operators.

The integration of these components produces trained models and comprehensive simulation data that serve as inputs for the hardware integration and testing stage.

3) Hardware Integration & Testing: As illustrated in Fig. 2(c), the final stage bridges simulation and physical deployment through two key components: (i) While the full hardware-in-the-loop module is planned for future development, the current architecture already supports connections to external hardware through standardized ROS interfaces. The lab-developed Ford E-transit electric van serves as the primary testbed, equipped with dashboard monitors, power modules, and a computing rack for algorithm deployment. (ii) Sky-Drive

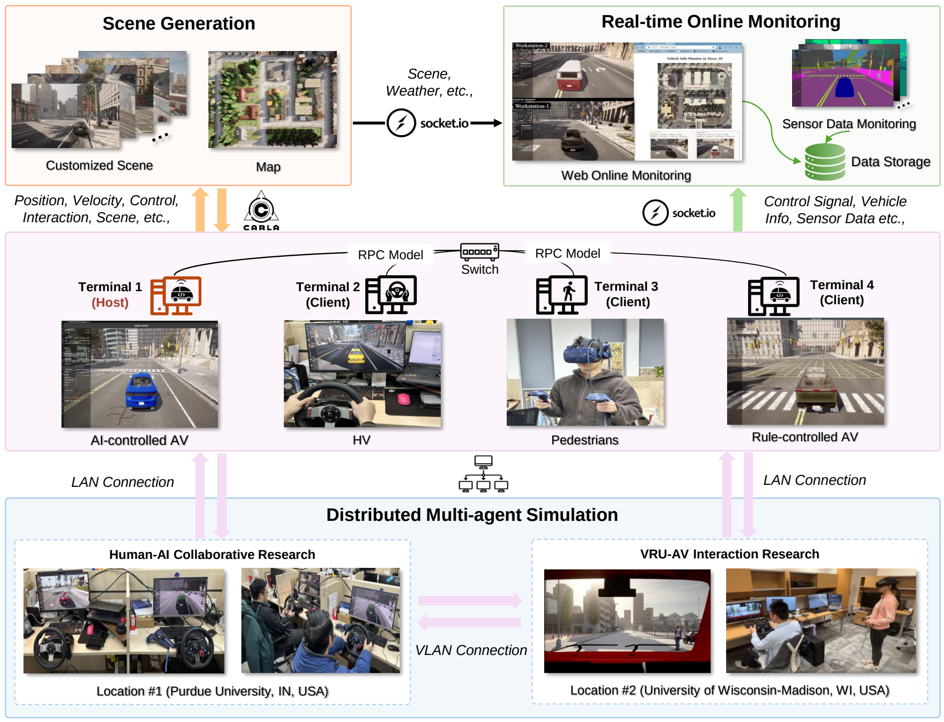

Fig. 3. Illustration of Sky-Drive's distributed multi-agent architecture. Sky-Drive enables synchronized simulation across multiple terminals while maintaining precise real-time interactions between AVs, HVs, and pedestrians through a sophisticated RPC networking model and Socket.IO-based communication platform, supporting comprehensive data collection and real-time analysis of multi-agent behaviors.

also supports the testing of vehicle-to-object (V2X) communication protocols, enabling evaluates cooperative perception and decision-making capabilities across multiple vehicles and infrastructure elements. This testing is essential for validating the performance of autonomous systems in complex traffic environments that require coordination between multiple agents.

This closed-loop workflow enables systematic development and validation of socially-aware autonomous driving systems and human-AI collaboration mechanisms, from initial concept testing through to real-world deployment, while maintaining safety and reliability throughout the process.

IV. Sky-Drive Features

A. Distributed Multi-agent Architecture

Sky-Drive introduces a novel distributed multi-agent architecture that enables synchronized simulation of multiple independently operating agents across different computing devices. As illustrated in Fig. 3, this architecture creates a comprehensive simulation environment where AVs, HVs, and pedestrians can interact in realistic traffic scenarios while being controlled from separate terminals.

1) System Architecture: Sky-Drive's core lies in a sophisticated RPC networking model built upon CARLA using the `rpclib` library. This implementation extends CARLA's proven vehicle control system while introducing crucial enhancements for distributed multi-agent simulation.

As shown in Fig. 3(c), Terminal 1 functions as the host (server) that maintains the global simulation environment, while Terminals 2-4 operate as clients controlling different agent types. Each terminal in the network can independently control its corresponding agent through various input devices while maintaining seamless interaction with other agents in the shared environment. The host terminal manages scene customization and map generation through CARLA and feeds this information to the distributed terminals. As illustrated in Fig. 3(a), the scene generation component creates detailed virtual environments with customizable traffic conditions, weather patterns, and road infrastructure. This architecture supports multiple agent types simultaneously, including AI-controlled AVs, HVs with steering wheel and keyboard interfaces, pedestrians controlled through VR systems, and rule-based AVs following predefined behaviors.

2) Communication Infrastructure: The communication infrastructure employs a dual-port TCP system on each terminal, enabling robust bidirectional data exchange between the host and clients. To achieve optimal performance in this distributed architecture, Sky-Drive implements a hybrid networking approach. For time-critical operations, we utilize a dedicated local area network (LAN) configured with high-performance switches and Ethernet connections. This setup achieves remarkably low latency, measured at 0.3 milliseconds, enabling smooth real-time interactions among agents. For scenarios requiring broader network coverage and geographically distributed research (as shown in Fig. 3(d) with locations at Purdue University and University of Wisconsin-Madison), the architecture employs virtual LAN (VLAN) configurations that maintain communication efficiency while extending the platform's reach.

3) Real-time Monitoring Platform: A key component of Sky-Drive's distributed architecture is its comprehensive monitoring and data management system. Complementing the core networking infrastructure, Sky-Drive developed a Socket.IO-based communication platform that monitors the real-time transmission of agent data, including position coordinates, velocity metrics, live video feeds, and sensor readings. As shown in Fig. 3(b), the platform features a web-based monitoring system that provides real-time visualization of all agent activities. This platform streams all data to a centralized system where agent interactions can be monitored and analyzed in real time. All simulation data, including agent states, environmental conditions, and interaction events, are automatically logged to a centralized database, facilitating comprehensive post-simulation analysis and scenario reproduction.

This distributed multi-agent architecture significantly advances autonomous driving research beyond existing platforms, such as Nocturne [35], MetaDrive [11], and Waymax [33], which primarily focus on simplified multi-agent interactions on a single machine. By enabling truly distributed control of multiple agents while maintaining precise synchronization across geographical locations, Sky-Drive creates a uniquely powerful environment for studying the complex social dynamics of future transportation systems.

B. Multi-modal Human-in-the-loop Framework

Sky-Drive provides a multi-modal human-in-the-loop framework that captures, synchronizes, and interprets rich behavioral signals from human participants.

1) Eye Tracking: Sky-Drive provides an immersive virtual reality experience through a custom-developed VR interface built on top of the Unreal Engine. Participants engage in the simulation using an HTC Vive Pro Eye headset, which supports full 6-DoF head tracking via SteamVR and integrated eye tracking via the SRanipal SDK. Our system captures high-frequency (up to 120 Hz) behavioral signals including 3D gaze vectors, pupil positions and diameters, eye openness, and fixation points. These data streams are essential for analyzing driver attention distribution, situational awareness, and cognitive state during complex driving tasks.

2) Voice Interaction: Sky-Drive supports voice commands as an explicit behavioral input modality. Spoken language is transcribed via Whisper, an OpenAI automatic speech recognition (ASR) model, and then interpreted by LLMs such as GPT-4. These models extract driver intent and sentiment from natural language input, whether structured ("slow down at the next intersection") or informal ("too fast"). LLMs map these expressions to semantic driving directives or policy preferences, which are then translated into model guidance signals or policy constraints

3) Facial Expression Recognition: A high-resolution in-cabin camera captures facial micro-expressions in real time. Sky-Drive employs a convolutional neural network-based expression classification model (e.g., EfficientFace or MobileNet-V2) trained on affective datasets to recognize expressions such as stress, confusion, or satisfaction. These cues serve as implicit indicators of driver state and comfort, which can modulate AI policy updates or trigger real-time interventions in the AIHM framework.

4) Physiological Signal Monitoring: Physiological states such as stress or alertness are inferred through biometric signals collected by wearable devices. Sky-Drive integrates the Garmin vívoactive 5 smartwatch to monitor heart rate and heart rate variability (HRV) in real time. This device provides continuous physiological data synchronized with other behavioral inputs, offering valuable indicators of driver arousal, cognitive workload, and fatigue. These signals are timestamp-aligned with other modalities and provide an additional implicit channel for modeling driver state and adapting AI behavior accordingly.

