J Intel Connect Veh

Research Article

Sky-Drive: A distributed multiagent simulation platform for human–AI collaborative and socially aware future transportation

Zilin Huang[1,†], Zihao Sheng[1,†], Zhengyang Wan[1,†], Yansong Qu[2], Yuhao Luo[1], Boyue Wang[1], Pei Li[1], Yen-Jung Chen[3], Jiancong Chen[2], Keke Long[1], Jiayi Meng[4], Yue Leng[5], Sikai Chen[1,✉]

Cite this article: Huang Z L, Sheng Z H, Wan Z Y, et al. *J Intel Connect Veh* 2025, **8**(4): 9210070. https://doi.org/10.26599/JICV.2026.9210070

ABSTRACT: Recent advances in autonomous system simulation platforms have significantly enhanced the safe and scalable testing of driving policies. Although existing simulators have greatly accelerated development by providing controlled testing environments, they face limitations in addressing the evolving needs of future transportation research, particularly in enabling effective human–artificial intelligence (human–AI) collaboration and modeling socially aware driving agents. This study introduces Sky-Drive, a novel distributed multiagent simulation platform that addresses these limitations through four key innovations: (1) a distributed architecture for synchronized simulation across multiple terminals; (2) a multimodal human-in-the-loop framework that integrates diverse sensors to collect rich behavioral data; (3) a human–AI collaboration mechanism that supports continuous and adaptive knowledge exchange; and (4) a digital twin framework for constructing high-fidelity virtual replicas of real-world transportation environments. Sky-Drive supports diverse applications, such as autonomous vehicle-human road user interaction modeling, human-in-the-loop training, socially aware reinforcement learning, personalized driving development, and customized scenario generation. Future extensions will incorporate foundation models for context-aware decision support and hardware-in-the-loop testing for real-world validation. By bridging scenario generation, data collection, algorithm training, and hardware integration, Sky-Drive has the potential to become a foundational platform for the next generation of human-centered and socially aware autonomous transportation system research.

KEYWORDS: driving simulator; autonomous vehicles; human–artificial intelligence (AI) collaboration; multiagent simulation; digital twin

1 Introduction

Autonomous systems and related technologies have made significant strides in recent years, demonstrating increasing maturity in terms of perception, decision-making, and control capabilities (Almaskati et al., 2024; Chen et al., 2023; Ma and Xue, 2024; Sheng et al., 2024a). As these technologies continue to advance, future transportation systems are expected to consist of human-driven vehicles (HVs) and diverse artificial intelligence-driven (AI-driven) intelligent agents, including autonomous vehicles (AVs), delivery robots, flying vehicles, and smart infrastructure, each capable of perceiving the environment, making decisions, and interacting with others in real time (Lv et al., 2024). For these agents to seamlessly integrate into traffic systems governed by humans, they must not only ensure their own safe and efficient operation but also exhibit acute perceptiveness in recognizing human intentions and preferences (e.g., comfort-oriented speed profiles) while skillfully demonstrating social behaviors comparable to those of their human counterparts (e.g., yielding behaviors and adherence to implicit right-of-way conventions) (Liu et al., 2024). This dual requirement, the ability to understand human expectations and to coordinate with other human road users (HRUs) (e.g., human drivers, pedestrians, and cyclists) to optimize overall system performance, constitutes what we refer to as human–AI collaboration and social awareness, respectively. Therefore, future research should move beyond the isolated validation of AV safety performance toward a more comprehensive investigation of human–AI collaboration and social awareness in mixed traffic environments, where multiple types of HRUs with varying levels of autonomy interact continuously.

Validating autonomous driving technologies in real-world settings presents considerable challenges due to safety risks, limited controllability, and the scale of testing required to

† Zilin Huang, Zihao Sheng, and Zhengyang Wan contributed equally to this work.

[1] Department of Civil and Environmental Engineering, University of Wisconsin–Madison, Madison 53706, USA. [2] Lyles School of Civil and Construction Engineering, Purdue University, West Lafayette 47907, USA. [3] Elmore Family School of Electrical and Computer Engineering, Purdue University, West Lafayette 47907, USA. [4] Department of Computer Science and Engineering, The University of Texas at Arlington, Arlington 76019, USA. [5] Google, Sunnyvale 94089, USA.

✉ Corresponding author. E-mail: sikai.chen@wisc.edu

Received: May 29, 2025; Revised: July 12, 2025; Accepted: August 25, 2025

demonstrate reliability (Feng et al., 2023; Huang et al., 2024c; Kalra and Paddock, 2016). To mitigate these barriers, the autonomous driving community has developed a variety of simulation platforms, including CARLA (Dosovitskiy et al., 2017), AirSim (Shah et al., 2017), SUMO (Lopez et al., 2018), Vissim (Fellendorf and Vortisch, 2010), Highway-Env (Leurent and Mercat, 2019), MetaDrive (Li et al., 2023b), SMARTS (Ramamohanarao et al., 2016), CarSim, and IPG CarMaker (Hong and Aparow, 2021; Li et al., 2019). These platforms have significantly accelerated development by providing controlled testing environments. However, they face key limitations in addressing the evolving needs of future transportation research. First, while existing simulation platforms can emulate multiple agents on a single machine via rule-based or pretrained learning-based methods, they generally do not support real-time participation of HRUs across multiple terminals. This limitation hinders the collection of authentic human behavior and HRU interaction data. Such data are particularly valuable for studying rare but safety-critical scenarios (e.g., interactions between AVs, HVs, and pedestrians), which pose significant safety and ethical risks when collected in the real world. A distributed simulation platform that enables participants to assume diverse roles across multiple terminals is urgently needed to collect such interaction data safely and to evaluate interaction algorithms in controlled environments.

Second, existing simulation platforms offer limited support for human-AI collaboration. While they can collect human inputs, these are often treated as low-level control signals rather than high-level feedback for improving autonomous driving algorithms (Gulino et al., 2023; Li et al., 2023a; Vinitsky et al., 2023). In contrast, effective human–AI collaboration refers to a bidirectional process in which humans provide feedback not only as commands but also as indications of preferences, situational understanding, and normative behaviors. AI systems, in turn, assist human drivers by offering real-time guidance, performance feedback, and personalized training. This bidirectional exchange enables AI systems to adapt continuously to human needs and expectations while simultaneously enhancing human driving performance through intelligent support. Additionally, the emergence of foundation models, such as large language models (LLMs) and vision-language models (VLMs), trained on large-scale, multimodal datasets and equipped with broad-world knowledge, offers new opportunities for capturing and utilizing human knowledge (Liao et al., 2024; Liu et al., 2023; Touvron et al., 2023; Wang et al., 2024a). However, in most simulators (Wei et al., 2024; Yang et al., 2024; Zhang et al., 2024), such models are used primarily for scenario generation rather than as active components in human-AI collaborative learning.

Third, although some simulation platforms have integrated reinforcement learning (RL) capabilities to improve autonomous driving policies (Berta et al., 2024; Li et al., 2023b; Zhou et al., 2021), they remain largely focused on optimizing vehicle-level metrics such as safety, efficiency, and route completion. However, advancing real-world deployment requires moving beyond individual vehicle performance to incorporate social awareness into the decision-making process. Social awareness refers to an autonomous system's ability to coordinate with other traffic participants to enhance overall system performance (Wang et al., 2022). This includes promoting traffic flow stability, mitigating congestion through cooperative maneuvers, enhancing the collective comfort of all road users, and enabling harmonious interactions between AVs and humans in mixed traffic settings. In this context, transportation science offers a valuable foundation. Decades of research have produced validated traffic flow theories and human behavioral models, such as the intelligent driver model (IDM) and the minimizing overall braking by lane changes (MOBIL) model, which can inform the design of socially aware autonomous systems capable of optimizing not only individual performance but also the efficiency and safety of the entire traffic ecosystem (Kesting et al., 2007; Treiber et al., 2000).

To address these challenges, this paper proposes Sky-Drive, an open-source simulation platform designed to advance research in socially aware autonomous driving and human-AI collaboration. Sky-Drive integrates scenario generation, data collection, algorithm training, and hardware integration into a unified platform, supporting distributed multiagent operation and multimodal human-in-loop interaction. As illustrated in Fig. 1, Sky-Drive introduces four key innovations:

1) Sky-Drive introduces a distributed multiagent architecture that enables synchronized simulation across multiple devices through remote procedure call (RPC). This design allows independent control of agents on separate terminals while maintaining shared environmental states, better reflecting future mixed traffic.

2) Sky-Drive provides a multimodal 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 multimodal streams, enabling detailed analysis of human driving patterns and responses to complex scenarios.

3) 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.

4) To bridge the gap between simulation and reality, Sky-Drive includes a digital twin (DT) framework that builds high-fidelity virtual replicas of transportation systems by integrating data collected from laboratory-developed AVs, roadside sensors, traffic cameras, and historical records.

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

1) Sky-Drive integrates foundation models (e.g., LLMs and VLMs) at both the system and agent levels. At the system level, foundation models provide global observations and feedback to optimize simulation dynamics. At the agent level, they enhance situational understanding and enable safer, more socially aware, and personalized decision-making.

2) Sky-Drive will incorporate a hardware-in-the-loop (HIL) framework via a robot operating system (ROS), enabling direct validation of autonomous driving algorithms on physical vehicles and safe evaluation of algorithms without exposing users to real-world risks.

The remainder of this paper is organized as follows: Section 2 reviews related work in driving simulators. Section 3 introduces the Sky-Drive workflow. Section 4 details Sky-Drive's features and technical implementation. Section 5 presents application examples. Section 6 discusses planned future enhancements. Finally, Section 7 concludes the paper and outlines future research directions.

