



 Latest updates: <https://dl.acm.org/doi/10.1145/3746449>

SURVEY

## Understanding World or Predicting Future? A Comprehensive Survey of World Models

JINGTAO DING, Tsinghua University, Beijing, China

YUNKE ZHANG, Tsinghua University, Beijing, China

YU SHANG, Tsinghua University, Beijing, China

YUHENG ZHANG, Tsinghua University, Beijing, China

ZEFANG ZONG, Tsinghua University, Beijing, China

JIE FENG, Tsinghua University, Beijing, China

[View all](#)

Open Access Support provided by:

[Tsinghua University](#)



PDF Download  
3746449.pdf  
15 February 2026  
Total Citations: 15  
Total Downloads:  
17762

Published: 09 September 2025

Online AM: 27 June 2025

Accepted: 13 June 2025

Revised: 11 June 2025

Received: 21 November 2024

[Citation in BibTeX format](#)

# Understanding World or Predicting Future? A Comprehensive Survey of World Models

JINGTAO DING, Tsinghua University, Beijing, China

YUNKE ZHANG, Tsinghua University, Beijing, China

YU SHANG, Tsinghua University, Beijing, China

YUHENG ZHANG, Tsinghua University, Beijing, China

ZEFANG ZONG, Tsinghua University, Beijing, China

JIE FENG, Tsinghua University, Beijing, China

YUAN YUAN, Electronic Engineering, Tsinghua University, Beijing, China

HONGYUAN SU, Tsinghua University, Beijing, China

NIAN LI, Tsinghua University, Beijing, China

NICHOLAS SUKIENNIK, Tsinghua University, Beijing, China

FENGLI XU, Tsinghua University, Beijing, China

YONG LI, Tsinghua University, Beijing, China

---

The concept of world models has garnered significant attention due to advancements in multimodal large language models such as GPT-4 and video generation models such as Sora, which are central to the pursuit of artificial general intelligence. This survey offers a comprehensive review of the literature on world models. Generally, world models are regarded as tools for either understanding the present state of the world or predicting its future dynamics. This review presents a systematic categorization of world models, emphasizing two primary functions: (1) constructing internal representations to understand the mechanisms of the world, and (2) predicting future states to simulate and guide decision-making. Initially, we examine the current progress in these two categories. We then explore the application of world models in key domains, including autonomous driving, robotics, and social simulacra, with a focus on how each domain utilizes these aspects. Finally, we outline key challenges and provide insights into potential future research directions. We summarize the representative articles along with their code repositories in <https://github.com/tsinghua-fib-lab/World-Model>.

---

Jingtao Ding and Yunke Zhang authors contributed equally.

Yu Shang, Yuheng Zhang, Zefang Zong, Jie Feng, Yuan Yuan, Hongyuan Su, and Nian Li contributed equally.

This work is supported by a grant from the National Key Research and Development Program of China under 2024YFC3307603, the National Natural Science Foundation of China under 62476152 and U24B20180.

Authors' Contact Information: Jingtao Ding (corresponding author), Tsinghua University, Beijing, China; e-mail: dingjt15@tsinghua.org.cn; Yunke Zhang, Tsinghua University, Beijing, Beijing, China; e-mail: zyk21@mails.tsinghua.edu.cn; Yu Shang, Tsinghua University, Beijing, Beijing, China; e-mail: shangy21@mails.tsinghua.edu.cn; Yuheng Zhang, Tsinghua University, Beijing, Beijing, China; e-mail: yuheng-z22@mails.tsinghua.edu.cn; Zefang Zong, Tsinghua University, Beijing, Beijing, China; email: zongzf19@mails.tsinghua.edu.cn; Jie Feng, Tsinghua University, Beijing, Beijing, China; e-mail: fengjie@tsinghua.edu.cn; Yuan Yuan, Electronic Engineering, Tsinghua University, Beijing, China; e-mail: y-yuan20@mails.tsinghua.edu.cn; Hongyuan Su, Tsinghua University, Beijing, Beijing, China; e-mail: suhy24@mails.tsinghua.edu.cn; Nian Li, Tsinghua University, Beijing, Beijing, China; e-mail: linian21@mails.tsinghua.edu.cn; Nicholas Sukiennik, Tsinghua University, Beijing, Beijing, China; e-mail: sukienikn10@mails.tsinghua.edu.cn; Fengli Xu, Tsinghua University, Beijing, Beijing, China; email: fenglixu@tsinghua.edu.cn; Yong Li (corresponding author), Tsinghua University, Beijing, China; e-mail: liyong07@tsinghua.edu.cn.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2025 Copyright held by the owner/author(s).

ACM 0360-0300/2025/09-ART57

<https://doi.org/10.1145/3746449>

CCS Concepts: • **Computing methodologies** → **Machine learning; Artificial intelligence; Modeling and simulation;**

Additional Key Words and Phrases: World model, model-based RL, video generation, embodied environment, autonomous driving, robots, social simulacra

#### ACM Reference Format:

Jingtao Ding, Yunke Zhang, Yu Shang, Yuheng Zhang, Zefang Zong, Jie Feng, Yuan Yuan, Hongyuan Su, Nian Li, Nicholas Sukiennik, Fengli Xu, and Yong Li. 2025. Understanding World or Predicting Future? A Comprehensive Survey of World Models. *ACM Comput. Surv.* 58, 3, Article 57 (September 2025), 38 pages. <https://doi.org/10.1145/3746449>

---

## 1 Introduction

The scientific community has long aspired to develop a unified model that can replicate its fundamental dynamics of the world in pursuit of **Artificial General Intelligence (AGI)** [109]. In 2024, the emergence of multimodal **large language models (LLMs)** and video generation models like Sora [146] has intensified discussions surrounding such **World Models**. While these models demonstrate an emerging capacity to capture aspects of world knowledge—such as Sora’s generated videos, which appear to perfectly adhere to physical laws—questions persist regarding whether they truly qualify as comprehensive world models. Therefore, a systematic review of recent advancements, applications, and future directions in world model research is both timely and essential as we look toward new breakthroughs in the era of **artificial intelligence (AI)**.

The definition of a world model remains a subject of ongoing debate, generally divided into two primary perspectives: *understanding the world* and *predicting the future*. As depicted in Figure 1, early work by Ha and Schmidhuber [66] focused on abstracting the external world to gain a deep understanding of its underlying mechanisms. In contrast, LeCun [109] argued that a world model should not only perceive and model the real-world but also possess the capacity to envision possible future states to inform decision-making. Video generation models such as Sora represent an approach that concentrates on simulating future world evolution and thus align more closely with the predictive aspect of world models. This raises the question of whether a world model should prioritize understanding the present or forecasting future states. In this article, we provide a comprehensive review of the literature from both perspectives, highlighting key approaches and challenges.

The potential applications of world models span a wide array of fields, each with distinct requirements for understanding and predictive capabilities. In autonomous driving, for example, world models need to perceive road conditions in real-time [198, 216] and accurately predict their evolution [143, 187, 264], with a particular focus on immediate environmental awareness and forecasting of complex trends. For robotics, world models are essential for tasks such as navigation [179], object detection [204], and task planning [69], requiring a precise understanding of external dynamics [51] and the ability to generate interactive and embodied environments [148]. In the realm of simulation of virtual social systems, world models must capture and predict more abstract behavioral dynamics, such as social interactions and human decision-making processes. Thus, a comprehensive review of advancements in these capabilities, alongside an exploration of future research directions and trends, is both timely and essential.

Existing surveys on world models can generally be classified into two categories, as shown in Table S1. The first category primarily focuses on describing the application of world models in specific fields such as video processing and generation [24, 266], autonomous driving [62, 112, 231], and agent-based applications [266]. The second category [130] concentrates on the technological transitions from multi-modal models, which are capable of processing data across various

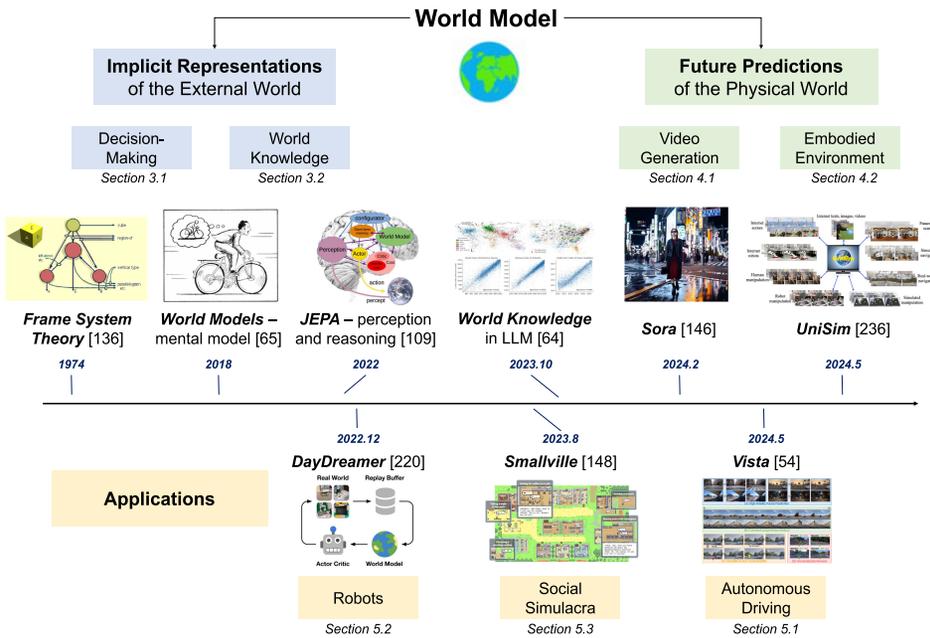


Fig. 1. The overall framework of this survey. We systematically define the essential purpose of a world model as understanding the dynamics of the external world and predicting future scenarios. The timeline illustrates the development of key definitions and applications.

modalities, to world models. However, these articles often lack a systematic examination of what precisely constitutes a world model and what different real-world applications require from these models. In this article, we aim at formally defining and categorizing world models, review recent technical progress, and explore their extensive applications.

The main contributions of this survey can be summarized as follows: (1) We present a novel categorization system for world models structured around two primary functions: *constructing implicit representations to understand the mechanism of the external world* and *predicting future states of the external world*. The first category focuses on the development of models that learn and internalize world knowledge to support subsequent decision-making, while the latter emphasizes enhancing predictive and simulative capabilities in the physical world from visual perceptions. (2) Based on this categorization, we classify how various key application areas, including autonomous driving, robots, and social simulacra, emphasize different aspects of world models. (3) We highlight future research directions and trends of world models that can adapt to a broader spectrum of practical applications.

The remainder of this article is organized as follows: In Section 2, we introduce the background of the world model and propose our categorization system. Section 3 and Section 4 elaborate on the details of current research progress on two categories of world models, respectively. Section 5 covers applications of the world model in three key research fields. Section 6 outlines open problems and future directions of world models.

## 2 Background and Categorization

In this section, we explore the evolving concepts of world models in the literature and categorize efforts to construct world models into two distinct branches: internal representation and future prediction.

The concept of building an internal model of the world has a long history in AI, dating back to foundational work such as Marvin Minsky’s frame representation in the 1960s [136], designed to systematically capture structured knowledge about the world. Ha et al. [65, 66] significantly revived and popularized the term “world model” in 2018 by proposing neural-network-based implicit models for learning latent representations. This line of research aligns with the psychological theory of “mental models” [94],<sup>1</sup> which holds that humans perceive the external world by abstracting it into simplified elements and relationships—an underlying philosophical root reflected alike in both frames and world models. This principle suggests that our descriptions of the world, when viewed from a deep, internal perspective, typically involve constructing an abstract representation that suffices without requiring detailed depiction. Building upon this conceptual framework, the authors introduce an agent model inspired by the human cognitive system, as illustrated in Figure 1. In this pioneering model, the agent receives feedback from the real-world environment, which is then transformed into a series of inputs that train the model. This model is adept at simulating potential outcomes following specific actions within the external environment. Essentially, it creates a mental simulation of potential future world evolutions, with decisions made based on the predicted outcomes of these states. This methodology closely mirrors the **Model-based Reinforcement Learning (MBRL)** method, where both strategies involve the model generating internal representations of the external world. These representations facilitate navigation through and resolution of various decision-making tasks in the real-world.

In the visionary article on the development of autonomous machine intelligence in 2022 [109], Yann LeCun introduced the system architecture for autonomous intelligence, a framework mirroring the human brain’s structure. As illustrated in Figure 1, it comprises a perception module that processes sensory data, followed by a cognitive module that evaluates this information, effectively embodying the world model. This model, designed as **Joint Embedding Predictive Architecture (JEPA)**, allows the brain to assess actions and determine the most suitable responses for real-world applications. LeCun’s framework is intriguing due to its incorporation of the dual-system concept, mirroring “fast” and “slow” thinking. System 1 involves intuitive, instinctive reactions: quick decisions made without a world model, such as instinctively dodging an oncoming person. In contrast, System 2 employs deliberate, calculated reasoning that considers the future state of the world. It extends beyond immediate sensory input, simulating potential future scenarios, like predicting events in a room over the next ten minutes and adjusting actions accordingly. This level of foresight requires constructing a world model to effectively guide decisions based on the anticipated dynamics and evolution of the environment. In this framework, the world model is essential for understanding and representing the external world. It models the state of the world using latent variables, which capture key information while filtering out redundancies. This approach allows for a highly efficient, minimalistic representation of the world, facilitating optimal decision-making and planning for future scenarios.

The ability of models to capture world knowledge is critical for their effective performance in a wide range of real-world tasks. In the recent wave of works on large language models starting from 2023, several have demonstrated the presence of latent world knowledge. In other words, these models capture intuitive knowledge, including spatial and temporal understanding, which enables them to make predictions about real-world scenarios [64, 133]. Furthermore, LLMs are capable of modeling the external world through cognitive maps, as indicated by recent research revealing the brain-like structures embedded within them [117]. These models can even learn to predict future events based on prior experiences, thereby enhancing their utility and applicability in real-world contexts.

---

<sup>1</sup><https://plato.stanford.edu/entries/mental-representation/>

The above world models primarily represent an implicit understanding of the external world. Powered by generative learning like diffusion modeling and model architecture like transformer, recent video generation models (e.g., Sora [146], Keling [104], Gen-2 [29]) take text instructions or real-world visual data as input and output high-quality video frames. Notably, these models demonstrate exceptional modeling capabilities, such as maintaining consistency in 3D video simulations, producing physically plausible outcomes, and simulating digital environments. These capabilities suggest that they not only mimic the appearance of but also model the real-world dynamics within simulation scenarios, focusing on realistically modeling dynamic world changes rather than merely representing static world states.

Whether focusing on learning internal representations of the external world or simulating its operational principles, these concepts coalesce into a shared consensus: the essential purpose of a world model is to understand the dynamics of the world and compute the next state with certainty (or with some guarantee) which empowers the model to extrapolate longer-horizon evolution and to support downstream decision-making and planning. From this perspective, we conduct a thorough examination of recent advancements in world models, analyzing them through the following lenses, as depicted in Figure 1.

- **Implicit representation of the external world** (Section 3): This research category constructs a model of environmental change to enable more informed decision-making, ultimately aiming at predicting the evolution of future states. It fosters an implicit comprehension by transforming external realities into a model that represents these elements as latent variables. Furthermore, with the advent of LLMs, efforts previously concentrated on traditional decision-making tasks have been significantly enhanced by the detailed descriptive power of these models regarding world knowledge. We further focus on the integration of world knowledge into existing models.
- **Future predictions of the external world** (Section 4): We initially explore generative models that simulate the external world, primarily using visual video data. These works emphasize the realness of generated videos that mirror future states of the physical world. As recent advancements shift focus toward developing a truly interactive physical world, we further investigate the transition from visual to spatial representations and from video to embodiment. This includes comprehensive coverage of studies related to the generation of embodied environments that mirror the external world.
- **Applications of world models** (Section 5): World models have a wide range of applications across various fields, including autonomous driving, robotics, and social simulacra. We explore how the integration of world models in these domains advances both theoretical research and practical implementations, emphasizing their transformative potential in real-world applications.

### 3 Implicit Representation of the External World

This section examines how world models enable informed decision-making by representing the environment as latent variables. Section 3.1 focuses on the world models in MBRL, while Section 3.2 explores the integration of world knowledge into advanced AI models, especially LLMs, enhancing real-world task performance.

#### 3.1 World Model in Decision-making

In decision-making tasks, understanding the environment is the major task in setting a foundation for optimized policy generation. As such, the world model in decision-making should include a comprehensive understanding of the environment. It enables us to take hypothetical actions

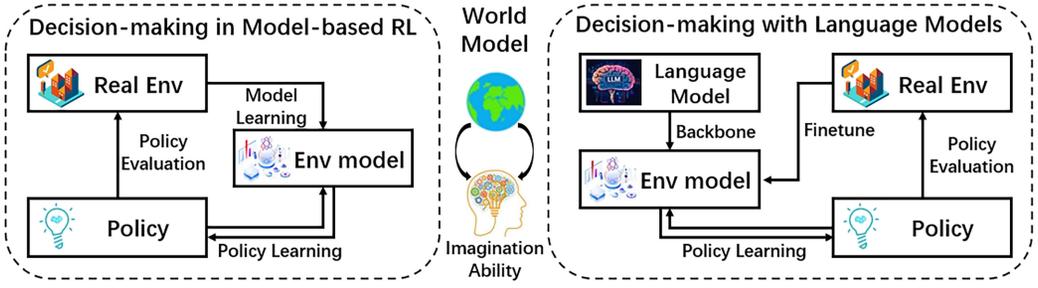


Fig. 2. Two schemes of utilizing world model in decision-making.

without affecting the real environment, facilitating a low trial-and-error cost. In literature, research on how to learn and utilize the world model was initially proposed in the field of model-based RL. Furthermore, recent progress on LLM and MLLM also provide comprehensive backbones for world model construction. With language serving as a more general representation, language-based world models can be adapted to more generalized tasks. The two schemes of leveraging world models in decision-making tasks are shown in Figure 2.

**3.1.1 World Model in Model-based RL.** In decision-making, the concept of the world model largely refers to the environment model in MBRL. A decision-making problem is typically formulated as a **Markov Decision Process (MDP)**, denoted with a tuple  $(S, A, M, R, \gamma)$ , where  $S, A, \gamma$  denotes the state space, action space and the discount factor each. The world model here consists of  $M$ , the state transition dynamics and  $R$ , the reward function. Since the reward function is defined in most cases, the key task of MBRL is to learn and utilize the transition dynamics, which can further support policy optimization.

**World Model Learning.** To learn an accurate world model, the most straightforward approach is to leverage the mean squared prediction error on each one-step transitions [89, 90, 108, 129, 163],

$$\min_{\theta} \mathbb{E}_{s' \sim M^*(\cdot|s, a)} [\|s' - M_{\theta}(s, a)\|_2^2], \quad (1)$$

where  $M^*$  is the real transition dynamics used to collect trajectory data and  $M_{\theta}$  is the parameterized transition to learn. Apart from directly utilizing the deterministic transition model, Chua et al. [27] further model the aleatoric uncertainty with the probabilistic transition model. The objective is to minimize the KL divergence between the transition models,

$$\min_{\theta} \mathbb{E}_{s' \sim M^*(\cdot|s, a)} \left[ \log \left( \frac{M^*(s'|s, a)}{M_{\theta}(s'|s, a)} \right) \right]. \quad (2)$$

In both settings, the phase of the world model learning task can be transformed into a supervised learning task. The learning labels are the trajectories derived from real interaction environments, also called the simulation data [128].

