Novelty Assessment Report

Paper: RAP: 3D Rasterization Augmented End-to-End Planning
PDF URL: https://openreview.net/pdf?id=a9bOgeqbdB
Venue: ICLR 2026 Conference Submission
Year: 2026
Report Generated: 2025-12-29

Abstract

Imitation learning for end-to-end driving trains policies only on expert demonstrations. Once deployed in a closed loop, such policies lack recovery data: small mistakes cannot be corrected and quickly compound into failures. A promising direction is to generate alternative viewpoints and trajectories beyond the logged path. Prior work explores photorealistic digital twins via neural rendering or game engines, but these methods are prohibitively slow and costly, and thus mainly used for evaluation. In this work, we argue that photorealism is unnecessary for training end-to-end planners. What matters is semantic fidelity and scalability: driving depends on geometry and dynamics, not textures or lighting. Motivated by this, we propose 3D Rasterization, which replaces costly rendering with lightweight rasterization of annotated primitives, enabling augmentations such as counterfactual recovery maneuvers and cross-agent view synthesis. To transfer these synthetic views effectively to real-world deployment, we introduce a Raster-to-Real (R2R) feature-space alignment that bridges the sim-to-real gap at the representation level. Together, these components form the Rasterization Augmented Planning (RAP) pipeline, a scalable data augmentation framework for planning. RAP achieves state-of-the-art closed-loop robustness and long-tail generalization, ranking 1st on four major benchmarks: NAVSIM v1/v2, Waymo Open Dataset Vision-based E2E Driving, and Bench2Drive. Our results demonstrate that lightweight rasterization with feature alignment suffices to scale end-to-end training, offering a practical alternative to photorealistic rendering. Code will be released.

Core Task Landscape

This paper addresses: **End-to-End Autonomous Driving Planning with Data Augmentation**

A total of **30 papers** were analyzed and organized into a taxonomy with **19 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Synthetic Data Generation and Augmentation Methods**
- **End-to-End Learning Architectures and Training Frameworks**
- **Imitation Learning and Behavioral Cloning**
- **Trajectory Selection and Multi-Modal Planning**
- **Reinforcement Learning for Autonomous Driving**

Complete Taxonomy Tree

- End-to-End Autonomous Driving Planning with Data Augmentation Survey Taxonomy
- Synthetic Data Generation and Augmentation Methods
  - Diffusion-Based Video and Image Generation (4 papers)
  - [2] Improving end-to-end autonomous driving with synthetic data from latent diffusion models (H Goel, 2024) View paper
  - [4] Unleashing Generalization of End-to-End Autonomous Driving with Controllable Long Video Generation (Ma, 2024) View paper
  - [5] Syndiff-ad: Improving semantic segmentation and end-to-end autonomous driving with synthetic data from latent diffusion models (Goel Harsh, 2024) View paper
  - [12] World model-based end-to-end scene generation for accident anticipation in autonomous driving (Yanchen Guan, 2025) View paper
  - Rasterization and Primitive-Based Rendering ★ (1 papers)
  - [0] RAP: 3D Rasterization Augmented End-to-End Planning (Anon et al., 2026) View paper
  - Style Transfer and Image Transformation (1 papers)
  - [20] Data augmentation technology driven by image style transfer in self-driving car based on end-to-end learning (Dongjie Liu, 2020) View paper
  - Simulation-Based Data Generation (2 papers)
  - [3] Learning robust control policies for end-to-end autonomous driving from data-driven simulation (Alexander Amini, 2020) View paper
  - [14] TrafficSim: Learning to Simulate Realistic Multi-Agent Behaviors (Simon Suo, 2021) View paper
- End-to-End Learning Architectures and Training Frameworks
  - Unsupervised and Self-Supervised Pretraining (1 papers)
  - [1] End-to-end autonomous driving without costly modularization and 3d manual annotation (Guo, 2025) View paper
  - Multi-Modal and Foundation Model Integration (2 papers)
  - [23] Drive Anywhere: Generalizable End-to-end Autonomous Driving with Multi-modal Foundation Models (Tsun-Hsuan Wang, 2023) View paper
  - [29] Exploring versatile neural architectures across modalities and perception tasks (Zhang, 2023) View paper
  - Knowledge Distillation and Multi-Mode Planning (1 papers)
  - [10] DistillDrive: End-to-End Multi-Mode Autonomous Driving Distillation by Isomorphic Hetero-Source Planning Model (Yu Rui, 2025) View paper