2 Related work

2.1 Driving simulators

Driving simulation platforms have evolved significantly to address

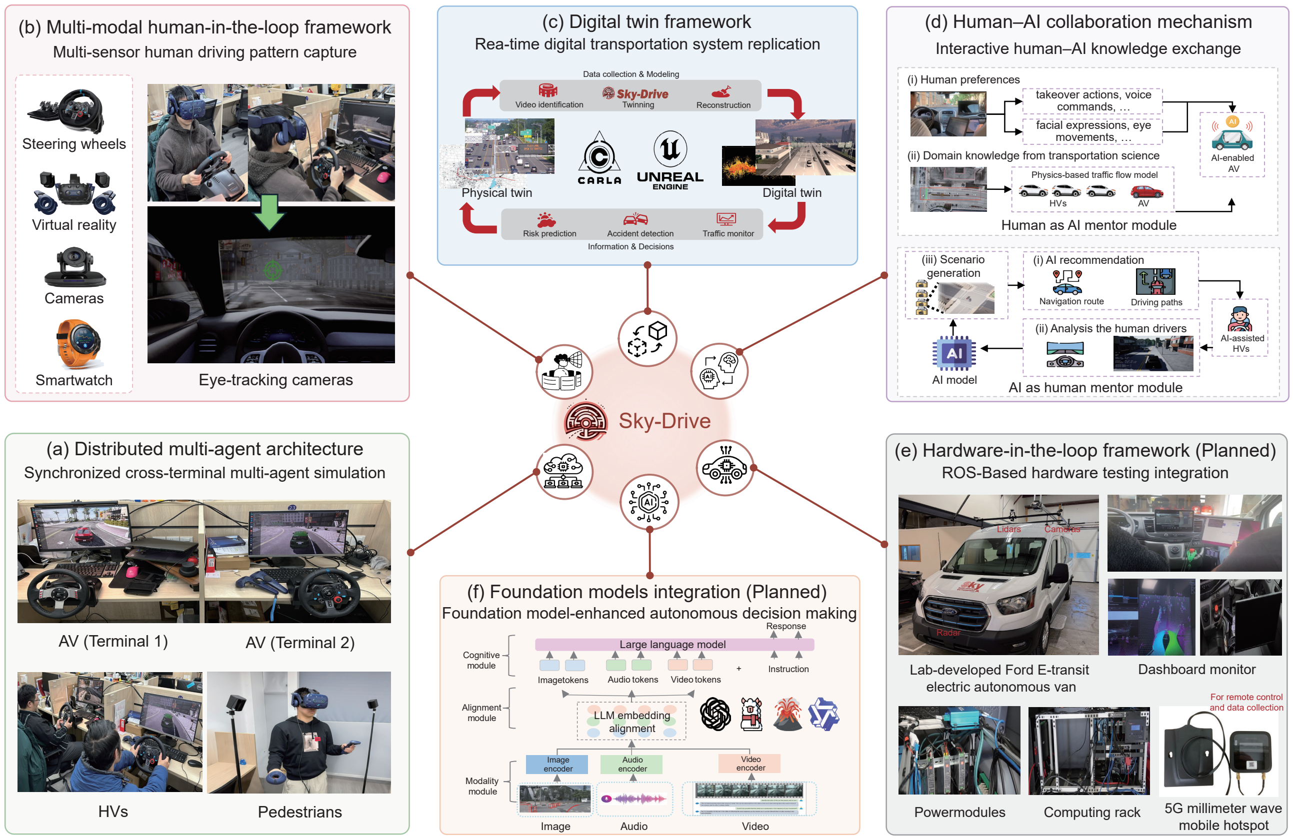

Fig. 1 Overview of Sky-Drive's key components and functionalities. (a) A distributed multiagent architecture enabling synchronized simulation across multiple terminals; (b) a multimodal human-in-the-loop framework capturing comprehensive behavioral data through integrated sensor systems; (c) a digital twin framework that creates high-fidelity virtual replicas of transportation systems through multisource data integration; (d) a human-AI collaboration mechanism facilitating knowledge exchange between humans and AI systems; (e) a hardware-loop framework, planned for future integration, enabling remote control and data collection and ensuring that algorithms are evaluated in real-world environments ; and (f) the planned integration of foundation models to enhance decision-making, enabling more adaptive and context-aware human-AI collaboration.

the growing needs of AV research. According to previous research (Li et al., 2024), these simulators can be categorized on the basis of 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 and LGSVL represent prominent open-source examples in this category (Dosovitskiy et al., 2017; Rong et al., 2020), offering rich environments for testing autonomous driving systems. Commercial solutions such as NVIDIA DRIVE Sim and rFpro, alongside academic developments such as DeepDrive and GarchingSim (Zhou et al., 2023), 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, Vissim, Flow (Wu et al., 2022), and CityFlow (Zhang et al., 2019). Recent developments have combined SUMO traffic modeling with three-dimensional (3D) simulators such as CARLA to merge scalability with realism.

Sensory data simulators, such as AirSim and Sim4CV (Müller et al., 2018), 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-EnvTORCS (Wymann et al., 2005), SUMMIT (Cai et al., 2020), MACAD (Palanisamy, 2020), SMARTS, and MetaDrive. Additionally, recent data-driven simulators such as Waymax, ScenarioNet, and Nocturne leverage real-world datasets to generate socially relevant traffic scenarios (Ramamohanarao et al., 2016; Vinitsky et al., 2023). Vehicle dynamics simulators, including CarSim, IPG CarMaker, and Gazebo (Koenig and Howard, 2004), 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 Table 1, most simulators run only on single devices, limiting their ability to model distributed multiagent scenarios. Additionally, current platforms provide insufficient support for socially aware algorithms that need to understand complex interactions with diverse road users. Considering CARLA's established strengths in sensor simulation and visualization, Sky-Drive leverages CARLA as its core engine while extending it with a distributed architecture for synchronized multiterminal simulation, immersive VR interfaces, and a DT framework. This integration creates an open-source platform specifically designed to support future transportation research.

2.2 Human–AI collaboration environments

Several simulation platforms have contributed to advancing human–AI collaboration in autonomous driving. For example, NVIDIA's DRIVE Sim and Omniverse platforms support collaboration by generating physics-based synthetic data for

Table 1 Comparison of representative simulators with sky drive

	Distributed multiagent simulation	Digital twin environment	Hardware-in-the-loop	Traffic flow	AI framework integration	Human-in-the-loop interface
Closed source						
NVIDIA DRIVE Sim	—	✓	✓	—	✓	✓
rFpro	—	✓	✓	—	—	✓
CarSim	—	—	✓	—	—	✓
Matlab	—	✓	✓	—	✓	✓
Open source						
DeepDrive 2.0	—	—	—	—	✓	—
GarchingSim	✓	—	✓	—	✓	✓
CARLA	—	✓	✓	—	✓	✓
SUMO	—	✓	✓	✓	—	—
Flow	—	—	—	✓	✓	—
CityFlow	—	—	—	✓	✓	—
TORCS	—	—	—	—	✓	—
SUMMIT	—	✓	—	—	✓	—
MACAD	—	—	—	—	✓	✓
MetaDrive	—	✓	—	—	✓	✓
SMARTS	—	—	—	—	✓	—
Nocturne	—	✓	—	✓	✓	—
Waymax	—	✓	—	✓	✓	—
Gazebo	—	✓	—	—	—	✓
Sky-Drive (ours)	✓	✓	✓	✓	✓	✓

Note: The "distributed multiagent simulation" functionality in this table refers to the ability 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.

training autonomous systems. However, their approach largely enables one-way knowledge transfer, where simulated data inform AI models without supporting real-time, bidirectional human-AI collaboration. Applied intuition incorporates human-in-the-loop testing to allow operators to validate autonomous decisions, yet its framework is tailored primarily for offline validation rather than for continuous learning. MORAI provides digital twin environments that visualize AI decision-making for human drivers, but its interaction remains limited to basic feedback collection without mechanisms for mutual adaptation or learning.

More specialized platforms have made progress toward collaborative learning. Massachusetts Institute of Technology (MIT)'s VISTA enables domain adaptation between real and virtual environments but focuses mainly on perception rather than interactive decision-making (Amini et al., 2022). The GAMMA framework introduces mixed-reality traffic that incorporates human behavior, although it lacks explicit mechanisms for integrating human experience into AI learning (Luo et al., 2022). Wayve's LINGO architecture enhances transparency by providing natural language explanations for AI decisions (Marcu et al., 2024), whereas SafeMod leverages LLMs for bidirectional planning, mimicking human reasoning in autonomous decision-making Similarly, SurrealDriver generates realistic driving behaviors that align with human expectations via LLMs (Jin et al., 2024), and DarwinAI's GenSynth facilitates collaboration between human designers and AI in developing neural networks for driving tasks (Shafiee et al., 2019).

Despite these advances, most platforms still fall short in enabling true human-AI knowledge exchange. They often lack mechanisms for the continuous integration of human feedback, resulting in open-loop rather than closed-loop learning processes. Moreover, few platforms support comprehensive multimodal sensing from humans, such as gaze, voice, physiological, and control inputs, which are critical for modeling and understanding driving behaviors. Sky-Drive addresses these limitations through its HAIM and AIHM modules, its multimodal human-in-the-loop framework, and its closed-loop learning architecture, which continuously integrates human experience into AI development.

3 Sky-Drive workflow

3.1 Overview

As shown in Fig. 1, the Sky-Drive architecture consists of six core components. Fig. 1 provides a high-level overview of the platform's design. The workflow begins with the DT framework, which creates high-fidelity virtual replicas of transportation systems through multisource data integration. These virtual environments feed into the distributed multiagent architecture, enabling synchronized simulations across multiple devices and supporting complex interactions between autonomous agents. Together, the DT framework and the multiagent architecture form the simulation environment that serves as the testing ground for the multimodal human-in-the-loop framework, which captures comprehensive behavioral data from human participants. The human-AI collaboration mechanism then processes and utilizes these data, facilitating knowledge exchange between humans and autonomous systems. Foundation model integration enhances system- and agent-level capabilities, enabling observations for performance feedback and assisting individual agents in better understanding human behavior patterns. Finally, the HIL framework connects with the DT framework, enabling real-world algorithm validation while real-world performance data are fed back into the simulation.

3.2 Details

Fig. 2 complements this view by illustrating the workflow pipeline, demonstrating how the components interact during scenario generation, simulation, training, and real-world validation. In detail, it consists of three primary stages that form a continuous feedback loop.

1) Scenario generation and data collection

As shown in Fig. 2a, this stage employs two complementary approaches to ensure comprehensive scenario coverage: (1) 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. (2) The DT framework imports real-world data through multisource integration, including high-definition maps collected by lab-developed AVs, open-source data, and real-world traffic data collection. These data undergo categorization and twinning processes to create digital replicas of physical environments.

2) Simulation and algorithm training

As shown in Fig. 2b, this stage processes the generated scenarios through an integrated learning pipeline with four interconnected components: (1) The distributed multiagent architecture enables the concurrent operation of multiple agents across different terminals, facilitating complex traffic interactions in a shared environment while maintaining synchronization. (2) The human-in-the-loop component integrates multiple human participants, capturing human behavior through immersive interfaces and allowing for real-time feedback collection. (3) Sky-Drive integrates LLMs to enhance simulation capabilities, facilitating natural communication between human participants and autonomous systems for more intuitive interaction and knowledge transfer. (4) The human–AI collaboration mechanism integrates human feedback and domain knowledge into AI training, creating a continuous learning loop in which humans inform AI systems and AI provides feedback to humans.

3) Hardware integration and testing

As illustrated in Fig. 2c, this stage bridges simulation and physical deployment through two key components: (1) While the full HIL framework is planned for future development, the current architecture already supports connections to external hardware through standardized ROS. The lab-developed Ford E-transit electric van serves as the primary testbed and is equipped with dashboard monitors, power modules, drive-by-wire systems, and a computing rack for algorithm deployment. Testing is conducted primarily at the Madison College Public Safety Training Center in Columbus, WI, and MGA Research Corporation in Burlington, WI, USA. (2) Sky-Drive also integrates a fifth-generation (5G) millimeter wave (mmWave) mobile hotspot to support low-latency teleoperation, enabling remote human operators to monitor and, when necessary, control physical vehicles in real time.

3.3 Case demonstration

To demonstrate the Sky-Drive workflow, we consider the case of personalized autonomous driving. In this case, Sky-Drive develops an autonomous system that tailors its behavior to the driver's unique preferences, learning from their driving styles and comfort levels.

