





3D and 4D World Modeling: A Survey

Lingdong Kong , Wesley Yang , Jianbiao Mei , Youquan Liu , Ao Liang , Dekai Zhu , Dongyue Lu , Wei Yin , Xiaotao Hu , Mingkai Jia , Junyuan Deng , Kaiwen Zhang , Yang Wu , Tianyi Yan , Shenyuan Gao , Song Wang , Linfeng Li , Liang Pan , Yong Liu , Jianke Zhu , Wei Tsang Ooi , Steven C. H. Hoi , Ziwei Liu 

 WorldBench Team

 Equal Contributions  Project Lead  Corresponding Author

World modeling has become a cornerstone in AI research, enabling agents to understand, represent, and predict the dynamic environments they inhabit. While prior work largely emphasizes generative methods for 2D image and video data, they overlook the rapidly growing body of work that leverages **native 3D and 4D representations** such as RGB-D imagery, occupancy grids, and LiDAR point clouds for large-scale scene modeling. At the same time, the absence of a standardized definition and taxonomy for “world models” has led to fragmented and sometimes inconsistent claims in the literature. This survey addresses these gaps by presenting the first comprehensive review explicitly dedicated to 3D and 4D world modeling and generation. We establish precise definitions, introduce a structured taxonomy spanning video-based (**VideoGen**), occupancy-based (**OccGen**), and LiDAR-based (**LiDARGen**) approaches, and systematically summarize datasets and evaluation metrics tailored to 3D/4D settings. We further discuss practical applications, identify open challenges, and highlight promising research directions, aiming to provide a coherent and foundational reference for advancing the field. A systematic summary of existing literature is available on our project page.

 Project Page: <https://worldbench.github.io/survey>

 GitHub Repo: <https://github.com/worldbench/awesome-3d-4d-world-models>

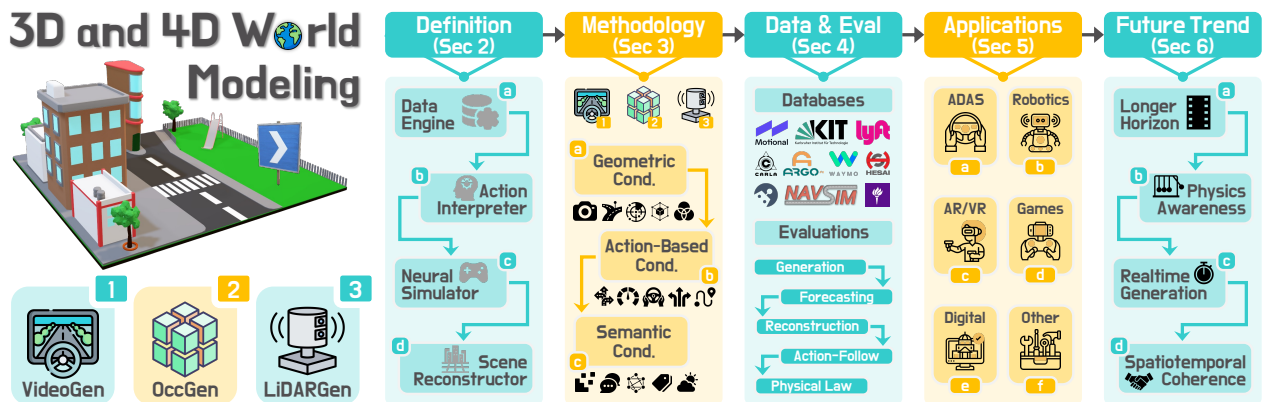


Figure 1 Outline of the survey. This work focuses on **native 3D and 4D representations**: video streams, occupancy grids, and LiDAR point clouds, guided by geometric (C_{geo}), action-based (C_{act}), and semantic (C_{sem}) conditions (Sec. 2). Methods are framed under two paradigms, *generative* (synthesis from observations and conditions) and *predictive* (forecasting from history and actions), and grouped into four functional types (Sec. 3). We cover three modality tracks and standardize evaluations (Sec. 4), practical applications (Sec. 5), and future trends (Sec. 6) across diverse generation, forecasting, and downstream task perspectives.

Why native 3D and 4D matters?

Unlike 2D projections, native 3D/4D signals encode metric geometry, visibility, and motion in coordinates where physics acts [17, 139]. This makes them *first-class carriers* of constraints needed for actionable modeling: multi-view and egocentric consistency, rigid-body and non-rigid kinematics, scene-scale occlusion reasoning, and map/topology adherence. In safety-critical settings, agents must not only produce photorealistic frames but also obey geometry, causality, and controllability; RGB-D, occupancy, and LiDAR provide the inductive bias to satisfy these requirements. Sec. 2 will formalize these representations and the conditioning signals ($\mathcal{C}_{\text{geo}}, \mathcal{C}_{\text{act}}, \mathcal{C}_{\text{sem}}$) we use throughout the survey.

Position in the broader landscape.

The adjacent lines – video/panorama/mesh world models [97] and object-centric 3D asset generators [277] – are complementary: they supply appearance, topology, and assets, while native 3D/4D world models supply geometry-grounded dynamics and interaction [296, 304]. Practical systems increasingly compose these capabilities: mesh/panorama worlds initialized from assets, then *driven* by occupancy- or LiDAR-based dynamics, or video models constrained by 3D priors for view and motion correctness. Our scope centers on the latter – native 3D/4D – while acknowledging and cross-referencing where cross-pollination occurs.

From conditions to functions.

A common pain point in the field is conflating “*what the model consumes*” (conditions) with “*what the model does*” (function). We therefore separate the roles of geometry/action/semantics conditions (Table 1) from functional types. Sec. 3 organizes methods by representation modality – **VideoGen**, **OccGen**, **LiDARGen** – and then by **four** functional roles: ¹*Data Engines* (diverse scene synthesis under $\mathcal{C}_{\text{geo}}, \mathcal{C}_{\text{sem}}, \mathcal{C}_{\text{act}}$), ²*Action Interpreters* (forecasting under \mathcal{C}_{act} with history), ³*Neural Simulators* (closed-loop rollouts with policy-in-the-loop), and ⁴*Scene Reconstructors* (completion/retargeting from partial 3D/4D observations). This decoupling lets us compare heterogeneous methods on common axes of fidelity, consistency, controllability, and scalability.

Contributions. To address the aforementioned gaps, this survey presents the first comprehensive review specifically dedicated to **3D and 4D world modeling and generation**. The primary contributions of this survey are **threefold**:

- We establish precise definitions for “world models” and “3D/4D world modeling”, providing the research community with consistent terminology and conceptual clarity.
- We propose a hierarchical taxonomy of methodologies, categorizing current approaches based on their representation modalities – namely, world modeling based on VideoGen, OccGen, and LiDARGen models.
- We provide extensive coverage of datasets and evaluation protocols specifically tailored for 3D and 4D scenarios, enabling a thorough benchmarking of existing and future world modeling and generation approaches.

Scope. Distinct from previous surveys, which predominantly focus on 2D generative models [144, 237, 322] or broadly define world modeling within limited contexts [55, 243, 258, 302, 303], this survey explicitly targets methodologies that utilize native 3D and 4D representations. This specialized focus includes approaches leveraging RGB-D, volumetric occupancy grids, LiDAR point clouds, and their spatiotemporal forms. By highlighting these modalities, our survey not only fills a critical knowledge gap but also serves as a foundational reference for researchers aiming to develop robust and generalizable 3D/4D generative models.

Organization. The remainder of this survey is organized as follows. Sec. 2 provides preliminaries, detailing fundamental concepts, definitions, and key generative paradigms relevant to world modeling. Sec. 3 introduces a new and hierarchical taxonomy, detailing VideoGen, OccGen, and LiDARGen methodologies, providing comparative analyses and insights into their respective strengths and limitations. Sec. 4 systematically summarizes and categorizes widely used datasets and evaluation metrics critical for world modeling tasks, as well as benchmarking recent methods in this related area. Sec. 5 reviews practical applications of 3D and 4D world models across autonomous driving, robotics, and simulation environments. Sec. 6 discusses major challenges and highlights promising future research directions, paving the way for continued innovation in the field. Finally, Sec. 7 concludes the key discussions drawn in this survey.

2 Preliminaries

In this section, we define critical concepts and establish unified mathematical notations essential for understanding 3D and 4D world modeling. This includes detailed descriptions of the key representations, definitions of generative and predictive world models, and model categorizations.

2.1 3D and 4D Representations

To systematically analyze 3D/4D world models, we first introduce the fundamental scene representations that serve as inputs, outputs, or intermediate states in generation and prediction. These representations differ in how they capture spatial geometry, temporal dynamics, and semantic context.

Video Streams. A video is denoted as $\mathbf{x}_v \in \mathbb{R}^{T \times H \times W \times C}$, where T is the number of frames, and H, W, C are the frame height, width, and channels. Unlike conventional 2D videos, 3D/4D modeling emphasizes *geometric coherence* and *temporal consistency* to ensure physically plausible simulations and accurate forecasting [86, 236, 295].

Occupancy Grids. A static occupancy grid is represented as $\mathbf{x}_o \in \{0, 1\}^{X \times Y \times Z}$, where each voxel indicates whether a location is occupied [165, 186]. Sequential occupancy grids $\mathbf{x}_o^t \in \{0, 1\}^{T \times X \times Y \times Z}$ extend this into 4D, capturing scene evolution over time. Such voxelized geometry enforces spatial constraints, making them well-suited for physics-consistent scene generation.

LiDAR Point Clouds. A LiDAR-acquired scan is expressed as $\mathbf{x}_l = \{(x_i, y_i, z_i)\}_{i=1}^N$, where (x_i, y_i, z_i) are the Cartesian coordinates in 3D space [190]. Sequential LiDAR $\mathbf{x}_l^t = \{(x_i, y_i, z_i, t_i)\}_{i=1}^{N_t}$ further records the timestamp t_i , enabling precise modeling of motion and interactions [113, 266]. Unlike RGB images, LiDAR captures geometry directly and remains robust to texture, lighting, or weather variations [118, 136].

Neural Representations. Implicit scene encodings, such as neural radiance fields (NeRF) and Gaussian splatting (GS), model continuous volumetric fields or explicit Gaussian primitives. NeRF maps a ray origin \mathbf{r} and direction \mathbf{d} to color \mathbf{c} and density σ , while GS represents the scene as a set of Gaussians parameterized by position, covariance, and color. Temporal extensions add dynamic components, enabling realistic 4D reconstructions and simulations.

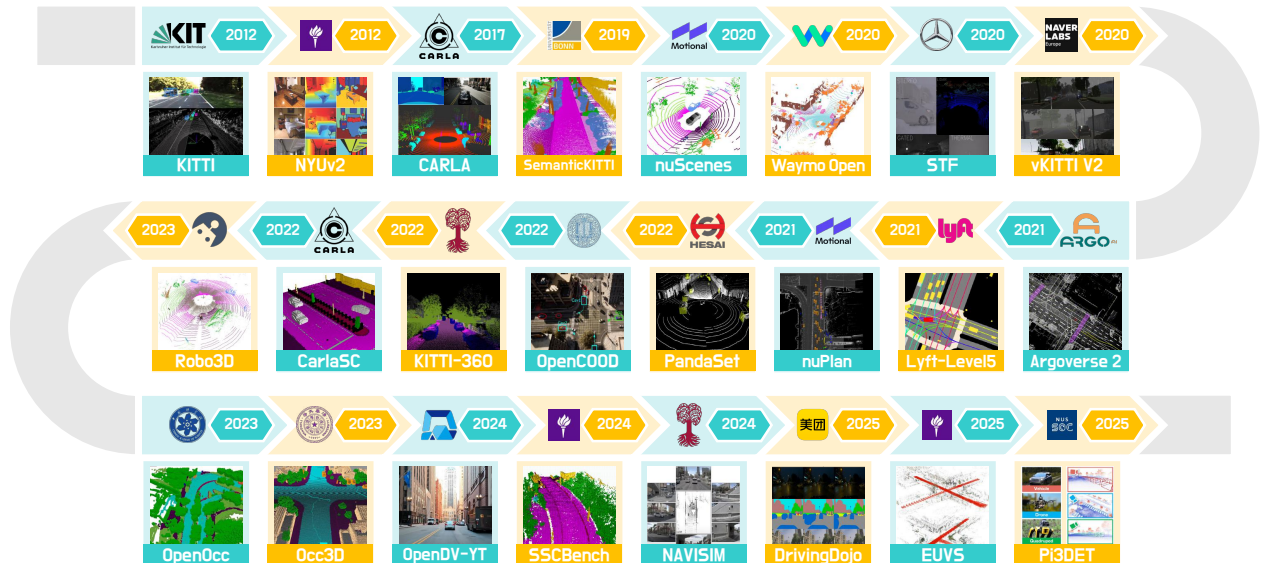


Figure 3 Summary of existing **datasets & benchmarks** used for training and evaluating **VideoGen**, **OccGen**, and **LiDARGen** models. For detailed configurations and statistics, kindly refer to Table 5. Images adopted from the original papers.

Table 1 Summary of the rich collection of conditions used by **VideoGen**, **OccGen**, and **LiDARGen** models. The conditions are categorized into **three** main groups: geometric conditions, action-based conditions, and semantic conditions. The tasks are **1** video generation (Sec. 3.1), **2** occupancy generation (Sec. 3.2), and **3** LiDAR generation (Sec. 3.3).

Group	Condition	Definition	Task
Geometry (\mathcal{C}_{geo})	C Camera Pose 📷	Position and orientation of the camera in world coordinates, controlling viewpoint	1
	D Depth Map 📏	Per-pixel depth values providing scene geometry constraints	1
	B BEV Map 🗺️	Bird’s-eye-view geometric representation of the scene	1 2 3
	H HD Map 🗺️	High-resolution semantic map with detailed road layout and traffic elements	1 3
	3 3D Bounding Box 📦	Object bounding boxes in 3D, defining positions, sizes, and orientations of objects	1 3
	F Flow Field 🌀	Optical or scene flow encoding per-pixel or per-point motion between frames	1 2
	P Past Occupancy 🗺️	Historical occupancy grids or voxel maps capturing prior scene geometry	2 3
	L LiDAR Pattern 🗺️	Sensor scan configuration including beam count, FOV, and resolution	3
	O Object Coordinate 📍	Set of Cartesian coordinates of instances from LiDAR point clouds	1 3
	P Partial Point Cloud 🗺️	Incomplete LiDAR point set capturing only a subset of the full 3D scene geometry	3
Action (\mathcal{C}_{act})	R RGB Frame 📷	Single color image frame from a monocular or multi-camera setup	1 2 3
	S Surface Mesh 🗺️	Triangular mesh or equivalent explicit geometry representation of the scene	1 2 3
	T Ego-Trajectory 📍	The planned or recorded path of the ego vehicle over time	1 2 3
	V Ego-Velocity 🚗	The speed and direction of the ego movement	1 2
	A Ego-Acceleration 📏	Rate of change of ego velocity, describing linear acceleration or deceleration	1
	S Ego-Steering 🚗	The steering angle or input controlling the ego direction	1
	c Ego-Command 📏	The control instructions given to the ego vehicle	1 2 3
Semantics (\mathcal{C}_{sem})	R Route Plan 🗺️	High-level navigation path through the environment, often from a planner	1
	A Action Token 🗺️	Encoded discrete actions or instructions influencing scene evolution	1
	S Scan Path 🗺️	Predefined movement or sweep pattern during LiDAR acquisition	3
	S Semantic Mask 🗺️	Pixel-/occupancy-/point-wise semantic categories	1 2 3
	T Text Prompt 🗺️	Natural language input specifying scene attributes, objects, or actions	1 2 3
	G Scene Graph 🗺️	Graph representation of scene entities and their spatial/semantic relationships	2 3
	O Object Label 🗺️	Class category annotation assigned to an object instance in the scene	1 2 3
	W Weather Tag 🌤️	Discrete label describing environmental conditions such as sunny, rainy, or foggy	1 2
	M Material Tag 🗺️	Classification of surface materials influencing appearance or LiDAR reflectance	1 3

2.2 Definition of World Modeling in 3D and 4D

The above scene representations form the structural backbone of 3D/4D world models. In practice, generating or forecasting them requires additional *conditions* – auxiliary signals that constrain spatial structure, describe agent behavior, or define high-level semantics. As summarized in Table 1, these conditions are typically grouped into:

- **geometric** \mathcal{C}_{geo} : specifying spatial layout such as camera pose, depth maps, or occupancy volumes;
- **action-based** \mathcal{C}_{act} : describing ego-vehicle or agent motion via trajectories, control commands, or navigation goals;
- **semantic** \mathcal{C}_{sem} : providing abstract scene intent such as textual prompts, scene graphs, or environment attributes.

These signals can be used independently or in combination, shaping the realism, controllability, and diversity of the generated or forecasted scenes in 3D and 4D.

2.2.1 Model Definitions

Depending on the modeling objective, 3D/4D world models generally fall into two complementary paradigms:

Generative World Models focus on synthesizing plausible scenes from scratch or from partial observations, guided by multimodal conditions. This process can be formulated as:

$$\mathcal{G}(\mathbf{x}_i, \mathcal{C}_{\text{geo}}, \mathcal{C}_{\text{act}}, \mathcal{C}_{\text{sem}}) \rightarrow \mathcal{S}_g, \quad (1)$$

where \mathbf{x}_i denotes the optional input representation, with $i \in \{\emptyset, v, o, l\}$, *e.g.*, noise, partial video, occupancy, or LiDAR data. \mathcal{C}_{geo} , \mathcal{C}_{act} , and \mathcal{C}_{sem} correspond to the geometric, action, and semantic conditions. The output \mathcal{S}_g is a generated 3D/4D scene, such as a video sequence, occupancy grid, or LiDAR sweep sequence.

Predictive World Models instead aim to forecast the future evolution of the scene based on historical observations, often under action conditions that describe planned or executed agent behavior. This process is formulated as:

$$\mathcal{P}(\mathbf{x}_i^{-t:0}, \mathcal{C}_{\text{act}}) \rightarrow \mathcal{S}_p^{1:k}, \quad (2)$$

where $\mathbf{x}_i^{-t:0}$ represents observations from the past t steps to the current step, and \mathcal{C}_{act} encodes agent actions (e.g., control commands or planned trajectories). The model outputs $\mathcal{S}_p^{1:k}$, the predicted scene representations over k future steps.

Together, these two paradigms capture the dual capability of world models: the ability to *imagine* diverse and controllable worlds (generative), and to *anticipate* their plausible future evolution under specific conditions (predictive).

2.2.2 Model Categorizations

Building on the generative and predictive paradigms, existing approaches can be further divided into **four functional types**. They differ in how they utilize historical observations, the nature of conditioning signals ($\mathcal{C}_{\text{geo}}, \mathcal{C}_{\text{act}}, \mathcal{C}_{\text{sem}}$), and whether they operate in an open-loop or closed-loop setting.

Type 1: Data Engines

Generate diverse 3D/4D scenes from geometric and semantic cues, optionally with action conditions.

- **Inputs:** \mathcal{C}_{geo} (geometric cond.), \mathcal{C}_{act} (action cond., optional), and \mathcal{C}_{sem} (semantic cond.)
- **Output:** \mathcal{S}_g (generated scene)

Focus on *plausibility* and *diversity* for large-scale data augmentation and scenario creation.

Type 2: Action Interpreters

Forecast future 3D/4D world states from historical observations under given action conditions.

- **Inputs:** $\mathbf{x}_i^{-t:0}$ (historical observations) and \mathcal{C}_{act} (action cond.)
- **Output:** $\mathcal{S}_p^{1:k}$ (predicted sequence)

Enable *action-aware forecasting* for trajectory planning, behavior prediction, and policy evaluation.

Type 3: Neural Simulators

Iteratively simulate closed-loop agent-environment interactions by generating successive scene states.

- **Inputs:** \mathcal{S}_g^t (current scene state) and π_{agent} (agent policy)
- **Output:** \mathcal{S}_g^{t+1} (next scene state)

Support *interactive simulation* for autonomous driving, robotics, and immersive XR training.

Type 4: Scene Reconstructors

Recover complete and coherent 3D/4D scenes from partial, sparse, or corrupted observations.

- **Inputs:** \mathbf{x}_i^p (partial observations) and \mathcal{C}_{geo} (optional geometric cond.)
- **Output:** $\hat{\mathcal{S}}_g$ (completed scene)

Facilitate *interactive tasks* on high-fidelity mapping, digital twin restoration, and post-event analysis.

Together, these four categories outline the functional landscape of 3D/4D world modeling. While all aim to produce physically and semantically coherent scenes, they differ in how they leverage past observations, conditioning signals, and interaction loops – serving applications ranging from large-scale data synthesis and policy evaluation to interactive simulation and scene restoration.

2.3 Generative Models

Generative models form the algorithmic core of 3D/4D world modeling, enabling agents to *learn*, *imagine*, and *forecast* sensory data under diverse conditions. They provide the mechanisms to synthesize realistic and physically plausible scenes, with different paradigms offering distinct trade-offs in quality, controllability, and efficiency. Representative families include variational autoencoders, generative adversarial networks, diffusion models, and autoregressive models.

Variational Autoencoders (VAEs) [115] learn a structured latent space via probabilistic encoding and decoding. Given input \mathbf{x} , the encoder defines a variational posterior $q_\phi(\mathbf{z}|\mathbf{x}) = \mathcal{N}(\mu_\phi(\mathbf{x}), \text{diag}(\sigma_\phi^2(\mathbf{x})))$ and samples \mathbf{z} using the reparameterization trick: $\mathbf{z} = \mu_\phi(\mathbf{x}) + \sigma_\phi(\mathbf{x}) \odot \epsilon$, where $\epsilon \sim \mathcal{N}(0, I)$. The decoder $p_\theta(\mathbf{x}|\mathbf{z})$ reconstructs the input, and the model is trained to maximize the variational lower bound that balances reconstruction fidelity and latent regularization:

$$\log p_\theta(\mathbf{x}) \geq \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})}[\log p_\theta(\mathbf{x}|\mathbf{z})] - D_{\text{KL}}(q_\phi(\mathbf{z}|\mathbf{x}) \parallel p(\mathbf{z})). \quad (3)$$

VAEs offer stable training and interpretable latent spaces, but may produce blurrier samples compared to other paradigms.

Generative Adversarial Networks (GANs) [70] generate data via a min-max game between a generator G_θ and discriminator D_ϕ . The generator maps latent variables $\mathbf{z} \sim p(\mathbf{z})$ to the data space, aiming to fool D_ϕ , while the discriminator distinguishes real from synthetic samples:

$$\min_G \max_D \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}}[\log D(\mathbf{x})] + \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z})}[\log(1 - D(G(\mathbf{z})))] \quad (4)$$

GANs can produce high-fidelity result samples but often suffer from training instability and mode collapse issues.

Diffusion Models (DMs) [82, 207] learn to reverse a gradual noising process. The forward process corrupts \mathbf{x}_0 into $\{\mathbf{x}_1, \dots, \mathbf{x}_T\}$ via $q(\mathbf{x}_t|\mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t}\mathbf{x}_{t-1}, \beta_t\mathbf{I})$, where β_t follows a variance schedule. The reverse process $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)$ is trained to denoise, minimizing:

$$\mathbb{E}_{\mathbf{x}, \epsilon, t}[\|\epsilon - \epsilon_\theta(\mathbf{x}_t, t)\|^2]. \quad (5)$$

DMs provide strong stability and sample quality, though inference can be slow due to iterative sampling.

Autoregressive Models (ARs) [215, 222] factorize the joint distribution as $p(\mathbf{x}) = \prod_{i=1}^n p(x_i | x_{<i})$, predicting each element conditioned on all previous ones. Transformer-based ARs offer exact likelihood estimation and flexible sequence modeling, but suffer from slow generation since samples are produced sequentially. Recent advances have adapted ARs to spatial and temporal tokens, making them well-suited for structured 3D scene generation and forecasting.

Summary. These paradigms form the algorithmic backbone for world models. Their differences in structure, training stability, and inference efficiency directly shape how 3D environments can be synthesized, forecasted, and controlled. As we move into native 3D/4D domains, these trade-offs are magnified, since scalability, controllability, and multi-modal integration are critical to constructing reliable world models for embodied AI and simulation.

3 Methods: A Hierarchical Taxonomy

In this section, we standardize and categorize existing 3D and 4D world modeling approaches based on their representation modalities. This includes descriptions and discussions of world modeling based on **Video Generation** (Sec. 3.1), **Occupancy Generation** (Sec. 3.2), and **LiDAR Generation** (Sec. 3.3) models, respectively.

3.1 World Modeling from Video Generation

Video-based generation has emerged as a new paradigm, offering visual cues and temporal dynamics to model complex real-world scenarios. By generating multi-view or egocentric video sequences, these models can

synthesize training data, predict future outcomes, and create interactive simulation environments. Based on their primary function, existing methods can be grouped into **three** categories: **1 Data Engines**, **2 Action Interpreters**, and **3 Neural Simulators**. Table 2 summarizes existing models under these domains.

3.1.1 Data Engines

Generative 3D data engines focus on generating diverse and controllable driving scenes to support perception, planning, and simulation [47, 63, 64, 104, 125, 131, 211, 275]. Research in this direction covers three major applications.

Perception Data Augmentation. Generative scene synthesis alleviates real-world data scarcity and addresses long-tail perception challenges. Early work focused on BEV-guided realistic street scenes. BEVGen [211] uses an autoregressive transformer and cross-view transformation to produce spatially consistent surrounding images aligned with a given BEV layout. BEVControl [275] centers on diffusion models to boost the quality of synthetic data, particularly for augmenting challenging long-tail scenarios. Subsequently, MagicDrive [64] made significant progress in driving scene generation and data augmentation, combining 3D geometry and semantic descriptions, and camera poses to generate high-fidelity images. Later work introduced finer conditioning. For instance, SyntheOcc [131] uses 3D semantic multi-plane images for comprehensive, spatially aligned conditioning, and PerLDiff [293] proposes perspective-layout diffusion models that fully leverage perspective 3D geometry to enhance realism and consistency. On the other hand, approaches such as Panacea [244], DrivingDiffusion [135], and SubjectDrive [93] introduce 4D attention, keyframes, and subject control to improve the temporal consistency and data diversity of 3D controllable multi-view videos. NoiseController [47] proposes multi-level noise decomposition and multi-frame collaborative denoising to enhance spatiotemporal coherence.

For long-horizon video generation, DiVE [104], MagicDrive-V2 [63], and Cosmos-Drive [194] leverage the flexibility and scalability of DiT to produce longer videos. Glad [256] uses latent-variable propagation, and STAGE [228] uses hierarchical temporal feature transfer to generate long videos in a streaming fashion. Others like UniScene [125] and BEVWorld [299] explore multi-modal data synthesis to broaden applications, supporting downstream perception tasks that leverage information from multiple modalities. These advances enable robust, scalable autonomous driving perception systems by delivering diverse, controllable, and long-horizon training data that capture real-world variability.