- Viewpoint Robustness and 3D Scene Reconstruction (1 papers)
- [17] VR-Drive: Viewpoint-Robust End-to-End Driving with Feed-Forward 3D Gaussian Splatting (Kang Jae Young, 2025) View paper
- World Model and Temporal Reasoning (1 papers)
- [16] KARNet: Kalman Filter Augmented Recurrent Neural Network for Learning World Models in Autonomous Driving Tasks (Manjunatha, 2023) View paper

- Imitation Learning and Behavioral Cloning
  - Conditional Imitation with Command Augmentation (2 papers)
  - [7] End-to-End Autonomous Guidance Method Integrated with Mixture-of-Experts for Intelligent Vehicles (Bowen Li, 2025) View paper
  - [26] Safer End-to-End Autonomous Driving via Conditional Imitation Learning and Command Augmentation (Wang, 2022) View paper
  - Data Balancing and Augmentation for Imitation (3 papers)
  - [9] An imitation learning method with data augmentation and post processing for planning in autonomous driving (W Xi, 2023) View paper
  - [28] Robust Behavioral Cloning for Autonomous Vehicles using End-to-End Imitation Learning (Tanmay Vilas Samak, 2022) View paper
  - [30] Meaningful Data Augmentation under Unbalanced Data in End-to-end Learning for Autonomous Driving (HM Chen, n.d.) View paper
  - Multi-State and Scenario-Specific Imitation (2 papers)
  - [18] End-to-end models for self-driving cars on UPB campus roads (Andrei Mihalea, 2019) View paper
  - [24] Multi-State End-to-End Learning for Autonomous Vehicle Lateral Control (S. Mentasti, 2020) View paper
  - Continual and One-Shot Learning from Disengagements (1 papers)
  - [11] Autonomous Driving Policy Continual Learning With One-Shot Disengagement Case (Zhong Cao, 2023) View paper
  - Vision-Based End-to-End Imitation Advances (2 papers)
  - [13] End-to-End Learning for Autonomous Driving in Secured Smart Cities (Dapeng Guo, 2021) View paper
  - [25] Advancing Vision-based End-to-End Autonomous Driving (Yi, 2024) View paper
  - Steering Angle Prediction and Lateral Control (2 papers)
  - [22] Prediction of Steering Angle in Autonomous Vehicles Using Deep Learning Approach (Yogesh Karemore, 2025) View paper
  - [27] Classifying driver behaviors for autonomous vehicle navigation (E Cheung, 2018) View paper
- Trajectory Selection and Multi-Modal Planning
  - Trajectory Scoring and Safety Assessment (2 papers)
  - [8] DriveSuprim: Towards Precise Trajectory Selection for End-to-End Planning (Yao Wenhao, 2025) View paper
  - [19] Generalized Trajectory Scoring for End-to-end Multimodal Planning (Li Zhenxin, 2025) View paper
- Reinforcement Learning for Autonomous Driving
  - Stability-Constrained Reinforcement Learning (1 papers)
  - [6] Stability-Aware Reinforcement Learning for Autonomous Driving With Dynamics-Augmented State and Lyapunov Constraints (Yutao Luo, 2025) View paper
  - Reinforcement Learning with Expert Demonstrations (1 papers)
  - [15] Reinforcement Learning-based Autonomous Parking with Expert Demonstrations (Yao Wu, 2023) View paper
  - Generalizable End-to-End Reinforcement Learning (1 papers)
  - [21] ColorDynamic: Generalizable, Scalable, Real-time, End-to-end Local Planner for Unstructured and Dynamic Environments (Xin, 2025) View paper