The workflow begins with the DT framework, which creates high-fidelity virtual replicas of real-world traffic environments via data such as high-definition maps. These environments are then input into the distributed multiagent architecture, enabling simulations of complex interactions between AVs and humans. During the simulation and training stage, real-time feedback, such as "It's too fast." is processed to infer driver preferences in terms of acceleration. This feedback is integrated into the human-AI

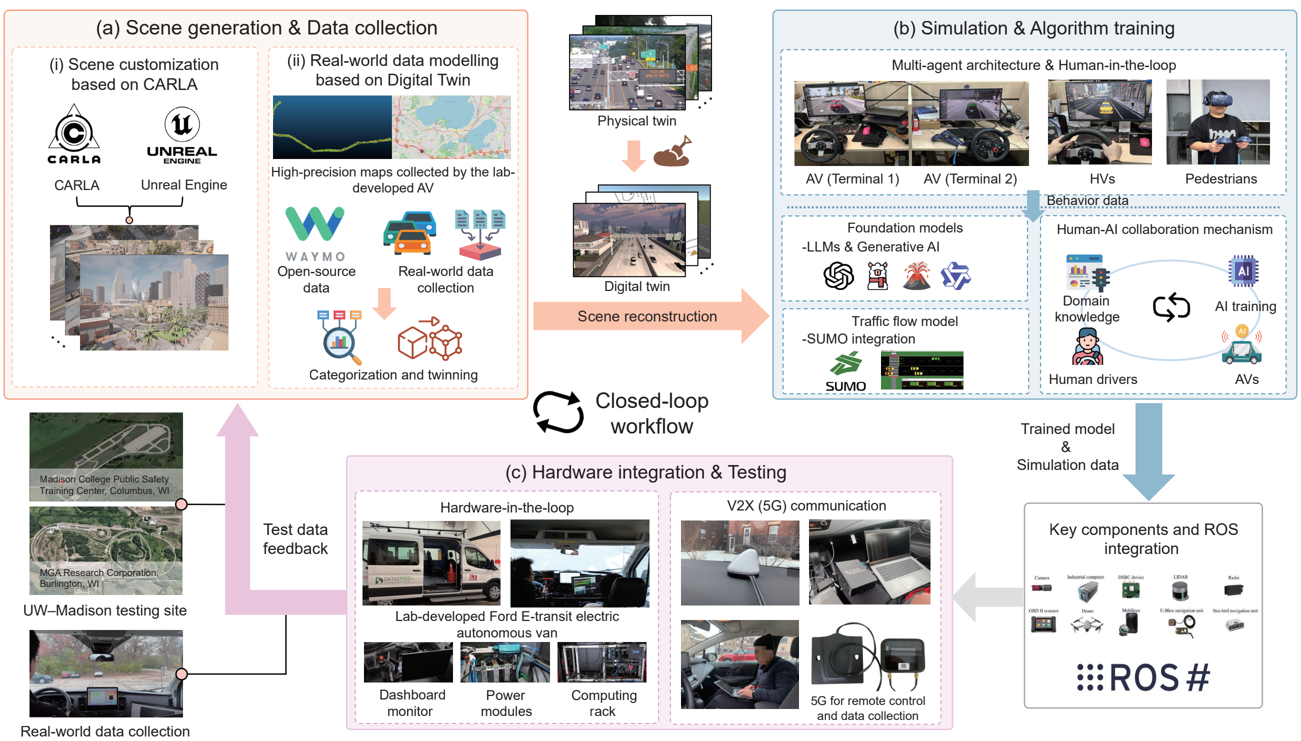

Fig. 2 Workflow of sky drive. (a) Scenario generation & data collection through CARLA-based synthetic environments and digital twin integration of real-world traffic data; (b) simulation & algorithm training enabled by distributed multiagent architecture and a human–AI collaboration mechanism; and (c) hardware integration & testing utilizing ROS compatibility for direct validation of autonomous driving algorithms on physical platforms.

collaboration mechanism, forming a continuous learning loop where the system adapts its driving strategies and provides more personalized guidance. The HIL framework connects the system to physical platforms, validating the personalized driving algorithm in real-world scenarios. This closed-loop workflow enables the development and validation of personalized autonomous systems, from concept testing to real-world deployment, ensuring safety and reliability.

4 Sky-Drive features

4.1 Distributed multiagent architecture

Future transportation systems will consist of multiple intelligent agents, such as AVs, HVs, and pedestrians, each operating independently, requiring simulation systems that can model and control these agents separately while ensuring seamless interaction between them. Existing platforms, such as Nocturne, MetaDrive, and Waymax, focus primarily on simplified multiagent interactions on a single machine, limiting their ability to model such complexity. To address this, Sky-Drive introduces a novel distributed multiagent architecture that enables the synchronized simulation of multiple independently operating agents across different computing devices.

1) System architecture: At the core of sky drive lies the RPC built upon CARLA, which uses the rpclib library. This extension enhances CARLA's vehicle control system, enabling crucial improvements for distributed multiagent simulation. As shown in Fig. 3c, Terminal 1 functions as the host (server) that maintains the global simulation environment, whereas Terminals 2–4 act as clients controlling different agent types. We designate the AI-controlled AV as the host because it is often the focus of training, testing, and evaluation. Hosting the simulation on this terminal minimizes latency for policy inference and synchronization and leverages its typically higher computational capacity. Each terminal independently controls its corresponding agent through various input devices while ensuring seamless interaction with other agents in the shared environment. The host terminal is responsible for scene customization and map generation, which are then distributed to the client terminals. The scene generation component (Fig. 3a) creates detailed virtual environments with customizable traffic conditions, weather patterns, and road infrastructure, supporting multiple agent types, including AI-controlled AVs, HVs, pedestrians controlled via VR, and rule-based AVs following predefined behaviors (Fig. 1a).

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 address challenges such as synchronization, latency, and data consistency, Sky-Drive adopts a centralized host-client model combined with time-aligned messaging protocols. For time-critical operations, a dedicated local area network (LAN) with high-performance switches and Ethernet connections achieves latency as low as 0.3 ms, ensuring real-time coherence among agents. For geographically distributed deployments, such as those

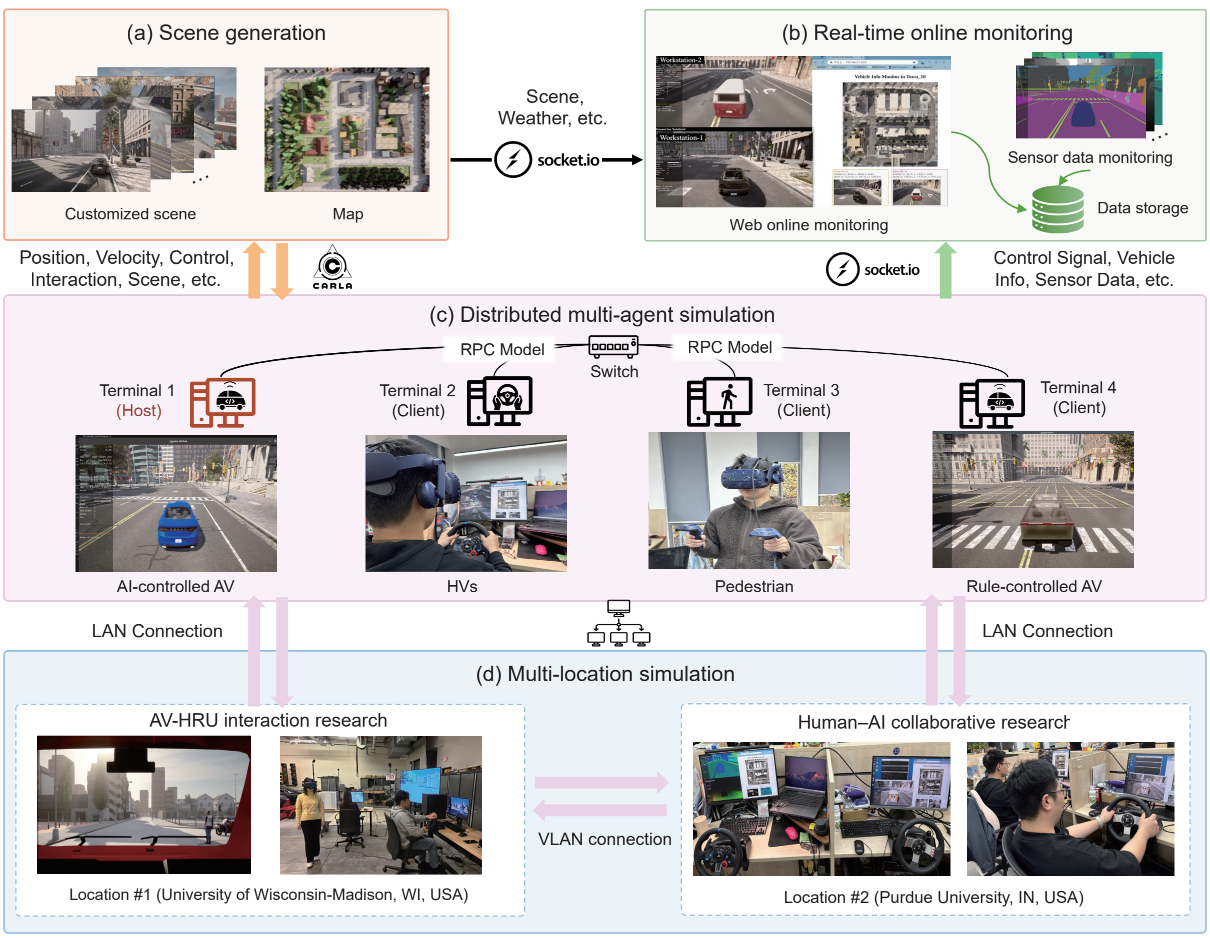

Fig. 3 Illustration of Sky-Drive's distributed multiagent architecture. Sky-Drive enables synchronized simulation across multiple terminals while maintaining precise real-time interactions between AVs, HVs, and pedestrians through RPC and Socket. IO-based communication platforms support comprehensive data collection and real-time analysis of multiagent behavior.

between Purdue University and the University of Wisconsin-Madison (Fig. 3d), virtual local area network (VLAN) configurations are used to extend reach while maintaining reliable communication. In practice, the Sky-Drive platform has been tested with up to 4 terminals and 3 simultaneous human participants and has maintained stable synchronization under both the LAN and the virtual private network (VPN) setups.

3) Real-time monitoring website: A key component of Sky-Drive's distributed architecture is its real-time monitoring and data management website. To complement the core communication infrastructure, Sky-Drive has developed a Socket. IO-based communication platforms track agent data, including position coordinates, velocity metrics, live video feeds, and sensor readings. As shown in Fig. 3b, the platform features a web-based system that provides real-time visualization of agent activities. It streams data to a centralized system where agent interactions are monitored and analyzed in real time. All simulation data, environmental conditions, and interaction events are synchronously logged to a centralized database, ensuring consistency across agents and enabling accurate scenario replay, analysis, and debugging. This design ensures data integrity across distributed nodes despite network delays or asynchronous actions.

4.2 Multimodal human-in-the-loop framework

To capture human preferences and cognitive states for adaptive AI behavior, Sky-Drive develops a multimodal human-in-the-loop framework, illustrated in Fig. 1b, that collects and synchronizes gaze patterns, voice commands, facial expressions, physiological signals, and control actions across multiple modalities.

1) Eye tracking: Sky-Drive provides an immersive experience through a custom-developed VR interface built on top of the Unreal Engine. The participants engage in the simulation via an HTC Vive Pro Eye headset, which supports full six degrees of freedom (6-DoF) head tracking via SteamVR and integrated eye tracking via the SRanipal SDK. The system captures high-frequency (up to 120 Hz) behavioral signals, including 3D gaze vectors, pupil positions and diameters, eye openness, and fixation points. These signals are critical for analyzing driver attention distribution, situational awareness, and cognitive workload 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 (Radford et al., 2023), and then interpreted by LLMs. The system extracts driver intent and sentiment from structured commands ("slow down at the next intersection") and informal feedback ("too fast"), mapping them into semantic driving directives or policy preferences to guide AI behavior.

3) Facial expression recognition: A high-resolution in-cabin camera captures facial microexpressions in real time. Sky-Drive employs expression classification models trained on affective datasets to recognize expressions such as stress, confusion, or satisfaction. These cues serve as implicit indicators of driver state and comfort, enabling real-time adaptation of AI behavior and intervention when necessary.

4) Physiological signal monitoring: Physiological states such as stress and alertness are inferred through biometric signals collected by wearable devices. Sky-Drive integrates the Garmin Vívoactive 5 smartwatch to continuously monitor heart rate and heart rate variability (HRV). These physiological signals are synchronized with other behavioral data streams, providing additional channels to model driver arousal, cognitive workload, and fatigue.

5) Steer wheel: The ego vehicle is equipped with a logitch G920 racing wheel and pedal system, with force feedback enabled through the open-source logitech wheel plugin. The steering, throttle, braking, and signaling inputs are logged in parallel with the gaze and head pose data. This setup supports realistic driving control and is fully compatible with CARLA's ScenarioRunner for scenario-based experiments.

While all the sensing components have been integrated and tested individually, the current case studies do not yet incorporate these multimodal inputs into active feedback loops. Their integration into human–AI collaboration and adaptive policy learning is planned for future stages.