Planning-Oriented Data Mining. Beyond perception, data engines also mine rare and safety-critical scenarios for planning. Delphi [157] employed a diffusion-based long video generation framework and a failure-case-driven approach utilizing pre-trained visual language models to synthesize data similar to failure scenarios, thereby enhancing sample efficiency and planning performance for end-to-end autonomous driving systems. DriveDreamer-2 [305] converted user queries into agent trajectories via a large language model, which are then used to produce traffic-compliant HDMaps for corner case generation. Nexus [320] simulated both regular and challenging scenarios from fine-grained tokens with independent noise states to improve reactivity and goal conditioning and collected a specialized corner-case dataset to complement challenging scenario generation. Challenger [267] exploited a physics-aware multi-round trajectory refinement to identify adversarial maneuvers

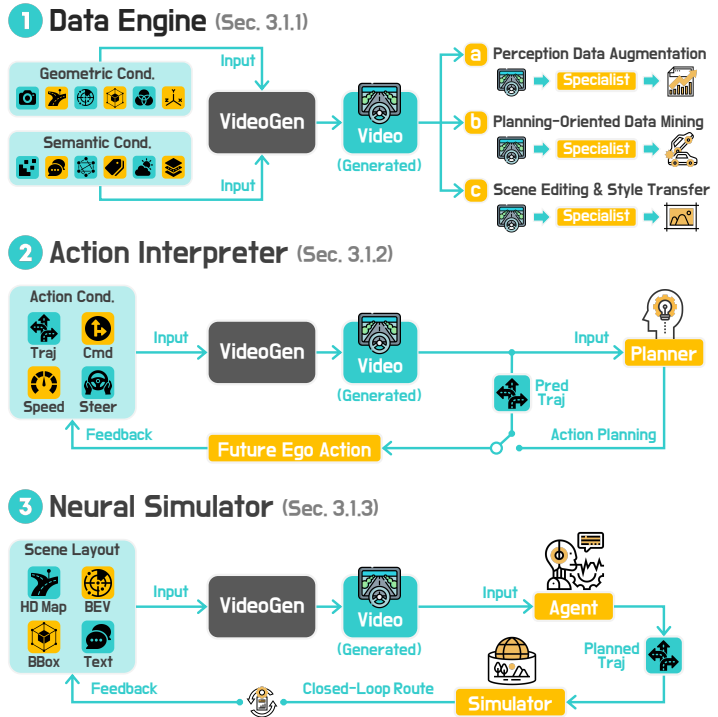


Figure 4 The categorization of **VideoGen** models based on functionalities, including data engines (Sec. 3.1.1), action interpreters (Sec. 3.1.2), and neural simulators (Sec. 3.1.3).

and a tailored scoring function to promote realistic yet challenging behaviors compatible with downstream video synthesis.

Scene Editing & Style Transfer. Many existing methods [155, 244, 319] also take world models for scene editing and style transfer to enrich the toolkit for autonomous driving simulation and data augmentation. Early methods primarily utilized scene descriptions [64] or reference images [34] for basic appearance modifications (*e.g.*, weather, lighting) and relied on bounding boxes or HD maps [275] for element-level adjustments. However, newer approaches explore richer representations for precise scene manipulation and diverse appearance control. WoVoGen [155] ensures cross-sensor consistency through world volume-aware synthesis, while SyntheOcc [131] employs occupancy grids for occlusion-aware scene editing. SimGen [319] bridges sim-to-real gaps via simulator-conditioned cascade diffusion, and DrivePhysica [286] simulates complex driving scenarios (*e.g.*, cut-ins) using CARLA and introduces motion representation learning and instance flow guidance for temporal consistency. Complementing these, GeoDrive [31] integrates explicit 3D geometry conditions and dynamic editing to enable interactive trajectory and object manipulation.

3.1.2 Action Interpreters

Action-driven generation models bridge agent intentions and environmental dynamics through action-guided world generation and forecast-driven action planning, enabling outcome anticipation and unifying low-level maneuvers and reasoning by mapping controls to plausible futures.




Action-Guided Video Generation. Action-conditioned generation models empower agents to predict future outcomes based on intended maneuvers, effectively bridging low-level control inputs with high-fidelity video rollouts of plausible futures. GAIA-1 [86] pioneered a generative model that fuses video, text, and action inputs to synthesize realistic driving scenarios with detailed control over ego-vehicle behavior and scene attributes. GAIA-2 [198] expanded this framework to include agent configurations, environmental factors, and road semantics. GenAD [274] further enhanced generalization by releasing the OpenDV dataset alongside a predictive model that supports zero-shot, language- and action-conditioned predictions. Vista [65] applies robust action conditioning across diverse scenarios, while GEM [80] delivers multimodal outputs with precise ego-motion control, and MaskGWM [178] boosts fidelity and long-horizon predictions using mask-based diffusion.



































































































































































To address error accumulation in long video synthesis, InfinityDrive [77] and Epona [295] proposed memory injection and a chain-of-forward training strategy, respectively. In addition, DrivingWorld [91] generates scenarios from predefined trajectories, functioning as a neural driving simulator. Other approaches, such as DriVerse [134], MiLA [226], PosePilot [107], and LongDWM [234], focus on trajectory alignment, temporal stability, pose controllability, and depth-free guidance. Collectively, these advances drive action-conditioned generation toward better precision, temporal coherence, and robustness.

Forecasting-Driven Action Planning. Another line of work forecasts future states from current observations and ego actions, letting planners evaluate outcomes before committing [132, 227, 311]. Different from purely reactive schemes, these approaches emphasize *anticipatory decision-making*, allowing the agent to virtually “test” multiple futures and avoid unsafe trial-and-error in the real world. Drive-WM [239] generates video rollouts of candidate maneuvers, scoring them with image-based rewards for trajectory selection. DriveDreamer [236] proposed the ActionFormer to predict future states and ego-environment interactions. ADriver-I [102] combines multimodal LLMs with autoregressive control signals and world evolution prediction. Vista [65] incorporates uncertainty-aware reward modules for robust action evaluation.

GPT-style designs such as DrivingGPT [35] and DrivingWorld [91] model visual and action tokens jointly for planning via next-token prediction. Integrated frameworks like Doe-1 [309] unify perception, prediction, and planning for closed-loop autonomous driving, while VaVAM [12] bridges video diffusion and an action expert for decision-making. ProphetDWM [233] further couples latent action learning with state forecasting for long-term planning. Overall, by simulating diverse futures and leveraging feedback, forecast-driven models enhance generalization and safety in end-to-end autonomous driving.

Table 2 Summary of video-based generation (**VideoGen**) models.

- **Datasets:** **N** nuScenes [24], **K** KITTI [69], **W** Waymo Open [210], **Y** OpenDV-YouTube [274], **A** Argoverse 2 [247], **N** nuPlan [25], **N** NAVSIM [41], **C** CARLA [48], and **P** Private (Internal) Data.
- **Input & Output:**  Noise Latent,  Video (Single-View and/or Multi-View), and  Ego-Action.
- **Architectures (Arch.):** **AR**: Autoregressive Models, **MLLM**: Multimodal Large Language Models, **SD**: Stable Diffusion Models, **DiT**: Diffusion Transformer, **GPT**: Generative Pre-trained Transformer.
- **Tasks:** **VG**: Video Generation, **E2E**: End-to-End Planning, and **3SR**: 3D Scene Reconstruction.
- **Categories:** **1** Data Engine (Sec. 3.1.1), **2** Action Interpreter (Sec. 3.1.2), and **3** Neural Simulator (Sec. 3.1.3).

#	Model	Venue	Dataset	Input	Output	Condition	Len.	Freq.	Arch.	Task	Cat.	URL
1	GAIA-1 [86]	arXiv'23	P			T C T	L	25Hz	AR	VG	2	
2	ADriver-I [102]	arXiv'23	N P			T	S	2Hz	MLLM	VG	2	
3	BEVControl [275]	arXiv'23	N			3 H	S	-	SD	VG	1	
4	BEVGen [211]	RA-L'24	N A			B	S	-	AR	VG	1	
5	MagicDrive [64]	ICLR'24	N			3 B C T	S	12Hz	SD	VG	1	
6	Panacea [244]	CVPR'24	N			3 H C T	S	2Hz	SD	VG	1	
7	Drive-WM [239]	CVPR'24	N			3 H T T	S	2Hz	SD	VG, E2E	2	
8	GenAD [274]	CVPR'24	Y			T T	S	2Hz	SD	VG, E2E	2	
9	DriveDreamer [236]	ECCV'24	N			3 H T T	S	2Hz	SD	VG, E2E	2	
10	DrivingDiffusion [135]	ECCV'24	N			3 T	S	2Hz	SD	VG	1	
11	WoVoGen [155]	ECCV'24	N			B T C T P	S	2Hz	SD	VG	1	
12	Vista [65]	NeurIPS'24	Y N			T	L	10Hz	SD	VG, E2E	2	
13	SimGen [319]	NeurIPS'24	N			3 B	S	2Hz	SD	VG	1	
14	MagicDrive3D [62]	arXiv'24	N			3 B C T	S	12Hz	SD	VG, 3SR	1 3	
15	Delphi [157]	arXiv'24	N			3 H T	L	2Hz	SD	VG	1	
16	BEVWorld [299]	arXiv'24	N C			T	S	12Hz	AR	VG	1	
17	Panacea+ [245]	arXiv'24	N A			3 H C T	S	2Hz	SD	VG	1	
18	DiVE [104]	arXiv'24	N			3 H T	L	12Hz	DiT	VG	1	
19	DreamForge [163]	arXiv'24	N			3 H C T	L	12Hz	SD, DiT	VG	3	
20	SyntheOcc [131]	arXiv'24	N			P	S	2Hz	SD	VG	1	
21	HoloDrive [253]	arXiv'24	N			3	S	-	SD	VG	1	
22	InfinityDrive [77]	arXiv'24	Y N			T T	L	10Hz	AR	VG	2	
23	CogDriving [153]	arXiv'24	N			3 B	S	2Hz	DiT	VG	1	
24	UniMLVG [34]	arXiv'24	Y N W A			3 B C T	L	12Hz	DiT	VG	1	
25	DrivePhysica [286]	arXiv'24	N			3 H C	L	12Hz	DiT	VG	1	
26	Doe-1 [309]	arXiv'24	N			P	S	2Hz	MLLM	VG, E2E	2 3	
27	OccScene [126]	arXiv'24	N K			T T	-	2Hz	SD	VG	1	
28	DrivingGPT [35]	arXiv'24	N N			T	L	10Hz	GPT	VG, E2E	2	
29	DrivingWorld [91]	arXiv'24	P N			T	L	10Hz	GPT	VG, E2E	2	
30	DriveDreamer-2 [305]	AAAI'25	N			3 H T	S	12Hz	SD	VG	1	
31	SubjectDrive [93]	AAAI'25	N			3 H T	S	2Hz	SD	VG	1	
32	Glad [256]	ICLR'25	N			3 H T	S	2Hz	SD	VG	1	
33	DualDiff [129]	ICRA'25	N			3 H C T P	S	-	SD	VG	1	
34	DriveScape [249]	CVPR'25	N			3 H T C T	S	10Hz	SD	VG	1	
35	DriveDreamer4D [304]	CVPR'25	W			3 H	S	-	SD	VG, 3SR	1 2	
36	DrivingSphere [268]	CVPR'25	N			P	L	12Hz	DiT	VG	3	
37	UniScene [125]	CVPR'25	N			T P	L	12Hz	SD	VG	1	
38	GEM [80]	CVPR'25	Y N			T	L	10Hz	SD	VG	2	
39	MaskGWM [178]	CVPR'25	Y N W			T T	L	10Hz	DiT	VG	2	
40	UMGen [252]	CVPR'25	W N			B T	L	2Hz	AR	VG, E2E	3	
41	PerLDiff [293]	ICCV'25	N K			3 H	S	-	SD	VG	1	
42	DriveArena [278]	ICCV'25	N			3 H C T	L	12Hz	SD, DiT	VG	1 3	
43	MagicDrive-V2 [63]	ICCV'25	N W			3 B T C T	L	12Hz	DiT	VG	1	
44	InfiniCube [156]	ICCV'25	W			H P	L	10Hz	SD	VG, 3SR	3	
45	DiST-4D [76]	ICCV'25	N			3 B T C	L	12Hz	DiT	VG, 3SR	3	
46	Epona [295]	ICCV'25	N N			T	L	5Hz	DiT	VG, E2E	2	
47	VaViM [12]	arXiv'25	Y			N/A	L	2Hz	MLLM	VG	2	
48	VaVAM [12]	arXiv'25	Y N N			T	L	2Hz	MLLM	VG, E2E	2	
49	DualDiff+ [285]	arXiv'25	N			3 H C T P	S	12Hz	SD	VG	1	
50	UniFuture [140]	arXiv'25	N			N/A	S	12Hz	SD	VG, 3SR	3	
51	MiLA [226]	arXiv'25	N			T C T	L	12Hz	DiT	VG	2	
52	GAIA-2 [198]	arXiv'25	P			T C T	L	30Hz	DiT	VG	2	
53	CoGen [101]	arXiv'25	N			P	L	12Hz	DiT	VG	1	
54	Nexus [320]	arXiv'25	N W P			B	S	2Hz	DiT	VG	1 3	
55	NoiseController [47]	arXiv'25	N			3 B C T	S	12Hz	SD	VG	1	
56	DriVerse [134]	arXiv'25	N W			T	L	12Hz	DiT	VG	2	
57	PosePilot [107]	arXiv'25	N			C	L	2Hz	SD, DiT, AR	VG	2	
58	GeoDrive [31]	arXiv'25	N			C	L	12Hz	DiT	VG, 3SR	1 3	
59	Challenger [267]	arXiv'25	N			3 B T	L	12Hz	DiT	VG	1	
60	ProphetDWM [233]	arXiv'25	N			T	L	2Hz	SD	VG, E2E	2	
61	LongDWM [234]	arXiv'25	N			T	L	10Hz	DiT	VG	2	
62	Cosmos-Drive [194]	arXiv'25	P			3 H T	L	-	DiT	VG	1	
63	STAGE [228]	arXiv'25	N			3 H T	L	12Hz	SD	VG	1	

3.1.3 Neural Simulators

Closed-loop simulators produce realistic virtual worlds that support effective planning, decision-making, and interaction. Regarding the difference in scene modeling, recent methods can be broadly categorized into two main approaches.

Generation-Driven Simulation. Recent advances in generative simulators for autonomous driving leverage conditional generative frameworks [20, 184, 197] to create interactive high-fidelity environments. DriveArena [278] establishes the first closed-loop framework through two core components: TrafficManager for scalable traffic synthesis and WorldDreamer for autoregressive scene generation. Building on this foundation, DreamForge [163] enhances long-term scenario modeling by integrating object-wise position encoding, supported by a novel temporal attention mechanism. Further extending these capabilities, DrivingSphere [268] introduces 4D semantic occupancy modeling that unifies static environments and dynamic objects, coupled with a visual synthesis module ensuring spatiotemporal consistency in multiview video generation. UMGGen [252] simulates behavioral interactions between ego-vehicles and user-defined agents, while Nexus [320] dynamically updates environments based on agent decisions, rigorously validated through nuPlan closed-loop benchmarks. GeoDrive [31] advances trajectory optimization for VLA systems via geometry-aware scene modeling and precision control modules. Collectively, these developments transition generative simulation from passive environment rendering to closed-loop systems capable of agent interaction and feedback-driven adaptation.

Reconstruction-Centric Simulation. Reconstruction-based simulators employ neural scene reconstruction techniques such as NeRF [166] and 3D GS [111] to convert driving logs into interactive neural environments [36, 53, 57, 67, 95, 105, 126, 128, 161, 172, 185, 195, 290, 298, 314, 323]. StreetGaussian [271] represented dynamic urban street as a set of point clouds equipped with semantic logits and 3D Gaussians, each associated with either a foreground vehicle or the background. Other key implementations include HUGSIM [313], which integrates physical constraints with 3D GS for aggressive behavior synthesis, alongside frameworks like UniSim [125] and Uni-Gaussians [291] that generate synchronized multi-modal sensor outputs through Gaussian primitive distillation. OmniRe [38] further enhances dynamic entity modeling via neural scene graph representations. While conventional 3D GS methods [62, 140, 154, 156, 271] struggle with viewpoint extrapolation artifacts, emerging solutions integrate 3D scene generation models as the data foundation to improve reconstruction robustness. ReconDreamer [177] applies progressive refinement to eliminate ghosting effects in dynamic scenes, while Stage-1 [230] achieves controllable 4D synthesis through multiview point cloud completion. These modeling methods enhanced approaches [76, 177, 272, 304, 306] demonstrate significant improvements in handling novel viewpoints, effectively bridging the fidelity gap between simulated and real-world environments.

3.2 World Modeling from Occupancy Generation

Generation models based on occupancy grids tailored to offer a geometry-centric representation that encodes both semantic and structural details of the 3D world. By generating, forecasting, or simulating occupancy in 3D/4D space, these models provide a geometry-consistent scaffold for perception, enable action-contingent future prediction, and support realistic large-scale simulation. Based on their primary function, existing methods can be grouped into **three** categories: ❶ **Scene Representors**, ❷ **Occupancy Forecasters**, and ❸ **Autoregressive Simulators**. Table 3 summarizes existing models under these domains.

3.2.1 Scene Representors

Occupancy-based 3D and 4D generation models, designed for learning structured 3D scene representations, treat the occupancy grid as a geometry-consistent intermediate for downstream tasks. Such a paradigm enhances perception robustness and provides structural guidance for 3D scene generation across two main applications.

3D Perception Robustness Enhancement. Occupancy-based representations have emerged as a powerful intermediate modality for enhancing perception robustness through generative modeling techniques. SSD [123] pioneered this direction by employing discrete [7] and latent diffusion [197] models for scene-level 3D categorical data generation, learning to map sparse occupancy inputs into dense semantic reconstructions. SemCity [124]

Table 3 Summary of occupancy-based generation (**OccGen**) models.

- **Datasets:** **S** SemanticKITTI [14], **C** CarlaSC [248], **N** Occ3D-nuScenes [216], **W** Waymo Open [210], **L** Lyft-Level5 [84], **A** Argoverse 2 [247], **3** KITTI-360 [141], **U** NYUv2 [206], and **O** OpenCOOD [264].
- **Input & Output:** \mathcal{A} Noise Latent, \mathcal{B} Latent Codebook, \mathcal{I} Images, \mathcal{C} 3D Occ, \mathcal{D} 4D Occ, and \mathcal{E} Ego-Action.
- **Architectures (Arch.):** *Enc-Dec*: Encoder-Decoder, *LDM*: Latent Diffusion Model, *MSSM*: Memory State-Space Model, *AR*: Autoregressive Model, *DiT*: Diffusion Transformer, *LLM*: Large Language Model.
- **Tasks:** *O3G*: 3D Occupancy Generation, *O4G*: 4D Occupancy Generation, *OF*: 4D Occupancy Forecasting, *PT*: Pre-Training, *SSC*: Semantic Scene Completion, and *E2E*: End-to-End Planning.
- **Categories:** **1** Scene Representer (Sec. 3.2.1), **2** Occ Forecaster (Sec. 3.2.2), and **3** AR Simulator (Sec. 3.2.3).

#	Model	Venue	Dataset	Input	Output	Condition	Len.	Arch.	Task	Cat.	URL
1	Emergent-Occ [112]	ECCV'22	N	\mathcal{C}	\mathcal{C} \mathcal{E}	N/A	7	Enc-Dec	OF, E2E	2	🔗
2	FF4D [113]	CVPR'23	S N A	\mathcal{C}	\mathcal{C}	N/A	5	Enc-Dec	OF	2	🔗
3	SSD [123]	arXiv'23	C	\mathcal{B}	\mathcal{C}	N/A	1	LDM	O3G	1	🔗
4	UniWorld [170]	arXiv'23	N	\mathcal{I}	\mathcal{C}	N/A	-	Enc-Dec	PT	2	🔗
5	UniScene [168]	RA-L'24	N	\mathcal{I}	\mathcal{C}	N/A	-	Enc-Dec	PT	2	🔗
6	Cam4DOcc [158]	CVPR'24	N L	\mathcal{I}	\mathcal{C}	T	4	Enc-Dec	OF	2	🔗
7	XCube [193]	CVPR'24	W	\mathcal{A}	\mathcal{C}	N/A	1	LDM	O3G	3	🔗
8	SemCity [124]	CVPR'24	S C	\mathcal{A}	\mathcal{C}	N/A	1	LDM	O3G, SSC	1 3	🔗
9	DriveWorld [169]	CVPR'24	N	\mathcal{I}	\mathcal{C}	V S	4	MSSM	OF, PT	2	🔗
10	UnO [2]	CVPR'24	S N A	\mathcal{C}	\mathcal{C}	N/A	6	Enc-Dec	OF	2	🔗
11	PDD [149]	ECCV'24	S C	\mathcal{B}	\mathcal{C}	N/A	1	LDM	O3G	3	🔗
12	OccWorld [308]	ECCV'24	N	\mathcal{C}	\mathcal{C} \mathcal{E}	T	6	AR	OF, E2E	2	🔗
13	WoVoGen [155]	ECCV'24	N	\mathcal{C}	\mathcal{C}	B	3	LDM	O4G	1 2	🔗
14	UrbanDiff [294]	arXiv'24	N	\mathcal{B}	\mathcal{C}	B	1	LDM	O3G	1	🔗
15	OccSora [229]	arXiv'24	C N W	\mathcal{B}	\mathcal{C}	T	32	DiT	O4G	2 3	🔗
16	LOPR [122]	arXiv'24	N W	\mathcal{I} \mathcal{I}	\mathcal{C}	T	15	Enc-Dec	OF	2	🔗
17	OccLLaMA [242]	arXiv'24	N	\mathcal{B}	\mathcal{C}	T	6	LLM	O4G	2	🔗
18	FSF-Net [75]	arXiv'24	N	\mathcal{C}	\mathcal{C} \mathcal{E}	N/A	4	Enc-Dec	OF, E2E	2	🔗
19	DOME [72]	arXiv'24	N	\mathcal{A}	\mathcal{C}	T	11	DiT	OF	2	🔗
20	GaussianAD [310]	arXiv'24	N	\mathcal{C}	\mathcal{C} \mathcal{E}	N/A	6	Enc-Dec	OF, E2E	2	🔗
21	OccScene [126]	arXiv'24	N S U	\mathcal{I}	\mathcal{C}	N/A	1	SD	O3G	1	🔗
22	DFIT-OccWorld [292]	arXiv'24	N	\mathcal{I} \mathcal{I}	\mathcal{C} \mathcal{E}	T	6	Enc-Dec	OF, E2E	2	🔗
23	Drive-OccWorld [281]	AAAI'25	N L	\mathcal{I}	\mathcal{C} \mathcal{E}	C T V S	4	AR	OF, E2E	2	🔗
24	DynamicCity [17]	ICLR'25	C N W	\mathcal{A}	\mathcal{C}	S C T	16	DiT	O4G	3	🔗
25	PreWorld [133]	ICLR'25	N	\mathcal{I}	\mathcal{C} \mathcal{E}	T	6	Enc-Dec	OF, E2E	2	🔗
26	OccProphet [33]	ICLR'25	N L	\mathcal{I}	\mathcal{C}	T	4	Enc-Dec	OF	2	🔗
27	RenderWorld [273]	ICRA'25	N	\mathcal{I} \mathcal{I}	\mathcal{C} \mathcal{E}	T	6	AR	OF, E2E	2	🔗
28	Occ-LLM [265]	ICRA'25	N	\mathcal{I} \mathcal{I}	\mathcal{C} \mathcal{E}	T	6	LLM	OF, E2E	2	🔗
29	DrivingSphere [268]	CVPR'25	N	\mathcal{B}	\mathcal{C}	B	-	LDM	O4G	1 3	🔗
30	EfficientOCF [263]	CVPR'25	N L	\mathcal{C}	\mathcal{C}	T	4	Enc-Dec	OF	2	🔗
31	UniScene [125]	CVPR'25	N	\mathcal{A}	\mathcal{C}	B	6	DiT	O4G, OF	1 3	🔗
32	DIO [44]	CVPR'25	A	\mathcal{A}	\mathcal{C}	N/A	5	Enc-Dec	OF	2	🔗
33	InfiniCube [156]	ICCV'25	W	\mathcal{A}	\mathcal{C}	3 H	1	LDM	O3G	1 2	🔗
34	Control-3D-Scene [150]	ICCV'25	C	\mathcal{A}	\mathcal{C}	G	1	LDM	O3G	1	🔗
35	UniOcc [238]	ICCV'25	N C W O	\mathcal{I} \mathcal{I}	\mathcal{C}	N/A	6	N/A	OF	2	🔗
36	I ² World [142]	ICCV'25	N W	\mathcal{C}	\mathcal{C}	C T V S	6	AR	OF	2	🔗
37	\mathcal{X} -Scene [280]	NeurIPS'25	N	\mathcal{A}	\mathcal{C}	3 B H	1	LDM	O3G	1 3	🔗
38	T ³ Former [261]	arXiv'25	N	\mathcal{C}	\mathcal{C} \mathcal{E}	T	6	AR	OF, E2E	2	🔗
39	COME [204]	arXiv'25	N	\mathcal{C}	\mathcal{C}	T	6	DiT	OF	2	🔗
40	PrITTI [219]	arXiv'25	3	\mathcal{A}	\mathcal{C}	3 B	1	DiT	O3G	3	🔗

further improves geometric and semantic fidelity by conditioning diffusion on initial SSC outputs, reducing inconsistencies in reconstructed scenes.

Generation Consistency Guidance. Other works leverage occupancy to guide high-fidelity, temporally coherent scene synthesis. WoVoGen [155] proposes 4D temporal occupancy volumes to drive multi-view video generation with intra-world and inter-sensor consistency. UrbanDiff [294] uses semantic occupancy grids as geometric priors for 3D-aware image synthesis, while DrivingSphere [268] transforms dynamic 4D occupancy scenes into temporally consistent video via semantic rendering. UniScene [125] generalizes occupancy-based generation across modalities, combining Gaussian-based rendering [111] with prior-guided sparse modeling for unified video and LiDAR synthesis. Collectively, these methods highlight the role of occupancy grids as a unifying

structural prior for producing spatially and temporally consistent outputs with high structural fidelity.

3.2.2 Occupancy Forecasters

Models for 4D occupancy forecasting predict future occupancy from ego actions and past observations, allowing anticipation of environmental changes. This capability serves two purposes: as a self-supervised pretraining task for building generalizable 3D/4D models, and as a dynamic predictor for behavior-aware, controllable future scene generation.