Narrative

Core task: end-to-end autonomous driving planning with data augmentation. The field addresses the challenge of training robust driving policies by combining neural architectures that map sensor inputs directly to control outputs with techniques that expand or enrich training data. The taxonomy reveals five main branches. Synthetic Data Generation and Augmentation Methods explores how to create or transform driving scenarios—ranging from diffusion-based scene synthesis (Synthetic Diffusion Driving[2]) to rasterization and primitive-based rendering (RAP Rasterization Planning[0])—to overcome data scarcity and improve generalization. End-to-End Learning Architectures and Training Frameworks focuses on neural network designs and optimization strategies that enable direct sensor-to-action mappings (Modularization Free Driving[1]). Imitation Learning and Behavioral Cloning investigates how to distill expert demonstrations into policies, often addressing distribution shift and covariate mismatch (Robust Behavioral Cloning[28], Imitation Data Augmentation[9]). Trajectory Selection and Multi-Modal Planning deals with generating and scoring multiple candidate trajectories to handle uncertainty and diverse driving behaviors (Multimodal Trajectory Scoring[19], DistillDrive Multimode[10]). Finally, Reinforcement Learning for Autonomous Driving examines reward-driven policy optimization, including safe exploration and sample efficiency (Expert Parking RL[15]).

A particularly active line of work centers on how synthetic augmentation interacts with end-to-end architectures: some studies leverage diffusion models or style transfer (Style Transfer Augmentation[20]) to diversify training scenes, while others use lightweight rasterization to rapidly generate varied traffic configurations. RAP Rasterization Planning[0] sits within the rasterization and primitive-based rendering cluster, emphasizing efficient scene rendering to augment planning data. This contrasts with heavier generative approaches like Synthetic Diffusion Driving[2], which prioritizes photorealistic synthesis at higher computational cost, and with simulation-based methods (Robust Control Simulation[3]) that rely on physics engines. Meanwhile, imitation learning branches grapple with the trade-off between data quality and quantity: augmenting expert demonstrations can mitigate overfitting, but poorly designed augmentations risk introducing spurious correlations. Across these directions, open questions remain about the optimal balance between synthetic diversity and real-world fidelity, and how to integrate augmentation seamlessly into multi-modal planning frameworks that must reason over diverse candidate trajectories.

Related Works in Same Category

No sibling papers were found in the same taxonomy leaf. A taxonomy-subtopic-level comparison will be produced instead.

Taxonomy-Level Summary

The original leaf focuses on lightweight 3D rasterization of annotated primitives (e.g., bounding boxes, lane markers) to generate semantically consistent augmented views without photorealism. Sibling subtopics represent alternative synthetic data generation approaches: diffusion models for photorealistic video synthesis, simulation-based world models for trajectory rendering, and style transfer for texture/appearance modification. The key distinction is that rasterization prioritizes semantic fidelity and computational efficiency over photorealism, while siblings either pursue realism (diffusion, simulation) or appearance diversity (style transfer).

Similarities: - All subtopics aim to augment training data for end-to-end autonomous driving planning - All generate synthetic or modified visual data to improve model generalization - All address data scarcity or distribution shift challenges in driving scenarios

Differences: - Rasterization uses explicit 3D geometric primitives and rendering pipelines, while diffusion models learn implicit generative distributions from data - Rasterization produces semantically faithful but non-photorealistic outputs, whereas diffusion and simulation target photorealistic synthesis - Style transfer modifies existing images without generating new scenes or trajectories, unlike rasterization which can create novel viewpoints from primitives - Rasterization is computationally lightweight compared to diffusion models or complex simulation engines - Simulation-based methods render from learned world models or demonstrations, while rasterization directly renders from annotated scene primitives