5) Steer Wheel: To support realistic interaction, the ego vehicle is equipped with a Logitech G920 racing wheel and pedal system, with force feedback enabled through the open-source LogitechWheelPlugin. All inputs—steering, throttle, braking, and signaling—are logged in parallel with gaze and head pose data and are compatible with CARLA's ScenarioRunner for scenario-based experiments.

This framework provides the technical backbone for the human-AI collaboration mechanism, transforming traditional user input into actionable and structured knowledge for collaborative decision-making.

C. Human-AI Collaboration Mechanism

Sky-Drive implements an innovative human-AI collaboration mechanism that establishes effective knowledge exchange between humans and AI-enabled autonomous systems. This mechanism consists of two primary frameworks: Human as AI Mentor (HAIM) and AI as Human Mentor (AIHM).

1) Human as AI Mentor: In the HAIM framework, humans act as real-time mentors to AI-enabled autonomous vehicles, guiding AI behavior through rich, multi-source human knowledge. This knowledge consists of two complementary components: (i) individual behavioral knowledge, which includes both explicit behaviors (e.g., takeover actions, voice commands, touchscreen interactions) and implicit behaviors (e.g., facial expressions, eye movements, physiological signals) collected from human drivers through Sky-Drive's multi-modal human-in-the-loop framework; and (ii) domain knowledge from transportation science, which includes validated principles derived from decades of research, such as car-following models and lane-change behavior theories.

Sky-Drive adopts an RL framework enhanced by human preference modeling and physics-informed priors to integrate this dual-source human knowledge into the AI learning process. Rather than relying on manually crafted reward functions, Sky-Drive formulates the learning problem as preference-based policy optimization. Specifically, individual behavioral knowledge is incorporated as guidance signals during learning. For example, frequent human takeovers in specific contexts (e.g., near intersections or during aggressive merging) are interpreted as indicators of suboptimal AI behavior. These signals are used to define implicit cost functions or to prioritize policy adjustments via human-aligned trajectory comparisons. Meanwhile, domain knowledge from transportation models (e.g., IDM and MOBIL) is embedded into the learning framework as safety constraints and behavioral baselines, ensuring that learned policies remain physically plausible, stable, and socially compliant. This approach accelerates policy convergence, reduces unsafe exploration, and fosters trust between human drivers and autonomous systems.

2) AI as Human Mentor: The AIHM framework complements HAIM by allowing AI systems to act as real-time coaches and trainers for human drivers. It utilizes physics-enhanced residual learning (PERL) [47] to generate optimal driving paths that account for traffic dynamics, safety constraints, and personalized behavior models. These reference paths serve as a dynamic instructional baseline for the human driver, visualized in real time via an in-vehicle display or VR interface and updated continuously based on the driver's ongoing performance. AIHM evaluates human driver performance through multi-dimensional metrics, including path adherence, reaction time, control smoothness, and situational awareness. Personalized feedback is delivered via visualizations (e.g., 3D paths, heat maps), annotated replay, and AI-generated verbal summaries.

A key innovation of AIHM is the integration of generative AI for adaptive scenario creation. Based on the driver's recent performance, the framework identifies areas requiring improvement—such as emergency braking, lane discipline, or roundabout negotiation—and procedurally generates customized training scenarios to target those weaknesses. In addition, the AIHM framework uses real-time physiological and cognitive data to dynamically modulate the level of guidance. For example, if the system detects elevated heart rate and frequent steering correction—indicating stress or confusion—it can proactively reduce scenario complexity, simplify instructions, or offer reassurance. Conversely, when the driver exhibits proficiency, the system can gradually increase challenge levels to promote growth.

D. Digital Twin Framework

The digital twin framework creates high-fidelity virtual replicas of transportation systems through sophisticated multi-source data integration. This framework serves as the environmental backbone of the Sky-Drive platform, enabling data-driven scenario generation, human-in-the-loop simulation, and hybrid validation across both simulated and physical domains.

The framework comprises three primary components that work together to create accurate digital replicas. The multi-source data integration layer combines diverse data streams from roadside units, traffic cameras, historical traffic records, and high-definition mapping data collected by Lab-developed AVs equipped with LiDAR and cameras. These data sources undergo temporal alignment and spatial correlation to ensure consistency across all inputs. The virtual environment component, built on Unreal Engine, serves as a dynamic visualization platform that provides real-time 3D rendering of traffic conditions. This environment processes incoming sensor data using computer vision algorithms to detect and track road users for both real-time rendering and trajectory prediction [48]. The edge-computing architecture processes data locally through roadside units equipped with sensors that monitor traffic flow, vehicle behavior, and environmental conditions, performing preliminary data processing before transmitting relevant information to the central system.

Sky-Drive has successfully implemented a pilot deployment of this digital twin system on Wisconsin's Flex Lane along the Beltline in Dane County. This implementation demonstrates the framework's ability to process real-time traffic data, generate accurate state predictions, and provide decision support for traffic management by integrating data from Wisconsin DOT's traffic cameras, autonomous vehicle surveys, and historical databases.

V. Sky-Drive Application Case

A. VR-based AV-VRU Interaction

Sky-Drive provides a platform for studying complex interactions between AVs and vulnerable road users (VRUs) through its advanced VR-enabled simulation capabilities. The platform's distributed multi-agent architecture enables realistic modeling of dynamic traffic scenarios. It allows independent control of multiple agents across different devices while maintaining precise synchronization. This capability is particularly valuable for investigating safety-critical interactions between AVs and VRUs, which are challenging to study in real-world environments due to safety concerns.

As shown in Fig. 4(a), we conducted a case study focused on right-turn conflicts at unsignalized intersections—a scenario frequently associated with accidents in urban environments. This study leveraged Sky-Drive's synchronized multi-terminal architecture in a novel experimental setup where human participants experienced the scenario from the pedestrian's perspective through immersive VR, while researchers controlled an AV making right turns from a separate terminal. During each interaction, Sky-Drive captured multimodal behavioral data from both the AV and the pedestrian. The VR recorded 3D gaze vectors, eye fixations, and reaction times from the pedestrian, while simultaneously logging control signals, deceleration profiles, and trajectory predictions from the AV.

This integrated setup enables researchers to study not only the physical outcomes of AV–VRU interactions (e.g., successful yielding, near-misses), but also the cognitive and emotional dimensions of human response.

B. HAIM-based Deep Reinforcement Learning

Sky-Drive's human-AI collaboration mechanism provides a HAIM framework for learning autonomous driving policies directly from human feedback. To demonstrate the platform's capabilities, as shown in Fig. 4(b), we implemented HAIM-DRL [49], a reward-free RL framework that infers driver preferences from control takeovers. Unlike conventional RL methods that rely on handcrafted reward functions, HAIM-DRL leverages human interventions to guide policy optimization.

Sky-Drive implements HAIM-DRL by detecting and recording steering takeovers while synchronizing with vehicle state and scene context. The platform's multi-agent simulation environment allows HAIM-DRL to operate in traffic flow