Predictive Model Pretraining. Several methods explore occupancy forecasting as a pretext task to learn rich spatiotemporal features from LiDAR sequences, building generalizable generation models via self-supervised learning. Emergent-Occ [112, 113] introduces differentiable rendering to reconstruct point clouds from 4D occupancy predictions, enabling self-supervised training from raw sequences. UnO [2] models a continuous 4D occupancy field for joint perception and forecasting. Large-scale pretraining frameworks such as UniWorld [170], UniScene [168], and DriveWorld [169] combine image and LiDAR data to learn foundational occupancy models that can be finetuned for downstream tasks like detection and planning, reducing reliance on dense labels while improving generalization.

Ego-Conditioned Occupancy Forecasting. Other approaches forecast occupancy conditioned on both history and ego-agent actions, supporting behavior-aware and controllable prediction. OccWorld [308] jointly models ego motion and surrounding environment evolution in 3D occupancy space, while OccSora [229] generates trajectory-conditioned 4D occupancy over long horizons. Later works enhance controllability [72, 204], fidelity [292], temporal coherence [44, 75, 310], and efficiency [261]. Vision-centric pipelines like Cam4DOcc [158] and its successors [273, 281] integrate world models into end-to-end planning to empower their generative abilities. OccLLaMA [242] and Occ-LLM [265] unify vision, language, and action modalities with semantic occupancy as the shared representation to support embodied question answering, while UniOcc [238] establishes a benchmark combining real and simulated data for standardized evaluation. Together, these works position occupancy forecasting as both a powerful self-supervised learning objective and a key tool for modeling dynamic, action-contingent world states.

3.2.3 Autoregressive Simulators

The occupancy-based autoregressive simulators generate large-scale, temporally coherent 4D occupancy for realistic and interactive simulation. They serve as foundation simulators for perception, planning, and decision-making, with research focusing on two directions: generating scalable unbounded environments and modeling long-horizon dynamics for controllable closed-loop simulation.

Scalable Open-World Generation. Coarse-to-fine and outpainting strategies have been explored to construct large-scale, unbounded 3D occupancy environments. PDD [149] proposes a scale-varied diffusion framework that progressively generates outdoor scenes from coarse layouts to fine details, while XCube [193] adopts hierarchical voxel-based latent diffusion for multi-resolution generation. SemCity [124] adds manipulation functions for scene editing, and InfiniCube [156] and \mathcal{X} -Scene [280] integrate voxel-based occupancy with

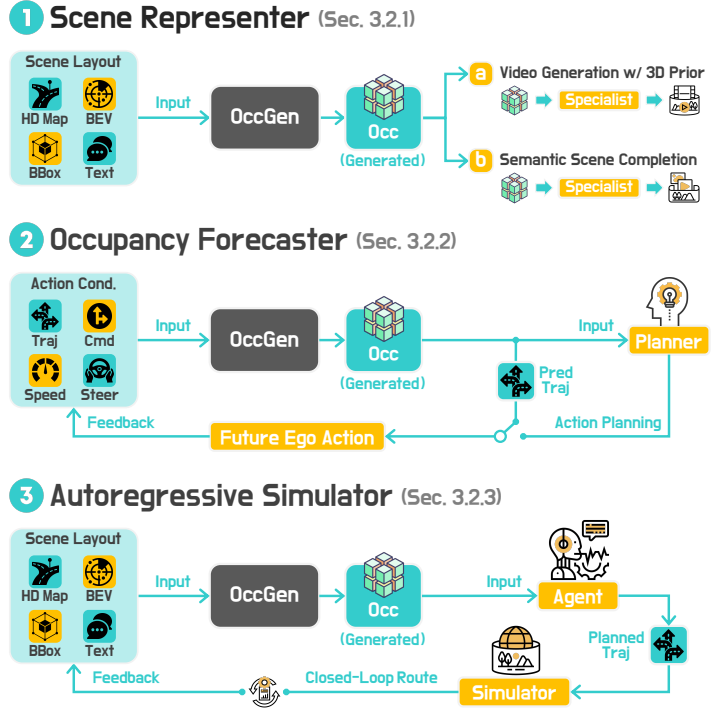


Figure 5 The categorization of **OccGen** models based on functionalities, including scene representors (Sec. 3.2.1), forecasters (Sec. 3.2.2), and autoregressive simulators (Sec. 3.2.3).

consistent visual synthesis for realistic, editable simulation worlds. Together, these works construct scalable occupancy-based representations that serve as interactive and extensible environments for embodied agents.

Long-Horizon Dynamic Simulation. Other works focus on autoregressive 4D occupancy generation to simulate dynamic world evolution. OccSora [229] produces trajectory-conditioned sequences over 16-second horizons, while DynamicCity [17] enables layout-aware and command-conditioned generation, supporting controllable scene synthesis and agent interaction. DrivingSphere [268] constructs a 4D world comprising static backgrounds and dynamic objects for closed-loop simulation, and UniScene [125] generates layout-conditioned 4D occupancy with rich semantic and geometric detail. These approaches integrate spatial structure and temporal coherence to create realistic, controllable environments for embodied agent simulation and decision-making.

3.3 World Modeling from LiDAR Generation

LiDAR-based generation models provide geometry-aware and appearance-invariant representations by modeling complex scenes from point clouds. They enable robust 3D scene understanding and high-fidelity geometric simulation, offering advantages over image- and occupancy-based approaches in both geometric fidelity and environmental robustness. Based on their primary function, these methods can be classified into **three** categories: ❶ **Data Engines**, ❷ **Action Interpreters**, and ❸ **Autoregressive Simulators**. Table 4 summarizes existing models under these domains.

3.3.1 Data Engines

LiDAR-based data engines mitigate the scarcity of large-scale LiDAR training data due to high acquisition costs and annotation challenges by generating diverse and controllable point clouds [117, 138]. Such models enhance perception robustness, enable geometrically accurate scene completion, and support the synthesis of rare or cross-modal scenarios [139]. Recent approaches focus on four major applications.

Perception Data Augmentation. LiDAR-based generative modeling supports data augmentation for core 3D perception tasks such as detection and segmentation, with an emphasis on geometric fidelity and sensor realism. Early approaches primarily focused on modeling uncertainty and spatial structure to synthesize realistic LiDAR scans. DUSTy [174] is a GAN-based framework that synthesizes realistic LiDAR scans by explicitly disentangling the underlying depth map from measurement uncertainty. DUSTy v2 [175] extends DUSTy by incorporating implicit neural representations, enabling the model to generate LiDAR range images at arbitrary resolutions. LiDARGen [324] pioneered the application of Langevin dynamics for LiDAR point cloud generation, achieving superior performance compared to GANs and VAEs. As the first work to adopt the denoising-diffusion paradigm in this domain, it has inspired numerous subsequent studies based on Denoising Diffusion Probabilistic Models (DDPMs) [82]. With explicit positional encoding, R2DM [173] achieves higher-precision LiDAR point cloud generation through a standardized DDPM process. Leveraging flow matching [146], R2Flow [176] significantly accelerates LiDAR point cloud generation.

LiDM [192], RangeLDM [89], and 3DiSS [180] adopt latent diffusion technology by first compressing raw-scale data into low-dimensional latent variables through a pretrained VAE, then training the diffusion model in this latent space. The generated outputs are reconstructed to the original resolution, substantially improving generation speed while preserving quality. LiDARGRIT [78] extends this paradigm by discretizing the latent space with VQ-VAE [221] and generating latent codes using an autoregressive transformer. LiDARGRIT [78] further introduces a raydrop estimation loss to explicitly enhance the raydrop noise modeling. SDS [54] proposes simultaneous diffusion sampling for multi-view LiDAR scene generation, producing all views together to achieve much better geometric consistency than generating each view separately. Recently, SPIRAL [321] pioneered the generation of segmentation-labeled LiDAR data and introduced a novel closed-loop inference strategy that enhances consistency between geometry and semantics. La La LiDAR [148] proposes a layout-guided generative framework that integrates scene graph-based layout diffusion with a foreground-aware control injector, enabling explicit modeling of object relations and controllable scene generation. Veila [147] introduces a conditional diffusion framework for panoramic LiDAR generation guided by a monocular RGB image. It addresses the challenges of reliable conditioning, cross-modal alignment, and maintaining structural coherence beyond the RGB field of view. These advances enhance LiDAR-based perception by generating diverse, controllable, and geometrically faithful training data that capture real-world sensing characteristics.

Table 4 Summary of LiDAR-based generation (LiDARGen) models.

- **Datasets:** **K** KITTI [69], **S** SemanticKITTI [14], **N** nuScenes [24], **3** KITTI-360 [141], **P** PandaSet [255] **C** Carla [48], **s** SeeingThroughFog [18], **W** Waymo [210], **N** NAVSIM [41], **A** Argoverse 2 [247] and **O** OmniDrive [232].
- **Input & Output:** Noisy Latent, Latent Codebook, Noisy LiDAR Point Cloud, LiDAR Point Cloud, LiDAR Sequence, and Images/Videos (Single-View and/or Multi-View).
- **Architectures (Arch.):** *GAN*: Generative Adversarial Network, *Enc-Dec*: Encoder-Decoder, *LDM*: Latent Diffusion Model, *AR*: Autoregressive Model, *DiT*: Diffusion Transformer, *LLM*: Large Language Model.
- **Tasks:** *LG*: LiDAR Generation, *L4G*: 4D LiDAR Generation, *SEG*: 3D Semantic Segmentation, *DET*: 3D Object Detection, *SC*: Scene Completion, *OP*: Occupancy Prediction, and *E2E*: End-to-End Planning.
- **Categories:** **1** Data Engine (Sec. 3.3.1), **2** Action Forecaster (Sec. 3.3.2), and **3** AR Simulator (Sec. 3.3.3).

#	Model	Venue	Dataset	Input	Output	Condition	Len.	Arch.	Task	Cat.	URL
1	DUSTy [174]	IROS'21	K			N/A	1	GAN	LG, SC	1	
2	LiDARGen [324]	ECCV'22	3 N			N/A	1	Enc-Dec	LG, SEG, SC	1	
3	DUSTy v2 [175]	WACV'23	K			L	1	GAN	LG, SEG	1	
4	UltraLiDAR [260]	CVPR'23	S K P			P	1	Enc-Dec	LG, DET, SC	1	
5	Copilot4D [296]	ICLR'24	A K N			T	6	LDM, AR	L4G	1 2	
6	R2DM [173]	ICRA'24	3 K			N/A	1	Enc-Dec	LG, SC	1	
7	ViDAR [284]	CVPR'24	N			N/A	6	Enc-Dec	L4G, DET, OP, E2E	1 2	
8	LiDiff [179]	CVPR'24	3 S			P	1	Enc-Dec	LG, SC	1	
9	LiDM [192]	CVPR'24	3 N S			S R T	1	LDM	LG	1	
10	RangeLDM [89]	ECCV'24	3 N			N/A	1	LDM	LG, SC	1	
11	Text2LiDAR [250]	ECCV'24	3 N			T	1	Enc-Dec	LG, SC	1	
12	LiDARGRIT [78]	arXiv'24	3 K			N/A	1	Enc-Dec, AR	LG	1	
13	BEVWorld [299]	arXiv'24	C N			T	6	LDM	L4G, DET	1 2	
14	SDS [54]	arXiv'24	3			P	1	Enc-Dec	LG, SC	1	
15	HoloDrive [253]	arXiv'24	N			3 T	8	LDM	L4G	1 3	
16	LOGen [116]	arXiv'24	N			3	N/A	DiT	LG	1	
17	OLiDM [269]	AAAI'25	3 N			T 3 O	1	Enc-Dec	LG, DET, SC	1	
18	X-Drive [259]	ICLR'25	N			S T	1	LDM	LG, DET	1	
19	LidarDM [325]	ICRA'25	3 W			H	N/A	LDM	L4G, DET	1 3	
20	LiDAR-EDIT [83]	ICRA'25	N			P	1	Enc-Dec	LG, SC, DET	1	
21	R2Flow [176]	ICRA'25	3 N			N/A	1	DiT	LG	1	
22	WeatherGen [251]	CVPR'25	3 S			P	1	Enc-Dec	LG, DET	1	
23	LiDPM [162]	IV'25	S			P	1	Enc-Dec	LG, SC	1	
24	DiffSSC [27]	IROS'25	S 3			P S	1	Enc-Dec	LG, SC	1	
25	HERMES [318]	ICCV'25	O N			T T	N/A	AR, LLM	L4G, E2E	1 2	
26	SuperPC [49]	CVPR'25	3			R P	1	Enc-Dec	LG, SC	1	
27	OpenDWM [200]	CVPR'25	A 3 N W			3 H	N/A	VQ-VAE	LG, L4G	1 3	
28	SPIRAL [321]	NeurIPS'25	S N			S	1	Enc-Dec	LG, SEG	1	
29	3DISS [180]	arXiv'25	S 3			N/A	1	LDM	LG, SEG	1	
30	Distill-DPO [301]	arXiv'25	S			P	1	Enc-Dec	LG, SC	1	
31	DriveX [203]	arXiv'25	N N			T	6	Enc-Dec	L4G, OP, E2E	1 2	
32	Veila [147]	arXiv'25	K S N			R	1	Enc-Dec	LG, SEG	1	
33	La La LiDAR [148]	AAAI'26	N W			O	1	Enc-Dec	LG, SEG, DET, SC	1	
34	LiDARCrafter [139]	AAAI'26	N			3 O	6	Enc-Dec	L4G	1 3	

Scene Completion. The completion of 3D scenes aims to reconstruct dense and coherent 3D geometry from sparse or occluded LiDAR scans, with recent generative methods improving geometric fidelity and controllability. UltraLiDAR [260] introduces a discrete voxel-based representation for LiDAR point clouds using a VQ-VAE [221], enabling efficient and controllable sparse-to-dense completion. LiDiff [179] and DiffSSC [27] utilize the denoising process of DDPM to reposition duplicated points, thereby densifying the LiDAR point cloud while simultaneously completing occluded areas. Building on UltraLiDAR [260] for background completion and AnchorFormer [37] for foreground object synthesis, LiDAR-EDIT [83] enables flexible editing of LiDAR scenes, including object removal and insertion. By enhancing the ability to denoise large-magnitude noise, LiDPM [162] extends LiDiff [179] to generate dense point clouds not only from sparse inputs but also from pure Gaussian noise, thus enabling the synthesis of entirely novel scenes. Similarly, Distillation-DPO [301] enhances both completion quality and inference efficiency of LiDiff [179] through the integration of Score Distillation [189] and Diffusion-DPO [224]. Recently, SuperPC [49] proposes a unified framework that transforms point clouds into representation features suitable for completion, upsampling, denoising, and colorization, thereby avoiding the error accumulation that can arise from sequentially applying separate models.

Rare Condition Modeling. To improve the robustness of 3D perception in adverse conditions, recent methods

explore controllable LiDAR generation for safety-critical scenarios. Text2LiDAR [250] presents a Transformer-based architecture that integrates textual information to enable text-controlled LiDAR point cloud generation. WeatherGen [251] targets rainy, snowy, and foggy conditions, generating high-quality LiDAR point clouds for these conditions within a unified controllable generative model. The practical utility of the generated point cloud data is validated through 3D object detection tasks in these adverse weather scenarios. OLiDM [269] addresses fidelity limitations at the object level via a two-stage pipeline: it first generates foreground objects, which are then used as conditions for scene generation, ensuring controllable and high-quality results at both object and scene levels. Meanwhile, LOGen [116] proposes an object-level point cloud generation model to synthesize traffic participants, conditioned on their relative orientation and distance to the sensor.

Multimodal Generation. Several recent methods [253, 299] investigate multimodal generation by synthesizing aligned LiDAR and image data. X-Drive [259] introduces a dual-branch diffusion architecture for jointly generating aligned LiDAR point clouds and multi-view camera images in driving scenarios. Its key innovation is the cross-modality epipolar condition module, which improves consistency between the point cloud and image modalities. Furthermore, X-Drive [259] supports controllable 3D scene generation conditioned on heterogeneous inputs, including text descriptions, object bounding boxes, and sensor data variants from the images or the LiDAR point clouds.

3.3.2 Action Forecasters

Based on past observations, the LiDAR-based world models functioning as action forecasters generate future LiDAR sequences conditioned on given future states.

Temporal Modeling. Copilot4D [296] proposes a scalable approach to building world models, primarily by (1) leveraging a VQ-VAE [221] model to tokenize complex, unstructured point cloud inputs, and (2) recasting the Masked Generative Image Transformer [30] as a discrete diffusion model to enable parallel denoising and decoding. Copilot4D takes as input 1–3 seconds of past LiDAR frames along with future ego actions (poses), and predicts high-quality LiDAR frames for the next 1–3 seconds. ViDAR [284] takes historical camera frames as input and predicts future LiDAR frames as output. This framework further enables pre-training for tasks such as perception, prediction, and planning.

Multi-Modal Action Forecasters. BEV-World [299] introduces a multi-modal tokenizer to extend the generative capability to both surround-view images and LiDAR point clouds. DriveX [203] supports multi-modal outputs, including point clouds, camera images, and semantic maps. By employing a decoupled latent world modeling strategy that separates world representation learning for spatial modeling from latent future decoding for future state prediction, DriveX effectively simplifies the modeling of complex dynamics in unstructured scenes. HERMES [318] integrates LLMs to generate textual descriptions of future frames in addition to LiDAR, thereby enhancing human-machine interaction.

3.3.3 Autoregressive Simulators

World models functioning as autoregressive simulators aim to generate temporally coherent LiDAR sequences for realistic and interactive simulation. These models serve as a foundation for perception, planning, and

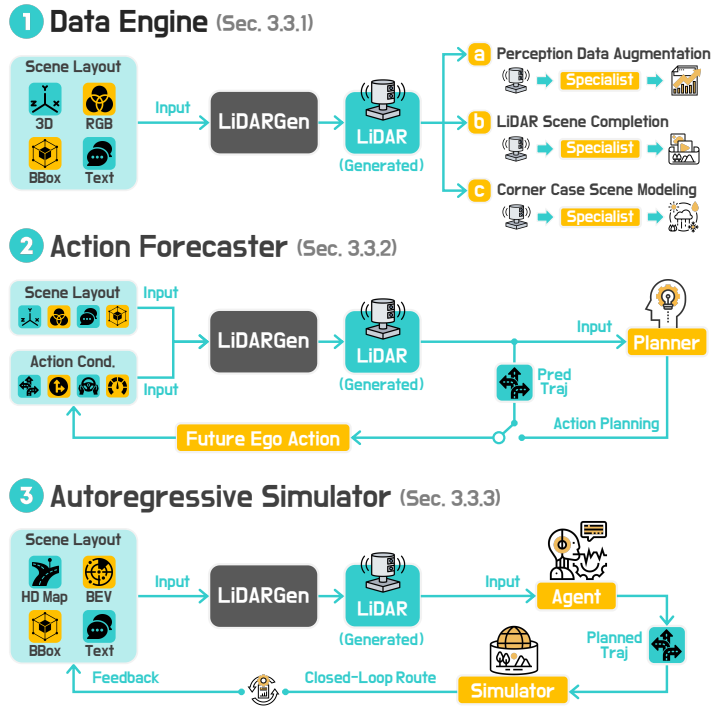


Figure 6 The categorization of **LiDARGen** models based on functionalities, including data engines (Sec. 3.3.1), action forecasters (Sec. 3.3.2), and autoregressive simulators (Sec. 3.3.3).

Table 5 Summary of datasets and benchmarks used for training **VideoGen**, **OccGen**, and **LiDARGen** models.

- **Column Keys:** # 📺 = Total number of frames; # 🏠 = Total number of occupancy scenes; # 🚗 = Total number of LiDAR scenes; **Freq** = Annotation frequency; Symbol “-” in a cell indicates the information is not provided.
- **Tasked by:** ① Video Gen. Models (**VideoGen**, cf. Sec. 3.1), ② Occupancy Gen. Models (**OccGen**, cf. Sec. 3.2), and ③ LiDAR Gen. Models (**LiDARGen**, cf. Sec. 3.3). Kindly refer to Table 1 for the definitions of conditions.

#	Dataset	Venue	# Scene	# 📺 (View)	# 🏠	# 🚗	Freq	Conditions	Tasked by	URL
K	KITTI [69]	CVPR'12	22	15k (×4)	-	15k	10	D 3 F	① ③	🔗
U	NYUv2 [206]	ECCV'12	464	1449 (×1)	1449	-	-	D S	②	🔗
C	CARLA [48]	CoRL'17	∞	∞	∞	∞	Free	3 T S	③	🔗
S	SemanticKITTI [14]	ICCV'19	22	-	43k	23k	10	S T	② ③	🔗
N	nuScenes [24]	CVPR'20	1000	1.4M (×6)	40k	400k	2	3 B H T V S	① ② ③	🔗
W	Waymo Open [210]	CVPR'20	1150	1M (×5)	-	230k	10	3 B H T V S	① ② ③	🔗
S	STF [18]	CVPR'20	-	1.4M (×2)	-	1.4M	0.1	3 T L P	⑤	🔗
V	Virtual KITTI 2 [23]	arXiv'20	5	40k (×2)	-	-	10	3 T	①	🔗
A	Argoverse 2 [247]	NeurIPS'21	1000	2.7M (×9)	-	150k	10	3 T H	② ③	🔗
L	Lyft-Level5 [84]	CoRL'21	170k	282M (×7)	42.5M	42.5M	10	T H	②	🔗
N	nuPlan [25]	CVPRW'21	-	24M (×6)	-	24M	-	T H	①	🔗
P	PandaSet [255]	ITSC'22	103	48k (×6)	-	16k	10	3 L P T S	③	🔗
O	OpenCOOD [264]	ICRA'22	73	11k (×4)	11k	11k	10	3 T	②	🔗
3	KITTI-360 [141]	TPAMI'22	379	150k (×4)	-	80k	10	3 L P T S H T	③	🔗
C	CarlaSC [248]	RA-L'22	24	-	43k	43k	10	T S	②	🔗
R	Robo3D [118]	ICCV'23	2194	-	-	476k	10	3 T S	③	🔗
O	OpenOccupancy [235]	ICCV'23	850	200k (×6)	34k	34k	2	3 T S	②	🔗
N	Occ3D-nuScenes [216]	NeurIPS'23	900	240k (×6)	40k	40k	2	3 T S	②	🔗
Y	OpenDV-YouTube [274]	CVPR'24	2139	60M (×1)	-	-	10	T C	①	🔗
S	SSCBench [137]	IROS'24	1859	404k (×6)	66k	66k	-	3 T S	②	🔗
N	NAVSIM [41]	NeurIPS'24	115k	920k (×8)	-	115k	2	3 T H	③	🔗
D	DrivingDojo [240]	NeurIPS'24	17.8k	1.7M (×1)	-	-	5	T T	①	🔗
O	OmniDrive [232]	CVPR'25	1000	1.4M (×6)	40k	-	2	3 T H T	③	🔗
E	EUVS [79]	ICCV'25	345	90k (×8)	-	-	-	T	①	🔗
P	Pi3DET [138]	ICCV'25	25	51k (×1)	-	51k	10	3 T	③	🔗

decision-making, with a focus on geometric fidelity and temporal consistency. Existing methods can be divided into two types based on their data generation paradigms.

Sequential Autoregressive LiDAR Generation. HoloDrive [253] presents an autoregressive framework for jointly generating multi-view camera images and LiDAR point clouds by introducing a depth prediction branch in the 2D generative model to improve alignment between 2D and 3D representations. More recently, LiDARcrafter [139] extends the layout-based two-stage framework of La La LiDAR [148] to the 4D domain, with an autoregressive LiDAR sequence generator, supporting fine-grained control, long-term temporal coherence, and diverse editing capabilities.

Scene-Scale Simulation from Meshes. LidarDM [325] constructs mesh grids from point clouds by removing dynamic objects across multiple frames. It then trains a diffusion model conditioned on the BEV layout, enabling it to generate a mesh world. By incorporating dynamic objects with motion trajectories into this mesh world and performing ray projection through the scene, LidarDM can synthesize long sequential LiDAR point clouds.

4 Datasets & Evaluations

In this section, we provide a comprehensive evaluation of 3D/4D world modeling across four aspects. ¹**Datasets** (Sec.4.1) introduce widely used benchmarks with multimodal inputs and annotations across video, occupancy, and LiDAR formats. ²**Metrics and Protocols** (Sec.4.2) define standardized criteria for assessing generation fidelity, forecasting accuracy, planning awareness, reconstruction quality, and downstream performance. ³**Quantitative Benchmarks** (Sec.4.3) report results of state-of-the-art models under these protocols. ⁴**Qualitative Analyses** (Sec.4.4) highlight strengths, limitations, and trade-offs across different modalities.

4.1 Datasets

In this survey, we discuss real, simulated, and augmented datasets that support research in 3D and 4D world modeling. These datasets span urban driving and related settings and provide rich annotations and conditions needed for **VideoGen**, **OccGen**, and **LiDARGen**. An overview of popular datasets and related benchmarks is illustrated in Figure 3. Additionally, Table 5 provides detailed statistics of each collection of the video, occupancy, LiDAR, and other relevant data formats from these mainstream datasets.

Among existing 3D/4D data collections, real-world datasets supply realism and multimodal context with reliable calibration. Recent web-scale corpora trade strict calibration for scale, diversity, and text supervision. Simulators contribute perfect labels, editable layouts, and rare or counterfactual scenarios. Together, these sources form a complementary foundation for training and evaluating controllable and planning-aware world models.

Video-based datasets provide long, coherent video sequences with reliable calibration, ego pose, and synchronized multi-view images. Conditions that aid controllability include action logs, HD maps, and language signals such as captions or driving commands. Real-world datasets, *e.g.*, nuScenes [24] and Waymo Open [210], provide surround-view imagery, accurate poses, and dense perception annotations, making them strong bases for video generation with map- or motion-conditioned control. Planning-aware datasets like NAVSIM [41] and nuPlan [25] pair short scenarios with ego motion, CAN signals, and maps to support policy-grounded video modeling. Web-scale video such as OpenDV-YouTube [274] contributes breadth and language supervision via captions and ego-action tags, trading off precise calibration for scale and diversity. Synthetic platforms like CARLA [48] offer poses and editable layouts for counterfactuals, rare events, and controlled ablations.

Occupancy-based datasets need voxelized 3D supervisions in a consistent coordinate frame, with semantic labels and tight alignment to the sensor rig. Conditions that stabilize learning include HD maps, ego trajectories, and either multi-view images or LiDAR to anchor the field over time. In driving settings, ready-to-use *real-world* benchmarks such as OpenOccupancy [235], Occ3D-nuScenes [216], NYUv2 [206], and SSCBench [137] provide standardized voxel grids and protocols for training and evaluation. *Simulated datasets* like CarlaSC [248] offer clean ground truth and full control of layout and motion, which is useful for ablations and stress tests. Semantic extensions like SemanticKITTI [14] couple point-wise labels with occupancy volumes and enable joint learning of geometry and semantics.

