

VGGT-World: Transforming VGGT into an Autoregressive Geometry World Model

Xiangyu Sun[†], Shijie Wang[†], Fengyi Zhang[†], Lin Liu[‡], Caiyan Jia[‡], Ziyang Song[‡], Zi Huang[†], and Yadan Luo^{†,*}

[†]UQMM Lab, The University of Queensland, Brisbane, Australia
{xiangyu.sun, shijie.wang, fengyi.zhang, y.luo}@uq.edu.au

[‡]Beijing Jiaotong University, Beijing, China
liulin010811@gmail.com, {cyjia, 22110110}@bjtu.edu.cn

Abstract. World models that forecast scene evolution by generating future video frames devote the bulk of their capacity to photometric details, yet the resulting predictions often remain geometrically inconsistent. We present VGGT-World, a *geometry world model* that side-steps video generation entirely and instead forecasts the temporal evolution of frozen geometry-foundation-model (GFM) features. Concretely, we repurpose the latent tokens of a frozen VGGT as the world state and train a lightweight temporal flow transformer to autoregressively predict their future trajectory. Two technical challenges arise in this high-dimensional ($d=1024$) feature space: (i) standard velocity-prediction flow matching collapses, and (ii) autoregressive rollout suffers from compounding exposure bias. We address the first with a clean-target (z -prediction) parameterization that yields a substantially higher signal-to-noise ratio, and the second with a two-stage latent flow-forcing curriculum that progressively conditions the model on its own partially denoised rollouts. Experiments on KITTI, Cityscapes, and TartanAir demonstrate that VGGT-World significantly outperforms the strongest baselines in depth forecasting while running $3.6\text{-}5\times$ faster with only 0.43B trainable parameters, establishing frozen GFM features as an effective and efficient predictive state for 3D world modeling. The code will be open-source at <https://github.com/SimonSun0810/VGGT-World>.

Keywords: Geometry World Model · Flow Matching · VGGT

1 Introduction

Recent 3D geometry foundation models (GFMs) [18, 24, 27, 35, 49, 50, 53], such as VGGT [49] have demonstrated that large transformers can recover holistic 3D scene structure from monocular RGB observations, including camera poses, depth maps, and point clouds. These models provide embodied agents with a strong geometric representation [28] of the observed world, especially benefiting tasks like collision-free navigation [20, 22, 55] and manipulation [51, 56]. However,

* Corresponding author.




Fig. 1: From video world models to geometry world models. Video world models predict future RGB in VAE latent space, coupling scene dynamics with appearance reconstruction. As a result, decoded predictions can remain geometrically invalid, with broken scene layout and mislocalized actors despite plausible video appearance. VGGT-World instead uses frozen geometry-foundation features as the latent state and models their temporal evolution directly, yielding a VAE-free, lightweight, and geometry-consistent alternative for future 3D forecasting.

they remain fundamentally *observational*: they infer geometry from *available views*, but do not predict how that geometry will evolve over time. For planning and control in dynamic environments, an agent needs not only perception, but also **foresight**: the ability to anticipate future scene states that remain consistent with underlying geometric and physical dynamics.

A natural way to equip agents with such foresight is to cast video generation models as *world models*. Representative systems such as Wan [47], MovieGen [37], and Cosmos [36] typically compress videos into VAE latents [25], and train diffusion or autoregressive predictors to reconstruct future RGB frames. While visually impressive, this video-centric paradigm is *poorly* suited to geometry forecasting: its objective prioritizes photometric details such as texture, lighting, and appearance changes, which is loosely coupled to 3D structure. As illustrated in Fig. 1, decoded future videos can still produce severely *distorted* geometry: the recovered structure may collapse into fragmented sheet-like surfaces, and dynamic objects can be mislocalized (*e.g.*, a cyclist may be shifted in front of the ego vehicle). Such failures are particularly concerning because they can corrupt scene understanding and potentially mislead downstream decisions. Even geometry-aware video generation variants that inject geometric priors, such as Gen3R [19], GeoWorld [48], GeometryForcing [55], and GeoVideo [5], remain cen-

tered on future RGB reconstruction rather than direct prediction of 3D geometric evolution. Beyond this representational mismatch, the paradigm is computationally *expensive*: Cosmos reportedly requires tens of millions of GPU-hours, while even Gen3R requires multi-day training on 24 H20 GPUs to finetune its VAE, making such pipelines impractical for most academic labs.