Suggested Search Directions: - Hybrid approaches combining rasterization with neural rendering for balancing semantic control and photorealism - Comparative studies on semantic consistency vs. photorealism trade-offs in augmentation effectiveness - Primitive-based rasterization for real-time online augmentation during training

Sibling Subtopics

- **Diffusion-Based Video and Image Generation** (leaves: 1, papers: 4)
- Scope: Papers using diffusion models to synthesize realistic multi-view or long-horizon driving videos for data augmentation.
- Exclude: Papers using rasterization or style transfer without diffusion models belong to other synthetic generation methods.
- **Simulation-Based Data Generation** (leaves: 1, papers: 2)
- Scope: Papers using data-driven simulation or world models to render novel trajectories and viewpoints from real human demonstrations.
- Exclude: Papers using diffusion models or style transfer belong to diffusion-based or style transfer categories.
- **Style Transfer and Image Transformation** (leaves: 1, papers: 1)
- Scope: Papers applying image style transfer to modify texture, contrast, or color for dataset diversity enhancement.
- Exclude: Papers generating entirely new scenes or trajectories belong to diffusion-based or simulation-based generation categories.

Contributions Analysis

Overall novelty summary. The paper proposes a rasterization-based augmentation framework for end-to-end driving, replacing photorealistic rendering with lightweight primitive rasterization to generate counterfactual recovery maneuvers and cross-agent viewpoints. Within the taxonomy, it occupies the 'Rasterization and Primitive-Based Rendering' leaf under 'Synthetic Data Generation and Augmentation Methods'. This leaf currently contains only the original paper itself, with no sibling papers identified. This isolation suggests the rasterization-based approach represents a relatively sparse research direction compared to neighboring leaves like 'Diffusion-Based Video and Image Generation' (four papers) or 'Simulation-Based Data Generation' (two papers).

The taxonomy reveals a crowded landscape of synthetic augmentation methods. Adjacent leaves include diffusion-based synthesis (Synthetic Diffusion Driving, four papers), style transfer techniques (one paper), and simulation-based generation (two papers). The scope notes clarify boundaries: diffusion methods prioritize photorealism via generative models, while simulation-based approaches use physics engines or world models. The original paper explicitly diverges by arguing photorealism is unnecessary—semantic fidelity and scalability matter more. This positions the work as a computational efficiency alternative to heavier generative pipelines, though it shares the broader goal of expanding training data beyond logged expert demonstrations.

Among eighteen candidates examined across three contributions, none were found to clearly refute the proposed methods. The '3D Rasterization Pipeline' examined six candidates with zero refutations; 'Raster-to-Real Feature Alignment' examined ten with zero refutations; 'RAP Framework with Counterfactual Augmentation' examined two with zero refutations. This absence of overlapping prior work within the limited search scope suggests the specific combination of lightweight rasterization, feature-space sim-to-real alignment, and counterfactual augmentation strategies has not been directly addressed in the top-eighteen semantically similar papers. However, the search scale is modest and does not cover the full breadth of autonomous driving or graphics literature.

Based on the limited search scope, the work appears to introduce a distinct technical approach within a broader augmentation landscape. The taxonomy structure shows active research in diffusion-based and simulation-based generation, but the rasterization-based direction remains sparsely populated. The contribution-level statistics indicate no direct prior work among examined candidates, though this reflects top-eighteen semantic matches rather than exhaustive coverage. The novelty assessment is thus conditional on the search boundaries and may shift with deeper exploration of graphics-oriented or real-time rendering communities.

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: 3D Rasterization Pipeline for Driving Scene Reconstruction

Description: The authors introduce a lightweight, training-free rasterization method that converts annotated driving logs (lane polylines, agent cuboids) into perspective camera views. This approach prioritizes semantic and geometric fidelity over photorealism, enabling fast and controllable scene generation for end-to-end planning.