LiDAR-based datasets require raw LiDAR-acquired sweeps with precise extrinsics, per-sweep ego poses, and object-level annotations. Additional 2D and 3D cues, such as HD maps, radar, and camera imagery, enable cross-modal conditioning, while coverage across weather conditions and sensor configurations improves robustness. Representative real-world sources include KITTI [69], nuScenes [24], Waymo Open [210], and Argoverse2[247]. NAVSIM [41] supplements these with short scenario snippets paired with control signals, supporting downstream planning tasks. For robustness testing, recent benchmarks [18, 118, 138] capture adverse weather, inject systematic corruptions, and cover multiple platforms to assess generalization. Synthetic platforms, such as CARLA [48], offer clean LiDAR simulations, editable environments, and controllable signals.

4.2 Evaluation Metrics & Protocols

Standardized evaluations lay the foundation for the development of generation models. However, existing literature has overlooked the importance of establishing a systematic protocol for evaluations in 3D and 4D.

Here, we organize evaluation metrics for world models into five perspectives:

- **¹Generation Quality** (Sec. 4.2.1) assesses the realism, coherence, and controllability of synthesized outputs.
- **²Forecasting Quality** (Sec. 4.2.2) evaluates future predictions given partial observations.
- **³Planning-Centric Quality** (Sec. 4.2.3) metrics measure safety and rule compliance in planning.
- **⁴Reconstruction-Centric Quality** (Sec. 4.2.4) examines the ability of generation models to reproduce or simulate novel views.

- ⁵ **Downstream Evaluation** (Sec. 4.2.5) tasks test how world models support tasks like detection, segmentation, and reasoning.

A comprehensive summary of evaluation metrics is provided in Table 14. Together, these metrics cover both perceptual fidelity and utility in embodied decision-making and beyond.

4.2.1 Generation Quality

Generation quality focuses on whether a world model can produce realistic and coherent outputs given a prompt or condition. This involves four dimensions: fidelity, consistency, controllability, and human preference.

Fidelity evaluates how closely a generator matches the real data distribution and is typically divided into two families. *Perceptual metrics* project samples into a feature space learned from human-labeled data, where distances align with human judgments of realism. The Fréchet family [81, 116, 149, 173, 192, 205, 220] encodes samples, fits Gaussians to real and generated features, and reports the Fréchet distance. Some variants differ by modality and encoder, while semantic versions [321] add labels to align categories. Other representative metrics include Inception Score [199], which uses Inception logits to reward confident and diverse predictions without real references. *Statistical metrics* operate directly on geometry or density. They ask whether the generated set covers the real set, stays within it, and matches the low-level structure. Some metrics [17, 116, 116] target the fidelity-coverage trade-off, probing set overlap by measuring whether generated samples stay on the real manifold while sufficiently covering it, while other metrics [78, 324, 324] quantify distributional discrepancy in geometry or density via different distance metrics.

Consistency evaluates whether a world model produces coherent outputs across space, time, and semantics. *Spatial Consistency* scores geometric alignment. Some [211, 239] quantify multi-view agreement by matching keypoints in overlapping regions, while others evaluate alignment by projecting the 3D outputs and comparing them with monocular depth estimates [259]. *Temporal Consistency* is measured by cosine similarity [104] between adjacent-frame embeddings from foundation models [181, 191], and *Subject Consistency* [98] tracks identity persistence by comparing subject-region features [181] across frames.

Controllability measures how well a model adheres to user-specified inputs, with metric design tailored to the conditioning modality. When the condition is reference frames, CLIP Similarity [157, 274] averages cosine similarity between CLIP embeddings of generated and reference frames to gauge semantic alignment.

Table 6 Benchmarking **VideoGen** models on the **Perceptual Fidelity** of generation quality evaluations. The reported metrics are FID and FVD scores on the official nuScenes [24] validation set. All metrics are the lower the better (\downarrow).

Method	Resolution	Freq	FID \downarrow	FVD \downarrow
Single-View Video Generation				
DriveDreamer [236]	128×192	2 Hz	14.90	340.80
GenAD [274]	256×448	2 Hz	15.40	184.00
ProphetDWM [233]	256×448	2 Hz	6.90	190.50
Epona [295]	512×1024	5 Hz	7.50	82.80
MaskGWM [178]	288×512	10 Hz	4.00	59.40
LongDWM [234]	480×720	10 Hz	12.30	102.90
DriVerse [134]	480×832	10 Hz	18.20	95.20
InfinityDrive [77]	576×1024	10 Hz	10.93	70.06
GEM [80]	576×1024	10 Hz	10.50	158.50
Vista [65]	576×1024	10 Hz	6.90	89.40
UniFuture [140]	320×576	12 Hz	11.80	99.90
MiLA [226]	360×640	12 Hz	8.90	89.30
GeoDrive [31]	480×720	12 Hz	4.10	61.60
STAGE [228]	512×768	12 Hz	11.04	242.79
Doe-1 [309]	384×672	-	15.90	-
Multi-View Video Generation				
Drive-WM [239]	192×384	2 Hz	15.80	122.70
WoVoGen [155]	256×448	2 Hz	27.60	417.70
Panacea [244]	256×512	2 Hz	16.96	139.00
SubjectDrive [93]	256×512	2 Hz	15.98	124.00
Glad [256]	256×512	2 Hz	11.18	188.00
SynthOcc [131]	448×800	2 Hz	14.75	-
CogDriving [153]	480×720	2 Hz	15.30	37.80
DrivingDiffusion [135]	512×512	2 Hz	15.83	332.00
Delphi [157]	512×512	2 Hz	15.08	113.50
MaskGWM [178]	288×512	10 Hz	8.90	65.40
DriveScape [249]	576×1024	10 Hz	8.34	76.39
MagicDrive3D [62]	224×400	12 Hz	20.67	164.72
MagicDrive [64]	224×400	12 Hz	16.20	218.12
DreamForge [163]	224×400	12 Hz	14.61	209.90
DrivePhysica [286]	256×448	12 Hz	3.96	38.06
UniScene [125]	256×512	12 Hz	6.45	71.94
MiLA [226]	360×640	12 Hz	4.90	36.30
CoGen [101]	360×640	12 Hz	10.15	68.43
DiST-4D [76]	424×800	12 Hz	6.83	22.67
DiVE [104]	480×854	12 Hz	-	94.60
DrivingSphere [268]	480×1080	12 Hz	-	103.40
DriveDreamer-2 [305]	512×512	12 Hz	11.20	55.70
NoiseController [47]	512×1024	12 Hz	13.72	87.23
MagicDrive-V2 [63]	848×1600	12 Hz	20.91	94.84
BEVWorld [299]	-	12 Hz	19.00	154.00
UniMLVG [34]	-	12 Hz	5.80	36.10
DualDiff [129]	224×400	-	10.99	160.00
BEVGen [211]	224×400	-	24.54	-
PerLDiff [293]	256×708	-	13.36	-
HoloDrive [253]	-	-	13.60	103.00
BEVControl [275]	-	-	24.85	-

Table 7 Benchmarking **VideoGen** models on the **Downstream Evaluation** tasks. The reported metrics are mAP and NDS for **3D Object Detection**, mIoU (Lanes, Drivable, Divider) for **BEV Map Segmentation**, L2 and Collision Rates at timestamps 1s, 2s, and 3s for **Open-Loop Planning**, and PDMS (P) and ADS (A) scores [278] for **Closed-Loop Planning**. All results are computed using the UniAD [92] implementation and checkpoints on official nuScenes [24] validation set.

Method	3D Det \uparrow		BEV Seg mIoU (%) \uparrow			Open-Loop Planning \downarrow						Closed-Loop Planning \uparrow			
	mAP	NDS	Lane	Dri	Div	L2@1s	L2@2s	L2@3s	CR@1s	CR@2s	CR@3s	P@SG	A@SG	P@BOS	A@BOS
Baseline [92]	37.98	49.85	31.31	69.14	25.93	0.51	0.98	1.65	0.10	0.15	0.61	-	-	-	-
MagicDrive [64]	12.92	28.36	21.95	51.46	17.10	0.57	1.14	1.95	0.10	0.25	0.70	-	-	-	-
Panacea [244]	13.72	27.73	18.23	52.37	17.21	0.58	1.14	1.95	-	-	-	-	-	-	-
DiST-4D [76]	15.63	32.44	26.80	60.32	21.69	0.56	1.11	1.91	-	-	-	-	-	-	-
DriveArena [278]	16.06	30.03	26.14	59.37	20.79	0.56	1.10	1.89	0.02	0.18	0.53	0.76	0.13	0.50	0.045
DreamForge [163]	16.63	30.57	26.16	58.98	20.22	0.55	1.08	1.85	0.08	0.27	0.81	0.81	0.12	0.74	0.076
DrivingSphere [268]	21.45	34.16	57.99	62.87	22.29	0.54	1.10	1.76	-	-	-	-	-	-	-

Beyond this, layout and object-level control is typically scored by agreement with detectors or segmentors on boxes and masks [319], scene-graph control by count errors and set overlap [150], and camera-pose control by trajectory rotation and translation errors [107],

Human Preference captures subjective qualities like realism and plausibility that automated scores may miss. Studies typically adopt either two-alternative forced choice [234] or mean opinion score [150] setups, involving both experts and lay users to provide human evaluation on world models.

4.2.2 Forecasting Quality

Forecasting quality extends beyond unconditional generation by evaluating how well the model predicts the future evolution of a scene given partial observations. Here, forecasting quality is evaluated in spatial and temporal domains.

Spatial Predictive Accuracy in forecasting measures how well predictions match the actual future in the spatial domain. For frames and videos, FID, FVD, and frame-level L1/L2 errors remain standard. IoU evaluates occupancy forecasts [158] at multiple horizons to separate near- and long-range correctness. Point-cloud forecasts [296] are evaluated by comparing the predicted and ground-truth sweeps in 3D space, using Chamfer distance for geometric overlap and depth-wise errors to quantify per-ray distance accuracy.

Temporal Predictive Accuracy in 4D forecasting assesses whether predictions remain temporally coherent, especially without full supervision [238]. Typical examples are Key Object Dimension Probability [238], which penalizes unlikely object sizes using category-specific priors, and Temporal Background Environment Consistency [238], which tracks static voxels under ego-motion to verify scene rigidity.

4.2.3 Planning-Centric Quality

Planning-centric metrics assess whether the model’s outputs result in safe, efficient, and rule-compliant decisions, and its evaluation falls into open-loop and closed-loop.

Open-Loop Planning assessment evaluates predictions that do not influence future inputs. nuPlan [25] compares predictions to expert demonstrations using waypoints and heading error, and a horizon-dependent Miss Rate, which thresholds trajectory and heading errors into bounded scores. To approximate behavioral quality without full interaction, NAVSIM [28, 41] introduces short non-reactive rollouts and aggregate safety, drivable-area

Table 8 Benchmarking **VideoGen** models on the **Downstream Evaluation** tasks. The reported metrics are mAP and NDS for **3D Object Detection** (*w/* BEVFusion [151] and StreamPETR [231]) and Road-wise mIoU scores (RmIoU) and Vehicle-wise mIoU scores (VmIoU) for **BEV Map Segmentation** (*w/* CVT [312]). The results are on the official nuScenes [24] validation set. All metrics are the higher the better (\uparrow).

Method	BEVFusion		StreamPETR		CVT	
	mAP	NDS	mAP	NDS	RmIoU	VmIoU
Baseline	35.54	41.21	34.50	46.90	73.67	34.82
BEVControl [275]	-	-	-	-	60.80	26.80
BEVGen [211]	-	-	-	-	50.20	5.89
Panacea [244]	-	-	22.50	36.10	-	-
DrivingDiffusion [135]	-	-	-	-	63.20	31.60
SimGen [319]	-	-	-	-	62.90	31.20
CogDriving [153]	-	-	-	-	65.70	32.10
UniMLVG [34]	-	-	-	-	70.81	29.12
DrivePhysica [286]	-	-	35.50	43.67	-	-
SubjectDrive [93]	-	-	28.00	41.10	-	-
Glad [256]	-	-	27.10	40.80	-	-
DriveScape [249]	-	36.50	-	-	64.43	28.86
MagicDrive [64]	12.30	23.32	-	-	61.05	27.01
DreamForge [163]	13.01	22.16	26.00	41.10	65.27	28.36
DualDiff [129]	13.99	24.98	-	-	62.75	30.22
PerLDiff [293]	15.24	24.05	-	-	61.26	27.13
MagicDrive-V2 [63]	17.65	-	-	-	59.79	32.73
NoiseController [47]	20.93	27.96	-	-	64.85	27.32
DrivingSphere [268]	22.71	31.79	-	-	-	-

compliance, progress, and comfort into a single policy score, using gating and weighted averaging to align with closed-loop outcomes.

Closed-Loop Planning evaluation executes the policy in an interactive simulator and scores observed behavior. CARLA [48] reports route or goal completion and infraction distance statistics for opposite-lane driving, sidewalk incursions, and collisions with other agents. nuPlan [25] provides a broader suite of closed-loop checks, including no at-fault collisions, drivable-area and direction compliance, time-to-collision bounds, speed-limit compliance, progress along route, capturing both traffic legality and human-likeness.

4.2.4 Reconstruction-Centric Quality

Reconstruction-centric neural simulators aim to reproject the past into interactive sensor views or novel viewpoints.

Photometric Fidelity captures low-level rendering quality when ground-truth images under known viewpoints are available. Following standard practices in neural rendering, metrics such as PSNR [100], SSIM [241], and LPIPS [297] remain foundational. PSNR quantifies pixel-level accuracy, SSIM evaluates structural consistency in luminance and texture, while LPIPS measures perceptual similarity in deep feature space aligned with human visual preferences.

View Changing Consistency evaluates the spatiotemporal plausibility of novel or counterfactual viewpoints where ground truth is unavailable [177, 304]. In such settings, photometric comparison is insufficient. Metrics like Novel Trajectory Agent IoU [304] assess whether foreground agents maintain geometrically plausible behavior, offering targeted signals for validating realism in 4D interactive simulations.

4.2.5 Downstream Evaluation

While the above evaluations assess a world model in isolation, downstream evaluations measure its utility when integrated into end-to-end perception and decision-making systems. Tasks span *object detection* (mAP [145], nuScenes Detection Score [24]), *multi-object tracking* (MOTA, MOTP [15]), *semantic and BEV map segmentation* (mIoU), *3D occupancy prediction and scene completion* (voxel-level IoU, Voxelized Panoptic Quality). In language-grounded settings such as *visual question answering*, models like OccLLaMA [242] report exact-match Top-1 accuracy across question types and difficulty levels. These evaluations reflect how well a learned world model supports downstream reasoning, representation, and control tasks effectively.

4.3 Quantitative Experiments & Analyses

In this section, we quantitatively evaluate world modeling approaches through ¹**VideoGen Benchmarks** (Sec.4.3.1), ²**OccGen Benchmarks** (Sec.4.3.2), and ³**LiDARGen Benchmarks** (Sec. 4.3.3). Models are assessed on standardized datasets using fidelity, consistency, and forecasting metrics, along with downstream perception and planning tasks. These evaluations reveal both the progress and limitations of current methods, highlighting key trade-offs between realism, geometric accuracy, temporal stability, and controllability.

Table 9 Benchmarking **OccGen** models on **Reconstruction Quality**. The reported metrics are mIoU (%) for **Semantic Occupancy Reconstruction** and IoU (%) for **Occupancy Reconstruction**. All results are on the official nuScenes [24] validation set. Both metrics are the higher the better (\uparrow).

Method	Type	Resolution	mIoU \uparrow	IoU \uparrow
OccSora [229]	VQVAE	$(\frac{T}{8}, 25, 25, 512)$	27.40	37.00
OccLLaMA [242]	VQVAE	(50, 50, 128)	65.93	57.66
OccWorld [308]	VQVAE	(50, 50, 128)	66.38	62.29
UrbanDiff [294]	VQVAE	(50, 50, 2048)	80.00	98.80
I ² World [142]	VQVAE	(50, 50, 128)	81.22	68.30
Occ-LLM [265]	VAE	(50, 50, 64)	71.08	62.74
UniScene [125]	VAE	(50, 50, 8)	72.90	64.10
DOMe [72]	VAE	(25, 25, 64)	83.08	77.25
UniScene [125]	VAE	(100, 100, 8)	92.10	87.00
T ³ Former [261]	Triplane-VAE	(100, 100, 16, 8)	85.50	72.07
X-Scene [280]	Triplane-VAE	(100, 100, 16, 8)	92.40	85.60

Table 10 Benchmarking **OccGen** models on **4D Occupancy Forecasting Quality**. The reported metrics are mIoU (%) for **Semantic Occupancy Reconstruction** and IoU (%) for **Occupancy Reconstruction**, respectively, at timestamps 1s, 2s, and 3s. All results are on the official nuScenes [24] validation set. Both metrics are the higher the better (\uparrow).

Method	mIoU (%) \uparrow			IoU (%) \uparrow		
	1s	2s	3s	1s	2s	3s
GaussianAD [310]	6.29	5.36	4.58	14.13	14.09	14.04
PreWorld [133]	12.27	9.24	7.15	23.62	21.62	19.63
Occ-LLM [265]	24.02	21.65	17.29	36.65	32.14	28.77
OccLLaMA [242]	25.05	19.49	15.26	34.56	28.53	24.41
OccWorld [308]	25.78	15.14	10.51	34.63	25.07	20.18
RenderWorld [273]	28.69	18.89	14.83	37.74	28.41	24.08
COMe [204]	30.57	19.91	13.38	36.96	28.26	21.86
DFIT-OccWorld [292]	31.68	21.29	15.18	40.28	31.24	25.29
DOMe [72]	35.11	25.89	20.29	43.99	35.36	29.74
UniScene [125]	35.37	29.59	25.08	38.34	32.70	29.09
T ³ Former [261]	46.32	33.23	28.73	77.00	75.89	76.32
I ² World [142]	47.62	38.58	32.98	54.29	49.43	45.69

4.3.1 Benchmarking Video Generation Models

Generation Fidelity. Table 6 reports FID and FVD results on the nuScenes validation set for both single-view and multi-view vision-based world models. Early baselines such as GenAD [274] and DriveDreamer [236] operate at relatively low resolutions and frame rates, achieving modest performance (FID \sim 15, FVD 180–340). Later single-view models improve visual quality. Vista [65] and InfinityDrive [77] leverage higher resolutions and frame rates, reducing FVD below 100. Recent works like MaskGWM [178] and GeoDrive [31] set new state-of-the-art, reaching FID around 4–5 and FVD near 60. In the multi-view setting, early BEV-based approaches (BEVControl [275], BEVGen [211]) yield high FID ($>$ 20). Subsequent models, including DriveWM [239], Panacea [244], and MagicDrive [64] reduce errors but struggle with temporal stability (FVD $>$ 120). Strong improvements come from models emphasizing geometric consistency and spatio-temporal alignment. UniScene [125], DriveScape [249], and DiST-4D [76] achieve the best balance, with FVD scores below 80 and DiST-4D [76] reaching as low as 22.67.

The comparison suggests resolution and frame rate strongly influence generation fidelity. Besides, explicit multi-view modeling is challenging; although many methods reduce FID, temporal coherence remains difficult, highlighting the importance of structured 4D representations. Finally, methods combining geometry-aware priors with temporal reasoning, such as DiST-4D and UniScene, demonstrate that enforcing spatial structure and temporal consistency jointly is crucial for scalable autonomous driving video generation.

Downstream Evaluations. Table 8 and Table 7 evaluate downstream perception and planning on generated scenes. Early generative baselines (BEVControl [275], BEVGen [211]) provide limited perception benefits, especially in vehicle segmentation ($<$ 27% mIoU). More advanced methods such as MagicDrive [64] and DreamForge [163] improve both detection (up to 26 mAP on StreamPETR [231]) and segmentation ($>$ 61% road mIoU), while DrivePhysica [286] and Glad [256] further push detection accuracy (35.5 mAP, 43.7 NDS). For segmentation, UniMLVG [34] and CogDriving [153] achieve the highest fidelity (70.8% road, 32.1% vehicle mIoU). Beyond perception, planning performance highlights the persistent gap between synthetic and real data. While real nuScenes provides the upper bound (37.9 mAP, 49.9 NDS, 1.05 Avg L2), generative methods lag significantly in detection and planning accuracy. Nevertheless, world models like DriveArena [278] and DreamForge demonstrate reduced planning errors and collision rates, enabling preliminary closed-loop driving with non-trivial success rates (*e.g.*, 0.81 PDMS). DrivingSphere [268] achieves the strongest drivable area segmentation ($>$ 58% mIoU), while DiST-4D [76] balances detection and segmentation performance but lacks closed-loop validation.

Overall, the results show that photorealistic generation alone is insufficient to improve downstream tasks; explicit modeling of geometry, temporal consistency, and motion dynamics is crucial. Models that incorporate such priors not only enhance detection and segmentation but also support safer planning by reducing collisions and trajectory errors. Strong segmentation fidelity further demonstrates the benefit of multi-view and structure-aware models in capturing global layouts, yet the performance gap to real data remains significant, underscoring the challenge of aligning generative fidelity with task-level utility.

4.3.2 Benchmarking Occupancy Generation Models

Occupancy Reconstruction Quality. Table 9 evaluates the reconstruction capability of occupancy world models under VAE-based formulations. Conventional VAEs such as DOME [72] already achieve strong results (83.08% mIoU, 77.25% IoU), outperforming most VQVAEs. While UrbanDiff [294] and I²World [142] show competitive IoU, other variants like OccSora [229] degrade significantly under coarse temporal-spatial compression.

Table 11 Benchmarking **OccGen** models on **Motion Planning Quality**. The reported metrics are L2 Error Rate (in meters) and Collision Rate (%), respectively, at timestamps 1s, 2s, and 3s. All results are on the official nuScenes [24] validation set. Both metrics are the lower the better (\downarrow).

Method	L2 Error (m) \downarrow			Collision Rate (%) \downarrow		
	1s	2s	3s	1s	2s	3s
ST-P3 [90]	1.33	2.11	2.90	0.23	0.62	1.27
OccNet [217]	1.29	2.13	2.99	0.21	0.59	1.37
FSF-Net [75]	0.54	1.09	-	0.01	0.01	-
UniAD [92]	0.48	0.96	1.65	0.05	0.17	0.71
OccWorld [308]	0.43	1.08	1.99	0.07	0.38	1.35
PreWorld [133]	0.41	1.16	2.32	0.50	0.88	2.42
GaussianAD [310]	0.40	0.64	0.88	0.09	0.38	0.81
DFIT-OccWorld [292]	0.38	0.96	1.75	0.07	0.39	0.90
Occ-LLaMA [242]	0.37	1.02	2.03	0.04	0.24	1.20
GenAD [274]	0.36	0.83	1.55	0.06	0.23	1.00
RenderWorld [273]	0.35	0.91	1.84	0.05	0.40	1.39
T ³ Former [261]	0.32	0.91	1.76	0.08	0.32	0.51
Drive-OccWorld [281]	0.32	0.75	1.49	0.05	0.17	0.64
Occ-LLM [265]	0.12	0.24	0.49	-	-	-

Triplane-based VAEs [124, 261, 280] bring the largest gains, with T³Former [261] reaching 85.50% mIoU and X-Scene [280] establishing a new state-of-the-art at 92.40% mIoU and 85.60% IoU.

These results underline that latent representation design is decisive for reconstruction fidelity. Triplane factorization enforces geometric consistency and enables finer spatial detail, while simply enlarging latent dimensionality (*e.g.*, UrbanDiff [294] with 2048 channels) yields limited returns. Compact VAEs such as UniScene [125] further show that well-regularized low-dimensional spaces can generalize effectively, whereas aggressive compression (*e.g.*, OccSora [229]) sacrifices accuracy. Overall, effective compression combined with explicit geometric priors is key to scalable and accurate 3D and 4D scene modeling.

4D Occupancy Forecasting Quality. Table 10 presents 4D occupancy forecasting results over the period of 1–3 seconds. Baselines such as OccWorld [308] and OccLLaMA [242] achieve moderate performance (17–20% mIoU), while DOME [72] and UniScene [125] improve temporal stability (27.10% and 31.76% mIoU). More recent models show further progress: I²World [142] reaches 39.73% mIoU with balanced IoU, and T³Former [261] excels in spatial coherence with 76.40% IoU. The comparisons reveal three insights. First, naive autoregressive or generative approaches deteriorate rapidly at longer horizons, highlighting the need for structured priors. Second, triplane factorization substantially improves spatial fidelity, as reflected in the performance of T³Former [261]. Third, I²World shows that coupling scalable latent reasoning with temporal modeling yields the best balance across horizons. Accurate 4D forecasting thus requires not only generative power but also structured representations that enforce geometric and temporal consistency.

End-to-End Planning. Table 11 reports the performance of end-to-end planning, measured by trajectory error (L2) and collision rate. Sequence-based planners like ST-P3 [90] perform poorly (2.11 meters in L2 error), while UniAD [92] and GenAD [274] achieve substantial gains, with UniAD+DriveWorld [169] further improving to 0.69 meter in L2 error and 0.19% collisions. Occupancy-based world models such as OccWorld [308] and OccLLaMA [242] reduce errors to around 1.15 meters. Structured refinements (*e.g.*, DFIT-OccWorld [292], RenderWorld [273], and Drive-OccWorld [281]) achieve stronger accuracy and safety, with Drive-OccWorld reaching 0.85 m in L2 error and 0.29% collisions. Notably, GaussianAD [310] and T³Former [261] balance error and safety, while Occ-LLM [265] reports extremely low error (*i.e.*, only 0.28 meter in L2 error). The results show that integrating occupancy world models into planning pipelines consistently outperforms pure trajectory-based methods. Hybrid designs that refine occupancy priors, such as Drive-OccWorld [281] and DFIT-OccWorld [292], bring joint improvements in accuracy and safety, demonstrating the downstream robustness of generative modeling. Overall, structured occupancy representations form a strong foundation for end-to-end autonomous driving, enabling reliable long-horizon planning in complex scenarios.

4.3.3 Benchmarking LiDAR Generation Models

Generation Fidelity. Table 12 reports the performance of recent LiDAR scene generation methods on SemanticKITTI [14] using four fidelity metrics (FRD, FPD, JSD, and MMD). Earlier methods such as LiDARGen [324] and LiDM [192] exhibit relatively large distributional discrepancies, as reflected by high FRD and FPD scores. In contrast, more recent approaches, including R2DM [173], Text2LiDAR [250], and WeatherGen [251], achieve substantially better results across most metrics, indicating a closer alignment between generated and real LiDAR distributions.

Table 12 Benchmarking LiDARGen models on the **Perceptual Fidelity** evaluations. The reported metrics are FRD, FPD, JSD and MMD scores on the official SemanticKITTI [14] dataset. All metrics are the lower the better (\downarrow).