These limitations suggest that the key issue is not the forecasting mechanism itself, but how the world state is represented. Instead of learning a video VAE space, we argue that GFM features that already encodes multi-view geometry and metric structure are a more suitable state space for predictive 3D modeling. We term this perspective **geometry world modeling**: world modeling is formulated as geometry-state evolution rather than future video construction. Building on this idea, we introduce **VGGT-World** (Fig. 1), which repurposes frozen VGGT features as the world state and learns only their temporal evolution. Concretely, we extract tokens from observed frames as conditions and train a temporal flow transformer to autoregressively predict future camera and patch tokens in a chunk-wise manner. The predicted tokens are then decoded using a frozen VGGT decoder and heads into 3D geometric outputs such as depth, point maps, and camera-motion cues, enabling arbitrary-horizon geometry forecasting without video VAE training or heavy video-backbone adaptation.

However, learning a geometry world model directly in frozen GFM feature space is *non-trivial*. Compared with standard video VAE latents ($d = 16$), VGGT states are weakly compressed and remain substantially higher-dimensional ($d = 1024$), making direct flow-based prediction difficult to optimize. In addition, autoregressive rollout introduces exposure bias, since each predicted chunk must condition on previously generated geometry states that may gradually drift away from the ground truth. To address these challenges, we adopt a z-prediction parameterization for stable learning in high-dimensional geometry space, together with a simple two-stage flow-forcing curriculum strategy that bridges teacher-forced training and inference-time rollout. These designs enable stable and effective learning of geometry dynamics in high-dimensional GFM feature space.

We evaluate VGGT-World on multiple datasets, including KITTI [16], Cityscapes [10], and TartanAir [52], across a range of geometry forecasting tasks such as future depth prediction, 3D point-map forecasting, camera trajectory preservation, and efficiency analysis. Across KITTI and Cityscapes, VGGT-World consistently surpasses prior world models, reducing mean AbsRel by up to 21% and 32% over the strongest prior baselines while improving δ_1 across both short- and mid-term horizons. It is also markedly more efficient, being 3.6-5 $\times$ faster than Cosmos (12B) and Gen3R (9B) with only 0.43B trainable parameters. Together, these results establish frozen geometry foundation features as an effective and efficient predictive state for 3D geometry world modeling.

2 Related Work

Geometry Foundation Model. 3D scene understanding has recently shifted from iterative geometric optimization pipelines, such as Structure-from-Motion

(SfM) [2, 14, 39, 42, 54] and Bundle Adjustment [34, 45], toward end-to-end learning. Early feed-forward multi-view reconstruction models such as DUST3R [50] demonstrate that camera poses, depth, and point clouds can be inferred directly from unposed image collections in a single forward pass. Building on this paradigm, recent geometry foundation models (GFMs), including VGGT [49], π^3 [53], and MapAnything [24, 35] scale geometry learning toward modality-general, representations with flexible input-output geometry prediction. However, these models remain limited to observational inputs: They excel at analyzing what is from available views but lack the ability to model temporal dynamics for forecasting future states.

Video Generation. Predictive modeling of scene evolution has largely been studied in the video generation community [17, 44]. Frontier systems [11, 59] such as SVD [7], Sora [33], Cosmos [36], Wan2.1 [47] learn powerful spatiotemporal priors via bidirectional denoising over full temporal windows, achieving high perceptual realism and motion coherence. Despite recent progress, these models operate primarily in pixel [57] or compressed video latent spaces [59] and thus lack explicit 3D structure, which limits geometric consistency over long horizons. To address this, a growing line of works introduces geometric priors into video generation. Representative models such as Geometry Forcing [55], GeoWorld [48], and Gen3R [19] align diffusion generation with multi-view or reconstruction-based 3D representations, improving view consistency and structural stability. This trend reflects an emerging consensus that predictive world modeling requires geometry-aware representations rather than purely appearance-based video latents. Nevertheless, most generative video models still rely on VAE-compressed latents that entangle geometry and appearance, limiting causal 3D dynamics.