This contribution was assessed against **6 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. 3D Vision-Language Gaussian Splatting

URL: View paper

Brief Assessment

Vision Language Gaussian[37] focuses on semantic segmentation in general 3D scenes using Gaussian splatting with vision-language models, not on driving scene reconstruction from geometric primitives for end-to-end planning.

---

2. LSD-3D: Large-Scale 3D Driving Scene Generation with Geometry Grounding

URL: View paper

Brief Assessment

LSD Large Scale[38] focuses on generating novel 3D driving scenes from scratch using diffusion-based distillation methods, not on reconstructing scenes from existing annotated driving logs. The candidate generates geometry via voxel diffusion and optimizes Gaussian representations, whereas the original paper converts logged annotations into camera views via lightweight rasterization for training augmentation.

---

3. EGSRAL: An Enhanced 3D Gaussian Splatting based Renderer with Automated Labeling for Large-Scale Driving Scene

URL: View paper

Brief Assessment

EGSRAL Automated Labeling[36] focuses on 3D Gaussian Splatting-based rendering for photorealistic novel view synthesis, not lightweight rasterization of geometric primitives. The candidate uses neural rendering techniques requiring training and optimization, whereas the original paper explicitly proposes training-free rasterization as an alternative to costly photorealistic methods.

4. Autosplat: Constrained gaussian splatting for autonomous driving scene reconstruction

URL: View paper

Brief Assessment

AutoSplat Gaussian[35] focuses on 3D Gaussian Splatting for photorealistic scene reconstruction from multi-view images, not on lightweight rasterization from annotated primitives for training augmentation. The candidate reconstructs scenes for simulation/ rendering quality, while the original paper uses rasterization as a fast, training-free augmentation method to generate synthetic training data.

5. Online map vectorization for autonomous driving: A rasterization perspective

URL: View paper

Brief Assessment

Map Vectorization Rasterization[34] focuses on converting vectorized map elements into rasterized HD maps for evaluation and training supervision in map construction tasks, not on reconstructing full driving scenes from geometric primitives for end-to-end planning as in the original paper.

6. Real-time neural rasterization for large scenes

URL: View paper

Brief Assessment

Neural Rasterization Realtime[33] focuses on real-time neural rendering for novel view synthesis using neural textures on meshes, not on converting annotated driving logs (lane polylines, agent cuboids) into perspective views for end-to-end planning training.

Contribution 2: Raster-to-Real (R2R) Feature-Space Alignment Module

Description: The authors propose a feature-space alignment technique that minimizes the domain gap between rasterized and real images at both spatial and global levels, using MSE loss and gradient reversal for domain confusion. This enables effective transfer from synthetic rasterized inputs to real-world deployment.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Industrial Image Anomaly Detection via Synthetic-Anomaly Contrastive Distillation

URL: View paper

Brief Assessment

Synthetic Anomaly Contrastive[40] focuses on industrial defect detection using teacher-student distillation with synthetic anomalies, not autonomous driving or sim-to-real transfer for planning tasks.

2. Improved distribution matching distillation for fast image synthesis

URL: View paper

Brief Assessment

Distribution Matching Distillation[39] focuses on distilling diffusion models for image synthesis using distribution matching between teacher and student models, not on bridging synthetic rasterized inputs to real-world deployment for autonomous driving planning tasks.

3. Brain tumor segmentation using synthetic MR images - A comparison of GANs and diffusion models

URL: View paper

Brief Assessment

Synthetic MR Segmentation[43] focuses on generating synthetic medical images using GANs and diffusion models for brain tumor segmentation training, not on feature-space alignment between synthetic rasterized driving scenes and real camera inputs for autonomous driving planning.

4. Dual teacher knowledge distillation with domain alignment for face anti-spoofing

URL: View paper

Brief Assessment

Dual Teacher Alignment[45] focuses on face anti-spoofing using domain adversarial attacks and knowledge distillation from face recognition/editing models. The original paper addresses autonomous driving with rasterized scene reconstruction and feature alignment between synthetic and real driving images—fundamentally different application domains and technical approaches.