Method	Resolution	FRD \downarrow	FPD \downarrow	JSD \downarrow	MMD \downarrow
LiDARGen [324]	64×1024	681.37	115.17	0.1323	2.19×10 ⁻³
LiDM [192]	64×1024	-	108.70	0.0456	2.90×10 ⁻⁴
R2DM [173]	64×1024	192.81	19.29	0.0373	1.60×10 ⁻⁴
Text2LiDAR [250]	64×1024	522.32	11.09	0.0750	4.29×10 ⁻⁴
WeatherGen [251]	64×1024	184.11	11.42	0.0290	3.80×10 ⁻⁵

Table 13 Benchmarking LiDARGen models on **4D LiDAR Generation Quality**. The reported metrics are TTCE and CTC. The numbers indicate frame intervals. All results are on nuScenes [24]. Both metrics are the lower the better (\downarrow).

Method	TTCE \downarrow		CTC \downarrow			
	3	4	1	2	3	4
UniScene [125]	2.74	3.69	0.90	1.84	3.64	3.90
OpenDWM [200]	2.68	3.65	1.02	2.02	3.37	5.05
OpenDWM-DiT [200]	2.71	3.66	0.89	1.79	3.06	4.64
LiDARCrafter [139]	2.65	3.56	1.12	2.38	3.02	4.81



Figure 7 Qualitative comparisons of state-of-the-art **VideoGen** models on the nuScenes [24] dataset. From top to bottom rows: Reference (from the dataset), MagicDrive [64], DreamForge [163], DriveDreamer-2 [305], and OpenDWM [200].

The results reveal a clear progression in LiDAR generation quality. Among evaluated methods, WeatherGen [251] achieves the best performance across all metrics by employing Mamba [71] as its backbone. Interestingly, Text2LiDAR [250], despite its strong conditioning on textual input, produces higher FRD, suggesting that aligning with semantic prompts may compromise geometric fidelity. These findings underscore the importance of balancing semantic controllability with distributional realism in future LiDAR scene generation research.

4D LiDAR Generation Quality. Table 13 benchmarks recent LiDAR-based 4D scene generation methods on temporal coherence, using TTCE (Temporal Transformation Consistency Error) and CTC (Chamfer Temporal Consistency) as evaluation metrics. Unlike video generation, which has been extensively studied with standardized benchmarks, temporal LiDAR generation remains relatively underexplored, and current metrics mainly focus on explicit geometric alignment across frames. The results reveal several observations. First, end-to-end autoregressive methods such as UniScene [125] and OpenDWM-DiT [200] demonstrate clear advantages in maintaining short-horizon geometric consistency, as reflected in lower TTCE and CTC at 1–2 frame intervals. However, their fixed-length generation limits broader applicability, as error accumulation grows at longer horizons. Second, incorporating strong vector quantization modules [200] facilitates better condition embedding and fine-grained reconstruction, leading to improved temporal stability. Third, modality choices introduce inherent trade-offs: BEV-based generation offers smoother temporal continuity but sacrifices fidelity to the raw point cloud pattern, while range-based [139] generation better preserves LiDAR-specific sensing characteristics but requires careful design to embed conditions and sustain long-term consistency.

4.4 Qualitative Experiments & Analyses

In this section, we qualitatively evaluate the 3D and 4D generation approaches through ¹**VideoGen Visualizations** (Sec.4.4.1), ²**OccGen Visualizations** (Sec.4.4.2), and ³**LiDARGen Visualizations** (Sec. 4.4.3). These evaluations highlight the strengths, limitations, and trade-offs of current methods, informing future advances in realism, consistency, and generalization.



Figure 8 Qualitative comparisons of state-of-the-art **VideoGen** models on the nuScenes [24] dataset. From top to bottom rows: Reference (from the dataset), MagicDrive [64], DreamForge [163], DriveDreamer-2 [305], and OpenDWM [200].

4.4.1 Qualitative Analyses of VideoGen Models

Visual Realism. Figure 7 and Figure 8 compare recent video generation world models, including MagicDrive [64], DreamForge [163], DriveDreamer-2 [236], and OpenDWM [200]. The generated scenes capture overall layouts and semantics close to real-world distributions, but fine-grained details often suffer from pixel misalignment, blurred textures, and structural discontinuities. Among the methods, OpenDWM [200] achieves the most realistic, consistent, and controllable results, owing to its training on diverse datasets (OpenDV [274], nuScenes [24], and Waymo Open [210]), while others rely on a single dataset. This underscores the role of dataset diversity in improving generalization and robustness.

Physical Plausibility. In the absence of explicit physics constraints, generated videos may exhibit violations of physical realism, such as vehicle-background interpenetration, incorrect shadows, or scale distortions. While such issues may appear subtle in static frames, they significantly reduce realism when viewed as continuous video, undermining temporal coherence and physical plausibility.

Controllability. Appearance-level controls (weather, time-of-day, style) can be reliably controlled via large-scale pre-trained video generation models with text conditioning. By contrast, precise geometric control over object position, orientation, and velocity remains challenging, typically requiring dedicated control embeddings or structured conditioning mechanisms.

Long-Tail Categories. Rare and small-scale classes (e.g., pedestrians, cyclists, traffic signs) remain hard to generate convincingly. Long-tail data imbalance often leads to unrealistic shapes, distorted geometry, or even omission of these critical objects.

Takeaways. The results suggest that future progress in video-based world models requires advances along five critical axes: (i) **realism**, reducing artifacts and enhancing detail fidelity; (ii) **consistency**, maintaining semantic and temporal coherence; (iii) **controllability**, unifying high-level appearance control with fine-grained geometric control; (iv) **physical plausibility**, incorporating physics priors to prevent unrealistic artifacts; and (v) **generalization**, leveraging diverse large-scale datasets to improve robustness.

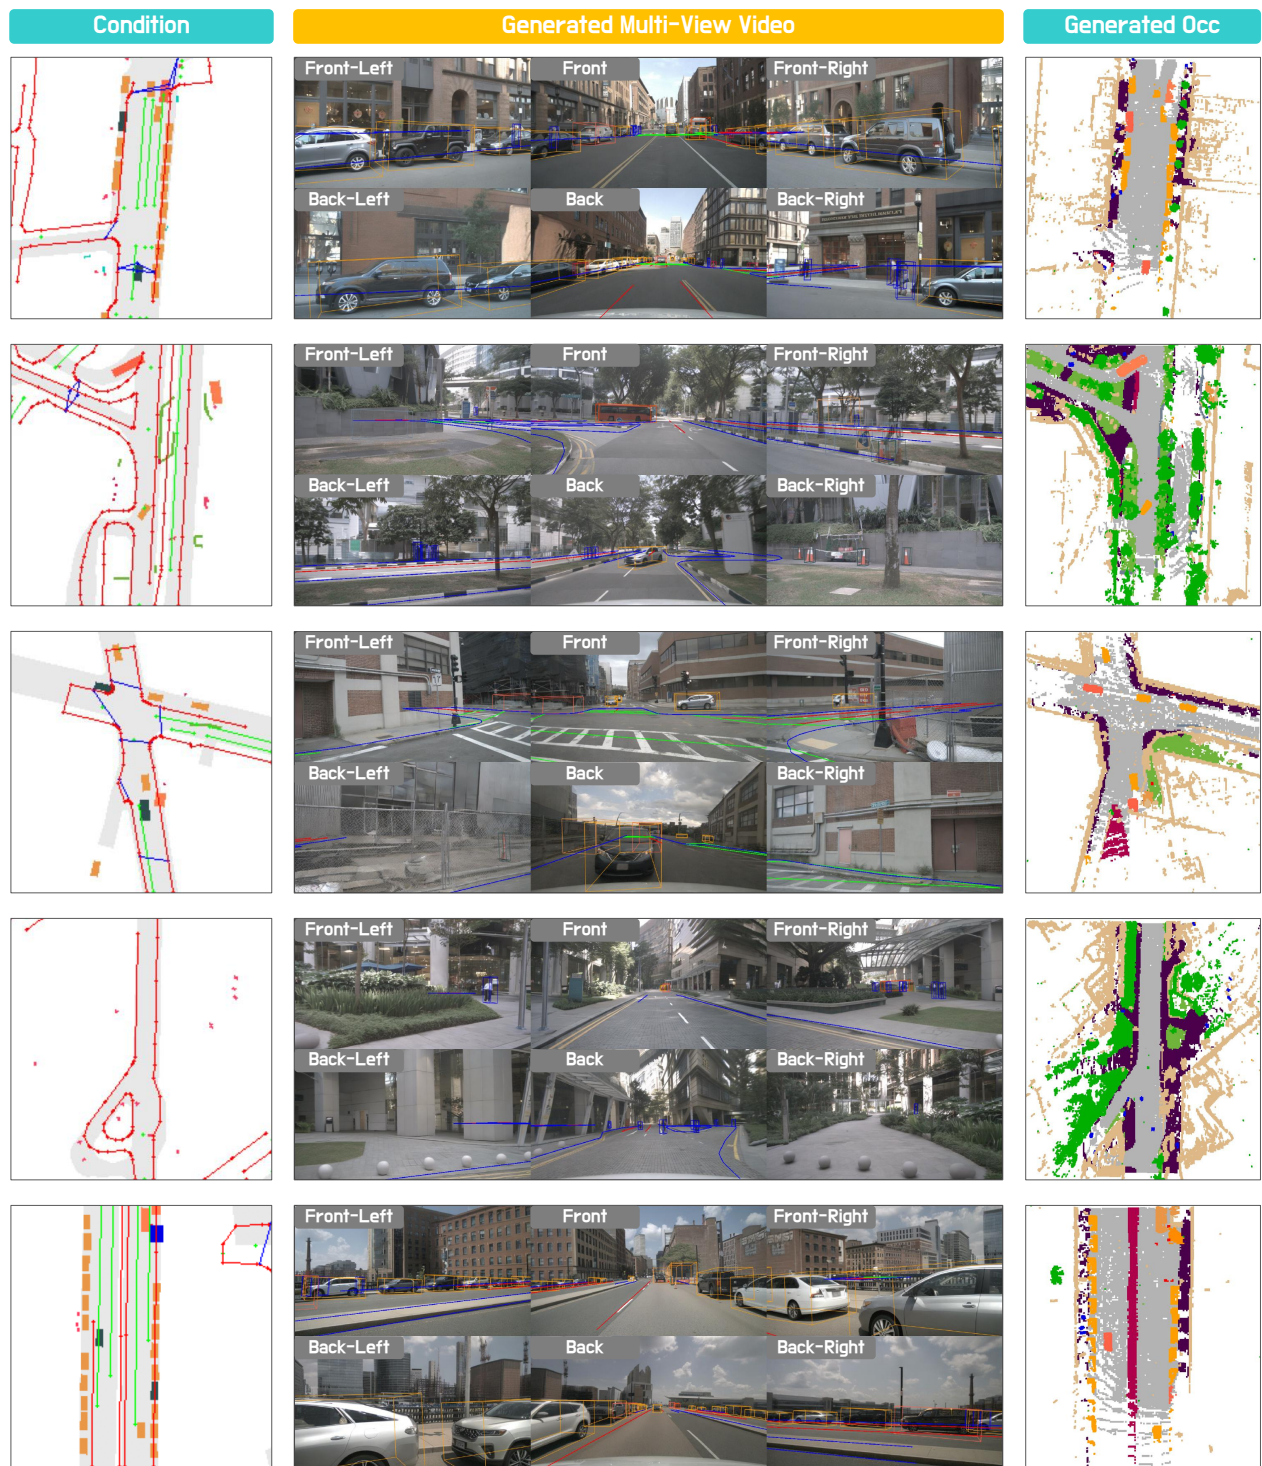


Figure 9 Qualitative examples of **OccGen** models on nuScenes [24]. From left to right columns: The input condition, the generated multi-view videos, and the generated occupancy grids. The results are generated using \mathcal{X} -Scene [280].

4.4.2 Qualitative Analyses of OccGen Models

3D Geometric Consistency. Figure 9 shows qualitative results of occupancy generation models conditioned on scene layouts. The generated videos and occupancies exhibit strong spatial alignment across perspectives. Such cross-view coherence is crucial for maintaining geometric plausibility in multi-camera settings.

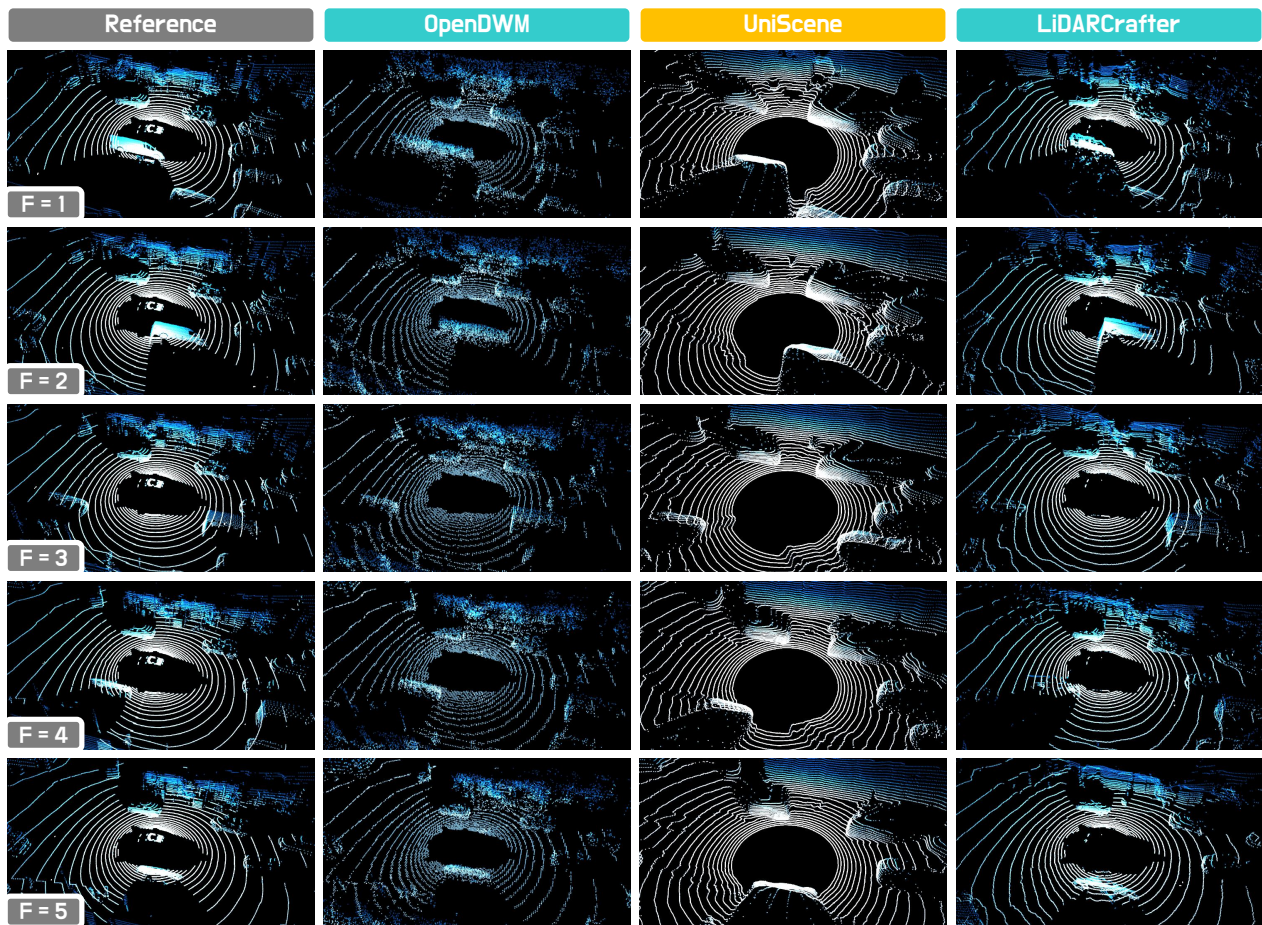


Figure 10 Qualitative comparisons of state-of-the-art **LiDARGen** models on the nuScenes [24] dataset. From left to right columns: Reference (from the dataset), OpenDWM [200], UniScene [125], and LiDARCrafter [139].

Occupancy Fidelity. The generated occupancies preserve key semantics, including drivable areas, sidewalks, and surrounding objects. While overall layouts are captured reliably for downstream perception, fine-grained geometry (e.g., thin lane boundaries, small dynamic agents) remains challenging, often leading to misalignment or incomplete reconstruction.

Controllability and Generalization. Conditioned on high-level scene priors, models can flexibly adapt to diverse intersection layouts and road structures, demonstrating promising controllability. However, rare structures and long-tail categories (e.g., bicycles, pedestrians) are often poorly represented, revealing limitations in data diversity and generalization capacity.

Takeaways. These results suggest that progress in occupancy generation hinges on three aspects: (i) **geometric consistency**, ensuring spatial coherence across 3D environments; (ii) **fine-grained fidelity**, particularly for small-scale and dynamic objects; and (iii) **generalization**, leveraging diverse datasets to handle rare layouts and long-tail classes. Advancing these aspects is essential for robust world models capable of supporting downstream tasks and closed-loop simulation.

4.4.3 Qualitative Analyses of LiDARGen Models

Global Patterns. Figure 10 compares representative LiDAR generation paradigms. The original scans exhibit dense rings with uniform angular spacing, faithfully capturing both static structures and dynamic objects. The voxel-based OpenDWM [200] emphasizes coherent scene geometry but often yields overly regularized patterns due to voxel-level modeling. The range-based LiDARCrafter [139] better preserves the native scanline structure with sharper rings, though it may introduce artifacts around occlusion boundaries. The occupancy-based

UniScene [125] reproduces global distributions but tends to oversmooth fine details, leading to discontinuities.

Point Cloud Sparsity. Given the inherent sparsity of LiDAR data, generation models must balance realistic density with structural consistency. OpenDWM [200] often produces overly sparse regions, especially at long ranges. LiDARcrafter [139] maintains more uniform angular density, closely following the sensor’s scanning characteristics. UniScene [125] provides globally complete coverage but sometimes introduces artificial filling inconsistent with real sensor patterns.

Object Completeness. Dynamic agents such as vehicles are particularly important for downstream perception and planning. OpenDWM [200] frequently underrepresents object contours, resulting in fragmented or partial shapes. LiDARcrafter [139] offers better surface completion, though finer details can be noisy. UniScene [125] reconstructs volumetrically plausible objects with consistent occupancy, but often lacks the sharp boundaries and crisp detail of real scans.

Takeaways. These results highlight three key attributes for LiDAR generation: (i) **global patterns**, ensuring coherent scene geometry while preserving sensor-specific scan structures; (ii) **point sparsity**, maintaining realistic density distributions that match LiDAR characteristics; and (iii) **object completeness**, accurately capturing dynamic agents with sharp contours and consistent surfaces. Future advances will require balancing these attributes to generate LiDAR sequences that are both perceptually realistic and physically faithful to sensor properties.

5 Applications

The versatility of 3D and 4D world models enables deployment across diverse domains. ¹**Autonomous Driving** (Sec. 5.1) supports simulation, evaluation, and scenario synthesis. ²**Robotics** (Sec. 5.2) leverages them for navigation, manipulation, and scalable simulation. ³**Video Games & XR** (Sec. 5.3) benefit from content generation, immersive rendering, and adaptive environments. ⁴**Digital Twins** (Sec. 5.4) enable city-scale reconstruction, event replay, and scene editing. ⁵**Emerging Applications** (Sec. 5.5) span scientific discovery, healthcare, industry, and disaster response. Together, these applications showcase the role of world models in unifying perception, prediction, and generation across domains.

5.1 Autonomous Driving

3D and 4D world models provide a principled foundation for autonomous driving, supporting simulation, evaluation, and scenario synthesis. They enable controllable, interactive, and safety-critical environments that cannot be easily reproduced in the real world. We outline three major applications.

Traffic Simulation. World models enable realistic traffic simulators with heterogeneous agents, diverse motion, and physics-compliant interactions. Compared with image-only platforms, volumetric representations such as occupancy grids [229, 292], multi-frame LiDAR point clouds [139], or scene-level meshes [165] provide richer geometry and temporal coherence [239, 283]. Modern systems further support controllable parameters (*e.g.*, traffic density, intent, weather) and stochastic perturbations, improving robustness and generalization for downstream policies [86, 198, 226, 236, 244, 245, 300, 305].

Closed-Loop Driving Evaluation. Beyond static benchmarks, closed-loop setups couple generative models with agents to assess *perception*→*planning*→*control* stacks over long horizons [283, 304]. By jointly modeling ego behavior and surrounding traffic dynamics, models create responsive environments that adapt to agent actions in real time [31, 139]. This allows scalable evaluation of robustness under distribution shifts, rare events, and recovery after failures [198, 300], while modular conditioning (*e.g.*, HD maps, text queries, and ego trajectories) enables targeted stress testing [31, 63].

Scenario Synthesis. World models can generate rare or safety-critical driving scenes that are underrepresented in real datasets, which is essential for evaluating robustness. Typical cases include severe occlusions, sudden intrusions, multi-agent conflicts, and adverse weather [96, 148, 321]. Controllable generation with HD maps, semantic masks, scene graphs, or textual prompts enables targeted testing [31, 139, 321]. Physics- and motion-aware models ensure dynamic feasibility [85, 315], while stochastic sampling improves coverage of rare

events. LiDAR-centric approaches such as LiDARcrafter [139] further extend this capability to 4D sequences with temporal coherence.

5.2 Robotics

3D and 4D world models have the potential to enhance robotic intelligence by supporting navigation, manipulation, and simulation. They provide spatial-temporal grounding, physical reasoning, and scalable synthetic environments, which are crucial for robust policy learning.

Embodied Navigation. Robots leverage world models to perceive and predict dynamic layouts, enabling long-horizon exploration, obstacle avoidance, and localization in both structured and unstructured settings [106, 212, 316]. Forecasting future states is critical in crowded or occluded scenes [106, 254], where multi-scan LiDAR, voxelized occupancy, and predictive dynamics provide reliable spatial-temporal cues [56, 283]. Recent studies also combine visual, topological, and linguistic signals for instruction following and adaptive decision-making [94, 316].

Object-Centric Manipulation. For this task, models capture object geometry and physical transitions, allowing robots to anticipate contact dynamics and plan stable grasps or rearrangements [13, 159, 317]. Representations such as meshes, keypoint graphs, and volumetric embeddings support fine-grained control and generalization to new objects [108, 171]. Integration of differentiable physics with generative models yields physically consistent predictions that can be optimized for various tasks [13, 51, 317].

Scene Generation for Simulation. Generative models create diverse synthetic environments, reducing manual design costs for training and evaluation [3, 106, 202]. Procedural variation in layout, semantics, and dynamics exposes robots to a wide range of scenarios, improving robustness and sim-to-real transfer [51, 212, 316, 317]. Flexible scene representations, from meshes to voxel grids and point clouds, further enable integration with both physics-based simulators and photorealistic renderers [13, 171].

5.3 Video Games & XR

World models transform gaming and XR by automating content creation, supporting immersive rendering, and enabling adaptive environments that respond to player actions.

Procedural World Generation. Generative models automate the design of expansive virtual worlds, supporting open exploration and emergent gameplay [42, 103, 164]. Procedural pipelines can incorporate maps, player states, or language prompts to scale content production beyond manual asset creation [88, 164]. Maintaining temporal and semantic coherence is key for believable dynamic evolution [68], while diverse scene representations such as point clouds, voxels, and neural radiance fields balance realism, style, and efficiency [111, 221].

Interactive Scene Rendering. Immersive XR requires real-time rendering of dynamic scenes where users move freely through evolving geometry and lighting [130, 282]. Neural representations including NeRF [221] and Gaussian Splatting [111] advance photorealistic synthesis, with temporal extensions modeling motion and state change [6, 61]. To ensure consistency and comfort, systems must maintain geometric fidelity under arbitrary viewpoints, adapt scene content to user actions [9, 289], and employ efficient pipelines to sustain high frame rates.

Playable Environment Adaptation. Adaptive worlds adjust geometry, layout, and agent behavior to sustain challenge and engagement [42, 66, 287, 288]. 3D/4D models support real-time transformations such as altering terrain, collapsing structures, or spawning entities based on player interactions [88, 288]. By leveraging priors or high-level instructions, these systems preserve style, physics, and narrative coherence [68, 139], thereby enhancing immersion, replayability, and personalized gameplay.

5.4 Digital Twins

3D and 4D world models underpin urban digital twins by enabling large-scale reconstruction, event replay, and interactive editing. These capabilities support planning, analysis, and simulation in smart city applications.

City-Scale Scene Modeling. Digital twins integrate multimodal sensing, including LiDAR, RGB-D, aerial photogrammetry, and drone surveys, to capture both static infrastructure and dynamic activities [17, 143, 257].

They enable applications such as traffic monitoring, infrastructure planning, and disaster response [14, 201], while dynamic modeling simulates pedestrian and vehicle flows for capacity planning [213, 214]. Recent advances in streaming pipelines and 4D compression maintain temporal consistency and allow metropolitan-scale deployment [43, 45].

Event Replay & Forecasting. World models reconstruct past or hypothetical events from sparse sensor logs, aiding analysis of incidents [40, 87], construction monitoring [225], or emergency response [73]. Replayable 4D scenes clarify causality, while predictive extensions enable what-if simulations for evaluating interventions. Alignment with sensor ground truth remains critical for reliability.

Scene Control & Editing. Interactive tools allow users to manipulate urban content for simulation and visualization, including vehicle removal, weather alteration, and layout modification [43, 68]. Such controllability improves planning workflows and supports immersive city-scale analysis.

5.5 Other Emerging Applications

Beyond autonomous driving and robotics, 3D and 4D world models are expanding into scientific, medical, industrial, and safety-critical domains. These applications highlight their versatility in modeling complex spatial-temporal systems.

Scientific Discovery and Environmental Modeling. World models capture natural dynamics from multimodal observations, supporting forecasting and exploratory simulation. Applications include climate and weather prediction [16, 121, 183], monitoring glacier retreat or floods, and simulating wildfire spread. By learning directly from data, they complement physics-based solvers with faster iteration.

Healthcare & Biomechanics. Generative 3D models reproduce anatomy deformation and tissue behavior for surgical training, planning, and guidance [279]. Predictive motion models aid rehabilitation, prosthetics, and injury prevention by anticipating joint trajectories [114], enhanced by multi-view capture and volumetric reconstruction.

Industrial Process & Manufacturing Simulation. Virtual prototyping with world models supports robotic assembly, material handling, and inspection [21, 38]. Temporal simulation of component interactions reduces costly trials and enables analysis of efficiency and fault recovery.

Security, Defense & Disaster Response. Synthetic environments simulate tactical operations, hazardous conditions, and evacuations [223]. Dynamic scene modeling further aids disaster preparedness by predicting structural collapse, fire spread, or chemical dispersion, and testing emergency response plans.

6 Challenges & Future Directions

In this section, we highlight key challenges of world models, including benchmarking, long-horizon fidelity, physical realism, efficiency, and cross-modal coherence, and outline directions for future research.