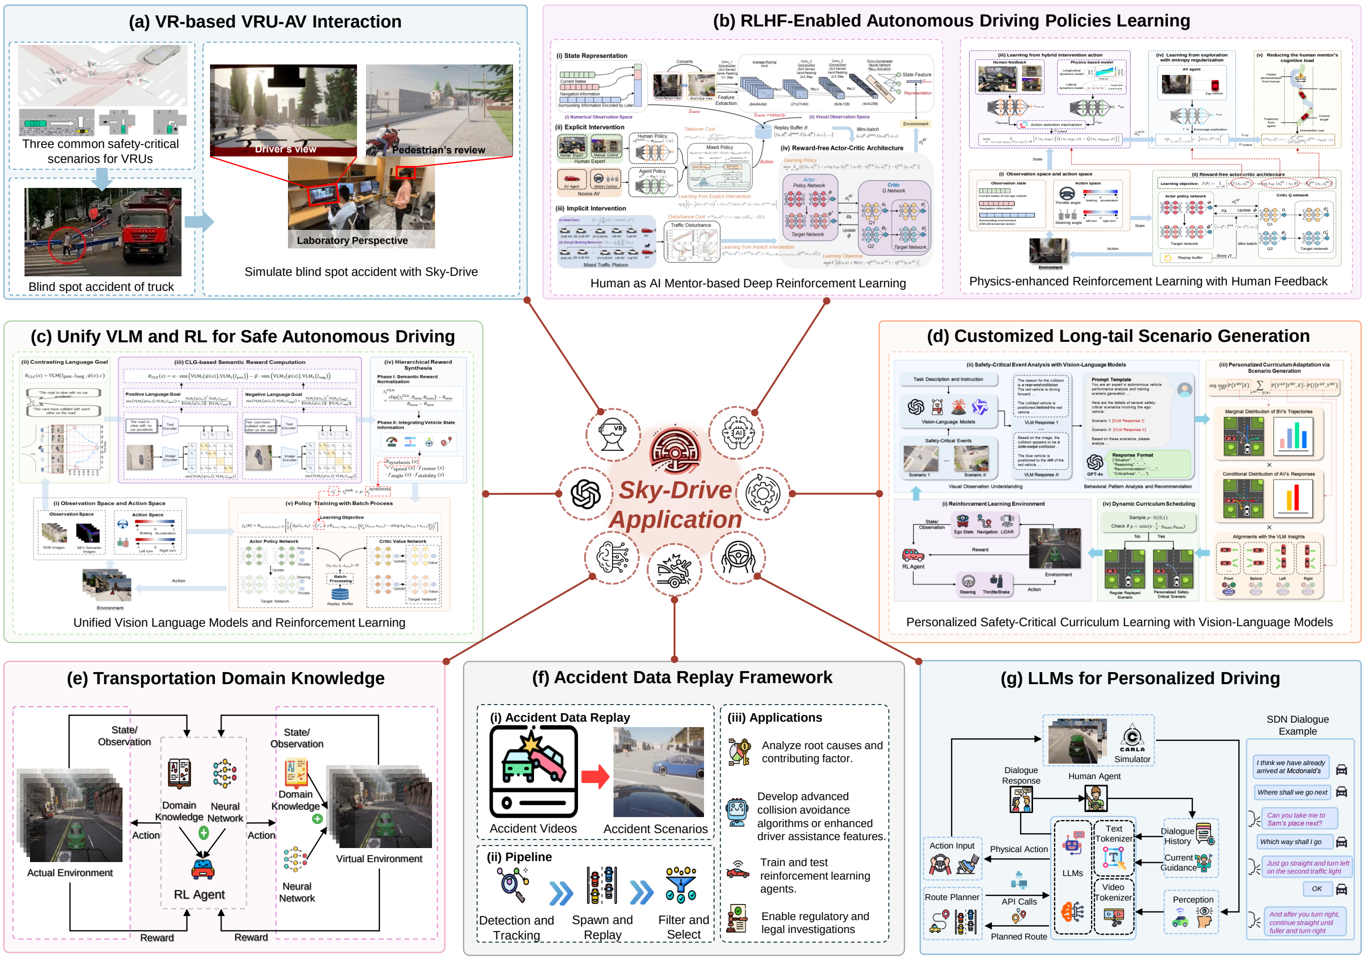

Fig. 4. Demonstration of Sky-Drive's key applications. (a) VR-enabled interaction studies between autonomous vehicles and vulnerable road users; (b) RLHF-enabled autonomous driving policy learning through HAIM-DRL and PE-RLHF frameworks; (c) VLM-RL framework integrating vision-language models with reinforcement learning for safe driving; (d) Customized long-tail scenario generation using CurricuVLM for personalized training; (e) Transportation domain knowledge integration through physics-enhanced reinforcement learning; (f) Accident data replay framework for systematic traffic incident analysis; (g) LLM-based system enabling personalized autonomous driving.

scenarios, where humans take over implicitly dissatisfied with agent behavior, such as when the AV merges too aggressively or follows too closely. By constructing human-aligned preference comparisons between pre- and post-takeover trajectories, the agent learns to avoid human-disapproved behaviors and improve its driving policy accordingly.

Mathematically, we can define the HAIM framework as follows: aligning AI behavior with human preferences as closely as possible:

$$\pi^*_{\text{AV}} = \arg\min_{\pi_{\text{AV}}} \mathbb{E}_{s_t \sim d_{\pi_{\text{AV}}}} \left[\mathcal{L}\left(\pi_{\text{AV}}(\cdot \mid s_t), \pi_{\text{human}}(\cdot \mid s_t)\right)\right], \quad (1)$$

where $d_{\pi_{\text{AV}}}$ represents the state distribution induced by the agent's policy π_{AV}, and $\mathcal{L}(\cdot,\cdot)$ is a measure of discrepancy (e.g., KL divergence). By minimizing this discrepancy over the state distribution, the AI agent is encouraged to learn from human knowledge and align its behavior with human preferences.

HAIM-DRL utilizes takeover actions as human knowledge input. A human expert acts as a mentor to the AI agent (AI-enabled AV), intervening and taking control of the vehicle in hazardous situations, demonstrating correct maneuvers to avoid potential accidents. To integrate data from agent exploration and human takeover, HAIM-DRL designs a switch function $\mathcal{T}$. Let $\mathcal{T}(a_t) = 1$ indicate that the human driver takes over the control, and $\mathcal{T}(a_t) = 0$ otherwise. We represent this process as follows:

$$\mathcal{T}(s_t, a_t, \pi_{\text{human}}) = \begin{cases} (a_t^{\text{AV}}, 0), & \text{if takeover;} \\ (a_t^{\text{human}} \sim \pi_{\text{human}}(\cdot \mid s_t), 1), & \text{otherwise.} \end{cases} \quad (2)$$

The actual trajectory during the training process is determined by the mixed behavior policy:

$$\pi_{\text{mix}}(a \mid s) = \pi_{\text{AV}}(a \mid s)(1 - I(s,a)) + \pi_{\text{human}}(a \mid s)F(s), \quad (3)$$

where $F(s) = \int_{a' \notin A_\eta(s)} \pi_{\text{AV}}(a' \mid s)\, da'$ represents the probability of the agent selecting an action that would be rejected by the human.

The overall learning objective of HAIM-DRL is specifically

TABLE II
THE PERFORMANCE OF PPO/HACO/HAIM-DRL METHODS IN THE CARLA SIMULATOR.

Method	Test Safety Violation	Test Return	Test Disturbance Rate	Test Success Rate	Train Samples
PPO	80.84	1591.00	-	0.35	500,000
HACO	12.14	1578.43	0.0137	0.35	8,000
HAIM-DRL	11.25	1590.85	0.0121	0.38	8,000

designed as:

$$\max_{\pi} \mathbb{E}\left[\psi \hat{Q}(s_t, a_t^{\mathrm{AV}}) - \alpha \log \pi_{\mathrm{AV}}(a_t^{\mathrm{AV}} \mid s_t; \theta) - \beta Q^{\mathrm{EX}}(s_t, a_t^{\mathrm{AV}}) - \varphi Q^{\mathrm{IM}}(s_t, a_t^{\mathrm{AV}})\right] \tag{4}$$

The first term is defined as

$$\hat{Q}(s_t, a_t^{\mathrm{AV}}) = \min_{\phi} \mathbb{E}_{(s_t, a_t^{\mathrm{AV}}, a_t^{\mathrm{human}}, I(s_t, a_t^{\mathrm{AV}}))\sim\mathcal{B}}\Big[I(s_t, a_t^{\mathrm{AV}})\big(\hat{Q}(s_t, a_t^{\mathrm{AV}}; \phi) - \hat{Q}(s_t, a_t^{\mathrm{human}}; \phi)\big)\Big]$$

which ensures that the AI agent mimics human-preferred behavior by minimizing the value discrepancy between its own actions and those demonstrated by a human mentor. The fourth term is defined as

$$Q^{\mathrm{IM}}(s_t, a_t^{\mathrm{AV}}) = C^{\mathrm{IM}}(s_t, a_t^{\mathrm{AV}}) + \gamma \mathbb{E}_{s_{t+1}\sim\mathcal{B}, a_{t+1}\sim\pi_{\mathrm{AV}}(\cdot|s_{t+1})}\left[Q^{\mathrm{IM}}(s_{t+1}, a_{t+1}^{\mathrm{AV}})\right]$$

where $C^{\mathrm{IM}}(s_t, a_t^{\mathrm{AV}})$ represents the traffic disturbance cost, and $I(s_t, a_t^{\mathrm{AV}})$ is an indicator function that equals 1 if the human rejects the action a_t^{AV}, and 0 otherwise.

As evidenced by Fig. **??** and Table II, HAIM-DRL method is successfully implemented in the Sky-Drive, where it exhibits superior performance, particularly in safety violation, success rate, and disturbance rate. The success of this implementation underscores Sky-Drive's ability to support closed-loop human-AI training and validate human-guided policies in simulated environments.

C. Physics-enhanced Reinforcement Learning with Human Feedback

Sky-Drive's human-AI collaboration mechanism also integrates the Physics-enhanced Reinforcement Learning with Human Feedback (PE-RLHF) framework for developing trustworthy autonomous driving policies. As shown in Fig. 4(c), we implemented PE-RLHF, a novel framework that synergistically combines human feedback with physics knowledge from traffic flow models to ensure safe and efficient driving decisions.

Unlike traditional RLHF methods that may falter with imperfect human feedback, PE-RLHF establishes a trustworthy safety performance lower bound through well-established traffic flow models. Sky-Drive implements PE-RLHF by incorporating a sophisticated Physics-enhanced Human-AI (PE-HAI) collaborative paradigm where three policies interact: a human policy (π_{human}), a physics-based policy (π_{phy}) derived from traffic flow models (IDM-MOBIL), and an AV policy (π_{AV}).