For high-dimensional environments, representation learning is essential for effective world-model training in MBRL. Early work by Ha and Schmidhuber [65] reconstructs images through an autoencoder–latent-state pipeline, whereas Hafner et al. [68, 70] couple a visual encoder with latent dynamics to master pixel-based control tasks. Their latest iteration, DreamerV3 [71], adds robust normalization and balancing techniques, solving over 150 tasks—including diamond collection in Minecraft—without human data or domain-specific tuning. Memory-centric extensions such as Recall-to-Imaging by Samsami et al. [171] further enhance long-horizon reasoning. A complementary trend is unified model learning via next-token prediction with transformer architectures, as

shown by Janner et al. [90] and expanded by Schubert et al. [177]. Further, Georgiev et al. [55] train a large off-policy multi-task world model whose smooth latent dynamics enable efficient per-task policy learning with first-order gradients, achieving strong scalability and performance without online planning. Recent work by Jonathan Richens et al. [167] further reinforces the necessity of world models, showing that any agent capable of generalizing to multi-step goal-directed tasks must have learned a predictive model of its environment, with the world model emerging from the agent’s policy. This insight aligns with the ongoing trend of incorporating predictive modeling into reinforcement learning to handle more complex and goal-oriented tasks.

**Policy Generation with World Model.** With an ideally optimized world model, one most straightforward way to generate a corresponding policy is **model predictive control (MPC)** [102]. MPC plans an optimized sequence of actions given the model as follows:

$$\max_{a_{t:t+\tau}} \mathbb{E}_{s_{t'+1} \sim p(s_{t'+1}|s_{t'}, a_{t'})} \left[ \sum_{t'=t}^{t+\tau} r(s_{t'}, a_{t'}) \right], \quad (3)$$

where  $\tau$  denotes the planning horizon. Nagabandi et al. [141] adopt a simple Monte Carlo method to sample action sequences. Rather than sampling actions uniformly, Chua et al. [27] propose a new probabilistic algorithm that ensembles with trajectory sampling. Further literature also improves the optimization efficiency by leveraging the world model usage [68, 80, 208, 245]. Hansen et al. [72] introduced an improved model-based RL algorithm called TD-MPC2 that integrates trajectory optimization within the latent space of a learned implicit world model. It achieves strong performance across diverse continuous control tasks and demonstrates scalability by training large agents with hundreds of millions of parameters across multiple domains.

Another popular approach to generating world model policies is the **Monte Carlo Tree Search (MCTS)**. By maintaining a search tree where each node refers to a state evaluated by a predefined value function, actions will be chosen such that the agent can be processed to a state with a higher value. AlphaGo and AlphaGo Zero are two significant applications using MCTS in discrete action space [189, 190]. Moerland et al. [138] extended MCTS to solve decision problems in continuous action space. Oh et al. [144] proposed a value prediction network that applies MCTS to the learned model to search for actions based on value and reward predictions.

**3.1.2 World Model with Language Backbone.** The rapid growth of language models, especially LLM and MLLM, benefits development in many related applications. With language serving as a universal representation backbone, language-based world models have shown their potential in many decision-making tasks.

**Direct Action Generation via LLM World Models.** LLM is capable of directly generating actions in decision-making tasks based on corresponding constructed world models. For example, in the navigation scenarios, Yang et al. [234] transfer pre-trained text-to-video models to domain-specific tasks for robot control, successfully annotating robot manipulation with text instructions as LLM outputs. Zhou et al. [263] further learn a compositional world model by factorizing the video generation process. Such a method enables a strong few-shot transfer ability to unseen tasks.

Besides training or fine-tuning specialized language-based world models, LLMs and MLLMs can be directly deployed to understand the world environment in decision-making tasks. For example, Long et al. [126] propose a multi-expert scheme to handle visual language navigation tasks. They construct a standardized discussion process where eight LLM-based experts participate to generate the final movement decision. An abstract world model is constructed from the discussion and further imagination (of future states) of the experts to support action generation. Zhao et al. [255] further combine LLMs and open-vocabulary detection to construct the relationship between multi-modal signals and key information in navigation. They propose an omni-graph to capture the

structure of the local space as the world model for the navigation task. Meanwhile, Yang et al. [239] utilize an LLM-based imaginative assistant to infer the global semantic graph as the world model based on the environment perception, and another reflective planner to directly generate actions.

**Modular Usage of LLM World Models.** Although taking LLM outputs as actions directly is straightforward in application and deployment, the decision quality in such a scheme heavily relies on the reasoning ability of the LLM itself. Although this year has witnessed the large potential of LLM reasoning capability [225], it can be further improved by integrating LLM-based world models as modules with external model-based verifiers or other effective planning algorithms [96].

Guan et al. [61] extract explicit world models by prompting GPT-4 to generate and iteratively refine PDDL domain descriptions, then pair these models with off-the-shelf planners, yielding good planning performance with much less human intervention. Xiang et al. [224] deploys an embodied agent in a world model, the simulator of VirtualHome [156], where the corresponding embodied knowledge is injected into LLMs. To better plan and complete specific goals, they propose a goal-conditioned planning schema where MCTS is utilized to search for the true embodied task goal. Lin et al. [119] introduce an agent, Dynalang, which learns a multimodal world model to predict future text and image representations, and which learns to act from imagined model rollouts. The policy learning stage utilizes an actor-critic algorithm purely based on the previously generated multimodal representations. Liu et al. [125] further cast reasoning in LLMs as learning and planning in Bayesian adaptive **Markov decision processes (MDPs)**. LLMs, like the world model, perform in an in-context manner within the actor-critic updates of MDPs. The proposed RAFA framework shows significantly increased performance in multiple complex reasoning tasks and environments, such as ALFWorld [188].

### 3.2 World Knowledge Learned by Models

After pretraining on large-scale web text and books [145, 201], large language models attain extensive knowledge about the real-world and common sense relevant to daily life. This embedded knowledge is considered crucial for their remarkable ability to generalize and perform effectively in real-world tasks. For instance, researchers leverage the common sense of large language models for task planning [257], robot control [82], and image understanding [122]. Furthermore, Li et al. [117] discover brain-like structures of world knowledge embedded in the high-dimensional vectors that represent the universe of concepts in large language models. Also, Li et al. [111] demonstrate that language models partially converge towards representations isomorphic to those of vision models.

Unlike common sense and general knowledge, we focus on world knowledge within large language models from the perspective of a world model. As shown in Figure 3, based on objects and spatial scope, the world knowledge in the large language models can be categorized into three parts: (1) knowledge of the global physical world; (2) knowledge of the local physical world; and (3) knowledge of human society. We summarize recent works in Table 1.

*3.2.1 Knowledge of the Global Physical World.* We first introduce research focused on analyzing and understanding the knowledge of the global physical world. Gurnee et al. [64] present the first evidence that large language models genuinely acquire spatial and temporal knowledge of the world, rather than merely collecting superficial statistics. They identify distinct “spatial neurons” and “temporal neurons” in Llama2 [201], suggesting that the model learns linear representations of space and time across multiple scales. Distinct from previous observations focused on embedding space, Manvi et al. [132, 133] develop effective prompts about textual address to extract intuitive real-world knowledge about the geospatial space and successfully improve the performance of the model in various downstream geospatial prediction tasks.

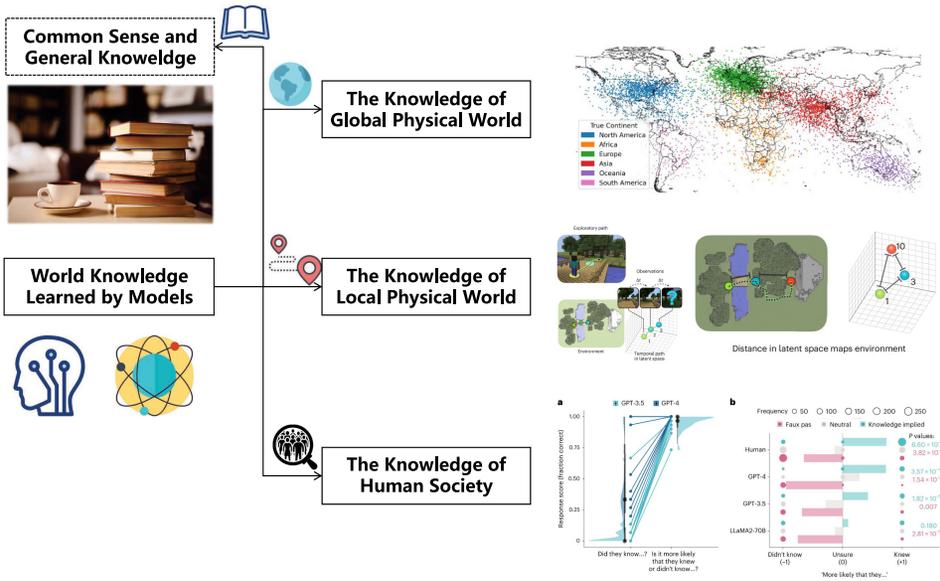


Fig. 3. World knowledge in large language models for world model.

Table 1. Overview of Recent Works in World Knowledge Learned by Models

Category	Methods/Model	YearandVenue	Modality	Content
Common Sense and General Knowledge	KoLA [242]	2024 ICLR	Language	Benchmark
	EWOK [86]	2024 arxiv	Language	Benchmark
	Geometry of Concepts [117]	2024 arxiv	Language	Analysis
Knowledge of Global Physical World	SpaceandTime [64]	2024 ICLR	Language	Analysis
	GeoLLM [133]	2024 ICLR	Language	Learning
	GeoLLM-Bias [132]	2024 ICML	Language	Learning
	GPT4GEO [169]	2023 NeurIPS(FMDM)	Language	Benchmark
	CityGPT [42]	2024 arxiv	Language	Learning
Knowledge of Local Physical World	CityBench [44]	2024 arxiv	LanguageandVision	Benchmark
	Predictive [59]	2024 NMI	Vision	Learning
	Emergent [93]	2024 ICML	Language	Learning
	E2WM [224]	2023 NeurIPS	Language	Learning
	Dynalang [119]	2024 ICML	LanguageandVision	Learning
Knowledge of Human Society	Testing ToM [194]	2024 NHB	Language	Benchmark
	High-order ToM [195]	2024 arxiv	Language	Benchmark
	COKE [218]	2024 ACL	Language	Learning
	MuMA-ToM [186]	2024 ACL	LanguageandVision	Benchmark
	SimToM [215]	2024 ACL	Language	Learning

While large language models do acquire some implicit knowledge of the real-world [64, 117], the quality of this knowledge remains questionable [42, 169]. For example, Feng et al. [42] find that the urban knowledge embedded in large language models is often coarse and inaccurate. To address this, they propose an effective framework to improve the acquisition of urban knowledge of specific cities in large language models. From these works, we can see that although large language models have demonstrated the ability to capture certain aspects of real-world knowledge [64, 117, 169], it is clear that further efforts are needed to enhance this knowledge to enable broader and more reliable real-world applications [43].

3.2.2 *Knowledge of the Local Physical World.* Unlike the knowledge of the global physical world, the local physical world represents the primary environment for human daily life and most real-world tasks. Therefore, understanding and modeling the local physical world is a more critical

topic for building a comprehensive world model. We first introduce the concept of the cognitive map [200], which refers to the mental representation that humans form to navigate and understand their environment, including spatial relationships and landmarks. Although initially developed to explain human learning processes, researchers have discovered similar structures in large language models [117] and have leveraged these insights to enhance the efficiency and performance of artificial models in learning and understanding the physical world.

Recent studies explore actively encouraging models to learn abstract knowledge through cognitive map-like processes across various environments. For example, Cornet et al. [59] show that in a simplified Minecraft world, visual predictive coding lets an agent build a spatial cognitive map purely from pixels. Once trained, the latent map encodes its metric distance to any target, enabling accurate rollout of future observations. Lin et al. [119] investigate teaching models to understand the game environments through a world model learning procedure, specifically by predicting the subsequent frame of the environment. In this way, the model can generate better actions in dynamic environments. Moreover, Jin et al. [93] find that language models can learn the emergent representations of program semantics by predicting the next token.

**3.2.3 Knowledge of the Human Society.** Beyond the physical world, understanding human society is another crucial aspect of world models. David Premack and Guy Woodruff proposed the Theory of Mind [155], which was later developed to explain how individuals infer the mental states of others around them. Recent works have extensively explored how large language models develop and demonstrate this social world model [173, 194]. Sap et al. [173] conduct an investigation focusing on evaluating the performance of large language models across various Theory of Mind tasks to determine whether their human-like behaviors reflect genuine comprehension of social rules and implicit knowledge. Strachan et al. [194] conduct a comparative analysis between human and LLM performance on diverse Theory of Mind abilities, such as understanding false beliefs and recognizing irony. While their findings demonstrate the potential of GPT-4 in these tasks, they also identify its limitations, particularly in detecting faux pas.

To address these limitations, researchers propose innovative methods to enhance the abilities of large language models in Theory of Mind for complex real-world applications. Wu et al. [218] introduce COKE, which constructs a knowledge graph to help large language models explicitly using theory in mind through cognitive chains. Additionally, Alex et al. [215] develop SimToM, a two-stage prompting framework, to enhance the performance of large language models in theory of mind tasks.

## 4 Future Prediction of the Physical World

### 4.1 World Model as Video Generation

The integration of video generation into world models marks a significant leap forward in the field of environment modeling [146]. Traditional world models primarily focused on predicting discrete or static future states [66, 109]. However, by generating video-like simulations that capture continuous spatial and temporal dynamics, world models [146, 233] have evolved to address more complex, dynamic environments. This breakthrough in video generation has pushed the capabilities of world models to a new level.

**4.1.1 Towards Video World Models.** A video world model is a computational framework designed to simulate and predict the future state of the world by processing past observations and potential actions within a visual context [146]. This concept builds on the broader idea of world models, which strive to capture the dynamics of an environment and enable machines to predict

how the world will evolve over time. In the case of a video world model, the focus is on generating sequences of visual frames that represent these evolving states.

**Sora as a World Simulator.** Sora [146] is a large-scale video generation model, which is designed to generate high-quality, temporally consistent video sequences, up to one minute long, based on various input modalities such as text, images, and videos. Sora leverages a combination of powerful neural network architectures, including encoder-decoder frameworks and transformers, to process multimodal inputs and generate visually coherent simulations. Sora's core capabilities lie in its ability to generate videos that align with real-world physical principles, such as the reflection of light on surfaces or the melting of candles. These properties suggest that Sora has the potential to act as a world simulator, predicting future states of the world based on its understanding of the initial conditions and simulation parameters.

**Sora's Limitations.** However, despite its impressive video generation abilities, Sora has several limitations in terms of understanding and simulating the external world. One key limitation concerns causal reasoning [24, 266], wherein the model is limited in simulating dynamic interactions within the environment. Thus, Sora can only passively generate video sequences based on an observed initial state, but cannot actively intervene or predict how changes in actions might alter the course of events. Another limitation is that it still fails to reproduce correct physical laws consistently [97]. While Sora can generate visually realistic scenes, it struggles with accurately simulating real-world physics, such as the behavior of objects under different forces, fluid dynamics, or the accurate depiction of light and shadow interactions.

**Other Video World Models.** Sora has undoubtedly catalyzed a significant wave of research into video world models, inspiring a surge of advancements in this field. Following Sora's success in generating high-quality video sequences, numerous subsequent models have been developed, each aiming at pushing the boundaries of what video world models can achieve. For example, some approaches have extended video lengths to enable long-form video simulation [77, 123, 241]. In addition to conventional language-guided video generation, more modalities are being integrated, such as images and actions [223, 258]. Researchers are also shifting their focus from basic video generation, which lacks user control, to interactive simulations that aim at replicating the decision space of the real-world and facilitate decision-making [87, 219, 223, 235, 237, 249]. Several studies have worked to enhance the smoothness of action transitions, improve the accuracy of physical laws, and maintain temporal consistency [17, 166, 229, 233]. Meanwhile, the concept of world models has evolved beyond imagination and is being applied in various scenario-specific simulations, including natural environments, games, and autonomous driving [12, 16, 77, 123, 134, 135, 209, 211, 261]. Table 2 summarizes the categorization of improvements in video world models across different aspects.

*4.1.2 Capabilities of Video World Models.* Despite the ongoing debate about whether models like Sora can be considered full-fledged world models, there is no doubt that video world models hold tremendous potential for advancing environment simulation and prediction [24, 97, 266]. These models can offer a powerful approach to understanding and interacting with complex environments by generating realistic, dynamic video sequences. To achieve this level of sophistication, this section outlines the key capabilities that video world models must possess to set them apart from traditional video generation models.

**Long-term Predictive Ability.** A robust video world model should be capable of making long-term predictions that adhere to the dynamic rules of the environment over an extended period. This capability allows the model to simulate how a scenario evolves, ensuring that the generated video sequences remain consistent with the temporal progression of the real-world. Although Sora has achieved the generation of minute-long video sequences with high-quality temporal coherence, it

Table 2. Overview of Recent Models in Video Generation Across Various Categories, which Summarizes Key Models in Long-term Video Generation, Multi-modal Learning, Interactive Video Generation, Temporal Consistency, and Diverse Environment Modeling

Category	Model	Description	Technique
Long-term	NUWA-XL [241]	“Coarse-to-fine” Diffusion over Diffusion architecture for long video generation.	Diffusion
	LWM [123]	Training large transformers on long video and language sequences.	Transformer
	GAIA-1 [77]	Generative world model predicting driving scenarios for autonomous driving.	Transformer, Diffusion
Multimodal	3D-VLA [258]	Integrates 3D perception, reasoning, and action in a world model for embodied AI.	Diffusion
	Pandora [223]	World-state simulation and real-time control with free-text actions.	LLM
	Genie [16]	Generative model from text, images, and sketches.	Transformer
Interactive	UniSim [235]	Simulates real-world interactions for vision-language and RL training.	Diffusion, RL
	VideoDecision[237]	Extends video models to real-world tasks like planning and RL.	Transformer, Diffusion
	iVideoGPT [219]	Combines visual, action, and reward signals for interactive world modeling.	Transformer
	PhysDreamer [249]	Simulates 3D object dynamics to generate responses to novel interactions.	Diffusion
Consistency	PEEKABOO [87]	Enhances interactivity with spatiotemporal control without extra training.	Diffusion Transformer
	WorldGPT [233]	Improves temporal consistency and action smoothness with multimodal learning and refined key frame generation.	Diffusion
	DiffDreamer [17]	Long-range scene extrapolation with improved consistency.	Diffusion
Diverse environments	ConsistI2V [166]	Enhances visual consistency in image-to-video generation.	Diffusion
	WorldDreamer [211]	World model capturing dynamic elements across diverse scenarios.	Transformer
	Genie [16]	Unsupervised generative model for action-controllable virtual environments.	Transformer
	MUVO [12] UniWorld [135]	Multimodal world model using camera and lidar data. 3D detection and motion prediction in autonomous driving.	Transformer Transformer

is still far from being able to simulate complex, long-term dynamics found in real-world environments. Recent efforts have explored extending video lengths to capture longer-term dependencies and improve temporal consistency [77, 123, 241].