6.1 Standardized Benchmarking & Evaluations

A major barrier to progress in the driving world models is the lack of common, standardized benchmarks and evaluation protocols. Current studies often utilize different datasets or ad hoc metrics, which makes it difficult to meaningfully compare models and assess their true performance in diverse realistic settings [50, 98, 127, 307]. Establishing unified benchmarks can provide a comprehensive evaluation framework that captures key metrics such as physical plausibility, temporal consistency, and controllability. Moreover, standardized evaluations should encompass both closed-loop simulation tests and real-world scenarios to validate the model’s capabilities under varying traffic densities, weather conditions, and complex urban architectures [278]. Future work must focus on developing these benchmarks to ensure fair and transparent comparisons across different approaches.

6.2 High-Fidelity & Long-Horizon Generation

Another critical challenge in world models for autonomous driving is achieving high-fidelity generation over long time horizons [34, 163]. While short-term predictions may capture immediate interactions with

reasonable accuracy, small errors tend to accumulate over longer sequences, leading to unrealistic behaviors and degradation of scene consistency. The difficulty of maintaining both high visual fidelity and long-horizon coherence is compounded by the complexity of dynamic urban environments, where interactions between multiple agents and environmental factors evolve continuously. Addressing these issues requires advanced generative techniques that explore novel training paradigms [295] and memory mechanisms [77] that effectively penalize long-term divergences to enable reliable long-term simulation.

6.3 Physical Fidelity, Controllability & Generalizability

From the perspective of the generation capability, current world models for autonomous driving are critically limited by a failure to ensure *physical realism*, offer fine-grained controllability, and achieve robust generalization [63, 286]. They often produce physically implausible events, such as non-deforming collision impacts and objects that lack temporal consistency [109]. Furthermore, their editing capabilities remain coarse, typically confined to adjusting traffic agents’ positions or appearances while neglecting granular control over environmental elements like architecture or road signs. Most critically, these models tend to overfit their training data, failing to generalize to new urban environments and rare objects, thus limiting their real-world applicability. Future work must overcome these challenges to build more faithful, controllable, and generalizable world models.

6.4 Computational Efficiency & Real-Time Performance

Another pressing limitation of current world models for autonomous driving lies in computational efficiency and real-time responsiveness. Existing methods often depend on heavy architectures and multi-step sampling strategies, leading to substantial latency and memory overhead, which undermines their practicality for large-scale data generation and simulation. Moving forward, research should prioritize sparse computation [196] and inference acceleration techniques [32] in order to enable world models that are both accurate and responsive while remaining scalable.

6.5 Cross-Modal Generation Coherence

Current world models often struggle to achieve consistent cross-modal generation, wherein visual, geometric, and semantic modalities must jointly interact to form a coherent representation of the environment. Misalignment can result in generated imagery that conflicts with the underlying 3D structure, undermining reliability in downstream perception and planning tasks. Overcoming these issues requires integrated architectures that jointly learn from diverse sensor data while enforcing strict consistency constraints during generation [125, 259]. Furthermore, ensuring fine-grained spatial alignment and temporal synchronization is crucial for accurately modeling the dynamic interactions in realistic driving environments. Future research should target this fundamental challenge to harmonize diverse data streams.

7 Conclusion

This survey has presented the first systematic review of **3D and 4D world modeling and generation**, clarifying definitions, organizing methods into a hierarchical taxonomy across VideoGen, OccGen, and LiDARGen, and summarizing datasets, evaluations, and applications. By shifting focus from purely visual realism to geometry-grounded modeling, native 3D and 4D representations enable models to achieve plausibility, controllability, and physical consistency, serving roles as data engines, action interpreters, neural simulators, and scene reconstructors. Despite rapid progress, challenges remain in scaling to real-world complexity, aligning multimodal signals, and establishing standardized evaluation for controllability, safety, and generalization. Looking forward, unifying generative and predictive paradigms, integrating language and reasoning, and advancing simulation and digital twin ecosystems represent promising directions. Equally important will be community efforts in creating open benchmarks, reproducible codebases, and large-scale datasets tailored for 3D/4D world models, which can accelerate progress and ensure comparability across methods. We hope this survey provides both a coherent foundation and a forward-looking roadmap for advancing robust, interpretable, and generalizable world models to power the next generation of embodied AI.

Table 14 Summary of the evaluation metrics used for evaluating the quality of ¹Generation, ²Forecasting, ³Planning, ⁴Reconstruction, and ⁵Downstream Tasks for the VideoGen, OccGen, and LiDARGen models in 2D, 3D, and 4D tasks.

Abbr.	-	Full Name	Description	Ref.
🏗️ Generation Quality - Perceptual Fidelity				
FID	↓	Fréchet Inception Distance	Statistical distance between multivariate Gaussians fitted to Inception features of real and generated samples, measuring distributional similarity.	[81]
FVD	↓	Fréchet Video Distance	Statistical distance between multivariate Gaussians fitted to I3D features [29] of real and generated videos, capturing temporal coherence.	[220]
FRD	↓	Fréchet Range Distance	Statistical distance between Gaussians fitted to RangeNet++ features [167] extracted from LiDAR range images, assessing distributional fidelity.	[173]
FPD	↓	Fréchet Point Cloud Distance	Statistical distance between Gaussians fitted to PointNet features [190] of raw 3D point clouds, evaluating geometric realism.	[205]
FSVD	↓	Fréchet Sparse Volume Distance	Statistical distance between Gaussians fitted to volumetric encoder features of sparse voxel inputs, capturing volumetric structure.	[192]
FPVD	↓	Fréchet Point Volume Distance	Statistical distance between Gaussians fitted to volumetric encoder features of hybrid point-voxel representations, measuring fidelity.	[192]
F3D	↓	Fréchet 3D Distance	Statistical distance between Gaussians fitted to occupancy grid features, evaluating volumetric realism in generated 3D data.	[149]
S-FRD	↓	Semantic Fréchet Range Distance	Class-aware extension of FRD that incorporates semantic labels for improved alignment of LiDAR range distributions.	[321]
S-FPD	↓	Semantic Fréchet Point Distance	Class-aware extension of FPD that integrates semantic labels to assess alignment of 3D point cloud distributions.	[321]
KID	↓	Kernel Inception Distance	Maximum Mean Discrepancy between Inception features using a polynomial kernel, providing an unbiased distributional similarity measure.	[19]
IS	↑	Inception Score	Evaluates image realism by rewarding confident and diverse class predictions from a pretrained Inception classifier, without real reference data.	[199]
IQ	↑	Image Quality	Predicts perceptual image quality by estimating human opinion scores with a learned quality assessor [110], without ground truth references.	[104]
🏗️ Generation Quality - Statistical Fidelity				
PR	↑	Precision-Recall	Reports sample fidelity as precision and distributional coverage as recall, characterizing closeness to the real data manifold.	[17]
SWD	↓	Sliced Wasserstein Distance	Mean Wasserstein distance over multiple random 1D projections of image patches at different scales, reflecting distributional similarity.	[78]
JSD	↓	Jensen-Shannon Divergence	Symmetric divergence measuring similarity between occupancy histograms of real and generated scenes, lower indicating better alignment.	[324]
MMD	↓	Minimum Matching Distance	Average Chamfer distance from each real sample to its nearest generated neighbor, quantifying geometric fidelity.	[324]
COV	↑	Coverage	Fraction of real samples matched by at least one generated output, measuring generative diversity and recall.	[116]
1-NNA	-	1-Nearest-Neighbor Accuracy	Overlap test using a 1-NN classifier trained across sets, where accuracy near 50% indicates distributional equivalence.	[116]
Diversity	↑	-	Degree of variability across generated outputs for fixed prompts, often measured via pixel- or feature-wise variance.	[319]
🏗️ Generation Quality - Spatial Consistency				
VCS	↑	View Consistency Score	Summation of LoFTR keypoint confidences [209] across overlapping views, evaluating multi-view geometric consistency and alignment quality.	[211]

Note: Continued on next page

Abbr.	-	Full Name	Description	Ref.
KPM	↑	Key Points Matching	Ratio of successfully matched keypoints between adjacent generated and real views, reflecting geometric alignment quality.	[239]
DAS	↓	Depth Alignment Score	Statistical discrepancy between projected point clouds and estimated monocular depth [276], measuring scene-level depth consistency.	[259]
⚖️ Generation Quality - Temporal Consistency				
CTC	↑	CLIP Temporal Consistency	Cosine similarity of CLIP features [191] across consecutive frames, measuring temporal stability and smoothness in generated video sequences.	[104]
DTC	↑	DINO Temporal Consistency	Cosine similarity of DINO features [181] across adjacent frames, evaluating temporal coherence and smooth transitions in generated sequences.	[104]
TTCE	↓	Temporal Transformation Consistency Error	Registration error between temporally generated and ground-truth point clouds, evaluating motion alignment and temporal consistency.	[139]
CTC	↓	Chamfer Temporal Consistency	Chamfer distance between generated point clouds across different timestamps, quantifying temporal stability and geometric coherence.	[139]
ICP	↓	ICP Energy / Outlier	Registration residuals and outlier ratios from Iterative Closest Point alignment of LiDAR frames, detecting temporal jitter and misalignment.	[325]
⚖️ Generation Quality - Subject Consistency				
SC	↑	Subject Consistency	Cosine similarity of subject-region features [181] across frames, evaluating identity persistence and stability in generated video sequences.	[98]
FDC	↑	Foreground Detection Confidence	Confidence scores of detected foreground objects in generated samples using a pretrained detector, reflecting semantic plausibility and realism.	[139]
CFCA	↑	Conditional Foreground Classification Accuracy	Semantic consistency of generated objects evaluated by classification accuracy with a pretrained object classifier, conditioned on ground truth.	[139]
CFSC	↑	Conditional Foreground Spatial Consistency	Mean IoU between 3D boxes regressed by a conditional VAE and ground truth boxes, assessing geometric alignment under conditioning.	[139]
⚖️ Generation Quality - Controllability				
CDA	↑	Conditional Detection Accuracy	Standard detection accuracy from a pretrained 3D detector applied to generated point clouds with box conditioning, measuring semantic fidelity.	[139]
CLIP-Sim	↑	CLIP Similarity	Average cosine similarity between CLIP embeddings of generated and reference frames, reflecting semantic alignment across modalities.	[157]
MAE	↓	Mean Absolute Error	Difference in predicted versus reference object counts within scene graphs, assessing accuracy of graph-level controllability.	[150]
JI	↑	Jaccard Index	Overlap ratio of predicted and reference category sets within scene graphs, evaluating graph-level semantic consistency.	[150]
RotErr	↓	Rotation Error	Angular difference between recovered and target camera trajectories, quantifying rotational alignment error.	[107]
TransErr	↓	Translation Error	Euclidean distance between recovered and target camera trajectories, quantifying translational alignment error.	[107]
⚖️ Generation Quality - Human Preference				
VQ	↑	Visual Quality (2AFC)	Win rates for perceptual visual quality in two-alternative forced choice comparisons, reflecting human-preferred realism of generations.	[234]

Note: Continued on next page

Abbr.	-	Full Name	Description	Ref.
MR	↑	Motion Rationality (2AFC)	Win rates for perceived motion rationality in two-alternative forced choice settings, evaluating naturalness of temporal dynamics.	[234]
DMOS	↑	Differential Mean Opinion Score	Average human-rated alignment with conditioning constraints (e.g., scene graphs), providing a relative perceptual quality measure.	[150]
⚖️ Forecasting Quality - Spatial Predictive Accuracy				
L1 Error	↓	Frame Mean Absolute Error	Pixel- or depth-space L1 distance between predicted and ground-truth frames, quantifying reconstruction fidelity and short-horizon accuracy.	[253]
L2 Error	↓	Frame Mean Squared Error	Pixel- or depth-space L2 distance between predicted and ground-truth frames, reflecting average squared reconstruction deviation.	[253]
IoU_c	↑	IoU at Current Timestamp	Intersection-over-Union between predicted and reference occupancy maps at the current frame, assessing immediate prediction quality.	[158]
IoU_f	↑	IoU at Future Timestamp	Intersection-over-Union between predicted and reference occupancy maps at a fixed future horizon, capturing long-range prediction quality.	[158]
IoU_{wf}	↑	IoU at Weighted Future Timestamp	Weighted average Intersection-over-Union across multiple future frames, emphasizing near-term predictions for smoother accuracy assessment.	[158]
CD	↓	Chamfer Distance	Bidirectional nearest-neighbor distance between point clouds from ray-cast predicted and ground-truth occupancy, measuring geometric fidelity.	[296]
L1 Med	↓	Median L1 Depth Error	Median absolute depth error along LiDAR rays after projection, robustly quantifying accuracy against outliers.	[296]
AbsRel Med	↓	Median Absolute Relative Error	Median of relative depth errors across all LiDAR rays, providing a scale-aware and robust measure of accuracy.	[296]
L1 Mean	↓	Mean L1 Depth Error	Mean absolute depth error along projected LiDAR rays, reflecting average deviation in meters from reference.	[296]
AbsRel Mean	↓	Mean Absolute Relative Error	Mean of relative depth errors across all rays, capturing overall scale-consistent accuracy of depth predictions.	[296]
⚖️ Forecasting Quality - Temporal Predictive Accuracy				
KODP	↑	Key Object Dimension Probability	Probability-based measure penalizing implausible object dimensions using category priors, encouraging physically realistic and consistent generation.	[238]
TFSC	↑	Temporal Foreground Shape Consistency	Voxel-level Intersection-over-Union of dynamic object instances across consecutive frames, ensuring shape persistence and temporal stability.	[238]
TBEC	↑	Temporal Background Environment Consistency	Consistency of static voxels under ego-motion compensation, validating environmental rigidity and long-term background stability.	[238]
⚖️ Planning Quality - Open-Loop Planning				
ADE	↓	Average Displacement Error	Mean displacement error between predicted trajectories and expert waypoints across the horizon, reflecting overall trajectory accuracy.	[274]
FDE	↓	Final Displacement Error	Displacement error at the final predicted waypoint compared with expert trajectories, emphasizing long-term accuracy.	[274]
SLE	↓	Speed L1 Error	Mean absolute error of predicted speed control signals.	[102]
SALE	↓	Steer Angle L1 Error	Mean absolute error of predicted steering angle control signals.	[102]
CR	↓	Collision rate	Fraction of rollouts in which the controlled vehicle collides with surrounding agents or obstacles, indicating safety risk.	[239]

Note: Continued on next page

Abbr.	-	Full Name	Description	Ref.
PDMS	↑	Predictive Driver Model Score	Aggregate score combining progress, spacing, and comfort after discarding unsafe rollouts, approximating human-like driving quality.	[41]
EPDMS	↑	Extended Predictive Driver Model Score	Extended version of PDMS that includes nine additional factors to reflect rule adherence and recovery behaviors.	[28]
AHE	↓	Average Heading Error	Mean absolute angular deviation between predicted and expert heading over the trajectory horizon, measuring orientation accuracy.	[25]
FHE	↓	Final Heading Error	Absolute angular deviation of predicted heading from expert at the final timestep, reflecting terminal orientation accuracy.	[25]
MR	↓	Miss Rate	Fraction of prediction timesteps where displacement error exceeds horizon-specific thresholds, reflecting failure in trajectory coverage.	[25]

Planning Quality - Closed-Loop Planning

SR	↑	Success Rate	Percentage of navigation episodes that successfully reach the goal within a fixed time budget, indicating overall task completion.	[48]
ID	↑	Infraction Distance	Average driving distance between two infractions, with longer distances reflecting safer and more reliable policy behavior.	[48]
ADS	↑	Arena Driving Score	Composite score combining route completion metrics with PDMS to summarize closed-loop driving performance in Arena environments.	[278]
NAC	↑	No At-Fault Collisions	Fraction of scenarios without ego-fault collisions, focusing exclusively on responsibility-aware collision evaluation.	[25]
DAC	↑	Drivable-Area Compliance	Boolean evaluation that checks whether the ego vehicle remains inside drivable polygons throughout the rollout.	[25]
DDC	↑	Driving-Direction Compliance	Boolean evaluation verifying that ego motion remains aligned with the designated lane’s legal driving direction.	[25]
MP	↑	Making Progress	Boolean check confirming that the ego vehicle makes sufficient forward route progress within the evaluation horizon.	[25]
TTC	↑	Time-to-Collision	Boolean verification that the time-to-collision value exceeds safety thresholds, preventing imminent crashes.	[25]
PAR	↑	Progress Along Route	Ratio of ego-vehicle progress compared to expert trajectory progress along the same route, reflecting efficiency.	[25]
SLC	↑	Speed-Limit Compliance	Score penalizing magnitude and duration of speed-limit violations, higher values indicating safer speed adherence.	[25]
Comfort	↑	Driving Comfort	Penalization of excessive jerk, acceleration, or yaw-rate, reflecting ride quality and passenger comfort.	[25]

Reconstruction Quality - Photometric Fidelity

PSNR	↑	Peak Signal-to-Noise Ratio	Logarithmic ratio of maximum possible signal power to reconstruction error, measuring pixel-level fidelity of generated images	[100]
SSIM	↑	Structural Similarity Index Measure	Quality index considering structural similarity, luminance consistency, and contrast preservation between generated and reference images.	[241]
LPIPS	↓	Learned Perceptual Image Patch Similarity	Feature-space distance between deep network activations of image patches, quantifying perceptual realism beyond pixel fidelity.	[297]

Reconstruction Quality - View Changing Consistency

NTA-IoU	↑	Novel Trajectory Agent IoU	Intersection-over-Union between projected 3D bounding boxes of foreground agents and detected 2D boxes under novel viewpoints.	[304]
NTL-IoU	↑	Novel Trajectory Lane IoU	Intersection-over-Union between projected lane structures and detected lane markings from novel viewpoints, evaluating background alignment.	[304]

Downstream Evaluation - Detection

Note: Continued on next page

Abbr.	-	Full Name	Description	Ref.
mAP	↑	Mean Average Precision	Average precision computed over multiple IoU thresholds for 2D detection boxes on standard benchmarks, reflecting detection accuracy.	[145]
mAP-3D	↑	Mean Average Precision in 3D	Average precision for 3D bounding boxes, integrating precision-recall across multiple IoU thresholds in 3D space.	[24]
LET-3D-AP	↑	Average Precision with Longitudinal Error Tolerance	3D average precision using LET-IoU, allowing depth shifts along the camera ray within a tolerance margin.	[99]
LET-3D-APL	↑	Longitudinal Affinity Weighted LET-3D-AP	Weighted LET-3D-AP that penalizes larger longitudinal corrections, improving realism in depth-sensitive evaluation.	[99]
mATE	↓	Mean Average Translation Error	Mean L2 distance between predicted and ground-truth object centers for matched true positives, evaluating localization accuracy.	[24]
mASE	↓	Mean Average Scale Error	Mean scale discrepancy defined as one minus IoU for matched true positives, reflecting object size accuracy.	[24]
mAOE	↓	Mean Average Orientation Error	Mean absolute yaw error of matched true positives, quantifying accuracy of predicted object orientation.	[24]
mAVE	↓	Mean Average Velocity Error	Mean L2 difference between predicted and ground-truth object velocities for matched true positives, measuring motion accuracy.	[24]
mAAE	↓	Mean Average Attribute Error	Mean error in predicting semantic attributes for matched true positives, computed as one minus attribute accuracy.	[24]
NDS	↑	nuScenes Detection Score	Composite metric combining mAP with normalized true positive errors (mATE, mASE, mAOE, mAVE, mAAE), reflecting holistic detection quality.	[24]
⚖️ Downstream Evaluation - Segmentation				
mIoU	↑	Mean Intersection over Union	Average Intersection-over-Union across all semantic classes in 2D images or 3D point clouds, measuring segmentation accuracy.	[52]
BEV-Map-IoU	↑	Bird’s-Eye-View Map IoU	Class-wise Intersection-over-Union for freespace, lane, and dynamic agents in BEV compared with HD maps, evaluating scene consistency.	[188]
⚖️ Downstream Evaluation - Tracking				
MOTA	↑	Multi-Object Tracking Accuracy	Composite metric aggregating false positives, false negatives, and identity switches into a single score, evaluating overall tracking reliability.	[15]
MOTP	↑	Multi-Object Tracking Precision	Mean Intersection-over-Union between matched predictions and ground truth across frames, quantifying spatial localization precision.	[15]
3D-AMOTA	↑	Average Multi-Object Tracking Accuracy in 3D	3D extension of MOTA averaged across recall thresholds, reducing sensitivity to threshold choices while evaluating tracking accuracy.	[246]
3D-AMOTP	↑	Average Multi-Object Tracking Precision in 3D	3D extension of MOTP averaging localization precision (IoU or center distance) over recall thresholds, reflecting robustness.	[246]
⚖️ Downstream Evaluation - Occupancy Prediction				
Occupancy-IoU	↑	Occupancy Intersection over Union	Intersection-over-Union between predicted and labeled voxel occupancies at class- or scene-level granularity, reflecting occupancy accuracy.	[208]
VPQ	↑	Voxelized Panoptic Quality	Panoptic quality metric for voxelized outputs, combining semantic segmentation accuracy with instance detection recall into a unified score.	[158]
⚖️ Downstream Evaluation - VQA				
Top-1 Acc	↑	Visual Question Answering Top-1 Accuracy	Exact-match accuracy of predicted answers across diverse question categories in autonomous driving VQA tasks, measuring reasoning reliability.	[242]

References

- [1] Niket Agarwal et al. Cosmos world foundation model platform for physical AI. *arXiv preprint arXiv:2501.03575*, 2025.
- [2] Ben Agro, Quinlan Sykora, Sergio Casas, et al. UnO: Unsupervised occupancy fields for perception and forecasting. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 14487–14496, 2024.
- [3] Naveed Ahmed et al. A systemic survey of the omniverse platform and its applications in data generation, simulation and metaverse. *Frontiers Computer Sci.*, 6:1423129, 2024.
- [4] Hassan Abu Alhaija et al. Cosmos-Transfer1: Conditional world generation with adaptive multimodal control. *arXiv preprint arXiv:2503.14492*, 2025.
- [5] Mahmoud Assran, Adrien Bardes, David Fan, et al. V-JEPA 2: Self-supervised video models enable understanding, prediction and planning. *arXiv preprint arXiv:2506.09985*, 2025.
- [6] Benjamin Attal, Jia-Bin Huang, Christian Richardt, et al. HyperReel: High-fidelity 6-DoF video with ray-conditioned sampling. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 16610–16620, 2023.
- [7] Jacob Austin, Daniel D Johnson, Jonathan Ho, Daniel Tarlow, and Rianne Van Den Berg. Structured denoising diffusion models in discrete state-spaces. *Adv. Neural Inf. Process. Syst.*, 34:17981–17993, 2021.
- [8] Alisson Azzolini et al. Cosmos-Reason1: From physical common sense to embodied reasoning. *arXiv preprint arXiv:2503.15558*, 2025.
- [9] Jianhong Bai et al. RecamMaster: Camera-controlled generative rendering from a single video. *arXiv preprint arXiv:2503.11647*, 2025.
- [10] Philip J. Ball, Jakob Bauer, Frank Belletti, et al. Genie 3: A new frontier for world models. <https://deepmind.google/discover/blog/genie-3-a-new-frontier-for-world-models/>, 2025.
- [11] Adrien Bardes et al. Revisiting feature prediction for learning visual representations from video. *arXiv preprint arXiv:2404.08471*, 2024.
- [12] Florent Bartoccioni, Elias Ramzi, Victor Besnier, et al. VaViM and VaVAM: Autonomous driving through video generative modeling. *arXiv preprint arXiv:2502.15672*, 2025.
- [13] Daniel M Bear et al. Physion: Evaluating physical prediction from vision in humans and machines. *arXiv preprint arXiv:2106.08261*, 2021.
- [14] Jens Behley, Martin Garbade, Andres Milioto, et al. SemanticKITTI: A dataset for semantic scene understanding of LiDAR sequences. In *IEEE/CVF Int. Conf. Comput. Vis.*, pages 9297–9307, 2019.
- [15] Keni Bernardin, Alexander Elbs, and Rainer Stiefelhagen. Multiple object tracking performance metrics and evaluation in a smart room environment. In *Eur. Conf. Comput. Vis. Worksh.*, volume 90, 2006.
- [16] Kaifeng Bi, Lingxi Xie, Hengheng Zhang, et al. Accurate medium-range global weather forecasting with 3D neural networks. *Nature*, 619(7970):533–538, 2023.
- [17] Hengwei Bian et al. DynamicCity: Large-scale 4D occupancy generation from dynamic scenes. In *Int. Conf. Learn. Represent.*, 2025.
- [18] Mario Bijelic et al. Seeing through fog without seeing fog: Deep multimodal sensor fusion in unseen adverse weather. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 11682–11692, 2020.
- [19] Mikołaj Bińkowski et al. Demystifying MMD GANs. *arXiv preprint arXiv:1801.01401*, 2018.
- [20] Andreas Blattmann, Tim Dockhorn, Sumith Kulal, et al. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*, 2023.
- [21] Conrad Botton et al. 4D simulation research in construction: A systematic mapping study. *Archives of Computational Methods in Engineering*, 30(4):2451–2472, 2023.
- [22] Jake Bruce, Michael D Dennis, Ashley Edwards, et al. Genie: Generative interactive environments. In *Int. Conf. Learn. Represent.*, 2024.
- [23] Yohann Cabon et al. Virtual KITTI 2. *arXiv preprint arXiv:2001.10773*, 2020.