The platform detects human takeovers while simultaneously evaluating actions generated by both the human and physics-based models. During intervention, an action selection mechanism determines which action to execute based on estimated Q-values:

$$a_{\mathrm{hybrid}} = T_{\mathrm{select}}(s) = \begin{cases} a_{\mathrm{human}}, & \text{if Mean}\left[E_{a\sim\pi_{\mathrm{human}}(\cdot|s)} Q^{\phi}(s,a) - E_{a\sim\pi_{\mathrm{phy}}}\right. \\ a_{\mathrm{phy}}, & \text{otherwise} \end{cases} \tag{5}$$

This ensures that the system always executes the action with higher expected value, establishing a performance floor guaranteed by interpretable physics-based models, even when human feedback quality deteriorates.

The overall learning objective of PE-RLHF is formulated as:

$$\max_{\pi} \mathbb{E}\left[\psi \hat{Q}(s_t, a_t^{\mathrm{AV}}) - \alpha \log \pi_{\mathrm{AV}}(a_t^{\mathrm{AV}}|s_t; \theta) - \beta Q^{\mathrm{int}}(s_t, a_t^{\mathrm{AV}})\right] \tag{6}$$

where $\hat{Q}(s_t, a_t^{\mathrm{AV}})$ is a proxy value function representing human preferences, the entropy term encourages exploration, and $Q^{\mathrm{int}}(s_t, a_t^{\mathrm{AV}})$ minimizes the need for human intervention.

Sky-Drive's multi-agent simulation environment enables PE-RLHF to operate effectively in complex traffic scenarios. The platform's digital twin framework provides realistic environments to test the physics-based policies, while the multi-modal human-in-the-loop framework captures nuanced human feedback through various input channels.

PE-RLHF demonstrates exceptional performance across key metrics when implemented in Sky-Drive, achieving a 91% reduction in safety violations compared to traditional RL methods while maintaining high success rates and trajectory efficiency. Particularly impressive is its ability to perform complex maneuvers such as overtaking and navigating around obstacles, which traditional physics-based models struggle with.

Tab. III shows the performance comparison of PE-RLHF with different physics-based model combinations and the standalone IDM-MOBIL model. In Stage I, we observe that PE-RLHF consistently outperforms the standalone IDM-MOBIL model across all configurations. The full PE-RLHF (with IDM-MOBIL) achieves the highest episodic return of 391.48 and a success rate of 0.85, compared to 206.30 and 0.31 for the standalone IDM-MOBIL model, respectively. This substantial improvement demonstrates the effectiveness of integrating RL with physics-based models. Moving to Stage II, we note that all PE-RLHF variants exhibit lower safety violations compared to the standalone IDM-MOBIL model. The full PE-RLHF configuration achieves the lowest safety violation of 0.47, indicating superior safety performance. Additionally, PE-RLHF variants consistently achieve greater travel distances, with the full configuration reaching 177.00m compared to 108.56m for the standalone model. In Stage III, the full PE-RLHF achieves the highest travel velocity (21.85km/h) and

TABLE III
PERFORMANCE COMPARISON OF PE-RLHF WITH DIFFERENT PHYSICS-BASED MODEL COMBINATIONS.

Method	**Driving Operation**	**Training**	**Testing**					
			Stage I		Stage II		Stage III	
		Total Safety Violation ↓	Episodic Return ↑	Success Rate (%) ↑	Safety Violation ↓	Travel Distance ↑	Travel Velocity ↑	Total Overtake Count ↑
IDM-MOBIL	Longitudinal & Lateral	-	206.30 $\pm_{35.23}$	0.31 $\pm_{0.15}$	0.49 $\pm_{0.08}$	108.56 $\pm_{55.23}$	19.78 $\pm_{2.67}$	0 $\pm_{0}$
PE-RLHF (without)	-	39.45 $\pm_{12.32}$	302.67 $\pm_{21.88}$	0.73 $\pm_{0.05}$	1.48 $\pm_{0.43}$	138.23 $\pm_{4.28}$	16.58 $\pm_{0.96}$	6.14 $\pm_{1.12}$
PE-RLHF (with IDM)	Longitudinal only	28.79 $\pm_{9.97}$	348.52 $\pm_{19.67}$	0.79 $\pm_{0.03}$	0.98 $\pm_{0.29}$	149.87 $\pm_{4.10}$	18.92 $\pm_{0.94}$	7.83 $\pm_{1.03}$
PE-RLHF (with MOBIL)	Lateral only	21.56 $\pm_{8.54}$	368.11 $\pm_{18.45}$	0.81 $\pm_{0.04}$	0.74 $\pm_{0.19}$	159.34 $\pm_{3.14}$	20.43 $\pm_{0.51}$	9.76 $\pm_{1.17}$
PE-RLHF (with IDM-MOBIL)	Longitudinal & Lateral	**16.61** $\pm_{9.96}$	**391.48** $\pm_{\mathbf{20.47}}$	**0.85** $\pm_{0.04}$	**0.47** $\pm_{0.01}$	**177.00** $\pm_{3.74}$	**21.85** $\pm_{0.02}$	**16.33** $\pm_{4.61}$

total overtake count (16.33), significantly outperforming the standalone IDM-MOBIL model (19.78 and 0, respectively).

Interestingly, we observe that incorporating either longitudinal (IDM) or lateral (MOBIL) components of the physics-based model into PE-RLHF yields improvements over the variant without any physics-based model. Yet, the combination of both IDM and MOBIL produces the best results across all metrics, suggesting a synergistic effect when integrating both longitudinal and lateral control models. It is worth noting that while the standalone IDM-MOBIL model provides a baseline level of performance, it struggles with overtaking maneuvers, as evidenced by its zero overtake count. In contrast, all PE-RLHF variants demonstrate the ability to perform overtaking, with the full configuration showing the highest proficiency in this regard. The results demonstrate that the PE-RLHF framework not only leverages the safety guarantees provided by these models but also enhances their performance through learning.

The integration of physics knowledge into the learning process uniquely positions PE-RLHF to address the challenges of safety-critical autonomous driving scenarios, making it an invaluable addition to Sky-Drive's human-AI collaboration capabilities.

D. VLM-RL

Sky-Drive's human-AI collaboration mechanism also supports the VLM-RL framework, as shown in Fig. 4(c). VLM-RL integrates pre-trained Vision-Language Models (VLMs) with Reinforcement Learning (RL) to generate reward signals using image observation and natural language goals for safe autonomous driving.

The core of VLM-RL is the Contrasting Language Goal (CLG)-as-reward paradigm, which leverages pre-trained VLMs to measure semantic alignment between driving states and contrasting language descriptions. Specifically, positive language goals (e.g., "the road is clear with no car accidents") and negative language goals (e.g., "two cars have collided with each other on the road") are used to guide the learning process, providing more informative and context-aware rewards.

We implement VLM-RL by first encoding RGB images through the CLIP vision encoder and language goals through the text encoder to obtain their respective embeddings in a shared latent space. The reward is computed as:

$$R_{CLG}(s) = \alpha \cdot \mathrm{sim}(VLM_I(\psi(s)), VLM_L(l_{pos})) - \beta \cdot \mathrm{sim}(VLM_I(\psi(s)), V \tag{7}$$

where $\mathrm{sim}(\cdot,\cdot)$ denotes the cosine similarity between embeddings, and $\alpha, \beta > 0$ are weighting factors. This formulation encourages the agent to seek states similar to the positive goal while avoiding states similar to the negative goal.

To enhance learning stability, we introduce a hierarchical reward synthesis approach that combines CLG-based semantic rewards with vehicle state information, providing comprehensive and stable reward signals. Additionally, a batch-processing technique is employed to optimize computational efficiency during training, where batches of observations are periodically sampled from a replay buffer and processed through the pre-trained VLM.

Extensive experiments in the CARLA simulator demonstrate that VLM-RL outperforms state-of-the-art baselines, achieving a 10.5% reduction in collision rate, a 104.6% increase in route completion rate, and robust generalization to unseen driving scenarios. This approach provides a more balanced and comprehensive learning signal compared to existing methods that rely solely on positive or negative goals, enabling the agent to better navigate the complex trade-offs between safety and efficiency in autonomous driving.

Through Sky-Drive's human-AI collaboration mechanism, VLM-RL can seamlessly integrate human feedback to refine contrasting language goals and enhance semantic reward signals, thereby improving the safety performance of autonomous systems while maintaining their efficiency in complex traffic scenarios.

To further validate the effectiveness of VLM-RL, we conduct comprehensive testing evaluations across 10 predefined routes and compare the performance with baseline methods. The route completion metric represents the average route

TABLE IV
PERFORMANCE COMPARISON WITH BASELINES DURING TESTING. MEAN AND STANDARD DEVIATION OVER 3 SEEDS. THE BEST RESULTS ARE MARKED IN **BOLD**.