**Multi-modal Integration.** In addition to language-guided video generation, video world models are increasingly integrating other modalities, such as images and actions, to enhance realism and interactivity [223, 258]. The integration of multiple modalities allows for richer simulations that better capture the complexity of real-world environments, improving both the accuracy and diversity of generated scenarios.

**Interactivity.** Another critical capability of video world models is their potential for controllability and interactivity. An ideal model should not only generate realistic simulations but also allow for interaction with the environment. This interactivity involves simulating the consequences of different actions and providing feedback, enabling the model to be used in applications requiring

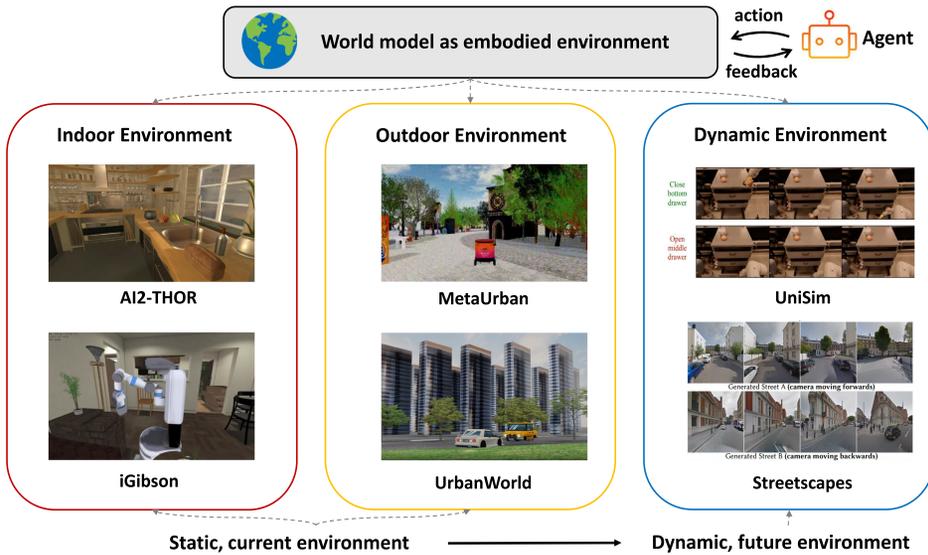


Fig. 4. Classification of world models as interactive embodied environments, including indoor, outdoor and dynamic environments. The modeling of the outside world is evolving from constructing static, current environments to predicting dynamic, future environments.

dynamic decision-making. Recent work is focusing on enhancing control over the simulations, allowing for more user-guided exploration of scenarios [219, 237].

**Diverse Environments.** Finally, video world models are being adapted to a variety of scenario-specific simulations, including natural environments, autonomous driving, and gaming. These models are evolving beyond basic video generation to replicate real-world dynamics and support a wide range of applications [16, 123, 211].

## 4.2 World Model as Embodied Environment

The development of world models for embodied environments is crucial for simulating and predicting how agents interact with and adapt to the external world. Initially, generative models focused on simulating visual aspects of the world, using video data to capture dynamic changes in the environment. More recently, the focus has shifted toward creating fully interactive and embodied simulations. These models not only represent the visual elements of the world but also incorporate spatial and physical interactions that more accurately reflect real-world dynamics. By integrating spatial representations and transitioning from video-based simulations to immersive, embodied environments, world models can now provide a more comprehensive platform for developing agents capable of interacting with complex real-world environments.

World models as embodied environments can be divided into three categories: indoor, outdoor, and dynamic environments, as shown in Figure 4, and the relevant works are summarized in Table 3. It can be summarized that most current works focus on developing static, existing indoor and outdoor embodied environments. An emerging trend is to predict the dynamic, future world through generative models producing first-person, dynamic video-based simulation environments. Such environments can offer flexible and realistic feedback for training embodied agents, enabling them to interact with ever-changing environments and improve their generalization ability.

**4.2.1 Indoor Environments.** Indoor environments offer controlled, structured scenarios where agents can perform detailed, task-specific actions such as object manipulation, navigation, and

Table 3. Comparison of Existing Works on World Models as Embodied Environments, Including Indoor, Outdoor, and Dynamic Environments

Type	Name	Environment	Year	Num of Scenes	Modality	Physics	3D Assets
Indoor	AI2-THOR [101]	Home	2017	120	V	✓	✓
Indoor	Matterport 3D [18]	Home	2018	90	V	✗	✗
Indoor	Virtual Home [156]	Home	2018	50	V	✓	✓
Indoor	Habitat [174]	Home	2019	-	V	✓	✓
Indoor	SAPIEN [222]	Home	2020	46	V	✓	✓
Indoor	iGibson [183]	Home	2021	15	V, L	✓	✓
Indoor	AVLEN [150]	Home	2022	85	V, T, A	✓	✓
Indoor	ProcTHOR [31]	Home	2022	Arbitrary	V	✓	✓
Indoor	Holodeck [238]	Home	2024	Arbitrary	V	✓	✓
Indoor	AnyHome [48]	Home	2024	Arbitrary	V	✓	✓
Indoor	LEGENT [21]	Home	2024	Arbitrary	V, T	✓	✓
Indoor	TDW [49]	Home	2021	-	V, A	✓	✓
In and Outdoor	GRUtopia [205]	Home, City	2024	100k	V, T	✓	✓
Outdoor	MineDOJO [41]	Game	2022	-	V	✗	✗
Outdoor	MetaUrban [221]	City	2024	13800	V, L	✓	✓
Outdoor	UrbanWorld [180]	City Building	2024	Arbitrary	V	✓	✓
Outdoor	EmbodiedCity [52]	City	2024	87.1k	V, T	✓	✓
Dynamic	UniSim [236]	Home, City, Simulation	2023	Arbitrary	V, T	✗	✗
Dynamic	Streetscapes [32]	Street View	2024	Arbitrary	V, T	✗	✗
Dynamic	AVID [168]	Home, Game	2024	Arbitrary	V, T	✗	✗
Dynamic	EVA [23]	Home, Simulation	2024	Arbitrary	V, T	✗	✗
Dynamic	Pandora [223]	Home, Game, Simulation, Street View	2024	Arbitrary	V, T	✗	✗

In the “Modality” column, “V” refers to vision, “L” refers to lidar, “T” refers to text, and “A” refers to audio. In the “Num of Scenes” column, “-” means no reported data, and “Arbitrary” means the method can support generating any number of scenes.

real-time interaction with users [18, 53, 101, 150, 156, 174, 183, 222]. Early works on establishing indoor environments like AI2-THOR [101] and Matterport 3D [18] focus on providing only visual information. These works build indoor environments by providing photorealistic settings where agents can practice visual navigation and engage in interactive tasks that mimic real-life home activities. These environments emphasize the importance of using visual-based reinforcement learning techniques that allow agents to optimize their decision-making based on environmental cues. By simulating real-world tasks like cooking or cleaning, these platforms assess an agent’s capacity to generalize learned behaviors across different types of spaces and objects. A line of further works contributes toward expanding the data modalities of the provided environments. Among these, iGibson [183] introduces Lidar observation as additional signal feedback, contributing to more accurate environment perception of agents. AVLEN [150] further supplements audio signals allowing agents to execute tasks such as object manipulation and navigation in household-like settings. The challenge here lies in enabling agents to understand and act on multimodal input including vision, language, and sound within a constrained space. Adding a social dimension, environments like GRUtopia [205] introduce agents to spaces where they must navigate and interact with both objects and NPCs. Here, agents need to understand social dynamics, such as positioning and task sharing, which requires more advanced forms of interaction modeling. The inclusion of social interaction modules in these settings demonstrates how agents can be trained to balance human-like social behaviors with task performance. More recently, with the development of LLMs, some works [21, 48, 238] seek to provide a flexible environment generation pipeline, supporting the generation of arbitrary indoor environments with language instructions.

**4.2.2 Outdoor Environments.** In contrast to indoor environments, creating outdoor environments [41, 49, 180, 205, 221] faces greater challenges due to their larger scale and increased variability. Some existing works focus on urban environments, such as MetaUrban [221], where agents are deployed to navigate in large-scale urban environments, where they encounter challenges like dynamically changing traffic, varied building structures, and social interactions

with other entities. These tasks often require the use of context-aware navigation algorithms that allow agents to adjust their trajectories and behaviors based on the layout and conditions of the environment. However, the environments in MetaUrban are created by retrieving and organizing 3D assets from existing libraries. Recently, utilizing advanced generative techniques, UrbanWorld [180] significantly enhances the scope of outdoor environments, using 3D generative models to create complex, customizable urban spaces that allow for more diverse urban scenarios. This shift from static asset-based environments to generative ones ensures that agents are exposed to a wider variety of tasks, from navigating unfamiliar street layouts to interacting with new types of objects or structures. In addition to the above real open-world generation works, there are also some virtual open-world platforms like MineDOJO [41] that extend these challenges even further by simulating procedurally generated, sandbox-like environments. These platforms, inspired by the open-ended world of Minecraft, push agents to engage in tasks like resource collection, construction, and survival, demanding continuous exploration and adaptive learning. In such environments, agents are motivated to seek out new information and adapt their behavior to finish given tasks. Training in such environments can help agents learn knowledge across a broad range of tasks and terrains, enabling them to operate effectively in various outdoor environments.

*4.2.3 Dynamic Environments.* Dynamic environments mark a significant evolution from traditional, static simulators by utilizing generative models to create flexible, real-time simulations. Unlike predefined environments that require manual adjustments, these models allow for the dynamic creation of a wide variety of scenarios, enabling agents to experience diverse, first-person perspectives. This shift provides agents with richer, more varied training experiences, improving their adaptability and generalization in complex, unpredictable real-world situations. A representative work is UniSim [236], which dynamically generates robot manipulation video sequences based on input conditions like spatial movements, textual commands, and camera parameters. Leveraging multimodal data from 3D simulations, real-world robot actions, and internet media, this system generates varied, realistic environments where agents can practice tasks like object manipulation and navigation. The key advantage of this approach is its flexibility, allowing agents to adapt to various scenarios without the limitations of static physical environments. Pandora [223] expands the dynamic environment generation from robot actions in Unisim to wider domains including human and robot actions in both indoor and outdoor scenes. Another subsequent work, AVID [168] builds on UniSim by conditioning on actions and modifying noise predictions from a pre-trained diffusion model to generate action-driven visual sequences for dynamic environment generation. Beyond the video diffusion-based framework of Unisim, EVA [23] introduces an additional vision-language model for embodied video anticipation, producing more consistent embodied video predictions. As for the generation of open-world dynamic environments, Streetscapes [32] employs autoregressive video diffusion models to simulate urban environments where agents must navigate dynamic challenges like changing weather and traffic. These environments offer consistently coherent, yet flexible, urban settings, exposing agents to real-world-like variability. The core trend in dynamic environments is the use of generative world models that provide scalable, adaptable simulations. This approach significantly reduces the manual effort required for environment setup, allowing agents to train across a diverse range of scenarios quickly. Moreover, the focus on first-person training closely mimics real-world decision-making, enhancing the agents' ability to adapt to evolving situations. These advances are key in developing embodied environments supporting agent learning in complex, dynamic scenarios.

Given the above developments, it is evident that world models as embodied environments have made significant advances in simulating and predicting how agents interact with dynamic, real-world scenarios. Current research predominantly focuses on developing indoor, static

environments, with notable efforts expanding to large-scale outdoor environments and dynamic simulation environments. A promising direction is to construct dynamic environments, which can provide first-person, action-conditioned future world prediction, enabling agents to better adapt to unseen conditions. These methods are promising to offer flexible, scalable environments for training embodied agents, enhancing their generalization capabilities for real-world tasks.

## 5 Application

### 5.1 Autonomous Driving

In recent years, with the rapid advancement of vision-based generative models [14, 75, 193] and multimodal large language models [1, 121], world models have attracted growing interest in the field of autonomous driving. The modern autonomous driving pipeline is typically divided into four key components: *perception*, *prediction*, *planning*, and *control*. Among these, the perception and prediction stages correspond to driving scene understanding—i.e., learning an implicit representation of the vehicle’s external environment. In parallel, recent surveys [62] highlight the emergence of end-to-end world simulators that learn to simulate realistic driving environments based on multimodal inputs—such as images, point clouds, trajectories, and language—and then generate future states to support downstream tasks like planning and decision-making. These two perspectives align well with our earlier categorization of world models, and in the following, we detail their applications and developments within the autonomous driving domain accordingly.

*5.1.1 Learning Implicit Representations.* Autonomous vehicles typically utilize cameras, radar, and lidar to perceive the real-world, gathering information through images, video data, and point cloud data. In the initial decision-making paradigm [20, 175], models often take perceptual data as input and directly output motion planning results for the autonomous vehicle. Conversely, when humans operate vehicles, they typically observe and predict the current and future states of other traffic participants to determine their own driving strategies [83]. Thus, learning the implicit representation of the world through perceptual data and predicting the future states of the surrounding environment is a crucial step in enhancing the decision-making reliability of autonomous vehicles. We consider this process as it manifests in how autonomous vehicles learn a world model in latent space.

As shown in the left half of Figure 5, before the advent of multimodal large models and end-to-end autonomous driving technologies [79], the perception and prediction tasks of autonomous vehicles were typically assigned to distinct modules, each trained on their respective tasks and datasets. The perception module processed data from images, point clouds, and other sources to accomplish tasks such as object detection and map segmentation, projecting the perceived world into an abstract geometric space. Furthermore, the prediction module would typically operate within these geometric spaces to forecast the future states of the surrounding environment, including the trajectories and motions of traffic participants.

The processing of perceptual data is closely tied to the evolution of deep learning technologies, as shown in Table 4. Pointnet [158], introduced in 2017, was the first to employ deep learning methods for processing point cloud data. As convolutional neural networks advanced, perception techniques based on image data, exemplified by YOLOP [216] and MultiNet [198], emerged and excelled in driving scene understanding tasks [74, 107, 203, 262]. In recent years, the transformer architecture has gained prominence in natural language processing, and this technology has also been applied to image data understanding. BEVFormer [118] utilizes the attention mechanism to integrate images from multiple camera angles, constructing an abstract geometric space from a bird’s-eye view, and achieving state-of-the-art results in various tasks, including object detection. Additionally, Transfusion [8] enhances perceptual accuracy by fusing lidar and camera data

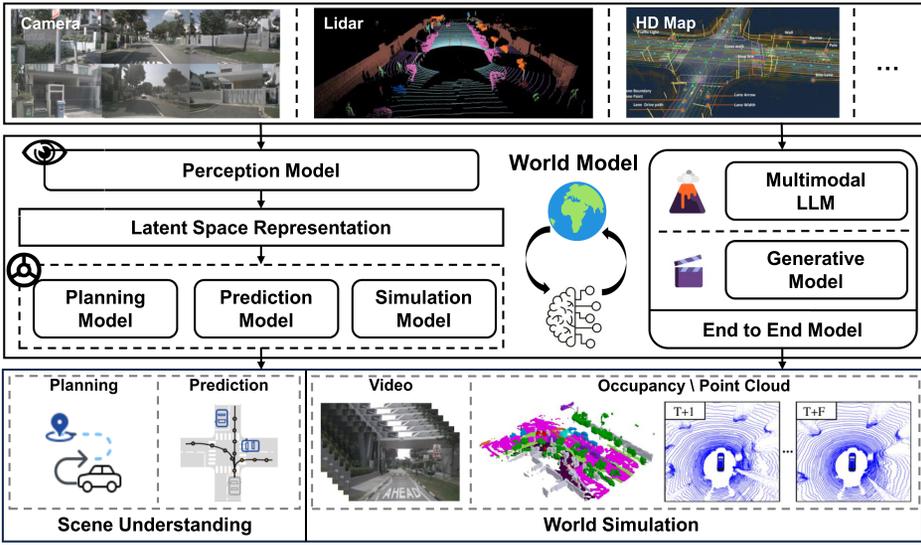


Fig. 5. Application of world model in autonomous driving.

Table 4. Comparison of Existing Works in Scene Understanding and World Simulation

	Task	Work	Year	Data Modality	Technique	Task Description
Driving Scene Understanding	Perception	Faster r-cnn [165]	2015	Camera	CNN	Object Detection
		Pointnet [158]	2017	Lidar	MLP	3D Classification
		MultiNet [198]	2018	Camera	CNN	Semantic Segmentation
		OmniDet [107]	2021	Camera	CNN and Attention	Multi-task Visual Perception
		YOLOP [216]	2022	Camera	CNN	Object Detection
		BEVFormer [118]	2022	Camera	Attention	3D Visual Perception
	Prediction	Transfusion [8]	2022	Camera and Lidar	Transformer	3D Object Detection
		Wayformer [142]	2022	Geometric Space	Attention	Trajectory Prediction
		MTR [187]	2022	Geometric Space	Transformer	Trajectory Prediction
		QCNet [264]	2023	Geometric Space	Transformer	Trajectory Prediction
End to End Scene Understanding	HPTR [253]	2023	Geometric Space	Transformer	Trajectory Prediction	
	Jiang et al. [92]	2023	Geometric Space	Diffusion	Trajectory Prediction	
	UniAD [79]	2023	Camera	Transformer	Motion Planning	
	TOKEN [199]	2024	Camera	MLLM	Motion Planning	
Driving World Simulation	Motion Simulation	OmniDrive [107]	2024	Camera	MLLM	Motion Planning
		SUMO [127]	2000	Geometric Space	Rule-based	Traffic Simulation
		Metadrive [116]	2022	Geometric Space	Data-driven	Traffic Simulation
		Trafficbots [252]	2023	Geometric Space	Transformer	Traffic Simulation
	End to End Sensor Simulation	Waymax [63]	2024	Geometric Space	Data-driven	Traffic Simulation
		GALA-1 [78]	2023	Camera	Transformer	Video Generation
		DriveDreamer [210]	2023	Camera	Diffusion	Video Generation
		Drive-WM [212]	2023	Camera	Diffusion	Video Generation
		OccWorld [260]	2023	Occupancy	Attention	Occupancy Generation
		OccSora [207]	2024	Occupancy	Diffusion	Occupancy Generation
Vista [54]	2024	Camera	Diffusion	Video Generation		
Copilot4D [248]	2024	Lidar	Diffusion	Point Cloud Generation		

through a cross-attention approach. Building on the perceptual results, a series of techniques such as RNNs [6, 98, 267], CNNs [25, 28, 152], and Transformers [84, 143, 187, 264] have been employed to encode historical scene information and predict the future behaviors of traffic participants.

With the emergence and rapid development of multimodal large language models in recent years, many efforts have sought to apply the general scene understanding capabilities of these models to the field of autonomous driving. TOKEN [199] tokenizes the whole traffic scene into object-level knowledge, using the reasoning ability of the language model to handle the long-tail prediction and planning problems, OmniDrive [107] sets up llm-based agents and covers multiple

tasks including scene description, counterfactual reasoning and decision making through visual question-answering.

**5.1.2 World Simulators.** As shown in Table 4, before the emergence of multimodal large models and vision-based generative models, traffic scenario simulations are often conducted in geometric spaces. The scene data on which these simulations rely is typically collected by the perception modules of autonomous vehicles or constructed manually. These simulations represent future states of the scenario in the form of geometric trajectories [63, 116, 127, 252], which require further modeling and rendering to produce outputs suitable for vehicle perception. The cascading of multiple modules often results in information loss and increases the complexity of simulations, making scenario control more challenging. Furthermore, realistic scene rendering typically requires substantial computational resources, which limits the efficiency of virtual traffic scenario generation.