Latent World Model. Recent latent world models [6, 13, 23, 30, 61] increasingly replace VAE latents [40] with the feature spaces of strong vision encoders such as DINOv2 [8], DINOv3 [41] and World4Drive [61]. DINO-Foresight [23] forecasts frozen DINO feature tokens with a masked feature transformer to enable future depth or segmentation prediction without pixels, while DINO-World [6] learns video dynamics directly in DINOv2 feature space, showing that DINO features alone can serve as a generalist world-model latent. LAW [30] and WOTE [31] performs future prediction directly in the BEV latent space. Together, these results suggest that predictive performance depends more on the choice of latent representation than on pixel-level reconstruction. However, DINO representations are primarily optimized for semantic invariance rather than explicit 3D geometry. In contrast, GFMs encode multi-view geometry, spatial consistency, and metric structure, making them inherently better suited as latent backbones for predictive 3D world modeling. We therefore advocate autoregressive world modeling directly in geometry foundation feature spaces.

3 Our Approach

Overview. As depicted in Fig. 1, we formulate world modeling as predicting the temporal evolution of *geometry states* in the latent space of a frozen geom-




Fig. 2: Comparison of v -prediction and z -prediction in frozen VGGT latent space. Left: z -prediction consistently achieves substantially higher signal-to-noise ratio than v -prediction in the layer-4 latent space. Middle: A target VGGT feature map and its corresponding RGB frame. Right: PCA visualizations of predicted latents during training. v -prediction remains noisy even after prolonged training, whereas z -prediction progressively recovers structured latent patterns aligned with the target VGGT feature.

etry foundation model, rather than reconstructing future RGB frames. VGGT-World separates spatial representation learning from temporal dynamics modeling: a frozen VGGT encoder extracts deterministic geometry states from observed frames (§3.1), a lightweight temporal flow transformer predicts their future evolution (§3.2), and the frozen VGGT decoder together with the original 3D task heads converts the predicted trajectory into geometric outputs. This design avoids training a video VAE or adapting a heavy video-generation backbone while preserving compatibility with the pretrained geometry decoder. To mitigate train-test mismatch during recursive rollout, we further introduce a two-stage latent flow-forcing curriculum that gradually exposes the model to its own generated histories (§3.3).

3.1 Geometry World Modeling via Autoregressive Flow

To model the temporal evolution of scene geometry, we first extract deterministic geometry states from observed history frames (k views), then iteratively forecast their future evolution using an autoregressive flow model, and finally decode the assembled trajectory into full 3D representations.

Geometry State Representation. Formally, given a video sequence of T frames $\mathbf{I}_{1:T} = (I_1, I_2, \dots, I_T)$, our goal is to model its temporal geometric evolution. We reuse the representation space of a frozen pretrained VGGT model $\Phi(\cdot)$, which consists of $L = 48$ (24 self-attention and 24 global-attention) transformer blocks. The frozen VGGT 3D heads $\mathcal{H}_{3D}(\cdot)$ consume a multi-scale feature hierarchy from layers $\{4, 11, 17, 23\}$ in both block types. Predicting deeper layers (*e.g.* L_{17}) would discard the required early layer (*e.g.* $L_{\{4,11\}}$) representation, which cannot be recovered by the forward network. We therefore use the L_4 output as the predictive state, as it is the earliest decoder-compatible representation. We

partition Φ into a deterministic geometry encoder Φ_{enc} (layer 0-4) and a feature propagator Φ_{dec} (layers 5-47) to produce tokenized clean geometry states:

$$\mathbf{z}_t = \Phi_{\text{enc}}(I_t) \in \mathbb{R}^{N \times d}, \quad (1)$$

where t is the discrete temporal index and $d = 1024$ is the embedding dimension. The sequence length $N = 5 + N_p$ comprises N_p spatial patch tokens and 5 special tokens (*e.g.*, camera and register tokens).

Chunk-wise Autoregressive Rollout. Rather than predicting one future frame at a time, we forecast geometry states in temporal chunks of length m . Given an observed history of k states $\mathbf{z}_{1:k}$, our goal is to forecast the geometric evolution up to a future horizon T . Assuming the future spans S chunks such that $T = k + S \cdot m$, we denote the i -th target chunk as $\mathbf{Z}_i = \mathbf{z}_{k+(i-1)m+1:k+im}$. Let $\mathbf{c}_i$ denote the context used to predict chunk $\mathbf{Z}_i$ formed by the most recent k geometry states in the trajectory. The rollout distribution factorizes as:

$$p(\mathbf{z}_{k+1:T} \mid \mathbf{z}_{1:k}) = \prod_{i=1}^S p(\mathbf{Z}_i \mid \mathbf{c}_i). \quad (2)$$

To maintain a constant context size and bound memory complexity, we do not accumulate a growing history of context during inference. Instead, the model acts as a sliding window: it predicts a full chunk of m states $\hat{\mathbf{Z}}_i$, and then extracts the latest k frames from the recently decoded sequence to serve as the fixed-size condition $\mathbf{c}_{i+1} = \hat{\mathbf{Z}}_i$ for the subsequent step. This allows the model to iteratively roll out chunk-by-chunk to construct the full sequence without context explosion.