4.3 Human–AI collaboration mechanism

Sky-Drive implements an adaptive human-AI collaboration mechanism that enables continuous, bidirectional knowledge exchange between human users and AI-enabled autonomous systems. As shown in Fig. 1d, this mechanism is built on two complementary modules: HAIM and AIHM.

1) Humans as AI mentors: In HAIM, humans serve as real-time mentors to AVs, guiding AI learning through two key sources of human knowledge: (1) Individual behavioral knowledge, encompassing both explicit behaviors (e.g., takeovers, voice commands, and touchscreen interactions) and implicit signals (e.g., facial expressions, eye movements, and physiological responses), captured via Sky-Drive's multimodal interface (Huang et al., 2024c, 2025); (2) domain knowledge from transportation science, including established models such as the IDM and MOBIL, which encode long-standing rules of human driving behavior (Huang et al., 2024b).

The HAIM adopts an RL paradigm enhanced by human preference modeling and physics-informed priors to incorporate this dual-source knowledge. Rather than relying on handcrafted reward functions, HAIM formulates learning as preference-based policy optimization. Frequent human takeovers in specific contexts (e.g., intersections, merges) are treated as implicit indicators of policy failure, shaping cost signals or trajectory ranking. Moreover, physics-based models act as behavioral constraints to ensure that learned policies remain safe, interpretable, and socially compliant. This design improves sample efficiency, reduces unsafe exploration, and fosters human trust in the AI system. In the current implementation, HAIM relies primarily on takeover signals (e.g., control interventions) to infer preferences and correct suboptimal behavior. The integration of multimodal sensing data, such as gaze, facial expression, and physiological state data, is part of our ongoing development roadmap.

2) AI as a human mentor: In parallel, AIHM enables AI to function as a real-time coach for human drivers. It leverages physics-enhanced residual learning (PERL) to generate optimal driving paths that consider vehicle dynamics (Long et al., 2025; Sheng et al., 2024b), safety margins, and individual driving styles (Sheng et al., 2024c). These reference trajectories are visualized in real time via VR or in-vehicle displays and are continuously updated on the basis of driver performance. The AIHM evaluates drivers via metrics such as path deviation, response latency, control stability, and situational awareness. Personalized feedback is delivered through annotated replays, visual heatmaps, and AI-generated verbal summaries.

A key innovation of AIHM is the use of generative AI for scenario customization. On the basis of performance analytics, the system dynamically generates targeted training tasks, such as emergency stops or lane changes, to address specific weaknesses.

The level of guidance is continuously adjusted via real-time physiological and behavioral signals: when elevated stress levels (e.g., increased heart rate, frequent steering corrections) are detected, the system reduces scenario complexity and provides calming feedback. Conversely, as the driver demonstrates improved performance, the system introduces more challenging conditions to encourage continued skill development (Sheng et al., 2025).

4.4 Digital twin framework

AI algorithms trained in simulations often fail to generalize to real-world traffic because of the lack of environmental fidelity. To address this sim-to-real gap and ensure practical applicability, Sky-Drive introduces a DT framework that creates dynamic, high-fidelity replicas of real transportation systems.

As illustrated in Fig. 2a, the DT framework consists of two core components: data integration and virtual environment construction. The multisource data integration layer fuses static and real-time inputs from traffic cameras, loop detectors, connected vehicle telemetry, Global Positioning System (GPS) traces, historical traffic records, and high-definition maps collected via laboratory-developed AVs equipped with LiDAR and radar. These inputs undergo temporal alignment, spatial correlation, and feature extraction to ensure semantic consistency across sources.

The virtual environment is built on CARLA and Unreal Engine and integrates real-time sensor data and computer vision models to detect and track road users for both rendering and trajectory prediction. By employing video recognition and object tracking models, the system reconstructs road user trajectories and maps them into the digital replica, enabling visual analytics, risk prediction, and event replay. Sky-Drive has implemented a pilot deployment of this framework along the Flex Lane on the Beltline in Dane County, Wisconsin. The DT uses real-time feeds from WisDOT 511 and historical records from WisTransPortal, enabling dynamic reconstruction of traffic states and generation of predictive insights.

5 Sky-Drive use cases

5.1 VR-based AV–HRU interaction

Studying the interactions between AVs and HRUs is critical for the safe deployment of AV technology in complex urban environments. Although incidents of AV–HRU conflicts have been reported, real-world crash data involving such cases remain scarce. More importantly, collecting such data in real traffic is unsafe, difficult to reproduce, and often restricted by ethical constraints. To address this challenge, Sky-Drive provides a VR-enabled platform for investigating AV–HRU interactions in a controlled, immersive, and data-rich environment. It is a distributed multiagent simulation architecture that enables synchronized control of multiple agents across separate terminals and devices while maintaining real-time coordination. This setup is particularly valuable for modeling high-risk scenarios that are difficult to observe or replicate in the physical world.

As shown in Fig. 4, we conducted a case study focused on right-turn conflicts at unsignalized intersections, which is a scenario frequently associated with accidents in urban environments. This study leveraged Sky-Drive's synchronized multiterminal architecture in a novel experimental setup where human participants experienced the scenario from the pedestrian's perspective through immersive VR, while researchers monitored the decision-making of an AV from a separate terminal and occasionally intervened when necessary. During each interaction, Sky-Drive captured multimodal behavioral data from both the AV and the pedestrian. The VR system 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 configuration allows researchers to analyze both the physical outcomes (e.g., successful yielding, near misses, and pedestrian hesitation) and the cognitive-emotional states of the human participant, offering insight into how HRUs perceive and respond to AV behavior.

5.2 HAIM-based deep reinforcement learning

To validate the HAIM module, we implemented and tested HAIM-DRL (Huang et al., 2024c), a reward-free RL approach that enables AI agents to learn driving behavior directly from human interventions. This demonstration serves as a proof-of-concept for the HAIM module's core functionality, which leverages real-time human feedback to guide policy learning within the multiagent, simulation-rich environment of sky drive.

Sky-Drive enables HAIM-DRL by detecting and recording

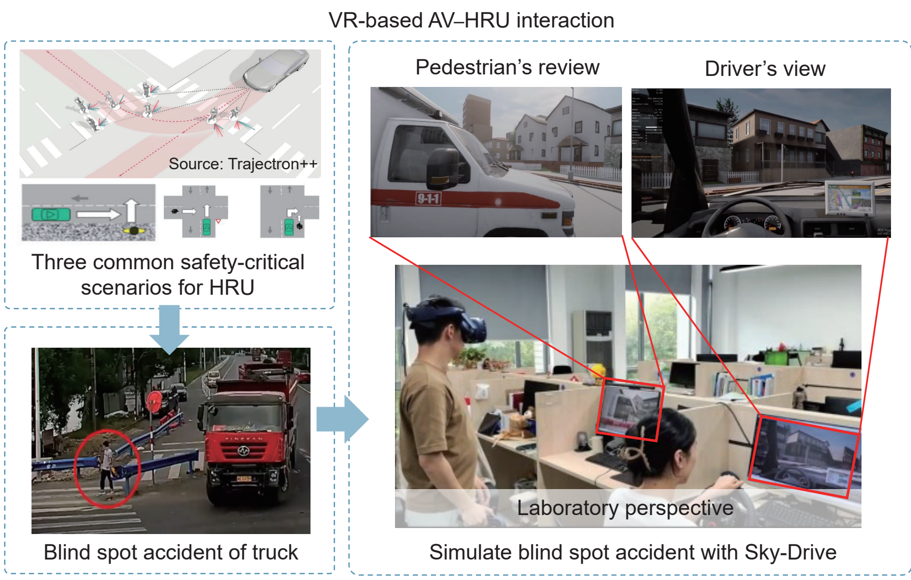

Fig. 4 VR-based experimental setup for studying AV–HRU interactions at unsignalized intersections.

steering takeovers synchronized with the vehicle state and surrounding scene context. Within its multiagent traffic simulation environment, human participants intervene when they are dissatisfied with the AV's behavior (e.g., aggressive merging, unsafe following), implicitly indicating suboptimal actions. These interventions are used to construct preference comparisons between pre- and post-takeover trajectories, allowing the agent to identify and avoid human-disapproved actions and iteratively refine its driving policy (Huang et al., 2024c).

Mathematically, the HAIM-DRL can be defined in the form of Eq. (1):

$$\pi_{\mathrm{AV}}^{*} = \mathrm{argmin}_{\pi_{\mathrm{AV}}} \mathbb{E}_{s_t \sim d_{\pi_{\mathrm{AV}}}} \left[\mathcal{L}(\pi_{\mathrm{AV}}(\cdot|s_t), \pi_{\mathrm{human}}(\cdot|s_t))\right] \quad (1)$$

where $d_{\pi_{\mathrm{AV}}}$ 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 align its behavior with human preferences.

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

$$\pi_{\mathrm{max}}(a|s) = \pi_{\mathrm{AV}}(a|s)(1 - I(s,a)) + \pi_{\mathrm{human}}(a|s)F(s) \quad (2)$$

where $F(s) = \int_{a' \notin A_\eta(s)} \pi_{\mathrm{AV}}(a'|s)\mathrm{d}a'$ represents the probability of the agent selecting an action that would be rejected by the human. $I(s,a)$ is an indicator function that equals 1 if the human rejects the agent action and 0 otherwise.

The overall learning objective of HAIM–DRL is specifically designed via Eq. (3):

$$\begin{aligned}\max_{\pi}\mathbb{E}\Big[&\psi\hat{Q}(s_t, a_t^{\mathrm{AV}}) - \alpha\log\pi_{\mathrm{AV}}(a_t^{\mathrm{AV}}|s_t;\theta) \\ &- \beta Q^{\mathrm{EX}}(s_t, a_t^{\mathrm{AV}}) - \varphi Q^{\mathrm{IM}}(s_t, a_t^{\mathrm{AV}})\Big]\end{aligned} \quad (3)$$

In Eq. (3), the first term $\hat{Q}\left(s_t, a_t^{\mathrm{AV}}\right)$ guides the agent to align with human-preferred behavior by minimizing the value discrepancy between its own actions and those demonstrated by the human mentor. The second term $\log\pi_{\mathrm{AV}}\left(a_t^{\mathrm{AV}}|s_t;\theta\right)$ introduces an entropy regularization factor that encourages the agent to explore diverse strategies. The third term $Q^{\mathrm{EX}}\left(s_t,a_t^{\mathrm{AV}}\right)$ penalizes actions that frequently trigger human takeovers. The fourth term $Q^{\mathrm{IM}}\left(s_t,a_t^{\mathrm{AV}}\right)$ constrains the agent to minimize disturbances to surrounding traffic.

As shown in Table 2, HAIM-DRL was successfully implemented and evaluated within the Sky-Drive platform, demonstrating clear advantages over conventional RL methods. Compared with PPO, HAIM-DRL achieves a drastic reduction in safety violations and eliminates the need for large-scale training data, reaching comparable or superior performance with only 8000 samples. Compared with HACO, which also leverages human interventions, HAIM-DRL further improves test returns, reduces the disturbance rate, and increases the success rate from 0.35 to 0.38. These results validate Sky-Drive's ability to support closed-loop human–AI training, enabling efficient, human-aligned policy learning through real-time feedback and preference-driven optimization.

5.3 Vision language model-enabled reinforcement learning

To validate Sky-Drive's ability to support VLM-enabled RL, we implemented VLM-RL (Huang et al., 2025), which integrates pretrained VLMs with RL to generate semantic reward signals from image observations and natural language goals. Specifically, we use the CLIP (ViT-bigG-14) model as the core VLM module, which is deployed via the 'open_clip' library. In each timestep, the agent's current observation is processed through the CLIP image encoder, whereas the language goals are encoded by the CLIP text encoder. The resulting semantic similarity scores are computed and normalized to serve as real-time reward feedback for policy learning. This demonstration highlights Sky-Drive's ability to enable high-level, human-interpretable guidance for safe and efficient autonomous driving.