Using diffusion-based video generation models as a world model partially addresses the aforementioned issues. By training on large-scale traffic scenario datasets, diffusion models can directly generate camera perception data that closely resembles reality. Additionally, the inherent controllability of diffusion models, combined with text-image alignment methods like CLIP [161], enables users to exert control over scenario generation in a straightforward manner. The GAIA-1 [78] and DriveDreamer series [210, 256] are among the first to employ this method for constructing world models. Building on this foundation, Drive-WM [212] introduces closed-loop control for planning tasks, while Vista [54] focuses on improving the resolution of generated results and extending prediction duration. In addition to methods that predict future states in video space, many other works have explored different forms of vehicle perception data. OccWorld [260] and OccSora [207] predict the future state of the world by forecasting 3D occupancy grids, whereas Copilot4D [248] constructs a world model by predicting changes in radar point cloud data. Compared to video data, these types of features better reflect the spatial characteristics of traffic scenarios.

## 5.2 Robots

World models have emerged as a transformative paradigm in robotics, empowering robots to perceive, predict, and act effectively in complex environments. This progress is driven in part by advances in neural architectures [75, 202] and learning algorithms [162, 178], which enable robots to build implicit representations that capture key aspects of the external world. Complementarily, prediction models [46, 47] offer the ability to forecast future environmental states, moving beyond static abstractions to support anticipatory and adaptive behavior. Together, these capabilities make it increasingly feasible for robots to learn directly from real-world interactions. In Table 5, we summarize the core learning tasks involved in constructing world models for robotics, categorized according to the three major perspectives outlined above (typical examples shown in Figure S1).

**5.2.1 Learning Implicit Representation.** Traditional robotic tasks (e.g., object grasping) are typically performed in highly structured environments where the critical components are explicitly modeled [38, 100], eliminating the need for the robot to independently learn or adapt its understanding of the world. However, when the robot is deployed in unfamiliar environments, especially those in which key features or dynamics have not been explicitly modeled, tasks that were previously successful may fail as the robot struggles to generalize to these unknown features [95, 137]. Thus, enabling a robot to learn an implicit representation of its environment is a crucial first step toward achieving intelligence.

To help a robot understand the objects in the world, visual models such as **convolutional neural networks (CNNs)** [57, 103, 110] and **vision transformers (ViT)** [34, 204] integrate visual characteristics of entities into representations, making it possible for robots to recognize critical

Table 5. Core Learning Tasks Involved in Constructing World Models for Robotics

	Task	Model	Year	Input	Backbone
Learning Inner Representation	Visual Representation	CNN [110]	1998	Image	CNN
		ViT [34]	2020	Image	Transformer
		RoboCraft [185]	2024	Image	GNN
	3D Representation	PointNet [157]	2017	3D point clouds	MLP
		Predictive Coding [59]	2024	Image	ResNet
		SpatialLM [197]	2025	3D point clouds	MLLM
Task Representation	BC-Z [88]	2022	Textand Video	LLMand ResNet	
	Text2Motion [120]	2023	Text	LLM	
	Gensim [206]	2023	Text	LLM	
Predicting Future Environment	Video Prediction	UniPi [36]	2024	Video	Diffusion
		VIPER [39]	2024	Video	Transformer
		GR-2 [19]	2024	Text and Video	Transformer
		IRASim [265]	2024	Trajectory	Diffusion
Real-world Planning	Real-world Adaptation	DayDreamer [220]	2023	Video	RSSM
		SWIM [134]	2023	Video	Transfer Learning
		CoSTAR [2]	2021	Multimodal	Belief Space
	Evaluation	OpenEQA [131]	2024	Imageand Text	LLM

objects for tasks. RoboCraft [185] transfers visual observation into particles and captures the structure of the underlying system through a graph neural network. Moreover, other attempts are made for the sensing of physical 3D space. PointNet [157, 159] first encodes the unstructured 3D point clouds with asymmetrical functions, capturing the spatial characteristics of the environment. A recent work [59] assembles observations acquired along local exploratory paths into a global representation of the physical space within its latent space, enabling robots to tail and approach specific targets. SpatialLM [197] further advances this direction by processing raw 3D point clouds into structured 3D scene representations with semantic labels, enhancing spatial reasoning for complex tasks in robotics and autonomous driving. With the advancement of language comprehension in LLMs [15, 35, 201], a novel paradigm for enabling robots to capture task intentions involves describing the task in textual form and then obtaining a textual representation through LLMs [56, 81, 140, 206]. BC-Z [88] utilizes language representations as task representations, enhancing the multi-task performance of robots. Text2Motion [120] splits the natural language instruction into task-level and motion-level plans with LLM to handle complex sequential manipulation tasks.

**5.2.2 Predicting Future States of the Environment.** Robotic tasks are inherently sequential and long-term, where early decisions significantly influence future outcomes [191]. Effective robotic planning thus relies heavily on accurately predicting future environmental states, enabling proactive decision-making and reducing costly errors. Classic closed-loop algorithms [10, 99], which select actions solely based on current observations, are typically short-sighted, risking irreversible mistakes. Additionally, approaches depending on explicit dynamic models crafted from expert knowledge tend to be limited in flexibility and robustness.

A key recent insight is the use of generative video models—particularly those leveraging diffusion [11, 22, 40, 73] and transformer architectures [230, 243]—to implicitly learn environmental dynamics directly from visual data. For instance, UniPi [36] frames action prediction explicitly as a video generation problem, conditioning a constrained diffusion model on the current state to visualize future scenarios. Similarly, VIPER [39] employs a pretrained autoregressive transformer to guide robotic actions, effectively leveraging rich representations learned from expert demonstration videos. IRASim [265] leverages diffusion models for trajectory-to-video generation tasks, starting from an initial given frame. Moreover, models like GR-2 [19, 217] benefit from the vast

scale of internet videos to establish robust priors, subsequently fine-tuning on specific robotics tasks to generate accurate image predictions and action trajectories.

Collectively, these methods demonstrate the promise of generative, vision-centric modeling as a foundation for anticipatory robotic control, significantly enhancing robots' capabilities to reason (to some extent) about future states and improve long-term task performance.

**5.2.3 From Simulation to Real-world.** Deep reinforcement learning has demonstrated remarkable capabilities in robotics, enabling autonomous performance in complex tasks such as stable locomotion [106, 192], precise object manipulation [33, 244], and intricate activities like tying shoelaces [5]. However, its practicality remains significantly limited by low sample efficiency. For instance, training a robot to solve a Rubik's Cube in the real-world can require tens of thousands of simulated years [3]. Consequently, most robot training is conducted within simulation environments, leveraging distributed training techniques to enhance efficiency [67, 170]. Unfortunately, due to discrepancies between simulations and real-world conditions, policies trained in simulation often fail when directly transferred to physical robots, particularly in complex or novel environments.

A pivotal recent insight is that world models can effectively bridge this simulation-to-reality gap by learning generalized representations of real-world dynamics. For example, NeBula [2] constructs a structured belief space that enables reasoning and rapid adaptation across diverse robot morphologies and unstructured environments. DayDreamer [220] further demonstrates the capability of generalized world models, allowing robots to directly learn locomotion in real-world environments within hours, significantly reducing the reliance on extensive simulations. Additionally, SWIM [134] highlights the power of human-video-based learning combined with minimal real-world fine-tuning, enabling task generalization with less than 30 minutes of interaction. These examples illustrate that by building robust, real-world-oriented internal representations, world models substantially narrow the gap between simulation and reality, facilitating rapid adaptation and generalization in robotics.

### 5.3 Social Simulacra

The concept of "social simulacra" was originally introduced as a prototyping technique in [149], aimed at helping designers create a virtual social computing system encompassing many diverse agents. The traditional methods of constructing agents based on expert-defined rules [13, 176] or reinforcement learning [259] face issues such as overly simplistic behaviors or a lack of interpretability. However, the emergence of LLMs provides a transformative tool for building more realistic social simulacra, achieving more convincing stylized facts [114] or accurate predictions. Social simulacra can be seen as a form of world model that mirrors real-world social computing systems. From another perspective, the agents within social simulacra also develop implicit representations of the external system; that is, they build an implicit world model that supports the generation of their social behaviors. The relationship between the world model and social simulacra is shown in Figure S2, and the summary of representative works is shown in Table 6.

**5.3.1 Building Social Simulacra Mirroring Real-world Society.** In the era of the rapid rise of LLM agents, building realistic social simulation systems becomes more practical. One of the most famous examples of social simulacra is AI Town [148], a world model composed of 25 generative agents, essentially forming a sandbox social environment. In this virtual community, the agents exhibit believable individual behaviors, and at the group level, emergent social behaviors similar to those that might appear in the real-world. Along these lines, more and more attempts are being made to replace humans in various social scenarios with LLM agents, in effect forming their

Table 6. Representative Works of Social Simulacra

	Advance	What to simulate	Effects of world model
World Model as Social Simulacra	AI Town [148]	Machine society	Stylized facts
	S3 [51]	Social network	Predictions
	Papachristou et al. [147]	Social network	Stylized facts and Predictions
	Xu et al. [228]	Games	Stylized facts
	EconAgent [114]	Macroeconomics	Stylized facts
	SRAP-Agent [91]	Resource allocation	Stylized facts
	Project Sid [4]	Collective rules (tax)	Stylized facts
World Model in Social Simulacra	AgentSociety [153]	Social life	Stylized facts
	Agent-Pro [250]	Games	Belief
	Zhang et al. [247]	Machine society	Reflection and Debate
	GovSim [154]	Resource sharing	Cognition
	AgentGroupChat [60]	Group chat	Belief and Memory

own scenario-specific social simulacra. These works have used the simulacra paradigm in such scenarios as social networks and cooperative or competitive games, among others [50].

Recent studies have extended social simulacra by demonstrating how LLM agents can reflect realistic social interactions and collective behaviors, effectively mirroring human societies. A critical insight is that LLM-driven agents, by modeling aspects such as emotions, attitudes, and decision-making patterns, can reproduce intricate social phenomena across multiple contexts. For example, S3 [51] leverages these human-like features to simulate realistic message propagation on social networks, successfully capturing the dynamics of real-world public events. Related studies [147] deepen this understanding by analyzing how network structures emerge organically among LLM agents, comparing these structures directly to those formed in human social networks. Beyond networks, the capacity of LLM agents to embody sophisticated strategic reasoning is highlighted in simulations of social games such as Werewolf [228], where the agent outputs exhibit patterns resembling human strategic behaviors such as deception and confrontation. Furthermore, in economic contexts, works such as EconAgent [114] reveal that collective behaviors emerging from individual economic reasoning by LLM agents can convincingly reproduce macroeconomic trends and regularities. Recently, AgentSociety [153] has further expanded this paradigm by creating a large-scale, LLM-powered societal simulation capable of modeling diverse social phenomena such as polarization and response to public policies, thus providing a versatile platform for computational social science research. These studies collectively illustrate the broad potential of social simulacra, demonstrating that agent-based world models powered by LLMs can capture diverse and realistic social, strategic, and economic interactions.

*5.3.2 Agent’s Understanding of External World in Social Simulacra.* LLM agents build their memory by storing observations obtained through interactions with the external environment [251], thereby forming implicit representations and basic cognition of the external world, especially in the context of simulating social scenarios. This cognition is stored in a memory bank in textual form for LLM agents to retrieve and use, enabling them to access useful information and fully leverage experiential knowledge from past interactions with the environment when making decisions.

Agent-Pro [250] transforms the memory of its interactions with the external environment (specifically with other agents in interactive tasks) into what are called “beliefs”. Based on these beliefs, it makes the next decision and updates its behavior strategy. These beliefs represent the agent’s social understanding of the environment and other agents within it, relating to Theory of Mind mentioned in Section 3.2. Other works on LLM agents have also adopted similar designs. For example, Zhang et al. [247] introduces mechanisms of reflection and debate from

a social psychology view for modeling multi-agent collaboration tasks. A more advanced study is GovSim [154], which explores whether cooperative behaviors aimed at sustainable resource development can emerge within a society composed of LLM agents. In this setup, each agent gathers information about the external world and other agents' behavioral strategies through multi-agent conversations and subsequently forms its own high-level insights, essentially creating an implicit representation of the world model. Another similar application scenario is Interactive Group Chat [60], where human-like behaviors and strategies emerge across four narrative scenarios, including Inheritance Disputes, Law Court Debates, and so on.

## 6 Open Problems and Future Directions

The recent advance of hyper-realistic generative AI has brought a lot of attention to development of the world model, with particular focus on the multi-modal big models like Sora [146]. Despite the rapid innovation, there are also a lot of important open problems that remain to be solved.

### 6.1 Physical Rules and Counterfactual Simulation

A key objective of world models is to capture the causal structure of their environments—especially the underlying physical rules—so they can reason about counterfactuals beyond the data distribution [151]. This capacity is crucial for handling rare, mission-critical events (e.g., autonomous-driving corner cases [45]) and for narrowing sim-to-real gaps. Recent progress raises the question of whether large-scale, purely data-driven generative models can acquire such rules from raw visual data alone. While transformer- and diffusion-based video generators such as Sora [146] produce strikingly realistic sequences, studies reveal persistent physical-law failures—e.g., inaccurate gravity, fluid, or thermal dynamics [213].

Hybrid approaches that explicitly embed physics are emerging as promising alternatives. Genesis [7] illustrates this direction by unifying fast, photo-realistic rendering with a re-engineered universal physics core, allowing language-conditioned data generation grounded in first-principles simulation. PhysGen [124] takes a similar stance at the image-to-video level: it couples a rigid-body simulator with a diffusion refiner, enabling controllable, physically plausible motion from a single image. Complementary diagnostic work underscores why such hybrids are needed. Kang et al. [97] show that scaling diffusion video models yields perfect in-distribution fidelity yet breaks down on out-of-distribution or combinatorial tests, indicating “case-based” rather than rule-based generalization. Motamed et al. [139] reach a similar conclusion with the Physics-IQ benchmark: current video generators achieve visual realism but largely fail on tasks requiring understanding of optics, fluid dynamics, or magnetism.

Taken together, the evidence suggests that data-driven scaling alone is insufficient to recover robust physical laws. Integrating explicit simulators—or otherwise enforcing physical priors [184]—remains a promising path toward world models that generalize to unseen counterfactual scenarios while retaining interpretability and transparency.

### 6.2 Enriching the Social Dimension

Simulating the physical elements alone is not sufficient for an advanced world model, since human behavior and social interaction also play a crucial role in many important scenarios [50, 58, 246]. For example, the behavior of urban dwellers is particularly important for building world models of the urban environment [9, 226]. Previous work shows that the human-like commonsense reasoning capabilities of LLMs provide a unique opportunity to simulate realistic human behavior with generative agents [148]. However, designing autonomous agents that can simulate realistic and comprehensive human behavior and social interactions remains an open problem. Recent studies suggest that theories of human behavior patterns and cognitive processes can inform the

design of agentic workflows, which in turn enhance the human behavior simulation capabilities of LLMs [148, 182], representing an important direction for future research. In addition, the evaluation of the realism of generated human behavior still largely relies on subjective human assessment, which is challenging to scale up to a large-scale world model. Developing a reliable and scalable evaluation scheme will be another future research direction that can enrich the social dimension of the world model.

### 6.3 Bridging Simulation and Reality with Embodied Intelligence

The world model has long been envisioned as a critical step toward developing embodied intelligence [174]. It can serve as a powerful simulator that creates comprehensive elements of the environment and models realistic relationships between them. Such an environment can facilitate embodied agents to learn through interaction with a simulated environment, reducing the need for supervision data. To achieve this goal, improving the multi-modal, multi-task, and 3D capacities of generative AI models has become an important research problem for developing general world models for embodied agents. Moreover, closing the simulation-to-reality gaps [76] has been a long-standing research problem for embodied environment simulators, and it is therefore important to transfer the trained embodied intelligence from the simulation environment to the physical world. Collecting more fine-grained sensory data is also a critical step toward this goal, which can be facilitated through the interface of embodied agents. Therefore, an interesting future research direction is to create self-reinforcing loops to harness the synergy power of generative world models and embodied agents.

### 6.4 Simulation Efficiency

Ensuring high simulation efficiency of world models is important for many applications. For example, number of frames per second is a key metric for high quality for learning sophisticated drone manipulating AIs. The popular transformer architecture of most big generative AIs poses a huge challenge for high-speed simulation because its autoregressive nature can only generate one token at a time. Several strategies are proposed to accelerate the inference of large generative models, such as incorporating big and small generative models [181] and distilling big models [182]. More holistic solutions include building a simulation platform that optimally schedule LLM requests [232]. High computation cost is also a problem for classic physics simulators when they are tasked to simulate large and complex systems. Previous research finds deep learning models like graph neural networks can be used to efficiently approximate physical systems [172]. Therefore, an important research direction will be to explore the synergy between smaller deep learning models and big generative AI models. Additionally, the overall improvement from underlying hardware to programming platform and AI models is also needed to achieve substantial speedup.

### 6.5 Ethical and Safety Concerns

**Data Privacy.** The recent trend of building world models with big generative AIs raises significant concerns of privacy risk, largely due to the massive and often opaque training data [240]. Extensive research effort is devoted to assessing the risk of inferring private information with big generative AIs like LLM [115], which could be especially sensitive in the context of video generation models. To be compliant with privacy regulations like GDPR [196], it is important to improve the transparency of the life cycle of generative AIs, helping the public understand how data is collected, stored, and used in these AI models.

**Simulating Unsafe Scenario.** The incredibly intelligent power of generative AIs makes safeguarding their access a paramount task. Previous studies on LLMs found they can be misled to generate unsafe content with adversarial prompting [85, 105]. The risk of unsafe use of world

models can be even larger. Adversarial users might leverage such techniques to simulate harmful scenarios, reducing the cost of planning illegal and unethical activities. Therefore, an important future research direction is to safeguard the usage of world models.

**Accountability.** The ability to generate hyper-realistic text, images, and videos has caused severe social problems of spreading misinformation and disinformation. For example, the emergence of deepfake technology gives rise to large-scale misuses that have widespread negative effects on social, economic, and political systems [214]. Thus, detecting AI-generated content has been a key research problem in addressing these risks [164]. However, this problem is becoming increasingly challenging due to the advance of generative AIs, and it will be even more difficult with the arrival of a world model that can generate consistent, multi-dimensional output. Technology like watermarking could help improve the accountability of world model usage [30]. More research attention, as well as legal solutions, are needed to improve the accountability of world model usage.

## 6.6 Benchmark

Benchmarking world models is both necessary and challenging. Because the community pursues divergent goals—*learning internal representations* vs. *predicting future worlds*—with heterogeneous technical approaches (e.g., LLM agents and video diffusion) and wide-ranging application domains (autonomous driving, robotics, social simulation), there is no single canonical task or metric. Nevertheless, several recent efforts illustrate how carefully designed testbeds can expose the specific gaps that prevent current models from becoming reliable world simulators.

**Video-centric world simulation.** Qin et al. propose WorldSimBench [160], a dual evaluation suite combining human-preference video judgments and action-level consistency across three embodied settings (open-ended sandboxes, autonomous driving, and robot manipulation). Complementarily, Duan et al. present WorldScore [37], which decomposes “world generation” into controllability, visual quality, and dynamics across 3,000 camera-specified scenes, enabling head-to-head comparisons of 3D, 4D, and video generators.