- [24] Holger Caesar, Varun Bankiti, Alex H Lang, et al. nuScenes: A multimodal dataset for autonomous driving. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 11621–11631, 2020.
- [25] Holger Caesar, Juraj Kabzan, Kok Seang Tan, et al. nuPlan: A closed-loop ML-based planning benchmark for autonomous vehicles. *arXiv preprint arXiv:2106.11810*, 2021.
- [26] Anh-Quan Cao and Raoul De Charette. MonoScene: Monocular 3D semantic scene completion. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 3991–4001, 2022.
- [27] Helin Cao and Sven Behnke. DiffSSC: Semantic LiDAR scan completion using denoising diffusion probabilistic models. *arXiv preprint arXiv:2409.18092*, 2024.
- [28] Wei Cao et al. Pseudo-simulation for autonomous driving. *arXiv preprint arXiv:2506.04218*, 2025.
- [29] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 6299–6308, 2017.
- [30] Huiwen Chang, Han Zhang, Lu Jiang, et al. MaskGIT: Masked generative image transformer. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 11315–11325, 2022.
- [31] Anthony Chen, Wenzhao Zheng, Yida Wang, et al. GeoDrive: 3D geometry-informed driving world model with precise action control. *arXiv preprint arXiv:2505.22421*, 2025.
- [32] Haoxuan Chen et al. Accelerating diffusion models with parallel sampling: Inference at sub-linear time complexity. In *Adv. Neural Inf. Process. Syst.*, volume 37, pages 133661–133709, 2024.
- [33] Junliang Chen et al. OccProphet: Pushing efficiency frontier of camera-only 4D occupancy forecasting with observer-forecaster-refiner framework. *arXiv preprint arXiv:2502.15180*, 2025.
- [34] Rui Chen et al. UniMLVG: Unified framework for multi-view long video generation with comprehensive control capabilities for autonomous driving. *arXiv preprint arXiv:2412.04842*, 2024.
- [35] Yuntao Chen, Yuqi Wang, and Zhaoxiang Zhang. DrivingGPT: Unifying driving world modeling and planning with multi-modal autoregressive transformers. *arXiv preprint arXiv:2412.18607*, 2024.
- [36] Yurui Chen, Junge Zhang, Ziyang Xie, et al. S-NeRF++: Autonomous driving simulation via neural reconstruction and generation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 47(6):4358–4376, 2025.
- [37] Zhikai Chen et al. AnchorFormer: Point cloud completion from discriminative nodes. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 13581–13590, 2023.
- [38] Ziyu Chen, Jiawei Yang, Jiahui Huang, et al. OmniRe: Omni urban scene reconstruction. In *Int. Conf. Learn. Represent.*, 2025.
- [39] Joseph Cho, Fachrina Dewi Puspitasari, Sheng Zheng, Jingyao Zheng, Lik-Hang Lee, Tae-Ho Kim, Choong Seon Hong, and Chaoning Zhang. Sora as an AGI world model? a complete survey on text-to-video generation. *arXiv preprint arXiv:2403.05131*, 2024.
- [40] Xingyuan Dai, Chen Zhao, Xiao Wang, et al. Image-based traffic signal control via world models. *Frontiers Info. Tech. Electro. Engineer.*, 23(12):1795–1813, 2022.
- [41] Daniel Dauner, Marcel Hallgarten, Tianyu Li, et al. NAVSIM: Data-driven non-reactive autonomous vehicle simulation and benchmarking. In *Adv. Neural Inf. Process. Syst.*, volume 37, pages 28706–28719, 2024.
- [42] Matt Deitke, Eli VanderBilt, Alvaro Herrasti, et al. ProcTHOR: Large-scale embodied AI using procedural generation. In *Adv. Neural Inf. Process. Syst.*, volume 35, pages 5982–5994, 2022.
- [43] Jie Deng et al. CityCraft: A real crafter for 3D city generation. *arXiv preprint arXiv:2406.04983*, 2024.
- [44] Christopher Diehl, Quinlan Sykora, Ben Agro, et al. DIO: Decomposable implicit 4D occupancy-flow world model. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 27456–27466, 2025.
- [45] Jingtao Ding, Yunke Zhang, Yu Shang, et al. Understanding world or predicting future? a comprehensive survey of world models. *ACM Computing Surveys*, 2024.
- [46] Jingtao Ding et al. Understanding world or predicting future? a comprehensive survey of world models. *arXiv preprint arXiv:2411.14499*, 2024.
- [47] Haotian Dong, Xin Wang, Di Lin, et al. NoiseController: Towards consistent multi-view video generation via noise decomposition and collaboration. *arXiv preprint arXiv:2504.18448*, 2025.

- [48] Alexey Dosovitskiy et al. CARLA: An open urban driving simulator. In *Conf. Robot Learn.*, pages 1–16. PMLR, 2017.
- [49] Yi Du et al. SuperPC: A single diffusion model for point cloud completion, upsampling, denoising, and colorization. *arXiv preprint arXiv:2503.14558*, 2025.
- [50] Haoyi Duan et al. WorldScore: A unified evaluation benchmark for world generation. *arXiv preprint arXiv:2504.00983*, 2025.
- [51] Kiana Ehsani, Winson Han, Alvaro Herrasti, et al. Manipulathor: A framework for visual object manipulation. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 4497–4506, 2021.
- [52] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The Pascal visual object classes (VOC) challenge. *Int. J. Comput. Vis.*, 88(2):303–338, 2010.
- [53] Lue Fan, Hao Zhang, Qitai Wang, et al. FreeSim: Toward free-viewpoint camera simulation in driving scenes. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 12004–12014, 2025.
- [54] Ryan Faulkner et al. Simultaneous diffusion sampling for conditional LiDAR generation. *arXiv preprint arXiv:2410.11628*, 2024.
- [55] Tuo Feng, Wenguan Wang, and Yi Yang. A survey of world models for autonomous driving. *arXiv preprint arXiv:2501.11260*, 2025.
- [56] Roya Firoozi, Johnathan Tucker, Stephen Tian, et al. Foundation models in robotics: Applications, challenges, and the future. *Int. J. Robot. Research*, 44(5):701–739, 2025.
- [57] Tobias Fischer, Jonas Kulhanek, Samuel Rota Bulò, et al. Dynamic 3D Gaussian fields for urban areas. In *Adv. Neural Inf. Process. Syst.*, volume 37, pages 80466–80494, 2024.
- [58] Whye Kit Fong, Rohit Mohan, Juana Valeria Hurtado, et al. Panoptic nuScenes: A large-scale benchmark for LiDAR panoptic segmentation and tracking. *IEEE Robot. Autom. Lett.*, 7:3795–3802, 2022.
- [59] Ao Fu et al. Exploring the interplay between video generation and world models in autonomous driving: A survey. *arXiv preprint arXiv:2411.02914*, 2024.
- [60] Pascale Fung, Yoram Bachrach, Asli Celikyilmaz, Kamalika Chaudhuri, DeLong Chen, Willy Chung, Emmanuel Dupoux, Hongyu Gong, Hervé Jégou, Alessandro Lazaric, et al. Embodied AI agents: Modeling the world. *arXiv preprint arXiv:2506.22355*, 2025.
- [61] Chen Gao et al. Dynamic view synthesis from dynamic monocular video. In *IEEE/CVF Int. Conf. Comput. Vis.*, pages 5712–5721, 2021.
- [62] Ruiyuan Gao, Kai Chen, Zhihao Li, et al. MagicDrive3D: Controllable 3D generation for any-view rendering in street scenes. *arXiv preprint arXiv:2405.14475*, 2024.
- [63] Ruiyuan Gao, Kai Chen, Bo Xiao, et al. MagicDrive-V2: High-resolution long video generation for autonomous driving with adaptive control. In *IEEE/CVF Int. Conf. Comput. Vis.*, 2025.
- [64] Ruiyuan Gao et al. MagicDrive: Street view generation with diverse 3D geometry control. In *Int. Conf. Learn. Represent.*, 2023.
- [65] Shenyuan Gao, Jiazhi Yang, Li Chen, et al. Vista: A generalizable driving world model with high fidelity and versatile controllability. In *Adv. Neural Inf. Process. Syst.*, volume 37, pages 91560–91596, 2024.
- [66] Shenyuan Gao et al. AdaWorld: Learning adaptable world models with latent actions. *arXiv preprint arXiv:2503.18938*, 2025.
- [67] Junhao Ge et al. Unraveling the effects of synthetic data on end-to-end autonomous driving. *arXiv preprint arXiv:2503.18108*, 2025.
- [68] Zhiqi Ge et al. WorldGPT: Empowering LLM as multimodal world model. In *ACM Int. Conf. Multimedia*, pages 7346–7355, 2024.
- [69] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the KITTI vision benchmark suite. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 3354–3361, 2012.
- [70] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Comm. of the ACM*, 63(11):139–144, 2020.

- [71] Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*, 2023.
- [72] Songen Gu, Wei Yin, Bu Jin, et al. DOME: Taming diffusion model into high-fidelity controllable occupancy world model. *arXiv preprint arXiv:2410.10429*, 2024.
- [73] Lin Guan et al. Leveraging pre-trained large language models to construct and utilize world models for model-based task planning. In *Adv. Neural Inf. Process. Syst.*, volume 36, pages 79081–79094, 2023.
- [74] Yanchen Guan, Haicheng Liao, Zhenning Li, et al. World models for autonomous driving: An initial survey. *IEEE Trans. Intell. Veh.*, pages 1–17, 2024.
- [75] Erxin Guo et al. FSF-Net: Enhance 4D occupancy forecasting with coarse BEV scene flow for autonomous driving. *arXiv preprint arXiv:2409.15841*, 2024.
- [76] Jiazhe Guo, Yikang Ding, Xiwu Chen, et al. DiST-4D: Disentangled spatiotemporal diffusion with metric depth for 4D driving scene generation. In *IEEE/CVF Int. Conf. Comput. Vis.*, 2025.
- [77] Xi Guo et al. InfinityDrive: Breaking time limits in driving world models. *arXiv preprint arXiv:2412.01522*, 2024.
- [78] Hamed Haghghi et al. Taming transformers for realistic LiDAR point cloud generation. *arXiv preprint arXiv:2404.05505*, 2024.
- [79] Xiangyu Han, Zhen Jia, Boyi Li, et al. Extrapolated urban view synthesis benchmark. *arXiv preprint arXiv:2412.05256*, 2024.
- [80] Mariam Hassan, Sebastian Stapf, Ahmad Rahimi, et al. GEM: A generalizable ego-vision multimodal world model for fine-grained ego-motion, object dynamics, and scene composition control. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 22404–22415, 2025.
- [81] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, et al. GANs trained by a two-time-scale update rule converge to a local Nash equilibrium. *Adv. Neural Inf. Process. Syst.*, 30, 2017.
- [82] Jonathan Ho et al. Denoising diffusion probabilistic models. In *Adv. Neural Inf. Process. Syst.*, volume 33, pages 6840–6851, 2020.
- [83] Shing-Hei Ho, Bao Thach, and Minghan Zhu. LiDAR-EDIT: LiDAR data generation by editing the object layouts in real-world scenes. In *IEEE Int. Conf. Robot. Autom.*, 2025.
- [84] John Houston, Guido Zuidhof, Luca Bergamini, et al. One thousand and one hours: Self-driving motion prediction dataset. In *Conf. Robot Learn.*, pages 409–418. PMLR, 2021.
- [85] Anthony Hu, Gianluca Corrado, Nicolas Griffiths, et al. Model-based imitation learning for urban driving. In *Adv. Neural Inf. Process. Syst.*, volume 35, pages 20703–20716, 2022.
- [86] Anthony Hu et al. GAIA-1: A generative world model for autonomous driving. *arXiv preprint arXiv:2309.17080*, 2023.
- [87] Junjun Hu, Xingyuan Dai, Xiaojun Li, et al. TrafficWise: Leveraging world models for generalized and interpretable traffic control. *IEEE Intell. Transport. Syst. Magazine*, pages 2–12, 2025.
- [88] Mengkang Hu et al. Text2World: Benchmarking large language models for symbolic world model generation. *arXiv preprint arXiv:2502.13092*, 2025.
- [89] Qianjiang Hu et al. RangeLDM: Fast realistic LiDAR point cloud generation. In *Eur. Conf. Comput. Vis.*, pages 115–135. Springer, 2024.
- [90] Shengchao Hu, Li Chen, Penghao Wu, et al. ST-P3: End-to-end vision-based autonomous driving via spatial-temporal feature learning. In *Eur. Conf. Comput. Vis.*, pages 533–549. Springer, 2022.
- [91] Xiaotao Hu, Wei Yin, Mingkai Jia, et al. DrivingWorld: Constructing world model for autonomous driving via video GPT. *arXiv preprint arXiv:2412.19505*, 2024.
- [92] Yihan Hu, Jiazhi Yang, Li Chen, et al. Planning-oriented autonomous driving. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 17853–17862, 2023.
- [93] Binyuan Huang, Yuqing Wen, Yucheng Zhao, et al. SubjectDrive: Scaling generative data in autonomous driving via subject control. In *AAAI Conf. Artif. Intell.*, volume 39, pages 3617–3625, 2025.

- [94] Chenguang Huang et al. Visual language maps for robot navigation. *arXiv preprint arXiv:2210.05714*, 2022.
- [95] Nan Huang et al. S3Gaussian: Self-supervised street Gaussians for autonomous driving. *arXiv preprint arXiv:2405.20323*, 2024.
- [96] Siyuan Huang, Zan Wang, Puhao Li, et al. Diffusion-based generation, optimization, and planning in 3D scenes. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 16750–16761, 2023.
- [97] Tianyu Huang et al. Voyager: Long-range and world-consistent video diffusion for explorable 3D scene generation. *arXiv preprint arXiv:2506.04225*, 2025.
- [98] Ziqi Huang, Yinan He, Jiashuo Yu, et al. VBench: Comprehensive benchmark suite for video generative models. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 21807–21818, 2024.
- [99] Wei-Chih Hung et al. LET-3D-AP: Longitudinal error tolerant 3D average precision for camera-only 3D detection. In *IEEE Int. Conf. Robot. Autom.*, pages 8272–8279, 2024.
- [100] Quan Huynh-Thu and Mohammed Ghanbari. Scope of validity of PSNR in image/video quality assessment. *Electronics Letters*, 44(13):800–801, 2008.
- [101] Yishen Ji, Ziyue Zhu, Zhenxin Zhu, et al. CoGen: 3D consistent video generation via adaptive conditioning for autonomous driving. *arXiv preprint arXiv:2503.22231*, 2025.
- [102] Fan Jia, Weixin Mao, Yingfei Liu, et al. ADriver-I: A general world model for autonomous driving. *arXiv preprint arXiv:2311.13549*, 2023.
- [103] Mingkai Jia et al. MGvQ: Could VQ-VAE beat VAE? a generalizable tokenizer with multi-group quantization. *arXiv preprint arXiv:2507.07997*, 2025.
- [104] Junpeng Jiang et al. DiVE: DiT-based video generation with enhanced control. *arXiv preprint arXiv:2409.01595*, 2024.
- [105] Junzhe Jiang et al. RealEngine: Simulating autonomous driving in realistic context. *arXiv preprint arXiv:2505.16902*, 2025.
- [106] Yunfan Jiang et al. Behavior robot suite: Streamlining real-world whole-body manipulation for everyday household activities. *arXiv preprint arXiv:2503.05652*, 2025.
- [107] Bu Jin, Weize Li, Baihan Yang, et al. PosePilot: Steering camera pose for generative world models with self-supervised depth. *arXiv preprint arXiv:2505.01729*, 2025.
- [108] R Kenny Jones, Theresa Barton, Xianghao Xu, et al. ShapeAssembly: Learning to generate programs for 3D shape structure synthesis. *ACM Trans. Graphics*, 39(6):1–20, 2020.
- [109] Bingyi Kang et al. How far is video generation from world model: A physical law perspective. In *Int. Conf. Mach. Learn.* PMLR, 2025.
- [110] Junjie Ke et al. MUSIQ: Multi-scale image quality transformer. In *IEEE/CVF Int. Conf. Comput. Vis.*, pages 5148–5157, 2021.
- [111] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3D gaussian splatting for real-time radiance field rendering. *ACM Trans. Graphics*, 42(4), 2023.
- [112] Tarasha Khurana, Peiyun Hu, Achal Dave, et al. Differentiable raycasting for self-supervised occupancy forecasting. In *Eur. Conf. Comput. Vis.*, pages 353–369. Springer, 2022.
- [113] Tarasha Khurana et al. Point cloud forecasting as a proxy for 4D occupancy forecasting. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 1116–1124, 2023.
- [114] Boah Kim and Jong Chul Ye. Diffusion deformable model for 4D temporal medical image generation. In *Int. Conf. Medical Image Comput. Computer-Assisted Intervention*, pages 539–548. Springer, 2022.
- [115] Diederik P Kingma, Max Welling, et al. Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [116] Ellington Kirby et al. LOGen: Toward LiDAR object generation by point diffusion. *arXiv preprint arXiv:2412.07385*, 2024.
- [117] Lingdong Kong, Youquan Liu, Runnan Chen, et al. Rethinking range view representation for lidar segmentation. In *IEEE/CVF Int. Conf. Comput. Vis.*, pages 228–240, 2023.

- [118] Lingdong Kong, Youquan Liu, Xin Li, et al. Robo3D: Towards robust and reliable 3D perception against corruptions. In *IEEE/CVF Int. Conf. Comput. Vis.*, pages 19994–20006, 2023.
- [119] Lingdong Kong, Xiang Xu, Jiawei Ren, et al. Multi-modal data-efficient 3D scene understanding for autonomous driving. *IEEE Trans. Pattern Anal. Mach. Intell.*, 47(5):3748–3765, 2025.
- [120] Lingdong Kong et al. LaserMix for semi-supervised LiDAR semantic segmentation. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 21705–21715, 2023.
- [121] Remi Lam, Alvaro Sanchez-Gonzalez, Matthew Willson, et al. Learning skillful medium-range global weather forecasting. *Science*, 382(6677):1416–1421, 2023.
- [122] Bernard Lange and others. Self-supervised multi-future occupancy forecasting for autonomous driving. *arXiv preprint arXiv:2407.21126*, 2024.
- [123] Jumin Lee, Woobin Im, Sebin Lee, and Sung-Eui Yoong. Diffusion probabilistic models for scene-scale 3D categorical data. *arXiv preprint arXiv:2301.00527*, 2023.
- [124] Jumin Lee, Sebin Lee, Changho Jo, et al. SemCity: Semantic scene generation with triplane diffusion. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 28337–28347, 2024.
- [125] Bohan Li, Jiazhe Guo, Hongsi Liu, et al. UniScene: Unified occupancy-centric driving scene generation. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 11971–11981, 2025.
- [126] Bohan Li et al. OccScene: Semantic occupancy-based cross-task mutual learning for 3D scene generation. *arXiv preprint arXiv:2412.11183*, 2024.
- [127] Dacheng Li et al. Worldmodelbench: Judging video generation models as world models. *arXiv preprint arXiv:2502.20694*, 2025.
- [128] Hao Li, Jingfeng Li, Dingwen Zhang, et al. VDG: Vision-only dynamic Gaussian for driving simulation. *arXiv preprint arXiv:2406.18198*, 2024.
- [129] Haoteng Li, Zhao Yang, Zezhong Qian, et al. DualDiff: Dual-branch diffusion model for autonomous driving with semantic fusion. In *IEEE Int. Conf. Robot. Autom.*, 2025.
- [130] Ke Li et al. Immersive neural graphics primitives. *arXiv preprint arXiv:2211.13494*, 2022.
- [131] Leheng Li, Weichao Qiu, Yingjie Cai, et al. SyntheOcc: Synthesize geometric-controlled street view images through 3D semantic mpis. *arXiv preprint arXiv:2410.00337*, 2024.
- [132] Peidong Li and Dixiao Cui. Navigation-guided sparse scene representation for end-to-end autonomous driving. *arXiv preprint arXiv:2409.18341*, 2024.
- [133] Xiang Li et al. Semi-supervised vision-centric 3D occupancy world model for autonomous driving. *arXiv preprint arXiv:2502.07309*, 2025.
- [134] Xiaofan Li, Chenming Wu, Zhao Yang, et al. DriVerse: Navigation world model for driving simulation via multimodal trajectory prompting and motion alignment. *arXiv preprint arXiv:2504.18576*, 2025.
- [135] Xiaofan Li et al. DrivingDiffusion: Layout-guided multi-view driving scenarios video generation with latent diffusion model. In *Eur. Conf. Comput. Vis.*, pages 469–485. Springer, 2024.
- [136] Ye Li et al. Is your LiDAR placement optimized for 3D scene understanding? In *Adv. Neural Inf. Process. Syst.*, volume 37, pages 34980–35017, 2024.
- [137] Yiming Li, Sihang Li, Xinhao Liu, et al. SSCBench: A large-scale 3D semantic scene completion benchmark for autonomous driving. In *IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2024.
- [138] Ao Liang, Lingdong Kong, Dongyue Lu, et al. Perspective-invariant 3d object detection. In *IEEE/CVF Int. Conf. Comput. Vis.*, 2025.
- [139] Ao Liang et al. LiDARCrafter: Dynamic 4D world modeling from LiDAR sequences. *arXiv preprint arXiv:2508.03692*, 2025.
- [140] Dingkan Liang, Dingyuan Zhang, Xin Zhou, et al. Seeing the future, perceiving the future: A unified driving world model for future generation and perception. *arXiv preprint arXiv:2503.13587*, 2025.
- [141] Yiyi Liao, Jun Xie, and Andreas Geiger. KITTI-360: A novel dataset and benchmarks for urban scene understanding in 2D and 3D. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(3):3292–3310, 2022.

- [142] Zhimin Liao, Ping Wei, Ruijie Zhang, et al. I2-World: Intra-inter tokenization for efficient dynamic 4D scene forecasting. *arXiv preprint arXiv:2507.09144*, 2025.
- [143] Liqiang Lin, Yilin Liu, Yue Hu, et al. Capturing, reconstructing, and simulating: the UrbanScene3D dataset. In *Eur. Conf. Comput. Vis.*, pages 93–109. Springer, 2022.
- [144] Minghui Lin et al. Exploring the evolution of physics cognition in video generation: A survey. *arXiv preprint arXiv:2503.21765*, 2025.
- [145] Tsung-Yi Lin et al. Microsoft COCO: Common objects in context. In *Eur. Conf. Comput. Vis.*, pages 740–755. Springer, 2014.
- [146] Yaron Lipman et al. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.
- [147] Youquan Liu, Lingdong Kong, Weidong Yang, et al. Veila: Panoramic LiDAR generation from a monocular RGB image. *arXiv preprint arXiv:2508.03690*, 2025.
- [148] Youquan Liu et al. La La LiDAR: Large-scale layout generation from LiDAR data. *arXiv preprint arXiv:2508.03691*, 2025.
- [149] Yuheng Liu, Xinke Li, Xueting Li, et al. Pyramid diffusion for fine 3D large scene generation. In *Eur. Conf. Comput. Vis.*, pages 71–87. Springer, 2024.
- [150] Yuheng Liu et al. Controllable 3D outdoor scene generation via scene graphs. *arXiv preprint arXiv:2503.07152*, 2025.
- [151] Zhijian Liu, Haotian Tang, Alexander Amini, Xinyu Yang, Huizi Mao, Daniela Rus, and Song Han. BEVFusion: Multi-task multi-sensor fusion with unified bird’s-eye view representation. In *IEEE Int. Conf. Robot. Autom.*, pages 2774–2781, 2023.
- [152] Xiaoxiao Long et al. A survey: Learning embodied intelligence from physical simulators and world models. *arXiv preprint arXiv:2507.00917*, 2025.
- [153] Hannan Lu, Xiaohe Wu, Shudong Wang, et al. Seeing beyond views: Multi-view driving scene video generation with holistic attention. *arXiv preprint arXiv:2412.03520*, 2024.
- [154] Hao Lu et al. DrivingRecon: Large 4D Gaussian reconstruction model for autonomous driving. *arXiv preprint arXiv:2412.09043*, 2024.
- [155] Jiachen Lu, Ze Huang, Zeyu Yang, et al. WoVoGen: World volume-aware diffusion for controllable multi-camera driving scene generation. In *Eur. Conf. Comput. Vis.*, pages 329–345. Springer, 2024.
- [156] Yifan Lu, Xuanchi Ren, Jiawei Yang, et al. InfiniCube: Unbounded and controllable dynamic 3D driving scene generation with world-guided video models. In *IEEE/CVF Int. Conf. Comput. Vis.*, 2025.
- [157] Enhui Ma, Lijun Zhou, Tao Tang, et al. Unleashing generalization of end-to-end autonomous driving with controllable long video generation. *arXiv preprint arXiv:2406.01349*, 2024.
- [158] Junyi Ma et al. Cam4DOcc: Benchmark for camera-only 4D occupancy forecasting in autonomous driving applications. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 21486–21495, 2024.
- [159] Xiao Ma, David Hsu, and Wee Sun Lee. Learning latent graph dynamics for visual manipulation of deformable objects. In *IEEE Int. Conf. Robot. Autom.*, pages 8266–8273, 2022.
- [160] Xinji Mai et al. From efficient multimodal models to world models: A survey. *arXiv preprint arXiv:2407.00118*, 2024.
- [161] Jiageng Mao, Boyi Li, Boris Ivanovic, et al. DreamDrive: Generative 4D scene modeling from street view images. *arXiv preprint arXiv:2501.00601*, 2025.
- [162] Tetiana Martyniuk et al. LiDPM: Rethinking point diffusion for LiDAR scene completion. In *IEEE Intell. Veh. Symp.*, 2025.
- [163] Jianbiao Mei, Tao Hu, Xuemeng Yang, et al. DreamForge: Motion-aware autoregressive video generation for multi-view driving scenes. *arXiv preprint arXiv:2409.04003*, 2024.
- [164] Timothy Merino, Megan Charity, and Julian Togelius. Interactive latent variable evolution for the generation of minecraft structures. In *Int. Conf. Foundation Digital Games*, pages 1–8, 2023.

- [165] Lars Mescheder, Michael Oechsle, Michael Niemeyer, et al. Occupancy networks: Learning 3D reconstruction in function space. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 4460–4470, 2019.
- [166] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. *Comm. of the ACM*, 65(1):99–106, 2021.
- [167] Andres Milioto et al. RangeNet++: Fast and accurate LiDAR semantic segmentation. In *IEEE/RSJ Int. Conf. Intell. Robots Syst.*, pages 4213–4220, 2019.
- [168] Chen Min, Liang Xiao, Dawei Zhao, et al. Multi-camera unified pre-training via 3D scene reconstruction. *IEEE Robot. Autom. Lett.*, 9(4):3243–3250, 2024.
- [169] Chen Min, Dawei Zhao, Liang Xiao, et al. DriveWorld: 4D pre-trained scene understanding via world models for autonomous driving. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 15522–15533, 2024.
- [170] Chen Min et al. UniWorld: Autonomous driving pre-training via world models. *arXiv preprint arXiv:2308.07234*, 2023.
- [171] Kaichun Mo et al. PartNet: A large-scale benchmark for fine-grained and hierarchical part-level 3D object understanding. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 909–918, 2019.
- [172] Sicheng Mo et al. Dreamland: Controllable world creation with simulator and generative models. *arXiv preprint arXiv:2506.08006*, 2025.
- [173] Kazuto Nakashima and Ryo Kurazume. LiDAR data synthesis with denoising diffusion probabilistic models. In *IEEE Int. Conf. Robot. Autom.*, pages 14724–1473, 2024.
- [174] Kazuto Nakashima and Ryo Kurazumeg. Learning to drop points for LiDAR scan synthesis. In *IEEE/RSJ Int. Conf. Intell. Robots Syst.*, pages 222–229, 2021.
- [175] Kazuto Nakashima, Yumi Iwashita, and Ryo Kurazume. Generative range imaging for learning scene priors of 3D LiDAR data. In *IEEE/CVF Winter Conf. Appl. Comput. Vis.*, pages 1256–1266, 2023.
- [176] Kazuto Nakashima et al. Fast LiDAR data generation with rectified flows. In *IEEE Int. Conf. Robot. Autom.*, 2025.
- [177] Chaojun Ni, Guosheng Zhao, Xiaofeng Wang, et al. ReconDreamer: Crafting world models for driving scene reconstruction via online restoration. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 1559–1569, 2025.
- [178] Jingcheng Ni, Yuxin Guo, Yichen Liu, et al. MaskGWM: A generalizable driving world model with video mask reconstruction. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 22381–22391, 2025.
- [179] Lucas Nunes et al. Scaling diffusion models to real-world 3D LiDAR scene completion. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 14770–14780, 2024.
- [180] Lucas Nunes et al. Towards generating realistic 3D semantic training data for autonomous driving. *arXiv preprint arXiv:2503.21449*, 2025.
- [181] Maxime Oquab et al. DINOv2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023.
- [182] Jack Parker-Holder et al. Genie 2: A large-scale foundation world model. <https://deepmind.google/discover/blog/genie-2-a-large-scale-foundation-world-model>, 2024.
- [183] Jaideep Pathak et al. FourcastNet: A global data-driven high-resolution weather model using adaptive fourier neural operators. *arXiv preprint arXiv:2202.11214*, 2022.
- [184] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *IEEE/CVF Int. Conf. Comput. Vis.*, pages 4195–4205, 2023.
- [185] Chensheng Peng, Chengwei Zhang, Yixiao Wang, et al. DeSiRe-GS: 4D street Gaussians for static-dynamic decomposition and surface reconstruction for urban driving scenes. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 6782–6791, 2025.
- [186] Songyou Peng et al. Convolutional occupancy networks. In *Eur. Conf. Comput. Vis.*, pages 523–540. Springer, 2020.
- [187] Jordan Peper et al. Four principles for physically interpretable world models. *arXiv preprint arXiv:2503.02143*, 2025.