At the core of VLM-RL is the contrast language goal (CLG)-as-reward paradigm, which uses pretrained VLMs to compute semantic similarity between driving states and paired language descriptions. Positive goals (e.g., "the road is clear with no accidents") and negative goals (e.g., "two cars have collided") are used to guide the agent's behavior by comparing how closely its current state aligns with each description. The reward is computed by encoding visual input via the CLIP image encoder and goals via its text encoder, both of which are mapped into a shared latent space, as shown in Eq. (4) (Huang et al., 2025):

$$\begin{aligned}R_{\mathrm{CLG}}(s) = \alpha \cdot \mathrm{sim}\left(\mathrm{VLM_I}(\psi(s)), \mathrm{VLM_L}(l_{\mathrm{pos}})\right) \\ - \beta \cdot \mathrm{sim}\left(\mathrm{VLM_I}(\psi(s)), \mathrm{VLM_L}(l_{\mathrm{neg}})\right)\end{aligned} \quad (4)$$

where l_{pos} and l_{neg} are the positive and negative language goals, respectively; $\mathrm{VLM_I}$ and $\mathrm{VLM_L}$ denote the image and language encoders of the pretrained VLM; $\psi(s)$ is the visual preprocessing function; and $\mathrm{sim}(\cdot,\cdot)$ represents cosine similarity. The weights α and β control the influence of the positive and negative goals, respectively.

To improve reward stability, VLM-RL introduces a hierarchical reward synthesis strategy that combines CLG-based semantic rewards with low-level vehicle state signals such as speed alignment, lane deviation, and directional consistency. The synthesized reward $r_t^{\mathrm{synthesis}}$ is as Eq. (5) (Huang et al., 2025):

$$r_t^{\mathrm{synthesis}} = f_t^{\mathrm{speed}} \cdot f_t^{\mathrm{center}} \cdot f_t^{\mathrm{angle}} \cdot f_t^{\mathrm{stability}} \quad (5)$$

Each term is bounded within [0, 1]. To improve readability, we explain each component below:

- Speed alignment $\left(f_t^{\mathrm{speed}}\right): f_t^{\mathrm{speed}} = 1 - \dfrac{\left|v_t - v_t^{\mathrm{target}}\right|}{v_{\mathrm{max}}}$ with $v_t^{\mathrm{target}} = r_t^{\prime\mathrm{CLG}} \cdot v_{\mathrm{max}}$. Here, $v_{\mathrm{max}} = 25\ \mathrm{km}\cdot\mathrm{h}^{-1}$, which simulates a low-speed urban environment. As it does not include lane-changing scenarios, the CLG reward is reflected primarily in the speed modulation term v_t^{target}, enabling the agent to learn cautious behaviors such as timely deceleration in response to potential collisions.

Table 2 Performance of the PPO, HACO, and HAIM-DRL methods

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	8000
HAIM–DRL (ours)	11.25	1590.85	0.0121	0.38	8000

Note: The results are based on data reported in our previous study (Huang et al., 2024c). For detailed definitions of the evaluation metrics and descriptions of the baseline methods, please refer to the original study.

• Lane alignment (f_t^{center}): This metric evaluates the vehicle's lateral position relative to the lane center.

• Heading alignment (f_t^{angle}): This measures the vehicle's orientation with respect to the road direction.

• Temporal stability ($f_t^{\text{stability}}$): accounts for the temporal consistency of the vehicle's lateral position relationships.

As shown in Table 3, VLM-RL significantly outperforms existing approaches across all key metrics. VLM-RL achieves the highest success rate and route completion while maintaining a low collision speed of 0.02 km·h^{-1}, which matches the safety level of the most conservative baselines. Unlike existing VLM-based methods, which suffer from overly cautious behavior and near-zero task success, VLM-RL balances safety with efficiency, reaching an average speed of 19.3 km·h^{-1} and a total driving distance of 2028.2 m. Compared with LLM-based methods such as Revolve, VLM-RL maintains comparable success and completion rates while drastically reducing the collision speed. The successful implementation of VLM-RL within the Sky-Drive platform validates its ability to support large-scale, multimodal policy learning grounded in human-understandable semantics.

5.4 Personalized safety-critical curriculum learning

To validate Sky-Drive's ability to support adaptive scenario generation and curriculum learning, we implement the CurricuVLM (Sheng et al., 2025). The CurricuVLM integrates VLMs to enable personalized, safety-critical training scenarios tailored to the evolving weaknesses of autonomous driving agents.

The core innovation of CurricuVLM lies in bridging the gap between scenario generation and policy learning. By continuously monitoring agent performance, the framework identifies failure patterns through a two-stage behavior analysis pipeline. Specifically, VLMs are first used to extract rich visual descriptions of unsafe events, which are then interpreted by a GPT-4o-based analyzer to uncover behavioral limitations. This process enables semantic understanding of critical driving mistakes without manual annotation.

On the basis of the analysis, scenario generation is formulated as a conditional trajectory generation problem:

$$P\left(Y^{\text{AV}}, Y^{\text{BV}} \mid I, X\right) \tag{6}$$

where X encodes the historical context (e.g., maps, past trajectories), I contains semantic insights from behavior analysis, and Y^{AV} and Y^{BV} denote future trajectories of the ego and background vehicles, respectively.

The framework optimizes Y^{BV} to generate targeted, informative interactions via Eq. (7):

$$Y^{\text{BV}*} = \text{argmax}_{Y^{\text{BV}}} P(Y^{\text{BV}}|X) \sum_{Y^{AV} \sim \mathcal{Y}(\pi)} P(Y^{\text{AV}}|Y^{\text{BV}}, X) \cdot P(I|Y^{\text{AV}}, Y^{\text{BV}}) \tag{7}$$

where $\mathcal{Y}(\pi)$ denotes the trajectory distribution induced by the current policy π and $P\left(I \mid Y^{\text{AV}}, Y^{\text{BV}}\right)$ represents how well the generated scenario aligns with the identified behavioral insight. This formulation encourages background vehicle behavior to induce targeted policy responses from the AV agent, forming the foundation for automated curriculum construction.

As shown in Table 4, CurricuVLM achieves the best overall performance across all key metrics, demonstrating both high safety and training effectiveness. In terms of task performance, CurricuVLM achieves the highest episode reward (48.9) and road completion rate (73.4%) while maintaining a low crash rate (25.1%), outperforming baselines such as CAT and CLIC. It also records the highest total driving distance (48.4 m) and failure-to-success rate (39.1%), indicating superior adaptability to previously failed scenarios. Moreover, its success-to-success rate (73.5%) reflects strong behavioral consistency and learning stability. These results validate that CurricuVLM not only enhances policy robustness under long-tail safety-critical scenarios but also

Table 3 Performance comparison with baselines during testing, with means and standard deviations over 3 seeds

Method	Average speed (km·h^{-1}) ↑	Route completion ↑	Traveled distance (m) ↑	Collision rate ↓	Success rate ↑
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**

Note: The best results are marked in bold. The results are based on data reported in our previous study (Huang et al., 2025). For detailed definitions of the evaluation metrics and descriptions of the baseline methods, please refer to the original study.

Table 4 Performance comparison with baselines in the safety-critical test scenarios

Model	Episode reward ↑	Road completion ↑	Total distance (m) ↑	Crash rate ↓	Average speed (km·h^{-1}) ↑	Failure-to-success rate ↑	Success-to-success rate ↑
SAC	38.4 ± 1.97	63.2 ± 1.21	40.9 ± 1.34	30.5 ± 2.33	9.25 ± 0.07	30.4 ± 7.00	56.9 ± 15.1
PPO	38.4 ± 0.86	62.7 ± 1.05	40.0 ± 0.70	32.0 ± 2.02	**9.94 ± 0.30**	26.7 ± 0.89	41.7 ± 8.33
TD3	42.4 ± 1.01	65.2 ± 1.40	42.6 ± 1.26	39.7 ± 1.04	8.02 ± 0.77	28.6 ± 2.79	64.3 ± 21.4
CAT	42.5 ± 3.95	66.6 ± 4.37	43.4 ± 3.48	32.1 ± 2.08	8.36 ± 1.17	35.2 ± 3.44	67.5 ± 7.50
CLIC	39.3 ± 0.72	64.3 ± 0.40	41.6 ± 0.78	26.2 ± 1.17	9.21 ± 0.26	34.7 ± 2.67	61.9 ± 26.9
CurricuVLM (ours)	**48.9 ± 1.53**	**73.4 ± 1.66**	**48.4 ± 1.31**	**25.1 ± 1.17**	9.45 ± 0.16	**39.1 ± 0.66**	**73.5 ± 21.1**

Note: The best results are marked in bold. The results are based on data reported in our previous study (Sheng et al., 2025). For detailed definitions of the evaluation metrics and descriptions of the baseline methods, please refer to the original study.

integrates seamlessly into Sky-Drive's human-AI collaboration mechanism. Some qualitative examples of generated scenarios are illustrated in Fig. 5.

5.5 Accident data replay

To validate Sky-Drive's ability to support real-world accident reconstruction and analysis, we implement an accident data replay framework that enables systematic reproduction of traffic collisions within the Sky-Drive environment. This framework addresses a fundamental challenge in autonomous driving development: understanding and learning from safety-critical real-world incidents in a controlled, repeatable, and risk-free setting.

The replay pipeline is built upon CenterTrack (Zhou et al., 2020), an advanced multiobject tracking algorithm used to extract 2D trajectories from raw accident video footage. These trajectories are then projected into 3D space through map matching, scene semantics inference, and spatial calibration. The reconstructed trajectories are then integrated into the Sky-Drive simulation engine, where they are rendered and executed via the platform's distributed multiagent architecture. Each agent involved in the accident (e.g., AV, HV, or pedestrian) is assigned to a separate terminal, allowing synchronized replay of the entire incident across multiple devices. For example, the motion of a vehicle from the original scene is reproduced on Terminal 1, whereas the actions of a pedestrian are replayed on Terminal 2. The global environment and timing are managed by the host terminal to maintain consistency across all agents. This design supports fine-grained multiagent analysis, including human-in-the-loop replay or intervention at any point in the incident.

As shown in Fig. 6, the key dynamics, such as vehicle positions, speeds, and interaction sequences, of the reconstructed scenes are preserved, along with contextual factors such as road layout and weather. Specifically, Sky-Drive incorporates a robust reconstruction validation process to ensure fidelity. A procedural matching algorithm selects the most appropriate simulation maps on the basis of road topology and scene semantics. A built-in quality assessment module scores the visual and kinematic

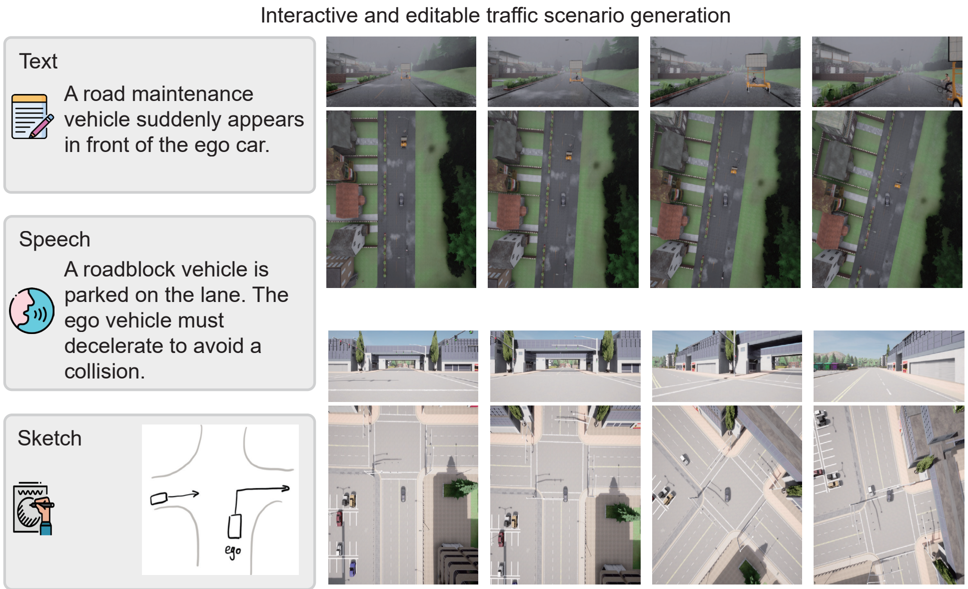

Fig. 5 Qualitative examples. Each scenario is downsampled to four frames for visualization.