**Physical and spatial reasoning.** PhysBench [26] (10k video–image–text triplets) and UrbanVideo-Bench [254] (5.2k drone clips) demonstrate that current VLMs lack basic physics understanding and urban navigation skills. Xu et al. translate psychometric testing into AI evaluation by defining five Basic Spatial Abilities and diagnosing geometry and rotation weaknesses across 13 VLMs [227].

**Embodied decision making.** The **Embodied Agent Interface (EAI)** [113] standardizes four LLM-based modules—goal parsing, sub-goal generation, action sequencing, and transition modeling—and reports fine-grained errors rather than a single success rate.

Despite these advances, benchmarking world models remains an open challenge. Future work should focus on building more diverse and realistic benchmarks to rigorously test generalization capabilities. Additionally, standardizing evaluation protocols will be key to improving comparability and robustness assessments across environments.

## 7 Conclusion

Understanding the world and predicting the future have been long-standing objectives for scientists developing artificial generative intelligence, underscoring the significance of constructing world models across various domains. This article presents the first comprehensive survey of world models that systematically explores their two primary functionalities: *implicit representations* and *future predictions* of the external world. We provide an extensive summary of existing research on these core functions, with particular emphasis on world models in decision-making, world knowledge learned by models, world models as video generation, and world models as embodied environments. Additionally, we review progress in key applications of world models,

including autonomous driving, robotics, and social simulation. Finally, recognizing the unresolved challenges in this rapidly evolving field, we highlight open problems and propose promising research directions with the hope of stimulating further investigation in this burgeoning area.

## References

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. arXiv:2303.08774. Retrieved from <https://arxiv.org/abs/2303.08774>
- [2] Ali Agha, Kyohei Otsu, Benjamin Morrell, David D Fan, Rohan Thakker, Angel Santamaria-Navarro, Sung-Kyun Kim, Amanda Bouman, Xianmei Lei, Jeffrey Edlund, et al. 2021. Nebula: Quest for robotic autonomy in challenging environments; team costar at the darpa subterranean challenge. arXiv:2103.11470. Retrieved from <https://arxiv.org/abs/2103.11470>
- [3] Ilge Akkaya, Marcin Andrychowicz, Maciek Chociej, Mateusz Litwin, Bob McGrew, Arthur Petron, Alex Paino, Matthias Plappert, Glenn Powell, Raphael Ribas, et al. 2019. Solving rubik’s cube with a robot hand. arXiv:1910.07113. Retrieved from <https://arxiv.org/abs/1910.07113>
- [4] Altera AL, Andrew Ahn, Nic Becker, Stephanie Carroll, Nico Christie, Manuel Cortes, Arda Demirci, Melissa Du, Frankie Li, Shuying Luo, et al. 2024. Project sid: Many-agent simulations toward AI civilization. arXiv:2411.00114. Retrieved from <https://arxiv.org/abs/2411.00114>
- [5] Jorge Aldaco, Travis Armstrong, Robert Baruch, Jeff Bingham, Sanky Chan, Kenneth Draper, Debidatta Dwibedi, Chelsea Finn, Pete Florence, Spencer Goodrich, et al. 2024. ALOHA 2: An enhanced low-cost hardware for bimanual teleoperation. arXiv:2405.02292. Retrieved from <https://arxiv.org/abs/2405.02292>
- [6] Florent Althé and Arnaud de La Fortelle. 2017. An LSTM network for highway trajectory prediction. In *Proceedings of the 2017 IEEE 20th International Conference on Intelligent Transportation Systems*. IEEE, 353–359.
- [7] Genesis Authors. 2024. Genesis: A Generative and Universal Physics Engine for Robotics and Beyond. Retrieved June 11, 2025 from <https://github.com/Genesis-Embodied-AI/Genesis>
- [8] Xuyang Bai, Zeyu Hu, Xinge Zhu, Qingqiu Huang, Yilun Chen, Hongbo Fu, and Chiew-Lan Tai. 2022. Transfusion: Robust lidar-camera fusion for 3d object detection with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 1090–1099.
- [9] Michael Batty. 2024. Digital twins in city planning. *Nature Computational Science* 4, 3 (2024), 192–199.
- [10] Dimitri Bertsekas. 2012. *Dynamic Programming and Optimal Control: Volume I*. Vol. 4. Athena scientific.
- [11] Kevin Black, Mitsuhiko Nakamoto, Pranav Atreya, Homer Walke, Chelsea Finn, Aviral Kumar, and Sergey Levine. 2023. Zero-shot robotic manipulation with pretrained image-editing diffusion models. arXiv:2310.10639. Retrieved from <https://arxiv.org/abs/2310.10639>
- [12] Daniel Bogdoll, Yitian Yang, and J. Marius Zöllner. 2023. Muvo: A multimodal generative world model for autonomous driving with geometric representations. *arXiv e-prints* (2023), arXiv–2311.
- [13] William A. Brock and Cars H. Hommes. 1998. Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control* 22, 8–9 (1998), 1235–1274.
- [14] Tim Brooks, Bill Peebles, Connor Holmes, Will DePue, Yufei Guo, Li Jing, David Schnurr, Joe Taylor, Troy Luhman, Eric Luhman, Clarence Ng, Ricky Wang, and Aditya Ramesh. 2024. Video generation models as world simulators. (2024). Retrieved June 11, 2025 from <https://openai.com/research/video-generation-models-as-world-simulators>
- [15] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in Neural Information Processing Systems* 33 (2020), 1877–1901.
- [16] Jake Bruce, Michael D. Dennis, Ashley Edwards, Jack Parker-Holder, Yuge Shi, Edward Hughes, Matthew Lai, Aditi Mavalankar, Richie Steigerwald, Chris Apps, et al. 2024. Genie: Generative interactive environments. In *Proceedings of the 41st International Conference on Machine Learning*.
- [17] Shengqu Cai, Eric Ryan Chan, Songyou Peng, Mohamad Shahbazi, Anton Obukhov, Luc Van Gool, and Gordon Wetzstein. 2023. Diffdreamer: Towards consistent unsupervised single-view scene extrapolation with conditional diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2139–2150.
- [18] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. 2017. Matterport3D: Learning from RGB-D data in indoor environments. *International Conference on 3D Vision* (2017). IEEE, 667–676.
- [19] Chi-Lam Cheang, Guangzeng Chen, Ya Jing, Tao Kong, Hang Li, Yifeng Li, Yuxiao Liu, Hongtao Wu, Jiafeng Xu, Yichu Yang, et al. 2024. GR-2: A generative video-language-action model with web-scale knowledge for robot manipulation. arXiv:2410.06158. Retrieved from <https://arxiv.org/abs/2410.06158>

- [20] Jianyu Chen, Bodi Yuan, and Masayoshi Tomizuka. 2019. Model-free Deep Reinforcement Learning for Urban Autonomous Driving. arXiv:1904.09503. Retrieved from <https://arxiv.org/abs/1904.09503>
- [21] Zhili Cheng, Zhitong Wang, Jinyi Hu, Shengding Hu, An Liu, Yuge Tu, Pengkai Li, Lei Shi, Zhiyuan Liu, and Maosong Sun. 2024. LEGENT: Open platform for embodied agents. arXiv:2404.18243. Retrieved from <https://arxiv.org/abs/2404.18243>
- [22] Cheng Chi, Zhenjia Xu, Siyuan Feng, Eric Cousineau, Yilun Du, Benjamin Burchfiel, Russ Tedrake, and Shuran Song. 2024. Diffusion policy: Visuomotor policy learning via action diffusion. *The International Journal of Robotics Research* 0, 0 (2024), 02783649241273668.
- [23] Xiaowei Chi, Hengyuan Zhang, Chun-Kai Fan, Xingqun Qi, Rongyu Zhang, Anthony Chen, Chi-min Chan, Wei Xue, Wenhan Luo, Shanghang Zhang, et al. 2024. EVA: An embodied world model for future video anticipation. arXiv:2410.15461. Retrieved from <https://arxiv.org/abs/2410.15461>
- [24] Joseph Cho, Fachrina Dewi Puspitasari, Sheng Zheng, Jingyao Zheng, Lik-Hang Lee, Tae-Ho Kim, Choong Seon Hong, and Chaoning Zhang. 2024. Sora as an agi world model? A complete survey on text-to-video generation. arXiv:2403.05131. Retrieved from <https://arxiv.org/abs/2403.05131>
- [25] Fang-Chieh Chou, Tsung-Han Lin, Henggang Cui, Vladan Radosavljevic, Thi Nguyen, Tzu-Kuo Huang, Matthew Niedoba, Jeff Schneider, and Nemanja Djuric. 2020. Predicting motion of vulnerable road users using high-definition maps and efficient convnets. In *Proceedings of the 2020 IEEE Intelligent Vehicles Symposium*. IEEE, 1655–1662.
- [26] Wei Chow, Jiageng Mao, Boyi Li, Daniel Seita, Vitor Guizilini, and Yue Wang. 2025. PhysBench: Benchmarking and enhancing vision-language models for physical world understanding. arXiv:2501.16411. Retrieved from <https://arxiv.org/abs/2501.16411>
- [27] Kurtland Chua, Roberto Calandra, Rowan McAllister, and Sergey Levine. 2018. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. *Advances in Neural Information Processing Systems* 31 (2018).
- [28] Henggang Cui, Vladan Radosavljevic, Fang-Chieh Chou, Tsung-Han Lin, Thi Nguyen, Tzu-Kuo Huang, Jeff Schneider, and Nemanja Djuric. 2019. Multimodal trajectory predictions for autonomous driving using deep convolutional networks. In *Proceedings of the 2019 International Conference on Robotics and Automation*. IEEE, 2090–2096.
- [29] Yifan Cui, Xinyi Shan, and Jeanhun Chung. 2024. A feasibility study on RUNWAY GEN-2 for generating realistic style images. *International Journal of Internet, Broadcasting and Communication* 16, 1 (2024), 99–105.
- [30] Sumanth Dathathri, Abigail See, Sumedh Ghaisas, Po-Sen Huang, Rob McAdam, Johannes Welbl, Vandana Bachani, Alex Kaskasoli, Robert Stanforth, Tatiana Matejovicova, et al. 2024. Scalable watermarking for identifying large language model outputs. *Nature* 634, 8035 (2024), 818–823.
- [31] Matt Deitke, Eli VanderBilt, Alvaro Herrasti, Luca Weihs, Kiana Ehsani, Jordi Salvador, Winson Han, Eric Kolve, Aniruddha Kembhavi, and Roozbeh Mottaghi. 2022. ProcTHOR: Large-scale embodied AI using procedural generation. *Advances in Neural Information Processing Systems* 35 (2022), 5982–5994.
- [32] Boyang Deng, Richard Tucker, Zhengqi Li, Leonidas Guibas, Noah Snavely, and Gordon Wetzstein. 2024. Streetscapes: Large-scale consistent street view generation using autoregressive video diffusion. In *Proceedings of the ACM SIGGRAPH 2024 Conference Papers*. 1–11.
- [33] Mehmet Dogar, Andrew Spielberg, Stuart Baker, and Daniela Rus. 2019. Multi-robot grasp planning for sequential assembly operations. *Autonomous Robots* 43 (2019), 649–664.
- [34] Alexey Dosovitskiy. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. arXiv:2010.11929. Retrieved from <https://arxiv.org/abs/2010.11929>
- [35] Nan Du, Yanping Huang, Andrew M. Dai, Simon Tong, Dmitry Lepikhin, Yuanzhong Xu, Maxim Krikun, Yanqi Zhou, Adams Wei Yu, Orhan Firat, et al. 2022. Glam: Efficient scaling of language models with mixture-of-experts. In *Proceedings of the International Conference on Machine Learning*. PMLR, 5547–5569.
- [36] Yilun Du, Sherry Yang, Bo Dai, Hanjun Dai, Ofir Nachum, Josh Tenenbaum, Dale Schuurmans, and Pieter Abbeel. 2024. Learning universal policies via text-guided video generation. *Advances in Neural Information Processing Systems* 36 (2024).
- [37] Haoyi Duan, Hong-Xing Yu, Sirui Chen, Li Fei-Fei, and Jiajun Wu. 2025. Worldscore: A unified evaluation benchmark for world generation. arXiv:2504.00983. Retrieved from <https://arxiv.org/abs/2504.00983>
- [38] Hugh Durrant-Whyte and Tim Bailey. 2006. Simultaneous localization and mapping: Part I. *IEEE Robotics and Automation Magazine* 13, 2 (2006), 99–110.
- [39] Alejandro Escontrela, Ademi Adeniji, Wilson Yan, Ajay Jain, Xue Bin Peng, Ken Goldberg, Youngwoon Lee, Danijar Hafner, and Pieter Abbeel. 2024. Video prediction models as rewards for reinforcement learning. *Advances in Neural Information Processing Systems* 36 (2024).
- [40] Patrick Esser, Johnathan Chiu, Parmida Atighehchian, Jonathan Granskog, and Anastasis Germanidis. 2023. Structure and content-guided video synthesis with diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 7346–7356.

- [41] Linxi Fan, Guanzhi Wang, Yunfan Jiang, Ajay Mandlekar, Yuncong Yang, Haoyi Zhu, Andrew Tang, De-An Huang, Yuke Zhu, and Anima Anandkumar. 2022. Minedojo: Building open-ended embodied agents with internet-scale knowledge. *Advances in Neural Information Processing Systems* 35 (2022), 18343–18362.
- [42] Jie Feng, Yuwei Du, Tianhui Liu, Siqi Guo, Yuming Lin, and Yong Li. 2024. CityGPT: Empowering urban spatial cognition of large language models. arXiv:2406.13948. Retrieved from <https://arxiv.org/abs/2406.13948>
- [43] Jie Feng, Jinwei Zeng, Qingyue Long, Hongyi Chen, Jie Zhao, Yanxin Xi, Zhilun Zhou, Yuan Yuan, Shengyuan Wang, Qingbin Zeng, et al. 2025. A survey of large language model-powered spatial intelligence across scales: Advances in embodied agents, smart cities, and earth science. arXiv:2504.09848. Retrieved from <https://arxiv.org/abs/2504.09848>
- [44] Jie Feng, Jun Zhang, Junbo Yan, Xin Zhang, Tianjian Ouyang, Tianhui Liu, Yuwei Du, Siqi Guo, and Yong Li. 2024. CityBench: Evaluating the capabilities of large language model as world model. arXiv:2406.13945. Retrieved from <https://arxiv.org/abs/2406.13945>
- [45] Shuo Feng, Haowei Sun, Xintao Yan, Haojie Zhu, Zhengxia Zou, Shengyin Shen, and Henry X Liu. 2023. Dense reinforcement learning for safety validation of autonomous vehicles. *Nature* 615, 7953 (2023), 620–627.
- [46] Chelsea Finn, Ian Goodfellow, and Sergey Levine. 2016. Unsupervised learning for physical interaction through video prediction. *Advances in Neural Information Processing Systems* 29 (2016).
- [47] Chelsea Finn and Sergey Levine. 2017. Deep visual foresight for planning robot motion. In *Proceedings of the 2017 IEEE International Conference on Robotics and Automation*. IEEE, 2786–2793.
- [48] Rao Fu, Zehao Wen, Zichen Liu, and Srinath Sridhar. 2025. Anyhome: Open-vocabulary generation of structured and textured 3d homes. In *Proceedings of the European Conference on Computer Vision*. Springer, 52–70.
- [49] Chuang Gan, Jeremy Schwartz, Seth Alter, Damian Mrowca, Martin Schrimpf, James Traer, Julian De Freitas, Jonas Kubilius, Abhishek Bhandwalder, Nick Haber, et al. 2020. Threedworld: A platform for interactive multi-modal physical simulation. arXiv:2007.04954. Retrieved from <https://arxiv.org/abs/2007.04954>
- [50] Chen Gao, Xiaochong Lan, Nian Li, Yuan Yuan, Jingtao Ding, Zhilun Zhou, Fengli Xu, and Yong Li. 2024. Large language models empowered agent-based modeling and simulation: A survey and perspectives. *Humanities and Social Sciences Communications* 11, 1 (2024), 1–24.
- [51] Chen Gao, Xiaochong Lan, Zhiheng Lu, Jinzhu Mao, Jinghua Piao, Huandong Wang, Depeng Jin, and Yong Li. 2023. S3: Social-network simulation system with large language model-empowered agents. arXiv:2307.14984. Retrieved from <https://arxiv.org/abs/2307.14984>
- [52] Chen Gao, Baining Zhao, Weichen Zhang, Jinzhu Mao, Jun Zhang, Zhiheng Zheng, Fanhang Man, Jianjie Fang, Zile Zhou, Jinqiang Cui, et al. 2024. EmbodiedCity: A benchmark platform for embodied agent in real-world city environment. arXiv:2410.09604. Retrieved from <https://arxiv.org/abs/2410.09604>
- [53] Qiaozhi Gao, Govind Thattai, Suhaila Shakiah, Xiaofeng Gao, Shreyas Pansare, Vasu Sharma, Gaurav Sukhatme, Hangjie Shi, Bofei Yang, Desheng Zhang, et al. 2024. Alexa arena: A user-centric interactive platform for embodied ai. *Advances in Neural Information Processing Systems* 36 (2024).
- [54] Shenyuan Gao, Jiazhi Yang, Li Chen, Kashyap Chitta, Yihang Qiu, Andreas Geiger, Jun Zhang, and Hongyang Li. 2024. Vista: A generalizable driving world model with high fidelity and versatile controllability. arXiv:2405.17398. Retrieved from <https://arxiv.org/abs/2405.17398>
- [55] Ignat Georgiev, Varun Giridhar, Nicklas Hansen, and Animesh Garg. 2025. PWM: Policy learning with multi-task world models. In *Proceedings of the 13th International Conference on Learning Representations*.
- [56] Elliot Gestrin, Marco Kuhlmann, and Jendrik Seipp. 2024. NL2Plan: Robust LLM-driven planning from minimal text descriptions. arXiv:2405.04215. Retrieved from <https://arxiv.org/abs/2405.04215>
- [57] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. 2014. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 580–587.
- [58] Jiahui Gong, Jingtao Ding, Fanjin Meng, Chen Yang, Hong Chen, Zuoqian Wang, Haisheng Lu, and Yong Li. 2025. BehaveGPT: A foundation model for large-scale user behavior modeling. arXiv:2505.17631. Retrieved from <https://arxiv.org/abs/2505.17631>
- [59] James Gornet and Matt Thomson. 2024. Automated construction of cognitive maps with visual predictive coding. *Nature Machine Intelligence* 6, 7 (2024), 820–833.
- [60] Zhouhong Gu, Xiaoxuan Zhu, Haoran Guo, Lin Zhang, Yin Cai, Hao Shen, Jiangjie Chen, Zheyu Ye, Yifei Dai, Yan Gao, et al. 2024. Agent group chat: An interactive group chat simulacra for better eliciting collective emergent behavior. arXiv:2403.13433. Retrieved from <https://arxiv.org/abs/2403.13433>
- [61] Lin Guan, Karthik Valmeekam, Sarath Sreedharan, and Subbarao Kambhampati. 2023. Leveraging pre-trained large language models to construct and utilize world models for model-based task planning. *Advances in Neural Information Processing Systems* 36 (2023), 79081–79094.
- [62] Yanchen Guan, Haicheng Liao, Zhenning Li, Jia Hu, Runze Yuan, Yunjian Li, Guohui Zhang, and Chengzhong Xu. 2024. World models for autonomous driving: An initial survey. *IEEE Transactions on Intelligent Vehicles* (2024).