- [188] Jonah Philion and Sanja Fidler. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3D. In *Eur. Conf. Comput. Vis.*, pages 194–210. Springer, 2020.
- [189] Ben Poole et al. DreamFusion: Text-to-3D using 2D diffusion. *arXiv preprint arXiv:2209.14988*, 2022.
- [190] Charles R Qi et al. PointNet: Deep learning on point sets for 3D classification and segmentation. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 652–660, 2017.
- [191] Alec Radford, Jong Wook Kim, Chris Hallacy, et al. Learning transferable visual models from natural language supervision. In *Int. Conf. Mach. Learn.*, pages 8748–8763. PMLR, 2021.
- [192] Haoxi Ran, Vitor Guizilini, and Yue Wang. Towards realistic scene generation with LiDAR diffusion models. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 14738–14748, 2024.
- [193] Xuanchi Ren, Jiahui Huang, Xiaohui Zeng, et al. XCube: Large-scale 3D generative modeling using sparse voxel hierarchies. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 4209–4219, 2024.
- [194] Xuanchi Ren, Yifan Lu, Tianshi Cao, et al. Cosmos-Drive-Dreams: Scalable synthetic driving data generation with world foundation models. *arXiv preprint arXiv:2506.09042*, 2025.
- [195] Yuan Ren et al. UniGaussian: Driving scene reconstruction from multiple camera models via unified Gaussian representations. *arXiv preprint arXiv:2411.15355*, 2024.
- [196] Carlos Riquelme, Joan Puigcerver, Basil Mustafa, et al. Scaling vision with sparse mixture of experts. In *Adv. Neural Inf. Process. Syst.*, volume 34, pages 8583–8595, 2021.
- [197] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 10684–10695, 2022.
- [198] Lloyd Russell, Anthony Hu, Lorenzo Bertoni, et al. GAIA-2: A controllable multi-view generative world model for autonomous driving. *arXiv preprint arXiv:2503.20523*, 2025.
- [199] Tim Salimans et al. Improved techniques for training GANs. *Adv. Neural Inf. Process. Syst.*, 29, 2016.
- [200] SenseTime-FVG. Open driving world models (OpenDWM). <https://github.com/SenseTime-FVG/OpenDWM>, 2025.
- [201] Yu Shang et al. UrbanWorld: An urban world model for 3D city generation. *arXiv preprint arXiv:2407.11965*, 2024.
- [202] Bokui Shen, Fei Xia, Chengshu Li, et al. iGibson 1.0: A simulation environment for interactive tasks in large realistic scenes. In *IEEE/RSJ Int. Conf. Intell. Robots Syst.*, pages 7520–7527, 2021.
- [203] Chen Shi et al. DriveX: Omni scene modeling for learning generalizable world knowledge in autonomous driving. *arXiv preprint arXiv:2505.19239*, 2025.
- [204] Yining Shi, Kun Jiang, Qiang Meng, et al. COME: Adding scene-centric forecasting control to occupancy world model. *arXiv preprint arXiv:2506.13260*, 2025.
- [205] Dong Wook Shu, Sung Woo Park, and Junseok Kwon. 3D point cloud generative adversarial network based on tree structured graph convolutions. In *IEEE/CVF Int. Conf. Comput. Vis.*, pages 3859–3868, 2019.
- [206] Nathan Silberman et al. Indoor segmentation and support inference from RGBD images. In *Eur. Conf. Comput. Vis.*, pages 746–760. Springer, 2012.
- [207] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.
- [208] Shuran Song, Fisher Yu, Andy Zeng, et al. Semantic scene completion from a single depth image. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 1746–1754, 2017.
- [209] Jiaming Sun, Zehong Shen, Yuang Wang, et al. LoFTR: Detector-free local feature matching with transformers. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 8922–8931, 2021.
- [210] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 2446–2454, 2020.
- [211] Alexander Swerdlow, Runsheng Xu, and Bolei Zhou. Street-view image generation from a bird’s-eye view layout. *IEEE Robot. Autom. Lett.*, 9(4):3578–3585, 2024.

- [212] Andrew Szot, Alexander Clegg, Eric Undersander, et al. Habitat 2.0: Training home assistants to rearrange their habitat. In *Adv. Neural Inf. Process. Syst.*, volume 34, pages 251–266, 2021.
- [213] Shuhan Tan, John Lambert, Hong Jeon, et al. SceneDiffuser++: City-scale traffic simulation via a generative world model. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 1570–1580, 2025.
- [214] Zheng Tang, Milind Naphade, Ming-Yu Liu, et al. CityFlow: A city-scale benchmark for multi-target multi-camera vehicle tracking and re-identification. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 8797–8806, 2019.
- [215] Keyu Tian et al. Visual autoregressive modeling: Scalable image generation via next-scale prediction. In *Adv. Neural Inf. Process. Syst.*, volume 37, pages 84839–84865, 2024.
- [216] Xiaoyu Tian, Tao Jiang, Longfei Yun, et al. Occ3D: A large-scale 3D occupancy prediction benchmark for autonomous driving. In *Adv. Neural Inf. Process. Syst.*, volume 36, pages 64318–64330, 2023.
- [217] Wenwen Tong et al. Scene as occupancy. In *IEEE/CVF Int. Conf. Comput. Vis.*, pages 8406–8415, 2023.
- [218] Sifan Tu et al. The role of world models in shaping autonomous driving: A comprehensive survey. *arXiv preprint arXiv:2502.10498*, 2025.
- [219] Christina Ourania Tze, Daniel Dauner, Yiyi Liao, et al. PrITTI: Primitive-based generation of controllable and editable 3D semantic scenes. *arXiv preprint arXiv:2506.19117*, 2025.
- [220] Thomas Unterthiner et al. Towards accurate generative models of video: A new metric & challenges. *arXiv preprint arXiv:1812.01717*, 2018.
- [221] Aaron Van Den Oord et al. Neural discrete representation learning. In *Adv. Neural Inf. Process. Syst.*, volume 30, pages 6309–6318, 2017.
- [222] Ashish Vaswani, Noam Shazeer, Niki Parmar, et al. Attention is all you need. In *Adv. Neural Inf. Process. Syst.*, volume 30, pages 6000–6010, 2017.
- [223] Styliani Verykokou et al. 3D reconstruction of disaster scenes for urban search and rescue. *Multimedia Tools Appl.*, 77(8):9691–9717, 2018.
- [224] Bram Wallace, Meihua Dang, Rafael Rafailov, et al. Diffusion model alignment using direct preference optimization. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 8228–8238, 2024.
- [225] Fei-Yue Wang. New control paradigm for industry 5.0: From big models to foundation control and management. *IEEE/CAA J. Autom. Sinica*, 10(8):1643–1646, 2023.
- [226] Haiguang Wang, Daqi Liu, Hongwei Xie, et al. MiLA: Multi-view intensive-fidelity long-term video generation world model for autonomous driving. *arXiv preprint arXiv:2503.15875*, 2025.
- [227] Hang Wang, Xin Ye, Feng Tao, et al. AdaWM: Adaptive world model-based planning for autonomous driving. *arXiv preprint arXiv:2501.13072*, 2025.
- [228] Jiamin Wang, Yichen Yao, Xiang Feng, et al. STAGE: A stream-centric generative world model for long-horizon driving-scene simulation. *arXiv preprint arXiv:2506.13138*, 2025.
- [229] Lening Wang, Wenzhao Zheng, Yilong Ren, et al. OccSora: 4D occupancy generation models as world simulators for autonomous driving. *arXiv preprint arXiv:2405.20337*, 2024.
- [230] Lening Wang et al. Stag-1: Towards realistic 4D driving simulation with video generation model. *arXiv preprint arXiv:2412.05280*, 2024.
- [231] Shihao Wang, Yingfei Liu, Tiancai Wang, et al. Exploring object-centric temporal modeling for efficient multi-view 3D object detection. In *IEEE/CVF Int. Conf. Comput. Vis.*, pages 3621–3631, 2023.
- [232] Shihao Wang, Zhiding Yu, Xiaohui Jiang, et al. OmniDrive: A holistic LLM-agent framework for autonomous driving with 3D perception, reasoning, and planning. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 22442–22452, 2025.
- [233] Xiaodong Wang and Peixi Peng. ProphetDWM: A driving world model for rolling out future actions and videos. *arXiv preprint arXiv:2505.18650*, 2025.
- [234] Xiaodong Wang, Zhirong Wu, and Peixi Peng. LongDWM: Cross-granularity distillation for building a long-term driving world model. *arXiv preprint arXiv:2506.01546*, 2025.

- [235] Xiaofeng Wang, Zheng Zhu, Wenbo Xu, et al. OpenOccupancy: A large-scale benchmark for surrounding semantic occupancy perception. In *IEEE/CVF Int. Conf. Comput. Vis.*, pages 17850–17859, 2023.
- [236] Xiaofeng Wang, Zheng Zhu, Guan Huang, et al. DriveDreamer: Towards real-world-drive world models for autonomous driving. In *Eur. Conf. Comput. Vis.*, pages 55–72. Springer, 2024.
- [237] Yuping Wang et al. Generative ai for autonomous driving: Frontiers and opportunities. *arXiv preprint arXiv:2505.08854*, 2025.
- [238] Yuping Wang et al. UniOcc: A unified benchmark for occupancy forecasting and prediction in autonomous driving. *arXiv preprint arXiv:2503.24381*, 2025.
- [239] Yuqi Wang et al. Driving into the future: Multiview visual forecasting and planning with world model for autonomous driving. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 14749–14759, 2024.
- [240] Yuqi Wang et al. DrivingDojo dataset: Advancing interactive and knowledge-enriched driving world model. *arXiv preprint arXiv:2410.10738*, 2024.
- [241] Zhou Wang et al. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.*, 13(4):600–612, 2004.
- [242] Julong Wei, Shanshuai Yuan, Pengfei Li, et al. OccLLaMA: An occupancy-language-action generative world model for autonomous driving. *arXiv preprint arXiv:2409.03272*, 2024.
- [243] Beichen Wen et al. 3d scene generation: A survey. *arXiv preprint arXiv:2505.05474*, 2025.
- [244] Yuqing Wen, Yucheng Zhao, Yingfei Liu, et al. Panacea: Panoramic and controllable video generation for autonomous driving. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 6902–6912, 2024.
- [245] Yuqing Wen, Yucheng Zhao, Yingfei Liu, et al. Panacea+: Panoramic and controllable video generation for autonomous driving. *arXiv preprint arXiv:2408.07605*, 2024.
- [246] Xinshuo Weng et al. 3D multi-object tracking: A baseline and new evaluation metrics. In *IEEE/RSJ Int. Conf. Intell. Robots Syst.*, pages 10359–10366, 2020.
- [247] Benjamin Wilson, William Qi, Tanmay Agarwal, et al. Argoverse 2: Next generation datasets for self-driving perception and forecasting. In *Adv. Neural Inf. Process. Syst.*, volume 34, 2021.
- [248] Joey Wilson, Jingyu Song, Yuewei Fu, et al. MotionSC: Data set and network for real-time semantic mapping in dynamic environments. *IEEE Robot. Autom. Lett.*, 7(3):8439–8446, 2022.
- [249] Wei Wu, Xi Guo, Weixuan Tang, et al. DriveScape: Towards high-resolution controllable multi-view driving video generation. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 17187–17196, 2025.
- [250] Yang Wu et al. Text2LiDAR: Text-fuied LiDAR point cloud generation via equirectangular transformer. In *Eur. Conf. Comput. Vis.*, pages 291–310. Springer, 2024.
- [251] Yang Wu et al. WeatherGen: A unified diverse weather generator for LiDAR point clouds via spider mamba diffusion. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 17019–17028, 2025.
- [252] Yanhao Wu, Haoyang Zhang, Tianwei Lin, et al. Generating multimodal driving scenes via next-scene prediction. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 6844–6853, 2025.
- [253] Zehuan Wu, Jingcheng Ni, Xiaodong Wang, et al. HoloDrive: Holistic 2D-3D multi-modal street scene generation for autonomous driving. *arXiv preprint arXiv:2412.01407*, 2024.
- [254] Fei Xia et al. Gibson env: real-world perception for embodied agents. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 9068–9079, 2018.
- [255] Pengchuan Xiao, Zhenlei Shao, Steven Hao, et al. PandaSet: Advanced sensor suite dataset for autonomous driving. In *IEEE Int. Conf. Intell. Transport. Syst.*, pages 3095–3101, 2021.
- [256] Bin Xie, Yingfei Liu, Tiancai Wang, et al. Glad: A streaming scene generator for autonomous driving. In *Int. Conf. Learn. Represent.*, 2025.
- [257] Haozhe Xie et al. Generative gaussian splatting for unbounded 3D city generation. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 6111–6120, 2025.
- [258] Ningwei Xie et al. From 2D to 3D cognition: A brief survey of general world models. *arXiv preprint arXiv:2506.20134*, 2025.

- [259] Yichen Xie, Chenfeng Xu, Chensheng Peng, et al. X-Drive: Cross-modality consistent multi-sensor data synthesis for driving scenarios. In *Int. Conf. Learn. Represent.*, 2025.
- [260] Yuwen Xiong et al. UltraLiDAR: Learning compact representations for LiDAR completion and generation. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 1074–1083, 2023.
- [261] Haoran Xu et al. Temporal triplane transformers as occupancy world models. *arXiv preprint arXiv:2503.07338*, 2025.
- [262] Huaiyuan Xu, Junliang Chen, Shiyu Meng, et al. A survey on occupancy perception for autonomous driving: The information fusion perspective. *Information Fusion*, 114:102671, 2025.
- [263] Jingyi Xu, Xieyuanli Chen, Junyi Ma, et al. Spatiotemporal decoupling for efficient vision-based occupancy forecasting. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 22338–22347, 2025.
- [264] Runsheng Xu et al. OPV2V: An open benchmark dataset and fusion pipeline for perception with vehicle-to-vehicle communication. In *IEEE Int. Conf. Robot. Autom.*, pages 2583–2589, 2022.
- [265] Tianshuo Xu, Hao Lu, Xu Yan, et al. Occ-LLM: Enhancing autonomous driving with occupancy-based large language models. In *IEEE Int. Conf. Robot. Autom.*, 2025.
- [266] Xiang Xu et al. 4D contrastive superflows are dense 3D representation learners. In *Eur. Conf. Comput. Vis.*, pages 58–80. Springer, 2024.
- [267] Zhiyuan Xu, Bohan Li, Huanang Gao, et al. Challenger: Affordable adversarial driving video generation. *arXiv preprint arXiv:2505.15880*, 2025.
- [268] Tianyi Yan, Dongming Wu, Wencheng Han, et al. DrivingSphere: Building a high-fidelity 4D world for closed-loop simulation. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 27531–27541, 2025.
- [269] Tianyi Yan et al. OLiDM: Object-aware LiDAR diffusion models for autonomous driving. In *AAAI Conf. Artif. Intell.*, volume 39, pages 9121–9129, 2025.
- [270] Xu Yan et al. Forging vision foundation models for autonomous driving: Challenges, methodologies, and opportunities. *arXiv preprint arXiv:2401.08045*, 2024.
- [271] Yunzhi Yan, Haotong Lin, Chenxu Zhou, et al. Street Gaussians: Modeling dynamic urban scenes with gaussian splatting. In *Eur. Conf. Comput. Vis.*, pages 156–173. Springer, 2024.
- [272] Yunzhi Yan, Zhen Xu, Haotong Lin, et al. StreetCrafter: Street view synthesis with controllable video diffusion models. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 822–832, 2025.
- [273] Ziyang Yan, Wenzhen Dong, Yihua Shao, et al. RenderWorld: World model with self-supervised 3D label. *arXiv preprint arXiv:2409.11356*, 2024.
- [274] Jiazhi Yang, Shenyuan Gao, Yihang Qiu, et al. Generalized predictive model for autonomous driving. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 14662–14672, 2024.
- [275] Kairui Yang, Enhui Ma, Jibin Peng, et al. BEVControl: Accurately controlling street-view elements with multi-perspective consistency via BEV sketch layout. *arXiv preprint arXiv:2308.01661*, 2023.
- [276] Lihe Yang, Bingyi Kang, Zilong Huang, et al. Depth anything v2. In *Adv. Neural Inf. Process. Syst.*, volume 37, pages 21875–21911, 2024.
- [277] Xianghui Yang et al. Hunyuan3D 1.0: A unified framework for text-to-3D and image-to-3D generation. *arXiv preprint arXiv:2411.02293*, 2024.
- [278] Xuemeng Yang, Licheng Wen, Yukai Ma, et al. DriveArena: A closed-loop generative simulation platform for autonomous driving. In *IEEE/CVF Int. Conf. Comput. Vis.*, 2025.
- [279] Yijun Yang et al. Medical world model: Generative simulation of tumor evolution for treatment planning. *arXiv preprint arXiv:2506.02327*, 2025.
- [280] Yu Yang, Alan Liang, Jianbiao Mei, et al. X-Scene: Large-scale driving scene generation with high fidelity and flexible controllability. *arXiv preprint arXiv:2506.13558*, 2025.
- [281] Yu Yang et al. Driving in the occupancy world: Vision-centric 4D occupancy forecasting and planning via world models for autonomous driving. In *AAAI Conf. Artif. Intell.*, volume 39, pages 9327–9335, 2025.

- [282] Yue Yang, Fan-Yun Sun, Luca Weihs, et al. HoloDeck: Language guided generation of 3D embodied AI environments. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 16227–16237, 2024.
- [283] Ze Yang et al. UniSim: A neural closed-loop sensor simulator. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 1389–1399, 2023.
- [284] Zetong Yang et al. Visual point cloud forecasting enables scalable autonomous driving. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 14673–14684, 2024.
- [285] Zhao Yang, Zezhong Qian, Xiaofan Li, et al. DualDiff+: Dual-branch diffusion for high-fidelity video generation with reward guidance. *arXiv preprint arXiv:2503.03689*, 2025.
- [286] Zhuoran Yang, Xi Guo, Chenjing Ding, et al. Physical informed driving world model. *arXiv preprint arXiv:2412.08410*, 2024.
- [287] Lance Ying et al. Assessing adaptive world models in machines with novel games. *arXiv preprint arXiv:2507.12821*, 2025.
- [288] Jiwen Yu et al. GameFactory: Creating new games with generative interactive videos. *arXiv preprint arXiv:2501.08325*, 2025.
- [289] Mark Yu et al. TrajectoryCrafter: Redirecting camera trajectory for monocular videos via diffusion models. *arXiv preprint arXiv:2503.05638*, 2025.
- [290] Zhongrui Yu, Haoran Wang, Jinze Yang, et al. SGD: Street view synthesis with Gaussian splatting and diffusion prior. In *IEEE/CVF Winter Conf. Appl. Comput. Vis.*, pages 3812–3822, 2025.
- [291] Zikang Yuan et al. Uni-Gaussians: Unifying camera and LiDAR simulation with Gaussians for dynamic driving scenarios. *arXiv preprint arXiv:2503.08317*, 2025.
- [292] Haiming Zhang, Ying Xue, Xu Yan, et al. An efficient occupancy world model via decoupled dynamic flow and image-assisted training. *arXiv preprint arXiv:2412.13772*, 2024.
- [293] Jinhua Zhang, Hualian Sheng, Sijia Cai, et al. PerLDiff: Controllable street view synthesis using perspective-layout diffusion models. In *IEEE/CVF Int. Conf. Comput. Vis.*, 2025.
- [294] Junge Zhang, Qihang Zhang, Li Zhang, et al. Urban scene diffusion through semantic occupancy map. *arXiv preprint arXiv:2403.11697*, 2024.
- [295] Kaiwen Zhang et al. Epona: Autoregressive diffusion world model for autonomous driving. In *IEEE/CVF Int. Conf. Comput. Vis.*, 2025.
- [296] Lunjun Zhang et al. Copilot4D: Learning unsupervised world models for autonomous driving via discrete diffusion. In *Int. Conf. Learn. Represent.*, 2024.
- [297] Richard Zhang, Phillip Isola, Alexei A Efros, et al. The unreasonable effectiveness of deep features as a perceptual metric. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 586–595, 2018.
- [298] Xiangwen Zhang et al. AccidentSim: Generating physically realistic vehicle collision videos from real-world accident reports. *arXiv preprint arXiv:2503.20654*, 2025.
- [299] Yumeng Zhang, Shi Gong, Kaixin Xiong, et al. BEVWorld: A multimodal world model for autonomous driving via unified BEV latent space. *arXiv preprint arXiv:2407.05679*, 2024.
- [300] Zhejun Zhang, Alexander Liniger, Dengxin Dai, et al. TrafficBots: Towards world models for autonomous driving simulation and motion prediction. In *IEEE Int. Conf. Robot. Autom.*, pages 1522–1529, 2023.
- [301] An Zhao et al. Diffusion distillation with direct preference optimization for efficient 3D LiDAR scene completion. *arXiv preprint arXiv:2504.11447*, 2025.
- [302] Changyuan Zhao et al. Edge general intelligence through world models and agentic AI: Fundamentals, solutions, and challenges. *arXiv preprint arXiv:2508.09561*, 2025.
- [303] Changyuan Zhao et al. World models for cognitive agents: Transforming edge intelligence in future networks. *arXiv preprint arXiv:2506.00417*, 2025.
- [304] Guosheng Zhao, Chaojun Ni, Xiaofeng Wang, et al. DriveDreamer4D: World models are effective data machines for 4D driving scene representation. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, 2025.

- [305] Guosheng Zhao, Xiaofeng Wang, Zheng Zhu, et al. DriveDreamer-2: LLM-enhanced world models for diverse driving video generation. In *AAAI Conf. Artif. Intell.*, volume 39, pages 10412–10420, 2025.
- [306] Guosheng Zhao et al. ReconDreamer++: Harmonizing generative and reconstructive models for driving scene representation. *arXiv preprint arXiv:2503.18438*, 2025.
- [307] Dian Zheng et al. VBench-2.0: Advancing video generation benchmark suite for intrinsic faithfulness. *arXiv preprint arXiv:2503.21755*, 2025.
- [308] Wenzhao Zheng, Weiliang Chen, Yuanhui Huang, et al. OccWorld: Learning a 3D occupancy world model for autonomous driving. In *Eur. Conf. Comput. Vis.*, pages 55–72. Springer, 2024.
- [309] Wenzhao Zheng et al. Doe-1: Closed-loop autonomous driving with large world model. *arXiv preprint arXiv:2412.09627*, 2024.
- [310] Wenzhao Zheng et al. GaussianAD: Gaussian-centric end-to-end autonomous driving. *arXiv preprint arXiv:2412.10371*, 2024.
- [311] Yupeng Zheng, Pengxuan Yang, Zebin Xing, et al. World4Drive: End-to-end autonomous driving via intention-aware physical latent world model. *arXiv preprint arXiv:2507.00603*, 2025.
- [312] Brady Zhou and Philipp Krähenbühl. Cross-view transformers for real-time map-view semantic segmentation. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, pages 13760–13769, 2022.
- [313] Hongyu Zhou et al. HUGSIM: A real-time, photo-realistic and closed-loop simulator for autonomous driving. *arXiv preprint arXiv:2412.01718*, 2024.
- [314] Jingqiu Zhou et al. FlexDrive: Toward trajectory flexibility in driving scene reconstruction and rendering. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, 2025.
- [315] Ming Zhou, Jun Luo, Julian Villella, et al. SMARTS: Scalable multi-agent reinforcement learning training school for autonomous driving. In *Conf. Robot Learn.*, 2020.
- [316] Siyuan Zhou et al. RoboDreamer: Learning compositional world models for robot imagination. *arXiv preprint arXiv:2404.12377*, 2024.
- [317] Xian Zhou et al. Genesis: A generative and universal physics engine for robotics and beyond. *arXiv preprint arXiv:2401.01454*, 2024.
- [318] Xin Zhou, Dingkan Liang, Sifan Tu, et al. HERMES: A unified self-driving world model for simultaneous 3D scene understanding and generation. In *IEEE/CVF Int. Conf. Comput. Vis.*, 2025.
- [319] Yunsong Zhou, Michael Simon, Zhenghao Mark Peng, et al. SimGen: Simulator-conditioned driving scene generation. In *Adv. Neural Inf. Process. Syst.*, volume 37, pages 48838–48874, 2024.
- [320] Yunsong Zhou, Naisheng Ye, William Ljungbergh, et al. Decoupled diffusion sparks adaptive scene generation. *arXiv preprint arXiv:2504.10485*, 2025.
- [321] Dekai Zhu, Yixuan Hu, Youquan Liu, et al. SPIRAL: Semantic-aware progressive LiDAR scene generation. *arXiv preprint arXiv:2505.22643*, 2025.
- [322] Zheng Zhu et al. Is Sora a world simulator? a comprehensive survey on general world models and beyond. *arXiv preprint arXiv:2405.03520*, 2024.
- [323] Yingshuang Zou et al. MuDG: Taming multi-modal diffusion with Gaussian splatting for urban scene reconstruction. *arXiv preprint arXiv:2503.10604*, 2025.
- [324] Vlas Zyrianov et al. Learning to generate realistic LiDAR point clouds. In *Eur. Conf. Comput. Vis.*, pages 17–35. Springer, 2022.
- [325] Vlas Zyrianov et al. LidarDM: Generative LiDAR simulation in a generated world. *arXiv preprint arXiv:2404.02903*, 2024.