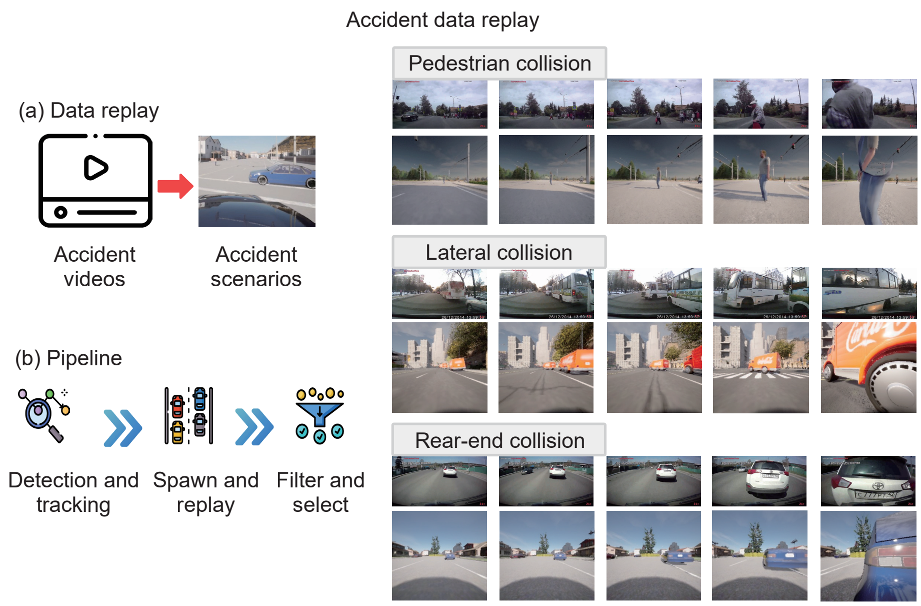

Fig. 6 Accident data replay framework for systematic traffic incident analysis.

https://doi.org/10.26599/JICV.2026.9210070

similarity between the replayed and original sequences, flagging low-fidelity cases for refinement. The framework also supports unsupervised domain adaptation to improve trajectory accuracy while offering manual editing tools when needed to ensure precise alignment with real-world footage.

This replay capability enables several downstream applications: (1) it provides a safe testbed for analyzing accident causation, enabling the development of improved safety mechanisms and behavior prediction models; (2) it allows RL agents to be trained and evaluated on real-world edge cases, significantly enhancing their robustness in critical scenarios; and (3) it supports regulatory compliance and postincident investigations by producing detailed, verifiable accident reconstructions. Through this integration, Sky-Drive enables scalable, high-fidelity replay of accident scenarios, positioning itself as a comprehensive platform for evaluating autonomous systems under rare, safety-critical conditions that are otherwise difficult or unsafe to replicate.

6 Future enhancements

6.1 Foundation model integration

1) Multimodal behavioral understanding: Interpreting human behavioral signals in a unified, context-aware manner remains an open challenge. Future iterations of sky drive leverage LLMs and VLMs to perform cross-modal reasoning across physiological, visual, verbal, and control-based modalities. For example, an elevated heart rate, downward gaze, and a quick verbal cue such as "too fast" may collectively indicate the driver's discomfort with vehicle acceleration. A more nuanced phrase such as "I feel a bit uneasy because the car accelerates too quickly." can be semantically aligned with facial tension and biometric signals such as heart rate variability. By combining these signals in the context of traffic density, road geometry, and interaction with nearby vehicles, Sky-Drive can construct rich behavioral profiles far beyond what single-modality systems can achieve. This capability will support personalized feedback generation, trust modeling, and adaptive control within the HAIM and AIHM modules.

2) Personalized autonomous driving: As shown in Fig. 7, Sky-Drive implements an LLM-based system that enables personalized autonomous driving through natural language interaction (Xu et al., 2025). Specifically, the system integrates three core modules: a visual encoder to process real-time camera feeds, an LLM to interpret language inputs, and a route planning module to generate executable commands on the basis of Sky-Drive maps. To ensure robustness, Sky-Drive uses a three-stage training pipeline. The first stage uses the BDD-X dataset to align visual and linguistic representations (Xu et al., 2017). The second stage fine-tunes language understanding via LoRA techniques on the SDN dataset (Ma et al., 2022). The final stage incorporates data generated within the Sky-Drive simulation environment to adapt model responses to realistic driving tasks. This integration will allow drivers to provide real-time feedback such as "slow down a bit here" or "take the next left" and have the vehicle respond accordingly. In the long term, this capability will support personalized, explainable, and user-aligned driving experiences.

3) Traffic brains: A Sky-Drive will position foundation models as intelligent "traffic brains" that govern decision-making in complex, multiagent traffic environments. While general-purpose models such as Qwen, GPT, and Llama exhibit strong language and reasoning abilities, they require domain-specific adaptation to meet the demands of autonomous driving. To address this challenge, Sky-Drive leverages transportation-specific datasets, including LMDrive (Shao et al., 2024), CCD (Bao et al., 2020), DoTA (Yao et al., 2023), and DriveCoT, to fine-tune pretrained foundation models (Wang et al., 2024b). This fine-tuning pipeline is designed to enhance the model's ability to handle dynamic scenario adaptation, hierarchical reasoning, and multitask decision-making, including generating safe control actions (e.g., steering, throttle, and braking) and predicting critical safety metrics such as time-to-collision (TTC) (Jiang et al., 2025). The refined models are deployed within the Sky-Drive simulation environment to enable coordinated behavior across AVs and other components, facilitating holistic control and system-level optimization.

6.2 Hardware-in-the-loop

1) Simulation-to-Reality integration: As shown in Fig. 1f, the center of the HIL framework is a Ford E-Transit electric van

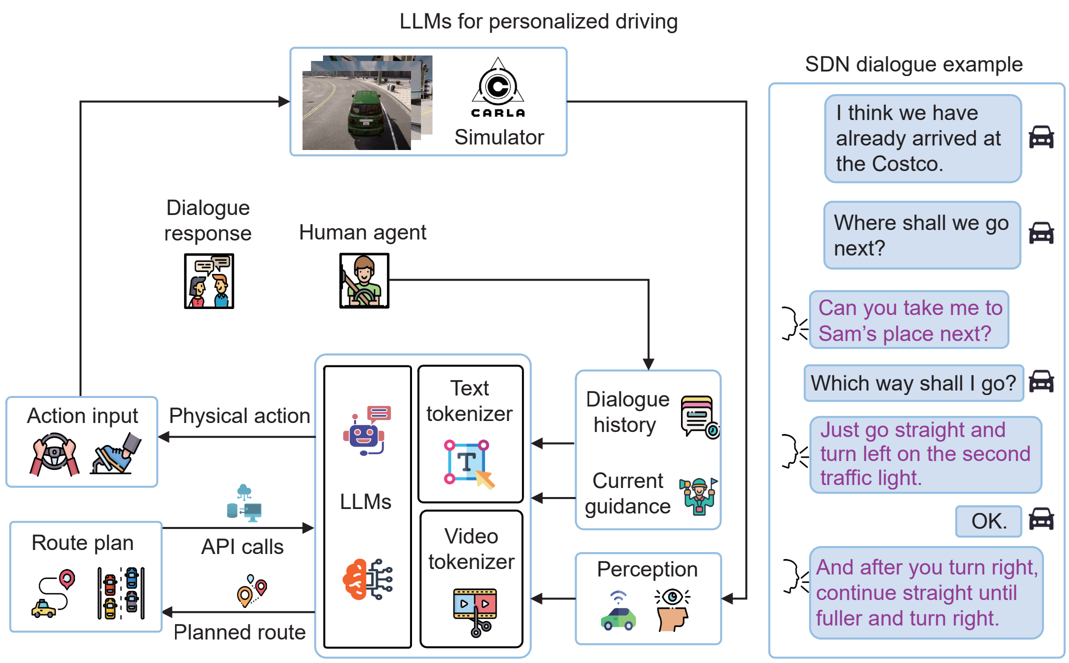

Fig. 7 LLM-based system enabling personalized autonomous driving.

retrofitted with fully automated driving capabilities. The vehicle is equipped with a comprehensive sensor suite, including LiDAR, radar, high-resolution cameras, and OxTS navigation units, and operates on a drive-by-wire system connected to an industrial-grade computing rack. The van will be tightly integrated with Sky-Drive's simulation environment via standardized ROS interfaces, allowing bidirectional data flow between simulation and physical execution. Simulated scenarios generated by Sky-Drive can be streamed as virtual sensor data (e.g., LiDAR point clouds, camera feeds) to the vehicle's onboard system, where perception and control algorithms interpret them in real time. Conversely, control outputs from the van (e.g., steering, throttle, braking) can be logged and compared against reference trajectories from the simulation. This closed-loop integration allows researchers to evaluate real hardware behavior under diverse and repeatable traffic scenarios, including safety-critical scenarios that are difficult or unsafe to test on public roads. By testing algorithms in simulations before deployment, Sky-Drive is expected to bridge the sim-to-real gap and ensure that experimental validation can be conducted safely without exposing operators or the public to real-world risks.

2) Remote driving and data collection: The HIL framework will establish 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-vehicle (V2X) communication, specified in the 3rd generation partnership project (3GPP) standards (Huang et al., 2024a; Wang et al., 2017; You et al., 2024). Sky-Drive's ROS integration enables wireless connectivity between the testbed vehicle and the human-in-the-loop framework via cellular or satellite networks. To support this, Sky-Drive integrates a 5G mmWave mobile hotspot (i.e., NETGEAR NighTalk M6 Pro) to enable remote control and real-time data collection. In human-in-the-loop experiments, a remote operator may control the van via the same multimodal input system used in the simulation (e.g., VR, logitch steering wheel). Physiological signals (e.g., heart rate), eye gaze, and intervention patterns are captured in parallel and analyzed to study driver behavior and human–AI collaboration under realistic conditions. Moreover, it will support remote collection of human behavioral data (e.g., steering patterns and pedal inputs), even when the vehicle and experiment coordinators are located in different cities or campuses. In addition, sky drives support the remote deployment of large foundation models (e.g., LLMs and VLMs). Sensor data from the vehicle are transmitted wirelessly via 5G to a nearby workstation for real-time inference, such as scene understanding or safety assessment (Yao and Chuah, 2025). The results are then sent back to the vehicle for decision-making. This edge-cloud architecture is expected to help overcome the onboard computational limitations of AVs.

7 Conclusions

This study presented Sky-Drive, a distributed multiagent simulation platform designed for socially aware autonomous driving and human-AI collaboration in future transportation systems. Unlike existing simulators that focus primarily on validating single-vehicle performance, Sky-Drive addresses the emerging need to explore complex interactions in mixed traffic environments where various intelligent agents must align with human preferences and societal norms.

Sky-Drive introduces several key innovations: (1) a distributed multiagent architecture enabling synchronized simulation across multiple terminals, allowing independent agent control while maintaining shared environmental states; (2) a multimodal human-in-the-loop framework integrating diverse sensors to capture comprehensive behavioral data; (3) a novel human-AI collaboration mechanism to facilitate bidirectional knowledge exchange; and (4) a digital twin framework creating high-fidelity virtual replicas of real transportation systems. The platform's effectiveness has been demonstrated through multiple application cases, including VR-based vulnerable road-user interactions, physics-enhanced reinforcement learning with human feedback, vision-language model-enabled reinforcement learning, personalized curriculum learning, and accident data replay. These applications show Sky-Drive's potential to advance autonomous driving research beyond traditional metrics of safety and efficiency toward more socially aware and human-aligned behavior.