- [63] Cole Gulino, Justin Fu, Wenjie Luo, George Tucker, Eli Bronstein, Yiren Lu, Jean Harb, Xinlei Pan, Yan Wang, Xianguy Chen, et al. 2024. Waymax: An accelerated, data-driven simulator for large-scale autonomous driving research. *Advances in Neural Information Processing Systems* 36 (2024).
- [64] Wes Gurnee and Max Tegmark. 2024. Language models represent space and time. In *Proceedings of the 12th International Conference on Learning Representations*.
- [65] David Ha and Jürgen Schmidhuber. 2018. Recurrent world models facilitate policy evolution. *Advances in Neural Information Processing Systems* 31 (2018).
- [66] David Ha and Jürgen Schmidhuber. 2018. World models. arXiv:1803.10122. Retrieved from <https://arxiv.org/abs/1803.10122>
- [67] Sehoon Ha, Peng Xu, Zhenyu Tan, Sergey Levine, and Jie Tan. 2020. Learning to walk in the real world with minimal human effort. arXiv:2002.08550. Retrieved from <https://arxiv.org/abs/2002.08550>
- [68] Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. 2019. Dream to control: Learning behaviors by latent imagination. arXiv:1912.01603. Retrieved from <https://arxiv.org/abs/1912.01603>
- [69] Danijar Hafner, Timothy Lillicrap, Ian Fischer, Ruben Villegas, David Ha, Honglak Lee, and James Davidson. 2019. Learning latent dynamics for planning from pixels. In *Proceedings of the International Conference on Machine Learning*. PMLR, 2555–2565.
- [70] Danijar Hafner, Timothy Lillicrap, Mohammad Norouzi, and Jimmy Ba. 2020. Mastering atari with discrete world models. arXiv:2010.02193. Retrieved from <https://arxiv.org/abs/2010.02193>
- [71] Danijar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. 2025. Mastering diverse control tasks through world models. *Nature* 640 (2025), 647–653.
- [72] Nicklas Hansen, Hao Su, and Xiaolong Wang. 2024. TD-MPC2: Scalable, robust world models for continuous control. In *Proceedings of the 12th International Conference on Learning Representations*.
- [73] Haoran He, Chenjia Bai, Ling Pan, Weinan Zhang, Bin Zhao, and Xuelong Li. 2024. Learning an actionable discrete diffusion policy via large-scale actionless video pre-training. In *Proceedings of the 38th Annual Conference on Neural Information Processing Systems*.
- [74] Juncai He and Jinchao Xu. 2019. MgNet: A unified framework of multigrid and convolutional neural network. *Science China Mathematics* 62, 7 (2019), 1331–1354. DOI : <https://doi.org/10.1007/s11425-019-9547-2>
- [75] Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems* 33 (2020), 6840–6851.
- [76] Sebastian Höfer, Kostas Bekris, Ankur Handa, Juan Camilo Gamboa, Melissa Mozifian, Florian Golemo, Chris Atkeson, Dieter Fox, Ken Goldberg, John Leonard, et al. 2021. Sim2real in robotics and automation: Applications and challenges. *IEEE Transactions on Automation Science and Engineering* 18, 2 (2021), 398–400.
- [77] Anthony Hu, Lloyd Russell, Hudson Yeo, Zak Murez, George Fedoseev, Alex Kendall, Jamie Shotton, and Gianluca Corrado. 2023. Gaia-1: A generative world model for autonomous driving. arXiv:2309.17080. Retrieved from <https://arxiv.org/abs/2309.17080>
- [78] Anthony Hu, Lloyd Russell, Hudson Yeo, Zak Murez, George Fedoseev, Alex Kendall, Jamie Shotton, and Gianluca Corrado. 2023. GAIA-1: A Generative World Model for Autonomous Driving. arXiv:2309.17080. Retrieved from <https://arxiv.org/abs/2309.17080>
- [79] Yihan Hu, Jiazhi Yang, Li Chen, Keyu Li, Chonghao Sima, Xizhou Zhu, Siqi Chai, Senyao Du, Tianwei Lin, Wenhai Wang, Lewei Lu, Xiaosong Jia, Qiang Liu, Jifeng Dai, Yu Qiao, and Hongyang Li. 2023. Planning-oriented Autonomous Driving. arXiv:2212.10156. Retrieved from <https://arxiv.org/abs/2212.10156>
- [80] Yi-Qi Hu, Hong Qian, and Yang Yu. 2017. Sequential classification-based optimization for direct policy search. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- [81] Pu Hua, Minghuan Liu, Annabella Macaluso, Yunfeng Lin, Weinan Zhang, Huazhe Xu, and Lirui Wang. 2024. GenSim2: Scaling robot data generation with multi-modal and reasoning LLMs. arXiv:2410.03645. Retrieved from <https://arxiv.org/abs/2410.03645>
- [82] Wenlong Huang, Fei Xia, Ted Xiao, Harris Chan, Jacky Liang, Pete Florence, Andy Zeng, Jonathan Tompson, Igor Mordatch, Yevgen Chebotar, et al. 2022. Inner monologue: Embodied reasoning through planning with language models. arXiv:2207.05608. Retrieved from <https://arxiv.org/abs/2207.05608>
- [83] Yanjun Huang, Jiatong Du, Ziru Yang, Zewei Zhou, Lin Zhang, and Hong Chen. 2022. A survey on trajectory-prediction methods for autonomous driving. *IEEE Transactions on Intelligent Vehicles* 7, 3 (2022), 652–674. DOI : <https://doi.org/10.1109/TIV.2022.3167103>
- [84] Zhiyu Huang, Xiaoyu Mo, and Chen Lv. 2022. Multi-modal motion prediction with transformer-based neural network for autonomous driving. In *Proceedings of the 2022 International Conference on Robotics and Automation*. IEEE, 2605–2611.
- [85] Hakan Inan, Kartikeya Upasani, Jianfeng Chi, Rashi Rungta, Krithika Iyer, Yuning Mao, Michael Tontchev, Qing Hu, Brian Fuller, Davide Testuggine, et al. 2023. Llama guard: Llm-based input-output safeguard for human-ai conversations. arXiv:2312.06674. Retrieved from <https://arxiv.org/abs/2312.06674>

- [86] Anna A Ivanova, Aalok Sathe, Benjamin Lipkin, Unnathi Kumar, Setayesh Radkani, Thomas H Clark, Carina Kauf, Jennifer Hu, RT Pramod, Gabriel Grand, et al. 2024. Elements of world knowledge (EWOK): A cognition-inspired framework for evaluating basic world knowledge in language models. arXiv:2405.09605. Retrieved from <https://arxiv.org/abs/2405.09605>
- [87] Yash Jain, Anshul Nasery, Vibhav Vineet, and Harkirat Behl. 2024. Peekaboo: Interactive video generation via masked-diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 8079–8088.
- [88] Eric Jang, Alex Irpan, Mohi Khansari, Daniel Kappler, Frederik Ebert, Corey Lynch, Sergey Levine, and Chelsea Finn. 2022. Bc-z: Zero-shot task generalization with robotic imitation learning. In *Proceedings of the Conference on Robot Learning*. PMLR, 991–1002.
- [89] Michael Janner, Justin Fu, Marvin Zhang, and Sergey Levine. 2019. When to trust your model: Model-based policy optimization. *Advances in Neural Information Processing Systems* 32 (2019).
- [90] Michael Janner, Qiyang Li, and Sergey Levine. 2021. Offline reinforcement learning as one big sequence modeling problem. *Advances in Neural Information Processing Systems* 34 (2021), 1273–1286.
- [91] Jiarui Ji, Yang Li, Hongtao Liu, Zhicheng Du, Zhewei Wei, Weiran Shen, Qi Qi, and Yankai Lin. 2024. SRAP-Agent: Simulating and optimizing scarce resource allocation policy with LLM-based agent. arXiv:2410.14152. Retrieved from <https://arxiv.org/abs/2410.14152>
- [92] Chiyu Max Jiang, Andre Cornman, Cheolho Park, Ben Sapp, Yin Zhou, and Dragomir Anguelov. 2023. MotionDiffuser: Controllable Multi-Agent Motion Prediction using Diffusion. arXiv:2306.03083. Retrieved from <https://arxiv.org/abs/2306.03083>
- [93] Charles Jin and Martin Rinard. 2024. Emergent representations of program semantics in language models trained on programs. In *Proceedings of the 41st International Conference on Machine Learning*.
- [94] Philip Nicholas Johnson-Laird. 1983. *Mental Models: Towards a Cognitive Science of Language, Inference, and Consciousness*. Number 6. Harvard University Press.
- [95] Gregory Kahn, Adam Villafior, Vitvich Pong, Pieter Abbeel, and Sergey Levine. 2017. Uncertainty-aware reinforcement learning for collision avoidance. arXiv:1702.01182. Retrieved from <https://arxiv.org/abs/1702.01182>
- [96] Subbarao Kambhampati, Karthik Valmeekam, Lin Guan, Mudit Verma, Kaya Stechly, Siddhant Bhambri, Lucas Paul Saldyt, and Anil B. Murthy. 2024. Position: LLMs can’t plan, but can help planning in LLM-modulo frameworks. In *Proceedings of the 41st International Conference on Machine Learning*.
- [97] Bingyi Kang, Yang Yue, Rui Lu, Zhijie Lin, Yang Zhao, Kaixin Wang, Gao Huang, and Jiashi Feng. 2024. How far is video generation from world model: A physical law perspective. arXiv:2411.02385. Retrieved from <https://arxiv.org/abs/2411.02385>
- [98] Atsushi Kawasaki and Akihito Seki. 2020. Multimodal trajectory predictions for urban environments using geometric relationships between a vehicle and lanes. In *Proceedings of the 2020 IEEE International Conference on Robotics and Automation*. IEEE, 9203–9209.
- [99] Oussama Khatib. 1987. A unified approach for motion and force control of robot manipulators: The operational space formulation. *IEEE Journal on Robotics and Automation* 3, 1 (1987), 43–53.
- [100] Kilian Kleiberger, Richard Borrmann, Werner Kraus, and Marco F Huber. 2020. A survey on learning-based robotic grasping. *Current Robotics Reports* 1 (2020), 239–249.
- [101] Eric Kolve, Roozbeh Mottaghi, Winson Han, Eli VanderBilt, Luca Weihs, Alvaro Herrasti, Matt Deitke, Kiana Ehsani, Daniel Gordon, Yuke Zhu, et al. 2017. Ai2-thor: An interactive 3d environment for visual ai. arXiv:1712.05474. Retrieved from <https://arxiv.org/abs/1712.05474>
- [102] Basil Kouvaritakis and Mark Cannon. 2016. Model predictive control. *Switzerland: Springer International Publishing* 38 (2016), 13–56.
- [103] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. 2012. Imagenet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems* 25 (2012).
- [104] Kuaishou. 2024. Kling AI: Next-Generation AI Creative Studio. Retrieved May 06, 2025 from <https://www.klingai.com/global/>.
- [105] Aounon Kumar, Chirag Agarwal, Suraj Srinivas, Soheil Feizi, and Hima Lakkaraju. 2023. Certifying llm safety against adversarial prompting. arXiv:2309.02705. Retrieved from <https://arxiv.org/abs/2309.02705>
- [106] Ashish Kumar, Zipeng Fu, Deepak Pathak, and Jitendra Malik. 2021. Rma: Rapid motor adaptation for legged robots. arXiv:2107.04034. Retrieved from <https://arxiv.org/abs/2107.04034>
- [107] Varun Ravi Kumar, Senthil Yogamani, Hazem Rashed, Ganesh Sistu, Christian Witt, Isabelle Leang, Stefan Milz, and Patrick Mäder. 2023. OmniDet: Surround View Cameras based Multi-task Visual Perception Network for Autonomous Driving. arXiv:2102.07448. Retrieved from <https://arxiv.org/abs/2102.07448>
- [108] Thanard Kurutach, Ignasi Clavera, Yan Duan, Aviv Tamar, and Pieter Abbeel. 2018. Model-ensemble trust-region policy optimization. arXiv:1802.10592. Retrieved from <https://arxiv.org/abs/1802.10592>

- [109] Yann LeCun. 2022. A path towards autonomous machine intelligence version 0.9. 2, 2022-06-27. *Open Review* 62, 1 (2022), 1–62.
- [110] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 11 (1998), 2278–2324.
- [111] Jiaang Li, Yova Kementchedjieva, Constanza Fierro, and Anders Søgaard. 2024. Do Vision and language models share concepts? A vector space alignment study. *Transactions of the Association for Computational Linguistics* 12 (2024), 1232–1249.
- [112] Lincan Li, Wei Shao, Wei Dong, Yijun Tian, Qiming Zhang, Kaixiang Yang, and Wenjie Zhang. 2024. Data-centric evolution in autonomous driving: A comprehensive survey of big data system, data mining, and closed-loop technologies. arXiv:2401.12888. Retrieved from <https://arxiv.org/abs/2401.12888>
- [113] Manling Li, Shiyu Zhao, Qineng Wang, Kangrui Wang, Yu Zhou, Sanjana Srivastava, Cem Gokmen, Tony Lee, Erran Li Li, Ruohan Zhang, et al. 2024. Embodied agent interface: Benchmarking llms for embodied decision making. *Advances in Neural Information Processing Systems* 37 (2024), 100428–100534.
- [114] Nian Li, Chen Gao, Mingyu Li, Yong Li, and Qingmin Liao. 2024. Econagent: Large language model-empowered agents for simulating macroeconomic activities. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 15523–15536.
- [115] Qinbin Li, Junyuan Hong, Chulin Xie, Jeffrey Tan, Rachel Xin, Junyi Hou, Xavier Yin, Zhun Wang, Dan Hendrycks, Zhangyang Wang, et al. 2024. Llm-pbe: Assessing data privacy in large language models. arXiv:2408.12787. Retrieved from <https://arxiv.org/abs/2408.12787>
- [116] Quanyi Li, Zhenghao Peng, Lan Feng, Qihang Zhang, Zhenghai Xue, and Bolei Zhou. 2022. Metadrive: Composing diverse driving scenarios for generalizable reinforcement learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 45, 3 (2022), 3461–3475.
- [117] Yuxiao Li, Eric J. Michaud, David D. Baek, Joshua Engels, Xiaoqing Sun, and Max Tegmark. 2024. The Geometry of Concepts: Sparse Autoencoder Feature Structure. arXiv:2410.19750. Retrieved from <https://arxiv.org/abs/2410.19750>
- [118] Zhiqi Li, Wenhai Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Yu Qiao, and Jifeng Dai. 2022. BEVFormer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers. arXiv:2203.17270. Retrieved from <https://arxiv.org/abs/2203.17270>
- [119] Jessy Lin, Yuqing Du, Olivia Watkins, Danijar Hafner, Pieter Abbeel, Dan Klein, and Anca Dragan. 2024. Learning to Model the World with Language. arXiv:2308.01399. Retrieved from <https://arxiv.org/abs/2308.01399>
- [120] Kevin Lin, Christopher Agia, Toki Migimatsu, Marco Pavone, and Jeannette Bohg. 2023. Text2motion: From natural language instructions to feasible plans. *Autonomous Robots* 47, 8 (2023), 1345–1365.
- [121] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. Visual Instruction Tuning. *Advances in Neural Information Processing Systems* 36 (2023), 34892–34916.
- [122] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2024. Visual instruction tuning. *Advances in Neural Information Processing Systems* 36 (2024).
- [123] Hao Liu, Wilson Yan, Matei Zaharia, and Pieter Abbeel. 2024. World model on million-length video and language with ringattention. arXiv:2402.08268. Retrieved from <https://arxiv.org/abs/2402.08268>
- [124] Shaowei Liu, Zhongzheng Ren, Saurabh Gupta, and Shenlong Wang. 2024. Physgen: Rigid-body physics-grounded image-to-video generation. In *Proceedings of the European Conference on Computer Vision*. Springer, 360–378.
- [125] Zhihan Liu, Hao Hu, Shenao Zhang, Hongyi Guo, Shuqi Ke, Boyi Liu, and Zhaoran Wang. 2024. Reason for Future, Act for Now: A Principled Framework for Autonomous LLM Agents with Provable Sample Efficiency. arXiv:2309.17382. Retrieved from <https://arxiv.org/abs/2309.17382>
- [126] Yuxing Long, Xiaoqi Li, Wenzhe Cai, and Hao Dong. 2024. Discuss before moving: Visual language navigation via multi-expert discussions. In *Proceedings of the 2024 IEEE International Conference on Robotics and Automation*. IEEE, 17380–17387.
- [127] Pablo Alvarez Lopez, Michael Behrisch, Laura Bieker-Walz, Jakob Erdmann, Yun-Pang Flötteröd, Robert Hilbrich, Leonhard Lücken, Johannes Rummel, Peter Wagner, and Evamarie Wießner. 2018. Microscopic traffic simulation using SUMO. In *Proceedings of the 21st IEEE Intelligent Transportation Systems Conference*. Retrieved from <https://elib.dlr.de/124092/>
- [128] Fan-Ming Luo, Tian Xu, Hang Lai, Xiong-Hui Chen, Weinan Zhang, and Yang Yu. 2024. A survey on model-based reinforcement learning. *Science China Information Sciences* 67, 2 (2024), 121101.
- [129] Yuping Luo, Huazhe Xu, Yuanzhi Li, Yuandong Tian, Trevor Darrell, and Tengyu Ma. 2018. Algorithmic framework for model-based deep reinforcement learning with theoretical guarantees. arXiv:1807.03858. Retrieved from <https://arxiv.org/abs/1807.03858>
- [130] Xinji Mai, Zeng Tao, Junxiong Lin, Haoran Wang, Yang Chang, Yanlan Kang, Yan Wang, and Wenqiang Zhang. 2024. From efficient multimodal models to world models: A survey. arXiv:2407.00118. Retrieved from <https://arxiv.org/abs/2407.00118>