Model	AS ↑	RC ↑	TD ↑	CS ↓	SR ↑
LLM-based Reward Methods					
Revolve	18.4 ± 0.03	0.92 ± 0.11	1915.3 ± 248.3	1.53 ± 2.16	0.83 ± 0.24
Revolve-auto	17.2 ± 0.76	0.80 ± 0.06	1539.6 ± 147.5	1.65 ± 0.28	0.63 ± 0.05
VLM-based Reward Methods					
VLM-SR	0.53 ± 0.27	0.02 ± 0.00	47.9 ± 9.2	0.18 ± 0.25	0.0 ± 0.0
RoboCLIP	0.44 ± 0.05	0.07 ± 0.03	146.3 ± 62.3	1.05 ± 0.58	0.0 ± 0.0
VLM-RM	0.20 ± 0.05	0.02 ± 0.01	35.9 ± 25.8	**0.003** ± 0.005	0.0 ± 0.0
LORD	0.17 ± 0.08	0.02 ± 0.02	45.1 ± 57.1	0.02 ± 0.02	0.0 ± 0.0
LORD-Speed	18.9 ± 0.36	0.87 ± 0.08	1783.4 ± 172.8	2.80 ± 1.16	0.67 ± 0.05
VLM-RL (ours)	19.3 ± 1.29	**0.97** ± 0.03	**2028.2** ± 96.6	0.02 ± 0.03	**0.93** ± 0.04

completion rates during each evaluation episode. The testing results in Tab. IV demonstrate significant advantages of our approach compared to the baselines.

LLM-based approaches demonstrate competitive performance during testing, with Revolve achieving a success rate of 0.83 and route completion of 0.92. However, their collision speeds of 1.53 km/h and 1.65 km/h indicate persistent safety issues. Most VLM-based methods, including VLM-SR, RoboCLIP, VLM-RM, and LORD, exhibit highly conservative behaviors with route completion rates below 0.07 and success rates of 0.0. LORD-Speed shows significantly improved efficiency metrics but records the highest collision speed at 2.80 km/h among all methods.

In contrast, VLM-RL achieves superior performance across all key metrics during testing. It maintains a high average speed of 19.3 km/h while recording a low collision speed of 0.02 km/h, matching the safety level of the most conservative approaches. Most notably, VLM-RL achieves the highest success rate of 0.93 and route completion of 0.97, along with the longest total driving distance of 2028.2m. These results demonstrate that our method not only learns more effective driving policies but also exhibits better generalization to testing scenarios. The significant improvements in both efficiency and safety metrics validate the effectiveness of our CLG-based and hierarchical reward design in providing comprehensive and well-balanced learning signals for safe driving tasks.

E. CurricuVLM

Sky-Drive's human-AI collaboration mechanism also supports the CurricuVLM framework, as shown in Fig. 4(d). CurricuVLM leverages Vision-Language Models (VLMs) to enable personalized safety-critical curriculum learning for autonomous driving agents.

The core innovation of CurricuVLM lies in its ability to bridge the gap between scenario generation and driving policy learning. By continuously monitoring agent behavior in various driving scenarios, CurricuVLM employs VLMs to analyze safety-critical events and identify recurring weaknesses. When unsafe situations occur, the framework generates comprehensive visual descriptions that are processed through a specialized GPT-4o-based analyzer. This two-stage behavior analysis pipeline combines VLMs' visual understanding capabilities with GPT-4o's reasoning abilities to assess the agent's current capabilities and limitations.

Based on this analysis, CurricuVLM formulates the scenario generation as a conditional trajectory generation problem, which can be expressed as:

$$P(Y^{AV}, Y^{BV} | I, X) \tag{8}$$

where X represents historical information including the HD map and past states of both AV and background vehicles, Y^{AV} and Y^{BV} denote their future trajectories, and I contains the behavioral insights from VLM analysis. The framework optimizes background vehicle trajectories by finding:

$$Y^{BV*} = \arg\max_{Y^{BV}} P(Y^{BV}|X) \sum_{Y^{AV} \sim \mathcal{Y}(\pi)} P(Y^{AV}|Y^{BV}, X) \cdot P(I|Y^{AV}, Y^{B} \tag{9}$$

This approach allows CurricuVLM to create diverse and realistic safety-critical scenarios that specifically target the agent's identified weaknesses, enabling more effective closed-loop training.

Experimental results in Sky-Drive demonstrate that CurricuVLM significantly outperforms state-of-the-art methods across both regular and safety-critical scenarios. In safety-critical testing, CurricuVLM achieves an episode reward of 48.9 compared to 42.5 for CAT and 39.3 for CLIC, while maintaining a lower crash rate of 25.1% versus 32.1% and 26.2% respectively. The framework shows strong compatibility with various RL algorithms including TD3, PPO, and SAC, demonstrating its potential as a general approach for enhancing autonomous driving systems.

Through Sky-Drive's integration, CurricuVLM enables more effective learning from safety-critical scenarios by dynamically generating personalized curricula that adapt to each agent's evolving capabilities, ultimately improving safety and performance in autonomous driving systems.

F. Talk2Traffic

Sky-Drive's human-AI collaboration mechanism also incorporates the Talk2Traffic framework, as shown in Fig. 4(f). Talk2Traffic leverages multimodal large language models (MLLMs) to enable intuitive and editable traffic scenario generation through natural language instructions, speech commands, and sketch-based inputs.

The core innovation of Talk2Traffic lies in its ability to bridge the gap between human designers' intuitive expressions and executable simulation scenarios. The framework processes diverse multimodal inputs through a specialized interpreter that extracts structured scene representations. These representations are then transformed into executable Scenic code using a retrieval-augmented generation (RAG) approach with a curated database of verified code snippets.

Talk2Traffic's multimodal instruction interpreter can be formulated as:

$$\mathbf{z} = \text{MLLM}(p, l, s), \tag{10}$$

where p is the task description, l represents textual or speech instructions, and s denotes sketch-based inputs. The extracted structured representation $\mathbf{z}$ captures essential scenario components including map configuration, weather conditions, and agent specifications.

To ensure high-quality code generation, Talk2Traffic employs a comprehensive database of description-snippet pairs:

$$\mathcal{D} = \{(d_j, c_j) | j \in \{1, \ldots, m\}\}, \tag{11}$$

where d_j represents natural language descriptions and c_j denotes corresponding Scenic code snippets. For each component of the structured representation, the framework retrieves relevant snippets based on semantic similarity computed through:

$$\text{sim}(\mathbf{s}_i^q, \mathbf{s}_j) = \frac{\mathbf{s}_i^q \cdot \mathbf{s}_j}{||\mathbf{s}_i^q|| \cdot ||\mathbf{s}_j||}, \tag{12}$$

where $\mathbf{s}_i^q$ and $\mathbf{s}_j$ are embeddings of the query and database descriptions respectively.

A distinguishing feature of Talk2Traffic is its human feedback guidance mechanism that enables iterative refinement of generated scenarios:

$$\mathcal{S}_{t+1} = \text{MLLM}(p_{\text{refine}}, \mathcal{H}_t, \mathcal{S}_t, f_t), \tag{13}$$

where $\mathcal{S}_t$ represents the scenario at iteration t, $\mathcal{H}_t$ is the conversation history, and f_t denotes user feedback. This interactive approach allows users to progressively align scenarios with specific testing objectives through natural language guidance.

Experimental results demonstrate Talk2Traffic's effectiveness in generating challenging traffic scenarios. The framework achieves state-of-the-art performance with an average collision rate of 0.877 across diverse scenario types, surpassing the second-best method by 4.6%. Talk2Traffic particularly excels in complex multi-agent interactions such as Red Light Running (0.900) and Unprotected Left Turn (0.833) scenarios, demonstrating its capability to create sophisticated testing environments for autonomous driving systems.

Through this integration, Sky-Drive's AIHM framework delivers on its promise of personalized driver training, using Talk2Traffic to dynamically generate the challenging scenarios needed to systematically improve driver performance across diverse traffic conditions and driving situations.

G. Accident Data Replay Framework

Sky-Drive implements an accident data replay framework that enables systematic reconstruction and analysis of real-world traffic accidents within its simulation environment. This capability addresses a critical need in autonomous driving development: understanding and learning from actual accident scenarios while maintaining safety and reproducibility. The framework employs a comprehensive technical pipeline centered around CenterTrack [50], an advanced detection and tracking algorithm that processes accident video footage to extract detailed object trajectories. Sky-Drive's integration with CARLA transforms these trajectories into precise 3D reconstructions that capture the relative positions and movements of all vehicles involved in the incident. The platform's high-fidelity simulation capabilities ensure accurate reproduction of critical environmental factors, including road conditions, vehicle dynamics, and interaction patterns.