5. Cad: Photorealistic 3d generation via adversarial distillation

URL: View paper

Brief Assessment

CAD Adversarial Distillation[44] focuses on adversarial distillation from 2D diffusion models to 3D GANs for photorealistic generation, not on bridging rasterized synthetic inputs to real-world deployment through feature-space alignment for autonomous driving planning.

6. Knowledge distillation-based domain generalization enabling invariant feature distributions for damage detection of rotating machines and structures

URL: View paper

Brief Assessment

Invariant Feature Distillation[47] focuses on domain generalization for damage detection in rotating machines using knowledge distillation, not on bridging synthetic rasterized and real driving images for autonomous planning.

7. Adversarial diffusion compression for real-world image super-resolution

URL: View paper

Brief Assessment

Adversarial Diffusion Compression[42] focuses on compressing diffusion models for image super-resolution through adversarial distillation in feature space, not on bridging synthetic rasterized inputs to real-world deployment for autonomous driving planning.

8. Align and distill: Unifying and improving domain adaptive object detection

URL: View paper

Brief Assessment

Align Distill Detection[46] focuses on domain adaptation for object detection across different visual domains (e.g., clear to foggy weather, synthetic to real cities), not on bridging rasterized synthetic driving scenes to real-world deployment for end-to-end planning.

9. Cycada: Cycle-consistent adversarial domain adaptation

URL: View paper

Brief Assessment

CyCADA Domain Adaptation[48] focuses on adversarial domain adaptation between image domains (e.g., synthetic to real photos, SVHN to MNIST) using cycle-consistency and semantic losses. The original paper's R2R module specifically bridges rasterized geometric primitives to real driving images, which is a distinct technical approach not addressed in CyCADA.

10. Dataset distillation via the wasserstein metric

URL: View paper

Brief Assessment

Wasserstein Dataset Distillation[41] focuses on dataset distillation via distribution matching in feature space for synthetic dataset generation, not on bridging rasterized and real driving images for autonomous driving deployment.

Contribution 3: RAP Framework with Counterfactual Augmentation Strategies

Description: The authors develop RAP, a complete data augmentation framework that combines 3D rasterization with recovery-oriented trajectory perturbations and cross-agent view synthesis. This framework addresses covariate shift in imitation learning by generating diverse training scenarios beyond logged expert demonstrations.

This contribution was assessed against **2 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Towards better efficiency and generalization in imitation learning: a causal perspective

URL: View paper

Brief Assessment

Causal Imitation Learning[32] is a doctoral thesis focused on causal perspectives in imitation learning efficiency and generalization. The provided context contains only the title page without technical content about data augmentation, counterfactual generation, or view synthesis methods that could challenge RAP's novelty claims.

2. Multi-Agent Imitation Learning: Value is Easy, Regret is Hard

URL: View paper

Brief Assessment

Multi Agent Regret[31] focuses on multi-agent coordination in Markov games with strategic agents and regret minimization, not on data augmentation for autonomous driving imitation learning.

Appendix: Text Similarity Detection

No high-similarity text segments were detected across any compared papers.