3D Joint Decoding with Forecasted Horizon. A unique property of VGGT-World is that forecasting the unseen future *improves* the geometric estimation of the seen present. Inferring holistic 3D geometry from a highly sparse context (*e.g.*, $k = 2$) is inherently ambiguous. Our framework naturally resolves this by leveraging the forecasted future as a temporal regularizer. After the rollout produces the future chunks $\hat{\mathbf{Z}}_{1:S}$, we aggregate them with the initial sparse context $\mathbf{z}_{1:k}$ to form a latent trajectory of the full scene:

$$\mathbf{z}_{1:T}^{\text{full}} = [\mathbf{z}_{1:k}; \hat{\mathbf{Z}}_{1:S}]. \quad (3)$$

This assembled sequence is then propagated jointly through the remaining frozen blocks Φ_{dec} , and the original VGGT 3D heads decode the final outputs (*e.g.*, point clouds $\mathcal{P}_{1:T}$, depth maps $\mathcal{D}_{1:T}$):

$$\mathcal{P}_{1:T}, \mathcal{D}_{1:T} = \mathcal{H}_{3D}(\Phi_{\text{dec}}(\mathbf{z}_{1:T}^{\text{full}})). \quad (4)$$

As a result, the geometric precision at the observed step k is significantly enhanced precisely because the model is forced to contextualize it within a consistent, long-horizon world model rollout.

3.2 Chunk Transition in High-Dimensional Geometry Space

To model the conditional transition of the next chunk $p(\mathbf{Z}_i \mid \mathbf{c}_i)$ in Eq. (2), we adopt continuous-time flow matching [32] in the geometry latent space. Standard

flow matching [32] defines a probability path $\mathbf{Z}_{i,\tau} = (1 - \tau)\mathbf{Z}_i + \tau\epsilon$ governed by the ODE:

$$\frac{d\mathbf{Z}_{i,\tau}}{d\tau} = v_\theta(\mathbf{Z}_{i,\tau}, \tau, \mathbf{c}_i), \quad (5)$$

where v_θ is typically trained to predict the vector field velocity (**v**-prediction).

The Curse of Dimensionality. However, we empirically found that standard **v**-prediction *fails* in our setting (Fig. 2): even after 4k training iterations, the denoised predictions remain noisy and weakly structured. We attribute this failure to the *high* dimensionality of our target space: unlike video VAEs operating in heavily compressed latents (*e.g.*, 16–32 channels), we model dynamics directly in the 1024-dimensional VGGT feature space where valid geometry states occupy only a structured subset of the ambient space. Isotropic Gaussian corruption therefore, mainly perturbs states along off-manifold directions, rendering velocity estimation ill-conditioned. In this regime, small velocity errors quickly drive trajectories away from the true geometry manifold and are further amplified by Φ_{dec} and $\mathcal{H}_{3\text{D}}$. While a common remedy is to learn an additional dimensionality reduction module [23], this would break alignment with the frozen VGGT heads and destroy their zero-shot geometric priors.

z-prediction Flow. We therefore directly predict the clean target latent $\mathbf{Z}_i$ instead of its velocity. In spirit, this is related to recent findings [3, 29] that direct clean-target prediction can provide stronger denoising signals than velocity prediction; however, unlike pixel-space *x*-prediction, our model predicts the clean geometry latent state. Fig. 2 confirms this choice empirically: **z**-prediction yields higher training SNR and progressively recovers structured target latents.