Sky-Drive is open source in its core design, including the distributed architecture, scenario generation tools, and HAIM and AIHM modules. It integrates open-source software such as CARLA and custom synchronization code. However, some components, such as the VR headset, logitch steering wheel, and Garmin smartwatch, are commercial hardware used for data collection and human-in-the-loop interaction. To further enhance Sky-Drive's capabilities, we have outlined two major planned functionalities: (1) the integration of foundation models to support multimodal behavior understanding, personalized driving, and system-level optimization via traffic brains; and (2) a hardware-in-the-loop framework via ROS integration to enable direct validation of algorithms on physical vehicles. These future enhancements will bridge the gap between simulation and reality, ensuring that algorithms are safely evaluated in real-world environments. As autonomous driving technology continues to evolve, sky drive will serve as a robust platform for ensuring that future transportation systems are not only safe and efficient but also socially aware of and aligned with human expectations.

Replication and data sharing

The data and codes that support the findings of this study are available at https://sky-lab-uw.github.io/Sky-Drive-website.

Acknowledgements

This work was funded by the U.S. Department of Transportation (No. #69A3552348305). In addition, this research was supported by grants from NVIDIA and utilized NVIDIA RTX PRO 6000 Blackwell Max-Q Workstation Edition GPUs and A100 GPU-Hours on Saturn Cloud.

Declaration of competing interest

The authors have no competing interests to declare that are relevant to the content of this article.

References

Almaskati, D., Kermanshachi, S., Pamidimukkala, A., 2024. Convergence of emerging transportation trends: A comprehensive review of shared autonomous vehicles. *J Intel Connect Veh*, **7**, 177–189.

Amini, A., Wang, T. H., Gilitschenski, I., Schwarting, W., Liu, Z., Han, S., et al., 2022. VISTA 2.0: An open, data-driven simulator for multi-modal sensing and policy learning for autonomous vehicles. In: 2022 International Conference on Robotics and Automation (ICRA), 2419–2426.

Bao, W., Yu, Q., Kong, Y., 2020. Uncertainty-based traffic accident anticipation with spatio-temporal relational learning. In: Proceedings of the 28th ACM International Conference on Multimedia, 2682–2690.

Berta, R., Lazzaroni, L., Capello, A., Cossu, M., Forneris, L., Pighetti, A., et al., 2024. Development of deep-learning-based autonomous agents for low-speed maneuvering in unity. *J Intel Connect Veh*, **7**, 229–244.

Cai, P., Lee, Y., Luo, Y., Hsu, D., 2020. SUMMIT: A simulator for urban driving in massive mixed traffic. In: 2020 IEEE International Conference on Robotics and Automation (ICRA), 4023–4029.

Chen, S., Zong, S., Chen, T., Huang, Z., Chen, Y., Labi, S., 2023. A taxonomy for autonomous vehicles considering ambient road infrastructure. *Sustainability*, **15**, 11258.

Dosovitskiy, A., Ros, G., Codevilla, F., Lopez, A., Koltun, V., 2017. CARLA: an open urban driving simulator. In: Proceedings of the 1st Annual Conference on Robot Learning, 1–16.

Fellendorf, M., Vortisch, P., 2010. *Microscopic Traffic Flow Simulator VISSIM*. In: Fundamentals of Traffic Simulation, Barceló, J., Eds. New York: Springer, 63–93.

Feng, S., Sun, H., Yan, X., Zhu, H., Zou, Z., Shen, S., et al., 2023. Dense reinforcement learning for safety validation of autonomous vehicles. *Nature*, **615**, 620–627.

Gulino, C., Fu, J., Luo, W., Tucker, G., Bronstein, E., Lu, Y., et, al., 2023. Waymax: an accelerated, data-driven simulator for large-scale autonomous driving research. In: Proceedings of the 37th International Conference on Neural Information Processing Systems, 7730–7742.

Hong, C. J., Aparow, V. R., 2021. System configuration of human-in-the-loop simulation for level 3 autonomous vehicle using IPG CarMaker. In: 2021 IEEE International Conference on Internet of Things and Intelligence Systems (IoTaIS), 215–221.

Huang, Z., Chen, S., Pian, Y., Sheng, Z., Ahn, S., Noyce, D. A., 2024a. Toward C-V2X enabled connected transportation system: RSU-based cooperative localization framework for autonomous vehicles. *IEEE Trans Intell Transp Syst*, **25**, 13417−13431.

Huang, Z., Sheng, Z., Chen, S., 2024b. Trustworthy human-AI collaboration: Reinforcement learning with human feedback and physics knowledge for safe autonomous driving. https://arxiv.org/abs/2409.00858

Huang, Z., Sheng, Z., Ma, C., Chen, S., 2024c. Human as AI mentor: Enhanced human-in-the-loop reinforcement learning for safe and efficient autonomous driving. *Commun Transp Res*, **4**, 100127.

Huang, Z., Sheng, Z., Qu, Y., You, J., Chen, S., 2025. VLM-RL: A unified vision language models and reinforcement learning framework for safe autonomous driving. *Transp Res Part C Emerg Technol*, **180**, 105321.

Jiang, S., Huang, Z., Qian, K., Luo, Z., Zhu, T., Zhong, Y., et al., 2025. A survey on vision-language-action models for autonomous driving. https://arxiv.org/abs/2506.24044

Jin, Y., Yang, R., Yi, Z., Shen, X., Peng, H., Liu, X. et al., 2024. SurrealDriver: Designing LLM-powered generative driver agent framework based on human drivers' driving-thinking data. In: 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 966–971.

Kalra, N., Paddock, S. M., 2016. Driving to safety: How many miles of driving would it take to demonstrate autonomous vehicle reliability? *Transp Res Part A Policy Pract*, **94**, 182–193.

Kesting, A., Treiber, M., Helbing, D., 2007. General lane-changing model MOBIL for car-following models. *Transp Res Rec J Transp Res Board*, **1999**, 86–94.

Koenig, N., Howard, A., 2004. Design and use paradigms for gazebo, an open-source multirobot simulator. In: 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2149–2154.

Leurent, E., Mercat, J., 2019. Social attention for autonomous decision-making in dense traffic. https://arxiv.org/abs/1911.12250

Li, Q., Peng, Z., Feng, L., Liu, Z., Duan, C., Mo, W., et al., 2023a. ScenarioNet: Open-source platform for large-scale traffic scenario simulation and modeling. https://arxiv.org/abs/2306.12241

Li, Q., Peng, Z., Feng, L., Zhang, Q., Xue, Z., Zhou, B., 2023b. MetaDrive: Composing diverse driving scenarios for generalizable reinforcement learning. *IEEE Trans Pattern Anal Mach Intell*, **45**, 3461–3475.

Li, Y., Deng, H., Xu, X., Wang, W., 2019. Modelling and testing of in-wheel motor drive intelligent electric vehicles based on co-simulation with Carsim/Simulink. *IET Intell Transp Syst*, **13**, 115−123.

Li, Y., Yuan, W., Zhang, S., Yan, W., Shen, Q., Wang, C., et al., 2024. Choose your simulator wisely: A review on open-source simulators for autonomous driving. *IEEE Trans Intell Veh*, **9**, 4861−4876.

Liao, H., Shen, H., Li, Z., Wang, C., Li, G., Bie, Y., et al., 2024. GPT-4 enhanced multimodal grounding for autonomous driving: Leveraging cross-modal attention with large language models. *Commun Transp Res*, **4**, 100116.

Liu, Y., Wu, F., Liu, Z., Wang, K., Wang, F., Qu, X., 2023. Can language models be used for real-world urban-delivery route optimization?, *Innovation*, **4**, 100520.

Liu, J., Zhou, D., Hang, P., Ni, Y., Sun, J., 2024. Towards socially responsive autonomous vehicles: A reinforcement learning framework with driving priors and coordination awareness. *IEEE Trans Intell Veh*, **9**, 827–838.

Long, K., Sheng, Z., Shi, H., Li, X., Chen, S., Ahn, S., 2025. Physical enhanced residual learning (PERL) framework for vehicle trajectory prediction. *Commun Transp Res*, **5**, 100166.

Lopez, P. A., Behrisch, M., Bieker-Walz, L., Erdmann, J., Flötteröd, Y. P., Hilbrich, R., et al., 2018. Microscopic traffic simulation using SUMO. In: 2018 21st International Conference on Intelligent Transportation Systems (ITSC), 2575–2582.

Luo, Y., Cai, P., Lee, Y., Hsu, D., 2022. GAMMA: A general agent motion model for autonomous driving. *IEEE Robot Autom Lett*, **7**, 3499–3506.

Lv, D., Wang, Y., Wang, L., Fei, Y., Wang, K., Qu, X., 2024. Modular flying vehicles: Scheduling modes, social benefits, and challenges. *Commun Transp Res*, **4**, 100144.

Koenig, Z., Sun, Q., Matsumaru, T., 2024. Bidirectional planning for autonomous driving framework with large language model. *Sensors*, **24**, 6723.

Ma, Z., VanDerPloeg, B., Bara, C.-P., Yidong, H., Kim, E.-I., Gervits, F., et al., 2022. DOROTHIE: Spoken dialog for handling unexpected situations in interactive autonomous driving agents. https://arxiv.org/abs/2210.12511

Marcu, A.-M., Chen, L., Hünermann, J., Karnsund, A., Hanotte, B., Chidananda, P., et al., 2024. LingoQA: Visual question answering for autonomous driving. https://doi.org/10.48550/arXiv.2312.14115

Müller, M., Casser, V., Lahoud, J., Smith, N., Ghanem, B., 2018. Sim4CV: A photo-realistic simulator for computer vision applications. *Int J Comput Vis*, **126**, 902–919.

Palanisamy, P., 2020. Multi-agent connected autonomous driving using deep reinforcement learning. In: 2020 International Joint Conference on Neural Networks (IJCNN), 1–7.

Radford, A., Kim, J.W., Xu, T., Brockman, G., McLeavey, C., Sutskever, I., et al., 2023. Robust speech recognition via large-scale weak supervision. In: Proceedings of the 40th International Conference on Machine Learning, 28492–28515.

Ramamohanarao, K., Xie, H., Kulik, L., Karunasekera, S., Tanin, E., Zhang, R., et al., 2016. SMARTS: Scalable microscopic adaptive road traffic simulator. *ACM Trans Intell Syst Technol*, **8**, 1–22.

Rong, G., Shin, B. H., Tabatabaee, H., Lu, Q., Lemke, S., Možeiko , M., et al., 2020. LGSVL simulator: A high fidelity simulator for autonomous driving. In: 2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC), 1–6.

Shafiee, M. J., Nentwig, M., Kassahun, Y., Li, F., Bochkarev, S., Kamal, A., et al., 2019. Human-machine collaborative design for accelerated design of compact deep neural networks for autonomous driving. https://arxiv.org/abs/1909.05587

Shah, S., Dey, D., Lovett, C., Kapoor, A., 2017. Airsim: High-fidelity visual and physical simulation for autonomous vehicles. In: Field and Service Robotics: Results of the 11th International Conference. Springer, 621–635.

Shao, H., Hu, Y., Wang, L., Song, G., Waslander, S. L., Liu, Y., et al., 2024.

LMDrive: Closed-loop end-to-end driving with large language models. In: 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 15120–15130.

Sheng, Z., Huang, Z., Chen, S., 2024a. Kinematics-aware multigraph attention network with residual learning for heterogeneous trajectory prediction. *J Intel Connect Veh*, **7**, 138–150.