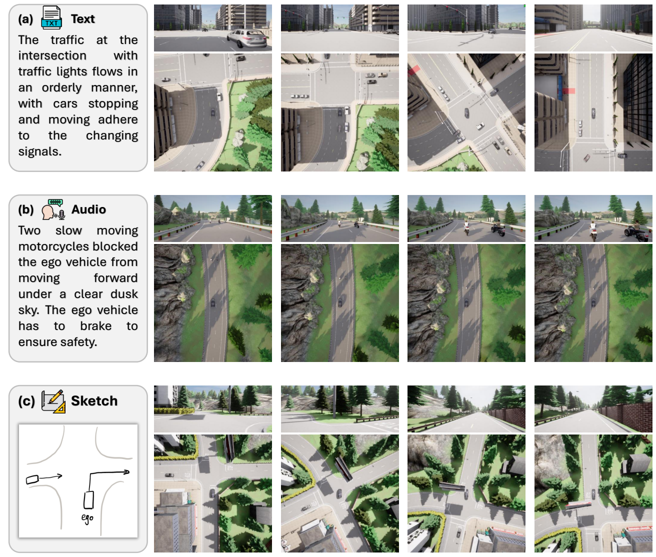

Fig. 5. Qualitative examples. Each scenario is downsampled to four frames for visualisation.

To ensure reconstruction accuracy, as shown in Fig. 4 (f), Sky-Drive implements a sophisticated validation process. The framework employs procedural matching algorithms to identify appropriate simulation maps that closely mirror the original accident conditions. A quality assessment module evaluates the fidelity of each reconstructed scenario, automatically flagging cases that require additional refinement. While the platform incorporates unsupervised domain adaptation techniques to enhance trajectory extraction accuracy, it also provides tools for manual refinement when needed, ensuring the highest possible reconstruction quality.

The implementation of accident data replay in Sky-Drive enables multiple critical applications in autonomous driving research and development. The framework provides a controlled environment for detailed analysis of accident causation, supporting the development of enhanced safety systems and collision avoidance algorithms. It serves as a valuable resource for training RL agents on real-world edge cases, significantly improving their ability to handle critical scenarios. Additionally, the platform's high-fidelity reconstructions support regulatory compliance and accident investigation processes. The framework's integration with Sky-Drive's broader simulation capabilities enables systematic testing of autonomous driving systems against a comprehensive database of real-world accident scenarios, advancing the development of safer and more robust autonomous vehicles.

VI. Future Enhancements

A. Foundation Models

Sky-Drive implements a comprehensive LLM-based system that enables personalized autonomous driving through natural language interactions. As shown in Fig. 4 (g), the platform integrates advanced language understanding with perception and navigation capabilities to enhance human-vehicle communication in complex driving scenarios.

The implementation leverages a multimodal architecture consisting of three core components integrated within Sky-Drive's simulation environment. A CLIP-based visual encoder processes real-time camera feeds to extract rich perceptual features. These features are seamlessly projected into the language decoder's embedding space, enabling a unified representation of visual and linguistic information [51]. A dedicated route planning module utilizes Sky-Drive's mapping capabilities to generate executable navigation commands, while a Vicuna-7B language model coordinates visual inputs, natural language instructions, and historical actions to produce appropriate driving responses.

To ensure robust performance, Sky-Drive employs a three-stage training pipeline for the LLM system. The first stage utilizes the BDD-X dataset to align visual and linguistic representations through careful tuning of LLaVA-7B's projection layers [52]. The second stage enhances command understanding through LoRA-based fine-tuning on the SDN dataset, enabling the system to interpret diverse human intentions. The final stage incorporates Sky-Drive's simulated perceptual data to refine real-world decision-making capabilities. Through this implementation, Sky-Drive establishes a foundation for personalized autonomous driving that seamlessly integrates natural language interaction with robust navigation capabilities.

Foundation models are transforming autonomous driving simulation by serving as sophisticated "traffic brains." These models enhance Sky-Drive's simulation capabilities through LLM/VLM-guided decision-making, personalized driving behavior, adaptive control, and predictive collision detection, leveraging their ability to process both visual and textual inputs contextually. While traditional driving simulators rely on deterministic rule-based approaches and manually defined policies, Sky-Drive's integration of foundation models provides the necessary adaptability and contextual awareness for modern autonomous driving systems. This advancement transforms Sky-Drive from a conventional testing environment into an advanced platform capable of equipping AVs with sophisticated LLM/VLM-based assistance and end-to-end driving capabilities.

Although general-purpose models like Qwen [53], GPT-4 [54], and Llama [55] demonstrate remarkable conversational abilities, they require specialized adaptation for driving applications. Sky-Drive addresses this limitation through targeted fine-tuning within the autonomous driving domain, emphasizing dynamic scenario adaptation, hierarchical reasoning in complex traffic situations, and multitask capabilities. These capabilities include generating safe driving actions (steering, throttle, and braking) and predicting critical safety metrics such as Time-to-Collision (TTC). The platform implements a comprehensive three-stage approach: initially leveraging pre-trained VLMs such as Qwen2-VL-7B [53] for visual-linguistic understanding, then fine-tuning these models on specialized autonomous driving datasets including LMDrive [56], CCD [57], DoTA [58], and DriveCoT [59]. These datasets provide diverse training signals across navigation states, crash scenarios, and multi-modal driving data. Finally, the refined models are integrated into Sky-Drive's simulation pipeline, incorporating advanced techniques such as safety-critical dataset balancing, computational efficiency through quantization, and RL for performance optimization.

This integration of foundation models establishes Sky-Drive as an intelligent and scalable platform for developing safe and efficient driving policies. By positioning these models at the core of its simulation systems, Sky-Drive creates a robust foundation for advancing autonomous driving technologies. Future developments will incorporate more sophisticated foundation models such as Qwen-QVQ-72B [60] to further enhance reasoning capabilities, safety features, and human-like decision-making processes.

1) Foundation Models Integration: Despite the integration of diverse sensing modalities—including VR-based eye tracking, voice interaction, facial expression recognition, physiological signal monitoring, and realistic steering wheel input—it must be acknowledged that Sky-Drive has not yet fully realized multimodal intent fusion and decision reasoning. One of our key future directions is to leverage large language models (LLMs) and vision-language models (VLMs) as core integration mechanisms. These models enable cross-modal reasoning by understanding the relationships between physiological signals, eye movement patterns, verbal expressions, and physical control inputs to interpret complex human behavior.

For instance, Sky-Drive can detect a pattern in which elevated heart rate, downward gaze direction, and a brief comment such as "too fast" collectively indicate discomfort with vehicle acceleration. In another case, a more nuanced verbal expression—"I feel a bit uneasy because the car accelerates too quickly"—can be semantically aligned with similar physiological and behavioral cues. LLMs reason over these varied expressions to infer underlying preferences, while simultaneously associating facial tension or stress with corresponding biometric signals.

By performing this type of integrated interpretation in environmental context—including factors such as traffic density, road geometry, and interactions with other road users—Sky-Drive can construct rich behavioral profiles that go far beyond what any single sensing modality could provide. This holistic understanding will empower future iterations of the HAIM and AIHM frameworks to adapt more intelligently and personally to each driver, enabling the development of socially-aware, trustworthy, and human-centered autonomous driving systems.

B. Hardware-in-the-Loop

Sky-Drive implements a HIL testing framework that seamlessly bridges simulation and real-world deployment through ROS integration. This framework enables rigorous validation of autonomous driving algorithms on physical platforms while maintaining stringent safety and reliability standards throughout the development process. At the core of the HIL system lies a testbed vehicle, a Ford E-Transit electric van retrofitted with complete autonomous driving capabilities. The vehicle features a comprehensive sensor array including LiDAR systems, radar units, high-resolution cameras, and OxTS navigation units. Advanced drive-by-wire control platforms integrate seamlessly with industrial-grade computing systems, facilitating precise vehicle control and real-time data acquisition. The vehicle software architecture leverages established ROS-based open-source packages enhanced with Sky-Drive's proprietary algorithms for expanded functionality.

To facilitate comprehensive testing capabilities, Sky-Drive has developed innovative portable roadside infrastructure units equipped with traffic signaling systems, regulatory signage, and advanced sensing capabilities including cameras and LiDAR arrays. These units serve dual purposes: enabling systematic validation of vehicle-to-infrastructure (V2I) communication protocols and supporting cooperative perception algorithm development across diverse environmental conditions [3]. The integration of these roadside units with the vehicle fleet establishes a comprehensive testing ecosystem for CAV technologies.