Implementation. Consequently, we explicitly parameterize our network F_θ to predict the clean latents of the target chunk directly (**z**-prediction). To enforce the separation of temporal reasoning and spatial refinement, we decouple F_θ into a composition of two distinct sub-networks: a dual-stream causal processor F_{dual} and a single-stream spatial denoiser F_{single} :

$$\mathbf{Z}_i \approx F_\theta(\mathbf{Z}_{i,\tau}, \tau, \mathbf{c}_i) = F_{\text{single}}\left(F_{\text{dual}}(\mathbf{Z}_{i,\tau}, \tau, \mathbf{c}_i), \tau\right). \quad (6)$$

This architectural design aligns with very recent analyses [4] showing that causality in video diffusers is separable from spatial denoising: temporal attention naturally dominates *early layers* to establish historical context, while *deeper layers* become highly diagonal, focusing almost entirely on intra-frame refinement. Specifically, continuous flow time τ is injected into all blocks via Adaptive Layer Normalization (adaLN). For early dual-stream blocks $l \in [1, L_d]$, the noisy m -frame target chunk $\mathbf{Z}_{i,\tau}$ asymmetrically cross-attends to the condition $\mathbf{c}_i$, absorbing temporal physics without corrupting the clean history representations:

$$\mathbf{Z}_{i,\tau}^{(l+1)} = \mathbf{Z}_{i,\tau}^{(l)} + \text{Attn}\left(\mathbf{Q} = \text{adaLN}(\mathbf{Z}_{i,\tau}^{(l)}, \tau), \mathbf{K}, \mathbf{V} = [\text{adaLN}(\mathbf{Z}_{i,\tau}^{(l)}, \tau); \mathbf{c}_i]\right). \quad (7)$$

The deeper layers then transition to single-stream blocks ($l \in [L_d + 1, L_d + L_s]$), operating exclusively on the target chunk via joint intra-chunk self-attention to

focus entirely on local spatiotemporal geometry refinement:

$$\mathbf{Z}_{i,\tau}^{(l+1)} = \mathbf{Z}_{i,\tau}^{(l)} + \text{Self-Attn}\left(\mathbf{Q}, \mathbf{K}, \mathbf{V} = \text{adaLN}(\mathbf{Z}_{i,\tau}^{(l)}, \tau)\right). \quad (8)$$

3.3 Latent Flow Forcing Curriculum

Autoregressive models trained exclusively on ground-truth histories suffer from severe exposure bias during inference. In our high-dimensional geometry space, predicting a sequence of chunks leads to compounding errors: a slight deviation in the predicted chunk $\hat{\mathbf{Z}}_i$ corrupts the subsequent condition $\mathbf{c}_{i+1}$, rapidly pushing the rollout trajectory off the true geometry manifold. To guarantee stable long-horizon rollouts, we propose a two-stage latent flow forcing curriculum.

Stage 1: Teaching-Forcing Training. In the first stage, we establish the base continuous-time vector field using oracle historical conditions. Specifically, the context is set as $\mathbf{c}_i = \mathbf{Z}_{i-1}$, where $\mathbf{Z}_{i-1}$ denotes the ground-truth VGGT latent extracted from previous frames. As established in Section 3.2, we optimize F_θ via a weighted z -prediction scheme to mitigate the extreme variance of the high-dimensional latent space. The Stage 1 objective is purely teacher-forced:

$$\mathcal{L}_{S1} = \mathbb{E}_{\mathbf{Z}_i, \epsilon, \tau} \left[\left\| \frac{F_\theta(\mathbf{Z}_{i,\tau}, \tau, \mathbf{c}_i) - \mathbf{Z}_{i,\tau}}{1 - \tau} - \frac{\mathbf{Z}_i - \mathbf{Z}_{i,\tau}}{1 - \tau} \right\|_2^2 \right], \quad (9)$$

However, at test time the model must condition on its own predicted history $\mathbf{c}_i = \hat{\mathbf{Z}}_{i-1}$ rather than oracle latents, creating a train-test mismatch. As a result, errors in the context compound across rollout, making the learned operator increasingly brittle over long horizons.

Stage 2: Trajectory-Consistent Flow Forcing. To mitigate this gap, in Stage 2, we expose the model to structured errors from its own rollout. Despite prior attempts [9, 21] to mitigate autoregressive drift, existing strategies are not ideal for our setting: analytic Gaussian corruption tends to move weakly compressed 1024-D VGGT latents off the geometry manifold, while fully self-generated histories are often slow in training. We therefore use partially denoised states from the model’s own ODE rollout as structured perturbations, yielding a smoother transition from oracle conditioning to self-conditioning.