References

- [0] RAP: 3D Rasterization Augmented End-to-End Planning View paper
- [1] End-to-end autonomous driving without costly modularization and 3d manual annotation View paper
- [2] Improving end-to-end autonomous driving with synthetic data from latent diffusion models View paper
- [3] Learning robust control policies for end-to-end autonomous driving from data-driven simulation View paper
- [4] Unleashing Generalization of End-to-End Autonomous Driving with Controllable Long Video Generation View paper
- [5] Syndiff-ad: Improving semantic segmentation and end-to-end autonomous driving with synthetic data from latent diffusion models View paper
- [6] Stability-Aware Reinforcement Learning for Autonomous Driving With Dynamics-Augmented State and Lyapunov Constraints View paper
- [7] End-to-End Autonomous Guidance Method Integrated with Mixture-of-Experts for Intelligent Vehicles View paper
- [8] DriveSuprim: Towards Precise Trajectory Selection for End-to-End Planning View paper
- [9] An imitation learning method with data augmentation and post processing for planning in autonomous driving View paper
- [10] DistillDrive: End-to-End Multi-Mode Autonomous Driving Distillation by Isomorphic Hetero-Source Planning Model View paper
- [11] Autonomous Driving Policy Continual Learning With One-Shot Disengagement Case View paper
- [12] World model-based end-to-end scene generation for accident anticipation in autonomous driving View paper
- [13] End-to-End Learning for Autonomous Driving in Secured Smart Cities View paper
- [14] TrafficSim: Learning to Simulate Realistic Multi-Agent Behaviors View paper
- [15] Reinforcement Learning-based Autonomous Parking with Expert Demonstrations View paper
- [16] KARNet: Kalman Filter Augmented Recurrent Neural Network for Learning World Models in Autonomous Driving Tasks View paper
- [17] VR-Drive: Viewpoint-Robust End-to-End Driving with Feed-Forward 3D Gaussian Splatting View paper
- [18] End-to-end models for self-driving cars on UPB campus roads View paper
- [19] Generalized Trajectory Scoring for End-to-end Multimodal Planning View paper
- [20] Data augmentation technology driven by image style transfer in self-driving car based on end-to-end learning View paper
- [21] ColorDynamic: Generalizable, Scalable, Real-time, End-to-end Local Planner for Unstructured and Dynamic Environments View paper
- [22] Prediction of Steering Angle in Autonomous Vehicles Using Deep Learning Approach View paper
- [23] Drive Anywhere: Generalizable End-to-end Autonomous Driving with Multi-modal Foundation Models View paper
- [24] Multi-State End-to-End Learning for Autonomous Vehicle Lateral Control View paper
- [25] Advancing Vision-based End-to-End Autonomous Driving View paper
- [26] Safer End-to-End Autonomous Driving via Conditional Imitation Learning and Command Augmentation View paper
- [27] Classifying driver behaviors for autonomous vehicle navigation View paper
- [28] Robust Behavioral Cloning for Autonomous Vehicles using End-to-End Imitation Learning View paper
- [29] Exploring versatile neural architectures across modalities and perception tasks View paper

- [30] Meaningful Data Augmentation under Unbalanced Data in End-to-end Learning for Autonomous Driving View paper
- [31] Multi-Agent Imitation Learning: Value is Easy, Regret is Hard View paper
- [32] Towards better efficiency and generalization in imitation learning: a causal perspective View paper
- [33] Real-time neural rasterization for large scenes View paper
- [34] Online map vectorization for autonomous driving: A rasterization perspective View paper
- [35] Autosplat: Constrained gaussian splatting for autonomous driving scene reconstruction View paper
- [36] EGSRAL: An Enhanced 3D Gaussian Splatting based Renderer with Automated Labeling for Large-Scale Driving Scene View paper
- [37] 3D Vision-Language Gaussian Splatting View paper
- [38] LSD-3D: Large-Scale 3D Driving Scene Generation with Geometry Grounding View paper
- [39] Improved distribution matching distillation for fast image synthesis View paper
- [40] Industrial Image Anomaly Detection via Synthetic-Anomaly Contrastive Distillation View paper
- [41] Dataset distillation via the wasserstein metric View paper
- [42] Adversarial diffusion compression for real-world image super-resolution View paper
- [43] Brain tumor segmentation using synthetic MR images - A comparison of GANs and diffusion models View paper
- [44] Cad: Photorealistic 3d generation via adversarial distillation View paper
- [45] Dual teacher knowledge distillation with domain alignment for face anti-spoofing View paper
- [46] Align and distill: Unifying and improving domain adaptive object detection View paper
- [47] Knowledge distillation-based domain generalization enabling invariant feature distributions for damage detection of rotating machines and structures View paper
- [48] Cycada: Cycle-consistent adversarial domain adaptation View paper