The HIL framework maintains tight integration with Sky-Drive's digital twin environment, enabling fluid transitions between simulation and physical testing phases. This integrated approach allows development teams to thoroughly validate algorithms in simulation before physical deployment, substantially reducing development cycles while maintaining safety standards. The system supports exhaustive testing of vehicle-to-vehicle (V2V) and V2I communication protocols, multi-sensor fusion algorithms, and autonomous driving capabilities, all while adhering to rigorous safety and reliability requirements throughout the development process.

The HIL framework also establishes a solid foundation for developing and testing teleoperated driving. Teleoperated driving allows humans (teleoperators) to remotely control vehicles, particularly in challenging scenarios, complementing fully/highly autonomous solutions. It is one of the important use cases of vehicle-to-everything (V2X) communication, specified in the 3GPP standards [61]. Sky-Drive's ROS integration enables wireless connectivity between its testbed vehicle and human-in-the-loop simulation platform—operated by a teleoperator—via cellular or satellite networks, e.g., 5G. Considering the wild fluctuations of network bandwidth, round-trip time (RTT), jitter time, and packet loss under driving conditions of 5G [62], Sky-Drive facilitates the collaboration between the vehicle and the simulation platform to dynamically decide what data (RGB images, LiDAR point cloud, and/or their pre-processed data) to transmit and how to transmit them to meet the end-to-end latency requirement for the teleoperation, i.e., below 100 milliseconds [63].

VII. Conclusions

This paper presented Sky-Drive, a comprehensive simulation platform that integrates virtual reality (VR), multi-agent interactions, and human-centered AI approaches for multi-agent traffic simulation and human-centered autonomous agent research. The key innovations include: a digital twin framework for high-fidelity virtual replication of transportation systems, a distributed multi-agent architecture enabling synchronized cross-terminal simulation, a multi-modal human-in-the-loop framework for capturing rich behavioral data, a novel human-AI bi-directional mentor mechanism for effective knowledge exchange, foundation models integration for enhanced human-machine collaboration, and a hardware-in-the-loop module for direct algorithm validation.

Looking forward, several promising directions exist for further development of Sky-Drive. First, improving computational efficiency remains a priority, particularly in processing real-time interactions among multiple agents and handling large-scale traffic scenarios. Second, while the current implementation successfully integrates foundation models, expanding their capabilities to handle more complex driving scenarios and environmental conditions could further enhance the platform's utility. Third, the platform could benefit from incorporating more sophisticated physics models and environmental simulations to improve fidelity in adverse weather conditions and complex urban environments.

Additionally, future work should focus on expanding Sky-Drive's capabilities in several key areas: enhancing the platform's ability to generate and validate edge cases for autonomous driving, developing more sophisticated methods for knowledge transfer between human drivers and AI systems, and improving the integration of real-world traffic data into simulation scenarios. Through its innovative integration of multiple cutting-edge technologies, Sky-Drive provides a robust foundation for advancing the field of autonomous driving while maintaining a strong focus on human-centered design and safety considerations. As autonomous driving technology continues to evolve, Sky-Drive will play an increasingly crucial role in ensuring the safety, reliability, and effectiveness of next-generation transportation systems.

References

[1] D. Almaskati, S. Kermanshachi, and A. Pamidimukkala, "Convergence of emerging transportation trends: A comprehensive review of shared autonomous vehicles," *J. Intell. Connected Veh.*, vol. 7, no. 3, pp. 177–189, 2024.

[2] S. Feng, H. Sun, X. Yan, H. Zhu, Z. Zou, S. Shen, and H. X. Liu, "Dense reinforcement learning for safety validation of autonomous vehicles," *Nature*, vol. 615, no. 7953, pp. 620–627, 2023.

[3] Z. Huang, S. Chen, Y. Pian, Z. Sheng, S. Ahn, and D. A. Noyce, "Toward c-v2x enabled connected transportation system: Rsu-based cooperative localization framework for autonomous vehicles," *IEEE Trans. Intell. Transp. Syst.*, 2024.

[4] Y. Xu, Z. Ye, and C. Wang, "Modeling commercial vehicle drivers' acceptance of advanced driving assistance system (adas)," *J. Intell. Connected Veh.*, vol. 4, no. 3, pp. 125–135, 2021.

[5] N. K. Sekar, V. Malaghan, and D. S. Pawar, "Micro-simulation insights into the safety and operational benefits of autonomous vehicles," *J. Intell. Connected Veh.*, vol. 6, no. 4, pp. 202–210, 2023.

[6] S. Chen, S. Zong, T. Chen, Z. Huang, Y. Chen, and S. Labi, "A taxonomy for autonomous vehicles considering ambient road infrastructure," *Sustainability*, vol. 15, no. 14, p. 11258, 2023.

[7] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "Carla: An open urban driving simulator," in *Proc. Conf. Robot Learn.* PMLR, 2017, pp. 1–16.

[8] S. Shah, D. Dey, C. Lovett, and A. Kapoor, "Airsim: High-fidelity visual and physical simulation for autonomous vehicles," in *Proc. Int. Conf. Field Serv. Robot.* Springer, 2018, pp. 621–635.

[9] E. Andreotti, Selpi, and P. Boyraz, "Potential impact of autonomous vehicles in mixed traffic from simulation using real traffic flow," *J. Intell. Connected Veh.*, vol. 6, no. 1, pp. 1–15, 2023.

[10] E. Leurent, "An environment for autonomous driving decision-making," 2018. [Online]. Available: https://github.com/eleurent/highway-env

[11] Q. Li, Z. Peng, L. Feng, Q. Zhang, Z. Xue, and B. Zhou, "Metadrive: Composing diverse driving scenarios for generalizable reinforcement learning," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 45, no. 3, pp. 3461–3475, 2022.

[12] M. Zhou, J. Luo, J. Villella, Y. Yang, D. Rusu, J. Miao, W. Zhang, M. Alban, I. Fadakar, Z. Chen *et al.*, "Smarts: An open-source scalable multi-agent rl training school for autonomous driving," in *Proc. Conf. Robot Learn.* PMLR, 2021, pp. 264–285.

[13] Mechanical Simulation. (2025) Carsim. [Online]. Available: https://www.carsim.com

[14] I. Automotive. (2025) Carmaker. [Online]. Available: https://ipg-automotive.com/products-services/simulation-software/carmaker/

[15] R. Umeno, M. Itoh, and S. Kitazaki, "Influence of automated driving on driver's own localization: A driving simulator study," *J. Intell. Connected Veh.*, vol. 1, no. 3, pp. 99–106, 2018.

[16] M. Lukovics, S. Prónay, and B. Nagy, "Ethical decision-making in older drivers during critical driving situations: An online experiment," *J. Intell. Connected Veh.*, vol. 7, no. 1, pp. 30–37, 2024.

[17] X. Qu, H. Lin, and Y. Liu, "Envisioning the future of transportation: Inspiration of chatgpt and large models," *Commun. Transp. Res.*, vol. 3, p. 100103, 2023.

[18] S. Kuang, Y. Liu, X. Wang, X. Wu, and Y. Wei, "Harnessing multimodal large language models for traffic knowledge graph generation and decision-making," *Commun. Transp. Res.*, vol. 4, p. 100146, 2024.

[19] Y. Li, W. Yuan, S. Zhang, W. Yan, Q. Shen, C. Wang, and M. Yang, "Choose your simulator wisely: A review on open-source simulators for autonomous driving," *IEEE Trans. Intell. Veh.*, vol. 9, pp. 4861–4876, 2023.

[20] G. Rong, B. H. Shin, H. Tabatabaee, Q. Lu, S. Lemke, M. Možeiko, E. Boise, G. Uhm, M. Gerow, S. Mehta *et al.*, "Lgsvl simulator: A high fidelity simulator for autonomous driving," in *2020 IEEE 23rd International conference on intelligent transportation systems (ITSC)*. IEEE, 2020, pp. 1–6.

[21] Nvidia, "Nvidia Drive End-to-End Platform for Software-Defined Vehicles," Online, 2024, accessed: Mar. 12, 2024. [Online]. Available: https://www.nvidia.com/en-us/self-driving-cars/

[22] rFpro, "The World's Most Accurate Simulation Environment," Online, 2023, accessed: Mar. 12, 2024. [Online]. Available: https://rfpro.com

[23] D. Team, "Deepdrive: a simulator that allows anyone with a pc to push the state-of-the-art in self-driving," 2019.

[24] L. Zhou, Y. Song, Y. Gao, Z. Yu, M. Sodamin, H. Liu, L. Ma, L. Liu, H. Liu, Y. Liu, H. Li, G. Chen, and A. Knoll, "Garchingsim: An autonomous driving simulator with photorealistic scenes and minimalist workflow," *Proc. IEEE Int. Conf. Intell. Transp. Syst.*, pp. 4227–4232, 2023.

[25] P. Group, "Ptv vissim," https://www.ptvgroup.com/en/solutions/ptv-vissim/, 2024, accessed: 2024-11-11.