- [131] Arjun Majumdar, Anurag Ajay, Xiaohan Zhang, Pranav Putta, Sriram Yenamandra, Mikael Henaff, Sneha Silwal, Paul Mcvay, Oleksandr Maksymets, Sergio Arnaud, et al. 2024. Openeqa: Embodied question answering in the era of foundation models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 16488–16498.
- [132] Rohin Manvi, Samar Khanna, Marshall Burke, David Lobell, and Stefano Ermon. 2024. Large language models are geographically biased. arXiv:2402.02680. Retrieved from <https://arxiv.org/abs/2402.02680>
- [133] Rohin Manvi, Samar Khanna, Gengchen Mai, Marshall Burke, David Lobell, and Stefano Ermon. 2023. GeoLLM: Extracting geospatial knowledge from large language models. arXiv:2310.06213. Retrieved from <https://arxiv.org/abs/2310.06213>
- [134] Russell Mendonca, Shikhar Bahl, and Deepak Pathak. 2023. Structured world models from human videos. arXiv:2308.10901. Retrieved from <https://arxiv.org/abs/2308.10901>
- [135] Chen Min, Dawei Zhao, Liang Xiao, Yiming Nie, and Bin Dai. 2023. Uniworld: Autonomous driving pre-training via world models. arXiv:2308.07234. Retrieved from <https://arxiv.org/abs/2308.07234>
- [136] Marvin Minsky. 1974. A framework for representing knowledge. MIT, Cambridge.
- [137] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, et al. 2015. Human-level control through deep reinforcement learning. *Nature* 518, 7540 (2015), 529–533.
- [138] Thomas M. Moerland, Joost Broekens, Aske Plaat, and Catholijn M. Jonker. 2018. A0c: Alpha zero in continuous action space. arXiv:1805.09613. Retrieved from <https://arxiv.org/abs/1805.09613>
- [139] Saman Motamed, Laura Culp, Kevin Swersky, Priyank Jaini, and Robert Geirhos. 2025. Do generative video models learn physical principles from watching videos? arXiv:2501.09038. Retrieved from <https://arxiv.org/abs/2501.09038>
- [140] Fangwen Mu, Lin Shi, Song Wang, Zhuohao Yu, Binquan Zhang, Chenxue Wang, Shichao Liu, and Qing Wang. 2023. ClarifyGPT: Empowering LLM-based code generation with intention clarification. arXiv:2310.10996. Retrieved from <https://arxiv.org/abs/2310.10996>
- [141] Anusha Nagabandi, Gregory Kahn, Ronald S. Fearing, and Sergey Levine. 2018. Neural network dynamics for model-based deep reinforcement learning with model-free fine-tuning. In *Proceedings of the 2018 IEEE International Conference on Robotics and Automation*. IEEE, 7559–7566.
- [142] Nigamaa Nayakanti, Rami Al-Rfou, Aurick Zhou, Kratarth Goel, Khaled S. Refaat, and Benjamin Sapp. 2022. Wayformer: Motion Forecasting via Simple and Efficient Attention Networks. arXiv:2207.05844. Retrieved from <https://arxiv.org/abs/2207.05844>
- [143] Jiquan Ngiam, Benjamin Caine, Vijay Vasudevan, Zhengdong Zhang, Hao-Tien Lewis Chiang, Jeffrey Ling, Rebecca Roelofs, Alex Bewley, Chenxi Liu, Ashish Venugopal, et al. 2021. Scene transformer: A unified multi-task model for behavior prediction and planning. arXiv:2106.08417. Retrieved from <https://arxiv.org/abs/2106.08417>
- [144] Junhyuk Oh, Satinder Singh, and Honglak Lee. 2017. Value prediction network. *Advances in Neural Information Processing Systems* 30 (2017).
- [145] OpenAI. 2022. Introducing ChatGPT. Retrieved May 06, 2025 from <https://openai.com/blog/chatgpt>.
- [146] OpenAI. 2024. Sora: Creating video from text. Retrieved May 06, 2025 from <https://openai.com/sora>.
- [147] Marios Papachristou and Yuan Yuan. 2024. Network formation and dynamics among multi-LLMs. arXiv:2402.10659. Retrieved from <https://arxiv.org/abs/2402.10659>
- [148] Joon Sung Park, Joseph O’Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2023. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. 1–22.
- [149] Joon Sung Park, Lindsay Popowski, Carrie Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. 2022. Social simulacra: Creating populated prototypes for social computing systems. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*. 1–18.
- [150] Sudipta Paul, Amit Roy-Chowdhury, and Anoop Cherian. 2022. Avlen: Audio-visual-language embodied navigation in 3d environments. *Advances in Neural Information Processing Systems* 35 (2022), 6236–6249.
- [151] Judea Pearl. 2009. Causal inference in statistics: An overview. (2009).
- [152] Tung Phan-Minh, Elena Corina Grigore, Freddy A. Boulton, Oscar Beijbom, and Eric M. Wolff. 2020. Covernet: Multimodal behavior prediction using trajectory sets. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 14074–14083.
- [153] Jinghua Piao, Yuwei Yan, Jun Zhang, Nian Li, Junbo Yan, Xiaochong Lan, Zhihong Lu, Zhiheng Zheng, Jing Yi Wang, Di Zhou, et al. 2025. AgentSociety: Large-scale simulation of LLM-driven generative agents advances understanding of human behaviors and society. arXiv:2502.08691. Retrieved from <https://arxiv.org/abs/2502.08691>
- [154] Giorgio Piatti, Zhijiang Jin, Max Kleiman-Weiner, Bernhard Schölkopf, Mrinmaya Sachan, and Rada Mihalcea. 2024. Cooperate or collapse: Emergence of sustainability behaviors in a society of LLM agents. arXiv:2404.16698. Retrieved from <https://arxiv.org/abs/2404.16698>

- [155] David Premack and Guy Woodruff. 1978. Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences* 1, 4 (1978), 515–526.
- [156] Xavier Puig, Kevin Ra, Marko Boben, Jiaman Li, Tingwu Wang, Sanja Fidler, and Antonio Torralba. 2018. Virtualhome: Simulating household activities via programs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 8494–8502.
- [157] Charles R. Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. 2017. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 652–660.
- [158] Charles R. Qi, Li Yi, Hao Su, and Leonidas J. Guibas. 2017. PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space. arXiv:1706.02413. Retrieved from <https://arxiv.org/abs/1706.02413>
- [159] Guocheng Qian, Yuchen Li, Houwen Peng, Jinjie Mai, Hasan Hammoud, Mohamed Elhoseiny, and Bernard Ghanem. 2022. Pointnext: Revisiting pointnet++ with improved training and scaling strategies. *Advances in Neural Information Processing Systems* 35 (2022), 23192–23204.
- [160] Yiran Qin, Zhelun Shi, Jiwen Yu, Xijun Wang, Enshen Zhou, Lijun Li, Zhenfei Yin, Xihui Liu, Lu Sheng, Jing Shao, et al. 2024. Worldsmbench: Towards video generation models as world simulators. arXiv:2410.18072. Retrieved from <https://arxiv.org/abs/2410.18072>
- [161] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning Transferable Visual Models From Natural Language Supervision. arXiv:2103.00020. Retrieved from <https://arxiv.org/abs/2103.00020>
- [162] Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D. Manning, Stefano Ermon, and Chelsea Finn. 2024. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems* 36 (2024).
- [163] Aravind Rajeswaran, Igor Mordatch, and Vikash Kumar. 2020. A game theoretic framework for model based reinforcement learning. In *Proceedings of the International Conference on Machine Learning*. PMLR, 7953–7963.
- [164] Md Shohel Rana, Mohammad Nur Nobil, Beddhu Murali, and Andrew H. Sung. 2022. Deepfake detection: A systematic literature review. *IEEE Access* 10 (2022), 25494–25513.
- [165] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Proceedings of the Advances in Neural Information Processing Systems*. C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett (Eds.), Vol. 28, Curran Associates, Inc. Retrieved from [https://proceedings.neurips.cc/paper\\_files/paper/2015/file/14bfa6bb14875e45bba028a21ed38046-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2015/file/14bfa6bb14875e45bba028a21ed38046-Paper.pdf)
- [166] Weiming Ren, Huan Yang, Ge Zhang, Cong Wei, Xinrun Du, Wenhao Huang, and Wenhui Chen. 2024. Consist2v: Enhancing visual consistency for image-to-video generation. arXiv:2402.04324. Retrieved from <https://arxiv.org/abs/2402.04324>
- [167] Jonathan Richens, David Abel, Alexis Bellot, and Tom Everitt. 2025. General agents need world models. arXiv:2506.01622. Retrieved from <https://arxiv.org/abs/2506.01622>
- [168] Marc Rigger, Tarun Gupta, Agrim Hilmkil, and Chao Ma. 2024. AVID: Adapting video diffusion models to world models. arXiv:2410.12822. Retrieved from <https://arxiv.org/abs/2410.12822>
- [169] Jonathan Roberts, Timo Lüddecke, Sowmen Das, Kai Han, and Samuel Albanie. 2023. GPT4GEO: How a language model sees the world’s geography. arXiv:2306.00020. Retrieved from <https://arxiv.org/abs/2306.00020>
- [170] Nikita Rudin, David Hoeller, Philipp Reist, and Marco Hutter. 2022. Learning to walk in minutes using massively parallel deep reinforcement learning. In *Proceedings of the Conference on Robot Learning*. PMLR, 91–100.
- [171] Mohammad Reza Samsami, Artem Zholus, Janarthanan Rajendran, and Sarath Chandar. 2024. Mastering memory tasks with world models. arXiv:2403.04253. Retrieved from <https://arxiv.org/abs/2403.04253>
- [172] Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter Battaglia. 2020. Learning to simulate complex physics with graph networks. In *Proceedings of the International Conference on Machine Learning*. PMLR, 8459–8468.
- [173] Maarten Sap, Ronan LeBras, Daniel Fried, and Yejin Choi. 2022. Neural theory-of-mind? On the limits of social intelligence in large lms. arXiv:2210.13312. Retrieved from <https://arxiv.org/abs/2210.13312>
- [174] Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, et al. 2019. Habitat: A platform for embodied ai research. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 9339–9347.
- [175] Dhruv Mauria Saxena, Sangjae Bae, Alireza Nakhaei, Kikuo Fujimura, and Maxim Likhachev. 2020. Driving in dense traffic with model-free reinforcement learning. In *Proceedings of the 2020 IEEE International Conference on Robotics and Automation*. IEEE. DOI : <https://doi.org/10.1109/icra40945.2020.9197132>
- [176] Thomas C. Schelling. 1971. Dynamic models of segregation. *Journal of Mathematical Sociology* 1, 2 (1971), 143–186.
- [177] Ingmar Schubert, Jingwei Zhang, Jake Bruce, Sarah Bechtel, Emilio Parisotto, Martin Riedmiller, Jost Tobias Springenberg, Arunkumar Byravan, Leonard Hasenclever, and Nicolas Heess. 2023. A generalist dynamics model for control. arXiv:2305.10912. Retrieved from <https://arxiv.org/abs/2305.10912>

- [178] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. arXiv:1707.06347. Retrieved from <https://arxiv.org/abs/1707.06347>
- [179] Dhruv Shah, Ajay Sridhar, Arjun Bhorkar, Noriaki Hirose, and Sergey Levine. 2023. Gnm: A general navigation model to drive any robot. In *Proceedings of the 2023 IEEE International Conference on Robotics and Automation*. IEEE, 7226–7233.
- [180] Yu Shang, Jiansheng Chen, Hangyu Fan, Jingtao Ding, Jie Feng, and Yong Li. 2024. UrbanWorld: An urban world model for 3D city generation. arXiv:2407.11965. Retrieved from <https://arxiv.org/abs/2407.11965>
- [181] Yu Shang, Yu Li, Fengli Xu, and Yong Li. 2024. DefInt: A default-interventionist framework for efficient reasoning with hybrid large language models. arXiv:2402.02563. Retrieved from <https://arxiv.org/abs/2402.02563>
- [182] Chenyang Shao, Fengli Xu, Bingbing Fan, Jingtao Ding, Yuan Yuan, Meng Wang, and Yong Li. 2024. Beyond imitation: Generating human mobility from context-aware reasoning with large language models. arXiv:2402.09836. Retrieved from <https://arxiv.org/abs/2402.09836>
- [183] Bokui Shen, Fei Xia, Chengshu Li, Roberto Martín-Martín, Linxi Fan, GuanZhi Wang, Claudia Pérez-D’Arpino, Shyamal Buch, Sanjana Srivastava, Lyne Tchappin, et al. 2021. igibson 1.0: A simulation environment for interactive tasks in large realistic scenes. In *Proceedings of the 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 7520–7527.
- [184] Hongzhi Shi, Jingtao Ding, Yufan Cao, Li Liu, Yong Li, et al. 2022. Learning symbolic models for graph-structured physical mechanism. In *Proceedings of the 11th International Conference on Learning Representations*.
- [185] Haochen Shi, Huazhe Xu, Zhiao Huang, Yunzhu Li, and Jiajun Wu. 2024. RoboCraft: Learning to see, simulate, and shape elasto-plastic objects in 3D with graph networks. *The International Journal of Robotics Research* 43, 4 (2024), 533–549.
- [186] Haojun Shi, Suyu Ye, Xinyu Fang, Chuanyang Jin, Layla Isik, Yen-Ling Kuo, and Tianmin Shu. 2024. MuMA-ToM: Multi-modal multi-agent theory of mind. arXiv:2408.12574. Retrieved from <https://arxiv.org/abs/2408.12574>
- [187] Shaoshuai Shi, Li Jiang, Dengxin Dai, and Bernt Schiele. 2022. Motion transformer with global intention localization and local movement refinement. *Advances in Neural Information Processing Systems* 35 (2022), 6531–6543.
- [188] Mohit Shridhar, Xingdi Yuan, Marc-Alexandre Côté, Yonatan Bisk, Adam Trischler, and Matthew Hausknecht. 2020. Aleworld: Aligning text and embodied environments for interactive learning. arXiv:2010.03768. Retrieved from <https://arxiv.org/abs/2010.03768>
- [189] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. 2016. Mastering the game of go with deep neural networks and tree search. *Nature* 529, 7587 (2016), 484–489.
- [190] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. 2017. Mastering the game of go without human knowledge. *Nature* 550, 7676 (2017), 354–359.
- [191] Reid G. Simmons. 1994. Structured control for autonomous robots. *IEEE Transactions on Robotics and Automation* 10, 1 (1994), 34–43.
- [192] Laura Smith, Ilya Kostrikov, and Sergey Levine. 2022. A walk in the park: Learning to walk in 20 minutes with model-free reinforcement learning. arXiv:2208.07860. Retrieved from <https://arxiv.org/abs/2208.07860>
- [193] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. 2020. Score-based generative modeling through stochastic differential equations. arXiv:2011.13456. Retrieved from <https://arxiv.org/abs/2011.13456>
- [194] James WA Strachan, Dalila Albergo, Giulia Borghini, Oriana Pansardi, Eugenio Scaliti, Saurabh Gupta, Krati Saxena, Alessandro Rufo, Stefano Panzeri, Guido Manzi, et al. 2024. Testing theory of mind in large language models and humans. *Nature Human Behaviour* (2024), 1–11.
- [195] Winnie Street, John Oliver Siy, Geoff Keeling, Adrien Baranes, Benjamin Barnett, Michael McKibben, Tatenda Kanyere, Alison Lentz, Robin IM Dunbar, et al. 2024. LLMs achieve adult human performance on higher-order theory of mind tasks. arXiv:2405.18870. Retrieved from <https://arxiv.org/abs/2405.18870>
- [196] Damian A Tamburri. 2020. Design principles for the general data protection regulation (GDPR): A formal concept analysis and its evaluation. *Information Systems* 91 (2020), 101469.
- [197] ManyCore Research Team. 2025. SpatialLM: Large Language Model for Spatial Understanding. Retrieved June 11, 2025 from <https://github.com/manycore-research/SpatialLM>.
- [198] Marvin Teichmann, Michael Weber, Marius Zoellner, Roberto Cipolla, and Raquel Urtasun. 2018. MultiNet: Real-time Joint Semantic Reasoning for Autonomous Driving. arXiv:1612.07695. Retrieved from <https://arxiv.org/abs/1612.07695>
- [199] Ran Tian, Boyi Li, Xinshuo Weng, Yuxiao Chen, Edward Schmerling, Yue Wang, Boris Ivanovic, and Marco Pavone. 2024. Tokenize the World into Object-level Knowledge to Address Long-tail Events in Autonomous Driving. arXiv:2407.00959. Retrieved from <https://arxiv.org/abs/2407.00959>

- [200] Edward C. Tolman. 1948. Cognitive maps in rats and men. *Psychological Review* 55, 4 (1948), 189.
- [201] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. arXiv:2302.13971. Retrieved from <https://arxiv.org/abs/2302.13971>
- [202] A Vaswani. 2017. Attention is all you need. *Advances in Neural Information Processing Systems* 30 (2017).
- [203] Dat Vu, Bao Ngo, and Hung Phan. 2022. HybridNets: End-to-End Perception Network. arXiv:2203.09035. Retrieved from <https://arxiv.org/abs/2203.09035>
- [204] Ao Wang, Hui Chen, Zijia Lin, Jungong Han, and Guiguang Ding. 2024. Repvit: Revisiting mobile cnn from vit perspective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 15909–15920.
- [205] Hanqing Wang, Jiahe Chen, Wensi Huang, Qingwei Ben, Tai Wang, Boyu Mi, Tao Huang, Siheng Zhao, Yilun Chen, Sizhe Yang, et al. 2024. Grutopia: Dream general robots in a city at scale. arXiv:2407.10943. Retrieved from <https://arxiv.org/abs/2407.10943>
- [206] Lirui Wang, Yiyang Ling, Zhecheng Yuan, Mohit Shridhar, Chen Bao, Yuzhe Qin, Bailin Wang, Huazhe Xu, and Xiaolong Wang. 2023. Gensim: Generating robotic simulation tasks via large language models. arXiv:2310.01361. Retrieved from <https://arxiv.org/abs/2310.01361>
- [207] Lening Wang, Wenzhao Zheng, Yilong Ren, Han Jiang, Zhiyong Cui, Haiyang Yu, and Jiwen Lu. 2024. OccSora: 4D occupancy generation models as world simulators for autonomous driving. arXiv:2405.20337. Retrieved from <https://arxiv.org/abs/2405.20337>
- [208] Tingwu Wang and Jimmy Ba. 2019. Exploring model-based planning with policy networks. arXiv:1906.08649. Retrieved from <https://arxiv.org/abs/1906.08649>
- [209] Xiaofeng Wang, Zheng Zhu, Guan Huang, Xinze Chen, Jiagang Zhu, and Jiwen Lu. 2023. Drivedreamer: Towards real-world-driven world models for autonomous driving. arXiv:2309.09777. Retrieved from <https://arxiv.org/abs/2309.09777>
- [210] Xiaofeng Wang, Zheng Zhu, Guan Huang, Xinze Chen, Jiagang Zhu, and Jiwen Lu. 2023. DriveDreamer: Towards Real-world-driven World Models for Autonomous Driving. arXiv:2309.09777. Retrieved from <https://arxiv.org/abs/2309.09777>
- [211] Xiaofeng Wang, Zheng Zhu, Guan Huang, Boyuan Wang, Xinze Chen, and Jiwen Lu. 2024. Worlddreamer: Towards general world models for video generation via predicting masked tokens. arXiv:2401.09985. Retrieved from <https://arxiv.org/abs/2401.09985>
- [212] Yuqi Wang, Jiawei He, Lue Fan, Hongxin Li, Yuntao Chen, and Zhaoxiang Zhang. 2023. Driving into the Future: Multiview Visual Forecasting and Planning with World Model for Autonomous Driving. arXiv:2311.17918. Retrieved from <https://arxiv.org/abs/2311.17918>
- [213] Yi Ru Wang, Jiafei Duan, Dieter Fox, and Siddhartha Srinivasa. 2023. NEWTON: Are large language models capable of physical reasoning? arXiv:2310.07018. Retrieved from <https://arxiv.org/abs/2310.07018>
- [214] Mika Westerlund. 2019. The emergence of deepfake technology: A review. *Technology Innovation Management Review* 9, 11 (2019), 39–52.
- [215] Alex Wilf, Sihyun Shawn Lee, Paul Pu Liang, and Louis-Philippe Morency. 2023. Think Twice: Perspective-Taking Improves Large Language Models’ Theory-of-Mind Capabilities. arXiv:2311.10227. Retrieved from <https://arxiv.org/abs/2311.10227>
- [216] Dong Wu, Man-Wen Liao, Wei-Tian Zhang, Xing-Gang Wang, Xiang Bai, Wen-Qing Cheng, and Wen-Yu Liu. 2022. YOLOP: You only look once for panoptic driving perception. *Machine Intelligence Research* 19, 6 (2022), 550–562. DOI: <https://doi.org/10.1007/s11633-022-1339-y>
- [217] Hongtao Wu, Ya Jing, Chilam Cheang, Guangzeng Chen, Jiafeng Xu, Xinghang Li, Minghuan Liu, Hang Li, and Tao Kong. 2023. Unleashing large-scale video generative pre-training for visual robot manipulation. arXiv:2312.13139. Retrieved from <https://arxiv.org/abs/2312.13139>
- [218] Jincenzi Wu, Zhuang Chen, Jiawen Deng, Sahand Sabour, Helen Meng, and Minlie Huang. 2024. Coke: A cognitive knowledge graph for machine theory of mind. arXiv:2305.05390. Retrieved from <https://arxiv.org/abs/2305.05390>
- [219] Jialong Wu, Shaofeng Yin, Ningya Feng, Xu He, Dong Li, Jianye Hao, and Mingsheng Long. 2024. iVideoGPT: Interactive VideoGPTs are scalable world models. arXiv:2405.15223. Retrieved from <https://arxiv.org/abs/2405.15223>
- [220] Philipp Wu, Alejandro Escontrela, Danijar Hafner, Pieter Abbeel, and Ken Goldberg. 2023. Daydreamer: World models for physical robot learning. In *Proceedings of the Conference on Robot Learning*. PMLR, 2226–2240.
- [221] Wayne Wu, Honglin He, Yiran Wang, Chenda Duan, Jack He, Zhizheng Liu, Quanyi Li, and Bolei Zhou. 2024. MetaUrban: A simulation platform for embodied AI in urban spaces. arXiv:2407.08725. Retrieved from <https://arxiv.org/abs/2407.08725>
- [222] Fanbo Xiang, Yuzhe Qin, Kaichun Mo, Yikuan Xia, Hao Zhu, Fangchen Liu, Minghua Liu, Hanxiao Jiang, Yifu Yuan, He Wang, et al. 2020. Sapien: A simulated part-based interactive environment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 11097–11107.