Sheng, Z., Huang, Z., Chen, S., 2024b. Traffic expertise meets residual RL: Knowledge-informed model-based residual reinforcement learning for CAV trajectory control. *Commun Transp Res*, **4**, 100142.

Sheng, Z., Huang, Z., Chen, S., 2024c. Ego-planning-guided multi-graph convolutional network for heterogeneous agent trajectory prediction. *Comput Aided Civ Infrastruct Eng*, **39**, 3357–3374.

Sheng, Z., Huang, Z., Qu, Y., Leng, Y., Bhavanam, S., Chen, S., 2025. CurricuVLM: Towards safe autonomous driving via personalized safety-critical curriculum learning with vision-language models. https://arxiv.org/abs/2502.15119

Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M. A., Lacroix, T., et al., 2023. LLaMA: Open and efficient foundation language models. https://arxiv.org/abs/2302.13971

Treiber, M., Hennecke, A., Helbing, D., 2000. Congested traffic states in empirical observations and microscopic simulations. *Phys Rev E*, **62**, 1805–1824.

Vinitsky, E., Lichtlé, N., Yang, X., Amos, B., Foerster, J., 2023. Nocturne: A scalable driving benchmark for bringing multiagent learning one step closer to the real world. https://arxiv.org/abs/2206.09889

Wang, P., Bai, S., Tan, S., Wang, S., Fan, Z., Bai, J., et al., 2024a. Qwen2-VL: Enhancing vision-language model's perception of the world at any resolution. https://arxiv.org/abs/2409.12191

Wang, T., Xie, E., Chu, R., Li, Z., Luo, P., 2024b. DriveCoT: Integrating chain-of-thought reasoning with end-to-end driving. https://doi.org/10.48550/arXiv.2403.16996

Wang, W., Wang, L., Zhang, C., Liu, C., Sun, L., 2022. Social interactions for autonomous driving: A review and perspectives. *Found Trends Robot*, **10**, 198–376.

Wang, X., Mao, S., Gong, M. X., 2017. An overview of 3GPP cellular vehicle-to-everything standards. *GetMobile Mob Comput Commun*, **21**, 19–25.

Wei, Y., Wang, Z., Lu, Y., Xu, C., Liu, C., Zhao, H., et al., 2024. Editable scene simulation for autonomous driving *via* collaborative LLM-agents. In: 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 15077–15087.

Wu, C., Kreidieh, A. R., Parvate, K., Vinitsky, E., Bayen, A. M., 2022. Flow: A modular learning framework for mixed autonomy traffic. *IEEE Trans Robot*, **38**, 1270–1286.

Wymann, B., Dimitrakakis, C., Sumner, A., Espie, E., Guionneau, C., 2005. TORCS: the open racing car simulator. *Software*, **4**, 1–5.

Xu, H., Gao, Y., Yu, F., Darrell, T., 2017. End-to-end learning of driving models from large-scale video datasets. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 3530–3538.

Xu, Z., Chen, T., Huang, Z., Xing, Y., Chen, S., 2025. Personalizing driver agent using large language models for driving safety and smarter human–machine interactions. *IEEE Intell Transp Syst Mag*, **17**, 96–111.

Yang, X., Wen, L., Ma, Y., Mei, J., Li, X., Wei, T., et al., 2024. DriveArena: A closed-loop generative simulation platform for autonomous driving. https://arxiv.org/abs/2408.00415

Yao, Y., Wang, X., Xu, M., Pu, Z., Wang, Y., Atkins, E., et al., 2023. DoTA: Unsupervised detection of traffic anomaly in driving videos. *IEEE Trans Pattern Anal Mach Intell*, **45**, 444–459.

Yao, Z., Chuah, M.C., 2025. Event-guided low-light video semantic segmentation. In: 2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), 3330–3341.

You, J., Shi, H., Jiang, Z., Huang, Z., Gan, R., Wu, K., et al., 2024. V2X-VLM: End-to-end v2x cooperative autonomous driving through large vision-language models. https://arxiv.org/abs/2408.09251

Zhang, H., Feng, S., Liu, C., Ding, Y., Zhu, Y., Zhou, Z., et al., 2019. CityFlow: A multi-agent reinforcement learning environment for large scale city traffic scenario. In: The World Wide Web Conference, 3620–3624.

Zhang, J., Xu, C., Li, B., 2024. ChatScene: Knowledge-enabled safety-critical scenario generation for autonomous vehicles. In: 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 15459–15469.

Zhou, L., Song, Y., Gao, Y., Yu, Z., Sodamin, M., Liu, H., et al., 2023. GarchingSim: An autonomous driving simulator with photorealistic scenes and minimalist workflow. In: 2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC), 4227–4232.

Zhou, M., Luo, J., Villella, J., Yang, Y., Rusu, D., Miao, J., et al., 2021. SMARTS: An open-source scalable multiagent RL training school for autonomous driving, In: Proceedings of the 2020 Conference on Robot Learning, 264–285.

Zhou, X., Koltun, V., Krähenbühl, P., 2020. Tracking objects as points. In: 2020 European Conference on Computer Vision, 474–490.

Zilin Huang is a Ph.D. candidate in the Department of Civil and Environmental Engineering at the University of Wisconsin–Madison (UW–Madison), WI, USA. Prior to joining UW–Madison, he worked at the Center for Connected and Automated Transportation (CCAT) at Purdue University, IN, USA. He received the M.S. degree in communication and transportation engineering from the South China University of Technology in 2021, and the B.S. degree from the School of Electromechanical Engineering, Guangdong University of Technology in 2018. His research interests include human-centered AI, autonomous driving, human–AI collaboration, robotics, and intelligent transportation.

Zihao Sheng received the B.S. degree in automation from Xi'an Jiaotong University in 2019 and the M.S. degree in control engineering from Shanghai Jiao Tong University in 2022. He is currently pursuing the Ph.D. degree at the Department of Civil and Environmental Engineering, University of Wisconsin–Madison. His main research interests include human-centered AI, autonomous driving, reinforcement learning, and intelligent transportation.

Zhengyang Wan received the B.S. degree in vehicle engineering from Tongji University in 2022, and the M.S. degree in transportation engineering from Tongji University in 2025. He is currently pursuing the Ph.D. degree in civil and environmental engineering at the University of Wisconsin–Madison. His research interests include autonomous driving, foundation model, and human-centered AI.

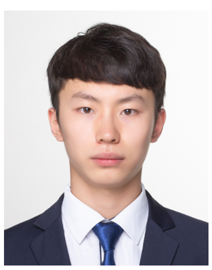

Yansong Qu received the B.S. degree in traffic engineering from Dalian Jiaotong University, Dalian, China, in 2021, and the M.S. degree from the College of Metropolitan Transportation at Beijing University of Technology, Beijing, China. He is currently pursuing the Ph.D. degree in civil engineering at Purdue University. His research interests include autonomous driving, reinforcement learning, and vision-language models.

清華大学出版社 Tsinghua University Press | Available on IEEE Xplore®

Yuhao Luo received the B.S. degree in vehicle engineering from Wuhan University of Technology, Wuhan, China, in 2023, and the M.S. degree with the Thrust of Intelligent Transportation, The Hong Kong University of Science and Technology (Guangzhou). He is currently pursuing the Ph.D. degree in civil and environmental engineering at the University of Wisconsin–Madison. His research interests include end-to-end autonomous driving, spatial-temporal data mining, and driving behavior modelling.

Boyue Wang received the B.S. degree in mechanical engineering from the University of Wisconsin–Madison, Madison, WI, USA, in 2025. He is currently pursuing the Ph.D. degree in civil and environmental engineering at the University of Wisconsin–Madison. His research interests include autonomous driving, reinforcement learning, and intelligent transportation systems.

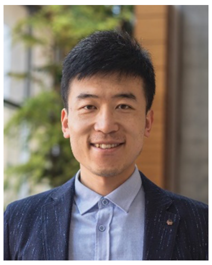

Pei Li received the Ph.D. in civil engineering from the University of Central Florida in 2021, after which he served as a Postdoctoral Research Fellow at the University of Michigan Transportation Research Institute. He is currently an Assistant Professor at the University of Wyoming. His research interests include smart mobility, machine learning, digital twins, and connected and automated vehicles.

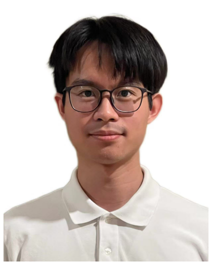

Yen-Jung Chen received the B.S. degree in computer engineering from Purdue University, West Lafayette, IN, USA, in 2024, where he is currently pursuing the M.S. degree in computer engineering. His research interests include computer vision, reinforcement learning, and vision-language models.

Jiancong Chen received the B.S. degree in transportation and communication from Chang'an University, Xi'an, China, in 2021 and M.S. degree in transportation engineering from Southeast University, Nanjing, China, in 2024. He is currently pursuing the Ph.D. degree at the Lyles School of Civil and Construction Engineering, Purdue University. His research interests lie at the intersection of artificial intelligence, statistics and robotics, with a focus on their integration in autonomous driving systems.

Keke Long received the Ph.D. degree in civil engineering from the University of Wisconsin–Madison, USA, in 2024, the M.S. degree in traffic engineering from Tongji University, China, in 2021, and the B.S. degree in traffic engineering from Chang'an University, China, in 2018. She is currently a Postdoc Research Associate at the University of Wisconsin–Madison. Her main research interests are connected and autonomous vehicles and artificial intelligence in transportation.

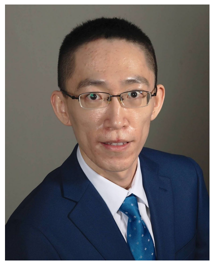

Jiayi Meng received the Ph.D. degree in computer science at Purdue University in 2023. His research focuses on immersive computing (e.g., VR/AR/MR and digital twins), mobile and edge computing, and cellular networks (4G/5G and beyond). He is currently an Assistant Professor in the Department of Computer Science and Engineering at the University of Texas at Arlington. His research aims to bridge the gap between the ever-increasing resource demands of immersive computing and the limited hardware capabilities of mobile devices over next-generation cellular networks. He has received the Rising STARs award from the University of Texas System (2023). His works have received ACM EdgeSys Best Paper Award (2022) and ACM IMC Best Paper Runner-Up Award (2023).

Yue Leng is a Senior Software Engineer at Google Cloud, where she works on integrating Gemini's generative AI to transform the full lifecycle of cloud application management. Before joining Google, she was a Quantitative Researcher at JPMorgan. She holds the Ph.D. degree in biomedical engineering and the M.S. degree in computer science from Purdue University. She has been driven by a passion for applying cutting-edge AI to solve tangible, real-world problems and deliver significant business value and innovation.

Sikai (Sky) Chen is an Assistant Professor in the Department of Civil and Environmental Engineering at the University of Wisconsin–Madison, with a courtesy appointment in the Department of Mechanical Engineering. He also serves as the Associate Director of the university's Smart Highway Research Center (SHRC), leading the Digital Infrastructure, Data Science, and Artificial Intelligence Program. He received the Ph.D. degree in civil engineering with a focus on Computational Science & Engineering from Purdue University. His research centers around three major themes: human users, AI, and transportation. He aims to innovate and develop safe, efficient, resilient, and human-centered transportation systems using cutting-edge methods and technologies. The focus is on incorporating human behaviors, interactive autonomy, digital infrastructure, and domain-specific AI. His Ph.D. dissertation, "Safety Implications of Roadway Design and Management: New Evidence and Insights in The Traditional and Emerging (Autonomous Vehicle) Operating Environments", received the 2019 Milton Pikarsky Memorial Award (a national award for best Ph.D. dissertation). He serves as PI and co-PI for research projects funded by USDOT, WisDOT, and industry partners. In addition, he is a Member of two ASCE national committees: Connected & Autonomous Vehicle Impacts, and Economics & Finance; and IEEE Emerging Transportation Technology Testing Technical Committee.