[26] C. Wu, A. R. Kreidieh, K. Parvate, E. Vinitsky, and A. M. Bayen, "Flow: A modular learning framework for mixed autonomy traffic," *IEEE Trans. Robot.*, vol. 38, no. 2, pp. 1270–1286, 2022.

[27] H. Zhang, S. Feng, C. Liu, Y. Ding, Y. Zhu, Z. Zhou, W. Zhang, Y. Yu, H. Jin, and Z. Li, "Cityflow: A multi-agent reinforcement learning environment for large scale city traffic scenario," in *Proc. World Wide Web Conf.*, 2019, pp. 3620–3624.

[28] M. Müller, V. Casser, J. Lahoud, N. G. Smith, and B. Ghanem, "Sim4cv: A photo-realistic simulator for computer vision applications," *Int. J. Comput. Vis.*, vol. 126, pp. 902 – 919, 2017.

[29] B. Wymann, E. Espié, C. Guionneau, C. Dimitrakakis, R. Coulom, and A. Sumner, "TORCS, the open racing car simulator," 2020, accessed: Mar. 12, 2024. [Online]. Available: https://sourceforge.net/projects/torcs/

[30] P. Cai, Y. Lee, Y. Luo, and D. Hsu, "Summit: A simulator for urban driving in massive mixed traffic," *Proc. IEEE Int. Conf. Robot. Autom.*, pp. 4023–4029, 2019.

[31] P. Palanisamy, "Multi-agent connected autonomous driving using deep reinforcement learning," *Proc. Int. Joint Conf. Neural Netw.*, pp. 1–7, 2019.

[32] Q. Li, Z. Peng, Z. Xue, Q. Zhang, and B. Zhou, "Metadrive: Composing diverse driving scenarios for generalizable reinforcement learning," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 45, pp. 3461–3475, 2021.

[33] C. Gulino, J. Fu, W. Luo, G. Tucker, E. Bronstein, Y. Lu, J. Harb, X. Pan, Y. Wang, X. Chen *et al.*, "Waymax: An accelerated, data-driven simulator for large-scale autonomous driving research," *Adv. Neural Inf. Process. Syst.*, vol. 36, 2024.

[34] Q. Li, Z. Peng, L. Feng, Z. Liu, C. Duan, W.-A. Mo, and B. Zhou, "Scenarionet: Open-source platform for large-scale traffic scenario simulation and modeling," *ArXiv*, vol. abs/2306.12241, 2023.

[35] E. Vinitsky, N. Lichtl'e, X. Yang, B. Amos, and J. N. Foerster, "Nocturne: a scalable driving benchmark for bringing multi-agent learning one step closer to the real world," *ArXiv*, vol. abs/2206.09889, 2022.

[36] O. Robotics. (2025) Gazebo. [Online]. Available: https://gazebosim.org/

[37] Mathworks. (2025) Vehicle dynamics blockset. [Online]. Available: https://www.mathworks.com/products/vehicle-dynamics.html

[38] NVIDIA, "Nvidia drive sim: Autonomous vehicle simulation platform," Retrieved from https://www.nvidia.com/en-us/self-driving-cars/drive-constellation/, 2025, accessed on March 20, 2025.

[39] Applied Intuition, "Vehicle software supplier for adas and autonomous vehicles," Retrieved from https://www.appliedintuition.com, 2023, accessed on March 20, 2025.

[40] MORAI, "Simulation platform for autonomous vehicle & urban air mobility," Retrieved from https://www.morai.ai, 2023, accessed on March 20, 2025.

[41] A. Amini, T.-H. Wang, I. Gilitschenski, W. Schwarting, Z. Liu, S. Han, S. Karaman, and D. Rus, "Vista 2.0: An open, data-driven simulator for multimodal sensing and policy learning for autonomous vehicles," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 2419–2426.

[42] Y. Luo, P. Cai, Y. Lee, and D. Hsu, "Gamma: A general agent motion model for autonomous driving," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 3499–3506, 2022.

[43] Wayve, "Lingo-1: Exploring natural language for autonomous driving," Retrieved from https://wayve.ai/thinking/lingo-natural-language-autonomous-driving/, September 2023, accessed on March 20, 2025.

[44] Z. Ma, Q. Sun, and T. Matsumaru, "Bidirectional planning for autonomous driving framework with large language model," *Sensors*, vol. 24, no. 20, p. 6723, 2024.

[45] Y. Jin, X. Shen, H. Peng, X. Liu, J. Qin, J. Li, J. Xie, P. Gao, G. Zhou, and J. Gong, "Surrealdriver: Designing generative driver agent simulation framework in urban contexts based on large language model," *arXiv preprint arXiv:2309.13193*, vol. 5, no. 7, p. 8, 2023.

[46] M. J. Shafiee, M. Nentwig, Y. Kassahun, F. Li, S. Bochkarev, A. Kamal, D. Dolson, S. Altintas, A. Virani, and A. Wong, "Human-machine collaborative design for accelerated design of compact deep neural networks for autonomous driving," *arXiv preprint arXiv:1909.05587*, 2019.

[47] K. Long, Z. Sheng, H. Shi, X. Li, S. Chen, and S. Ahn, "A physics enhanced residual learning (perl) framework for vehicle trajectory prediction," *Commun. Transp. Res.*, 2025.

[48] Z. Sheng, Z. Huang, and S. Chen, "Ego-planning-guided multi-graph convolutional network for heterogeneous agent trajectory prediction," *Comput.-Aided Civ. Infrastruct. Eng.*, vol. 39, no. 22, pp. 3357–3374, 2024.

[49] Z. Huang, Z. Sheng, C. Ma, and S. Chen, "Human as ai mentor: Enhanced human-in-the-loop reinforcement learning for safe and efficient autonomous driving," *Commun. Transp. Res.*, vol. 4, 2024.

[50] X. Zhou, V. Koltun, and P. Krähenbühl, "Tracking objects as points," in *Proc. Eur. Conf. Comput. Vis.* Springer, 2020, pp. 474–490.

[51] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nuscenes: A multimodal dataset for autonomous driving," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 11 621–11 631.

[52] J. Kim, A. Rohrbach, T. Darrell, J. Canny, and Z. Akata, "Textual explanations for self-driving vehicles," *Proc. Eur. Conf. Comput. Vis.*, 2018.

[53] P. Wang, S. Bai, S. Tan, S. Wang, Z. Fan, J. Bai, K. Chen, X. Liu, J. Wang, W. Ge, Y. Fan, K. Dang, M. Du, X. Ren, R. Men, D. Liu, C. Zhou, J. Zhou, and J. Lin, "Qwen2-vl: Enhancing vision-language model's perception of the world at any resolution," *arXiv preprint arXiv:2409.12191*, 2024.

[54] OpenAI, "Gpt-4 technical report," *arXiv preprint arXiv:2303.08774*, 2023.

[55] A. Dubey, A. Jauhri, A. Pandey, A. Kadian, A. Al-Dahle, A. Letman, A. Mathur, A. Schelten, A. Yang, A. Fan *et al.*, "The llama 3 herd of models," *arXiv preprint arXiv:2407.21783*, 2024.

[56] H. Shao, Y. Hu, L. Wang, G. Song, S. L. Waslander, Y. Liu, and H. Li, "Lmdrive: Closed-loop end-to-end driving with large language models," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2024, pp. 15 120–15 130.

[57] W. Bao, Q. Yu, and Y. Kong, "Uncertainty-based traffic accident anticipation with spatio-temporal relational learning," in *Proc. ACM Multimedia Conf.*, May 2020.

[58] Y. Yao, X. Wang, M. Xu, Z. Pu, Y. Wang, E. Atkins, and D. Crandall, "Dota: unsupervised detection of traffic anomaly in driving videos," *IEEE Trans. Pattern Anal. Mach. Intell.*, 2022.

[59] T. Wang, E. Xie, R. Chu, Z. Li, and P. Luo, "Drivecot: Integrating chain-of-thought reasoning with end-to-end driving," *arXiv preprint arXiv:2403.16996*, 2024.

[60] Q. Team, "Qvq: To see the world with wisdom," December 2024. [Online]. Available: https://qwenlm.github.io/blog/qvq-72b-preview/

[61] "Study on enhancement of 3GPP Support for 5G V2X Services," 3rd Generation Partnership Project (3GPP), Technical Report TR 22.886, 2020, available at: https://www.3gpp.org/ftp/Specs/archive/22_series/22.886/.

[62] M. Ghoshal, I. Khan, Z. J. Kong, P. Dinh, J. Meng, Y. C. Hu, and D. Koutsonikolas, "Performance of cellular networks on the wheels," in *Proceedings of the 2023 ACM on Internet Measurement Conference*, 2023, pp. 678–695.

[63] A. Podhurst, "Autonomous vehicle teleoperation: Is one network enough for remote driving?" 2025.