- [223] Jiannan Xiang, Guangyi Liu, Yi Gu, Qiyue Gao, Yuting Ning, Yuheng Zha, Zeyu Feng, Tianhua Tao, Shibo Hao, Yemin Shi, et al. 2024. Pandora: Towards general world model with natural language actions and video states. arXiv:2406.09455. Retrieved from <https://arxiv.org/abs/2406.09455>
- [224] Jiannan Xiang, Tianhua Tao, Yi Gu, Tianmin Shu, Zirui Wang, Zichao Yang, and Zhiting Hu. 2024. Language models meet world models: Embodied experiences enhance language models. *Advances in Neural Information Processing Systems* 36 (2024).
- [225] Fengli Xu, Qianyue Hao, Zefang Zong, Jingwei Wang, Yunke Zhang, Jingyi Wang, Xiaochong Lan, Jiahui Gong, Tianjian Ouyang, Fanjin Meng, et al. 2025. Towards large reasoning models: A survey of reinforced reasoning with large language models. arXiv:2501.09686. Retrieved from <https://arxiv.org/abs/2501.09686>
- [226] Fengli Xu, Jun Zhang, Chen Gao, Jie Feng, and Yong Li. 2023. Urban generative intelligence (ugi): A foundational platform for agents in embodied city environment. arXiv:2312.11813. Retrieved from <https://arxiv.org/abs/2312.11813>
- [227] Wenrui Xu, Dalin Lyu, Weihang Wang, Jie Feng, Chen Gao, and Yong Li. 2025. Defining and evaluating visual language models' basic spatial abilities: A perspective from psychometrics. arXiv:2502.11859. Retrieved from <https://arxiv.org/abs/2502.11859>
- [228] Yuzhuang Xu, Shuo Wang, Peng Li, Fuwen Luo, Xiaolong Wang, Weidong Liu, and Yang Liu. 2023. Exploring large language models for communication games: An empirical study on werewolf. arXiv:2309.04658. Retrieved from <https://arxiv.org/abs/2309.04658>
- [229] Wilson Yan, Danijar Hafner, Stephen James, and Pieter Abbeel. 2023. Temporally consistent transformers for video generation. In *Proceedings of the International Conference on Machine Learning*. PMLR, 39062–39098.
- [230] Wilson Yan, Yunzhi Zhang, Pieter Abbeel, and Aravind Srinivas. 2021. Videogpt: Video generation using vq-vae and transformers. arXiv:2104.10157. Retrieved from <https://arxiv.org/abs/2104.10157>
- [231] Xu Yan, Haiming Zhang, Yingjie Cai, Jingming Guo, Weichao Qiu, Bin Gao, Kaiqiang Zhou, Yue Zhao, Huan Jin, Jiantao Gao, et al. 2024. Forging vision foundation models for autonomous driving: Challenges, methodologies, and opportunities. arXiv:2401.08045. Retrieved from <https://arxiv.org/abs/2401.08045>
- [232] Yuwei Yan, Qingbin Zeng, Zhiheng Zheng, Jingzhe Yuan, Jie Feng, Jun Zhang, Fengli Xu, and Yong Li. 2024. OpenCity: A scalable platform to simulate urban activities with massive LLM agents. arXiv:2410.21286. Retrieved from <https://arxiv.org/abs/2410.21286>
- [233] Deshun Yang, Luhui Hu, Yu Tian, Zihao Li, Chris Kelly, Bang Yang, Cindy Yang, and Yuexian Zou. 2024. WorldGPT: A sora-inspired video AI agent as rich world models from text and image inputs. arXiv:2403.07944. Retrieved from <https://arxiv.org/abs/2403.07944>
- [234] Mengjiao Yang, Yilun Du, Bo Dai, Dale Schuurmans, Joshua B. Tenenbaum, and Pieter Abbeel. 2023. Probabilistic adaptation of text-to-video models. arXiv:2306.01872. Retrieved from <https://arxiv.org/abs/2306.01872>
- [235] Mengjiao Yang, Yilun Du, Kamyar Ghasemipour, Jonathan Tompson, Dale Schuurmans, and Pieter Abbeel. 2023. Learning interactive real-world simulators. arXiv:2310.06114. Retrieved from <https://arxiv.org/abs/2310.06114>
- [236] Sherry Yang, Yilun Du, Seyed Kamyar Seyed Ghasemipour, Jonathan Tompson, Leslie Pack Kaelbling, Dale Schuurmans, and Pieter Abbeel. 2024. Learning interactive real-world simulators. In *Proceedings of the 12th International Conference on Learning Representations*.
- [237] Sherry Yang, Jacob Walker, Jack Parker-Holder, Yilun Du, Jake Bruce, Andre Barreto, Pieter Abbeel, and Dale Schuurmans. 2024. Video as the new language for real-world decision making. arXiv:2402.17139. Retrieved from <https://arxiv.org/abs/2402.17139>
- [238] Yue Yang, Fan-Yun Sun, Luca Weihs, Eli VanderBilt, Alvaro Herrasti, Winson Han, Jiajun Wu, Nick Haber, Ranjay Krishna, Lingjie Liu, et al. 2024. Holodeck: Language guided generation of 3d embodied ai environments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 16227–16237.
- [239] Zeyuan Yang, Jiageng Liu, Peihao Chen, Anoop Cherian, Tim K. Marks, Jonathan Le Roux, and Chuang Gan. 2024. RILA: Reflective and imaginative language agent for zero-shot semantic audio-visual navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 16251–16261.
- [240] Yifan Yao, Jinhao Duan, Kaidi Xu, Yuanfang Cai, Zhibo Sun, and Yue Zhang. 2024. A survey on large language model (llm) security and privacy: The good, the bad, and the ugly. *High-Confidence Computing* 4, 2 (2024), 100211.
- [241] Shengming Yin, Chenfei Wu, Huan Yang, Jianfeng Wang, Xiaodong Wang, Minheng Ni, Zhengyuan Yang, Linjie Li, Shuguang Liu, Fan Yang, et al. 2023. Nuwa-xl: Diffusion over diffusion for extremely long video generation. arXiv:2303.12346. Retrieved from <https://arxiv.org/abs/2303.12346>
- [242] Jifan Yu, Xiaozhi Wang, Shangqing Tu, Shulin Cao, Daniel Zhang-Li, Xin Lv, Hao Peng, Zijun Yao, Xiaohan Zhang, Hanming Li, et al. 2023. Kola: Carefully benchmarking world knowledge of large language models. arXiv:2306.09296. Retrieved from <https://arxiv.org/abs/2306.09296>
- [243] Lijun Yu, Yong Cheng, Kihyuk Sohn, José Lezama, Han Zhang, Huiwen Chang, Alexander G. Hauptmann, Ming-Hsuan Yang, Yuan Hao, Irfan Essa, et al. 2023. Magvit: Masked generative video transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10459–10469.

- [244] Sheng Yu, Di-Hua Zhai, Yuanqing Xia, Haoran Wu, and Jun Liao. 2022. SE-ResUNet: A novel robotic grasp detection method. *IEEE Robotics and Automation Letters* 7, 2 (2022), 5238–5245.
- [245] Yang Yu, Hong Qian, and Yi-Qi Hu. 2016. Derivative-free optimization via classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- [246] Yuan Yuan, Jingtao Ding, Depeng Jin, and Yong Li. 2025. Learning the complexity of urban mobility with deep generative network. *PNAS Nexus* 4, 5 (2025), pga081.
- [247] Jintian Zhang, Xin Xu, and Shumin Deng. 2023. Exploring collaboration mechanisms for llm agents: A social psychology view. arXiv:2310.02124. Retrieved from <https://arxiv.org/abs/2310.02124>
- [248] Lunjun Zhang, Yuwen Xiong, Ze Yang, Sergio Casas, Rui Hu, and Raquel Urtasun. 2024. Copilot4D: Learning Unsupervised World Models for Autonomous Driving via Discrete Diffusion. arXiv:2311.01017. Retrieved from <https://arxiv.org/abs/2311.01017>
- [249] Tianyuan Zhang, Hong-Xing Yu, Rundi Wu, Brandon Y Feng, Changxi Zheng, Noah Snaveley, Jiajun Wu, and William T Freeman. 2025. Physdreamer: Physics-based interaction with 3d objects via video generation. In *Proceedings of the European Conference on Computer Vision*. Springer, 388–406.
- [250] Wenqi Zhang, Ke Tang, Hai Wu, Mengna Wang, Yongliang Shen, Guiyang Hou, Zeqi Tan, Peng Li, Yueting Zhuang, and Weiming Lu. 2024. Agent-pro: Learning to evolve via policy-level reflection and optimization. arXiv:2402.17574. Retrieved from <https://arxiv.org/abs/2402.17574>
- [251] Zeyu Zhang, Xiaohe Bo, Chen Ma, Rui Li, Xu Chen, Quanyu Dai, Jieming Zhu, Zhenhua Dong, and Ji-Rong Wen. 2024. A survey on the memory mechanism of large language model based agents. arXiv:2404.13501. Retrieved from <https://arxiv.org/abs/2404.13501>
- [252] Zhejun Zhang, Alexander Liniger, Dengxin Dai, Fisher Yu, and Luc Van Gool. 2023. Trafficbots: Towards world models for autonomous driving simulation and motion prediction. In *Proceedings of the 2023 IEEE International Conference on Robotics and Automation*. IEEE, 1522–1529.
- [253] Zhejun Zhang, Alexander Liniger, Christos Sakaridis, Fisher Yu, and Luc Van Gool. 2023. Real-Time Motion Prediction via Heterogeneous Polyline Transformer with Relative Pose Encoding. arXiv:2310.12970. Retrieved from <https://arxiv.org/abs/2310.12970>
- [254] Baining Zhao, Jianjie Fang, Zichao Dai, Ziyu Wang, Jirong Zha, Weichen Zhang, Chen Gao, Yue Wang, Jinqiang Cui, Xinlei Chen, et al. 2025. Urbanvideo-bench: Benchmarking vision-language models on embodied intelligence with video data in urban spaces. arXiv:2503.06157. Retrieved from <https://arxiv.org/abs/2503.06157>
- [255] Ganlong Zhao, Guanbin Li, Weikai Chen, and Yizhou Yu. 2024. OVER-NAV: Elevating iterative vision-and-language navigation with open-vocabulary detection and structured representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 16296–16306.
- [256] Guosheng Zhao, Xiaofeng Wang, Zheng Zhu, Xinze Chen, Guan Huang, Xiaoyi Bao, and Xingang Wang. 2024. DriveDreamer-2: LLM-Enhanced World Models for Diverse Driving Video Generation. arXiv:2403.06845. Retrieved from <https://arxiv.org/abs/2403.06845>
- [257] Zirui Zhao, Wee Sun Lee, and David Hsu. 2024. Large language models as commonsense knowledge for large-scale task planning. *Advances in Neural Information Processing Systems* 36 (2024).
- [258] Haoyu Zhen, Xiaowen Qiu, Peihao Chen, Jincheng Yang, Xin Yan, Yilun Du, Yining Hong, and Chuang Gan. 2024. 3d-vla: A 3d vision-language-action generative world model. arXiv:2403.09631. Retrieved from <https://arxiv.org/abs/2403.09631>
- [259] Stephan Zheng, Alexander Trott, Sunil Srinivasa, David C. Parkes, and Richard Socher. 2022. The AI economist: Taxation policy design via two-level deep multiagent reinforcement learning. *Science Advances* 8, 18 (2022), eabk2607.
- [260] Wenzhao Zheng, Weiliang Chen, Yuanhui Huang, Borui Zhang, Yueqi Duan, and Jiwen Lu. 2023. Occworld: Learning a 3d occupancy world model for autonomous driving. arXiv:2311.16038. Retrieved from <https://arxiv.org/abs/2311.16038>
- [261] Wenzhao Zheng, Weiliang Chen, Yuanhui Huang, Borui Zhang, Yueqi Duan, and Jiwen Lu. 2025. Occworld: Learning a 3d occupancy world model for autonomous driving. In *Proceedings of the European Conference on Computer Vision*. Springer, 55–72.
- [262] Hongyu Zhou, Zheng Ge, Zeming Li, and Xiangyu Zhang. 2022. MatrixVT: Efficient Multi-Camera to BEV Transformation for 3D Perception. arXiv:2211.10593. Retrieved from <https://arxiv.org/abs/2211.10593>
- [263] Siyuan Zhou, Yilun Du, Jiaben Chen, Yandong Li, Dit-Yan Yeung, and Chuang Gan. 2024. RoboDreamer: Learning compositional world models for robot imagination. arXiv:2404.12377. Retrieved from <https://arxiv.org/abs/2404.12377>
- [264] Zikang Zhou, Jianping Wang, Yung-Hui Li, and Yu-Kai Huang. 2023. Query-centric trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 17863–17873.
- [265] Fangqi Zhu, Hongtao Wu, Song Guo, Yuxiao Liu, Chilam Cheang, and Tao Kong. 2024. Irasim: Learning interactive real-robot action simulators. arXiv:2406.14540. Retrieved from <https://arxiv.org/abs/2406.14540>

[266] Zheng Zhu, Xiaofeng Wang, Wangbo Zhao, Chen Min, Nianchen Deng, Min Dou, Yuqi Wang, Botian Shi, Kai Wang, Chi Zhang, et al. 2024. Is sora a world simulator? A comprehensive survey on general world models and beyond. arXiv:2405.03520. Retrieved from <https://arxiv.org/abs/2405.03520>

[267] Alex Zyner, Stewart Worrall, and Eduardo Nebot. 2019. Naturalistic driver intention and path prediction using recurrent neural networks. *IEEE Transactions on Intelligent Transportation Systems* 21, 4 (2019), 1584–1594.

Appendices

A Related Survey

Table S1. Comparison with Existing Surveys

Survey	Venue and Year	Main Focus	Deficiency
[266]	Arxiv, 2024	General world model	Limited to discussion on applications
[130]	Arxiv, 2024	Efficient multimodal models	Limited to discussion on techniques
[24]	Arxiv, 2024	Text-to-video generation	Limited scope
[62]	IEEE T-IV, 2024	Autonomous driving	Limited scope
[112]	Arxiv, 2024	Autonomous driving	Limited scope
[231]	Arxiv, 2024	Autonomous driving	Limited scope

This paper focuses on a comprehensive overview of the systematic definition and the capabilities of world models.

B Figures and Tables

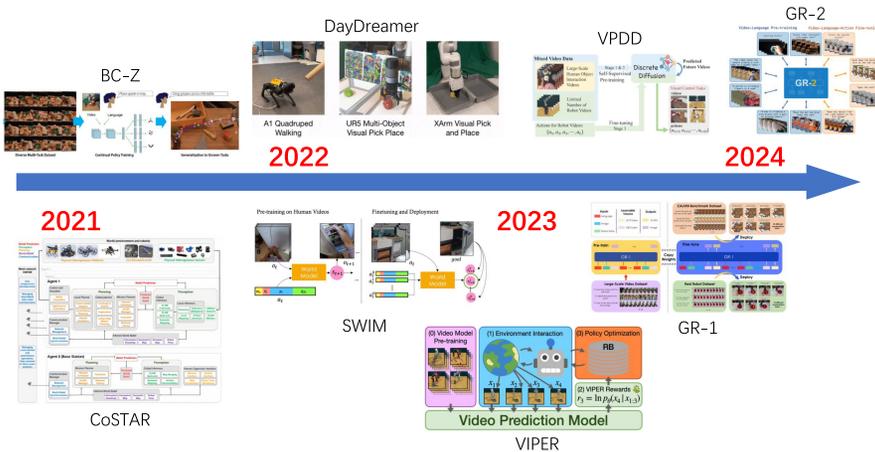


Fig. S1. The development of the robotic world model.

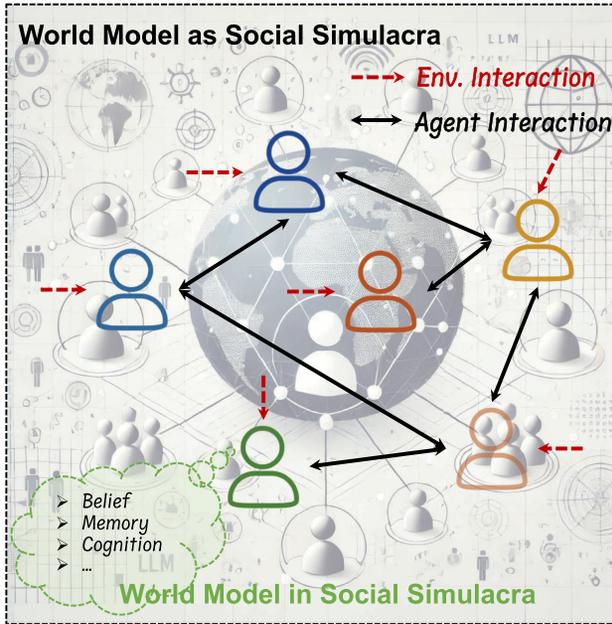


Fig. S2. World model and social simulacra.

Received 21 November 2024; revised 11 June 2025; accepted 13 June 2025