Partially-denoised Rollout State. To predict the target chunk $\mathbf{Z}_{i+1}$, we first roll out the preceding chunk $\mathbf{Z}_i$ from noise $\mathbf{Z}_{i,1} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ under its clean history $\mathbf{c}_i = \mathbf{Z}_{i-1}$, but stop the ODE integration using our Euler solver at an intermediate flow time $\tau_{\text{mid}} \in (0, 1)$:

$$\hat{\mathbf{Z}}_{i,\tau_{\text{mid}}} = \text{ODESolve}\left(F_\theta, \mathbf{Z}_{i,1}, \mathbf{c}_i, \tau \in [1 \rightarrow \tau_{\text{mid}}]\right). \quad (10)$$

This intermediate state $\hat{\mathbf{Z}}_{i,\tau_{\text{mid}}}$ is a partially denoised rollout. To formulate the condition for the next chunk $i+1$, we do not fully replace the ground truth immediately. We construct a mixed condition by interpolating this partially integrated trajectory with the clean ground-truth chunk $\mathbf{Z}_i$ using $\lambda \in [0, 1]$:

$$\mathbf{c}_{i+1}^{\text{mix}}(\lambda, \tau_{\text{mid}}) = (1 - \lambda)\mathbf{Z}_i + \lambda\hat{\mathbf{Z}}_{i,\tau_{\text{mid}}}. \quad (11)$$

Finally, we optimize F_θ to predict the clean target chunk $\mathbf{Z}_{i+1}$ conditioned on this mixed, partially denoised history. The target chunk itself is perturbed to a flow time τ , yielding the Stage 2 finetuning objective:

$$\mathcal{L}_{S2} = \mathbb{E}_{i, \lambda, \tau, \tau_{\text{mid}}} \left[\left\| \mathbf{Z}_{i+1} - F_\theta(\mathbf{Z}_{i+1, \tau}, \tau, \mathbf{c}_{i+1}^{\text{mix}}(\lambda, \tau_{\text{mid}})) \right\|_2^2 \right]. \quad (12)$$

Curriculum Stability. In Appendix, we further show theoretical analysis of the proposed flow forcing strategy establishes a sequential ELBO on the uncorrupted data likelihood (detailed in Theorem 1). However, the stability of this bound relies heavily on the variance of the rollout error, $\mathbf{E}_i = \hat{\mathbf{Z}}_{i, \tau_{\text{mid}}} - \mathbf{Z}_i$. Because our mixed condition expands to $\mathbf{c}_{i+1}^{\text{mix}} = \mathbf{Z}_i + \lambda \mathbf{E}_i$, the variance injected into the conditioning signal scales proportionally with $\lambda^2 \mathbb{E}[\|\mathbf{E}_i\|^2]$. Early in training, model can produce unstructured rollouts, meaning $\mathbb{E}[\|\mathbf{E}_i\|^2]$ is prohibitively large. Setting λ too high would inject massive variance, completely violating the step-wise KL penalty and causing the ELBO to collapse. To enforce convergence and maintain a tight lower bound, we introduce a *curriculum learning* schedule for the mixing parameter λ . We linearly anneal λ from 0 to 1 as finetuning progresses. As the intrinsic rollout error decays, smoothly increasing λ toward 1 progressively exposes the network to its inference-time distribution.

4 Experiments

Dataset. We assess our approach using the Cityscapes [10], KITTI [16] and TartanAir [52] datasets, offering video sequences from urban driving environments and diverse 6-DoF camera trajectories in photorealistic simulated scenes. Cityscapes includes 2,975 training sequences and 500 validation sequences, each containing 30 frames captured at 16 FPS with a resolution of 1024×2048 pixels. KITTI contains 143 training sequences and 13 validation sequences, with a native resolution of 512×1382 and a frame rate of 10 FPS. TartanAir is a large-scale synthetic dataset rendered in Unreal Engine, covering diverse indoor and outdoor environments, with approximately 700k frames in total. Following the split protocol of Gen3R [19], we randomly use 90% of the data for training and the remaining 10% for validation.

Implementation Details. For the two-stage training, we adopt AdamW with a learning rate of 2×10^{-4} and a weight decay of 0.05, and apply gradient clipping with ℓ_2 norm of 1.0. Training is conducted on 1 RTX Pro 6000 (96 GB) GPU. The batch size is set to 64 for Teaching-Forcing Training and 48 for Trajectory-Consistent Flow Forcing. We set the initial chunk size to 4 and perform autoregressive rollout with a stride of 1. The flow transformer is alike to [26], comprising $L_d = 8$ double-stream Transformer blocks and $L_s = 8$ single-stream Transformer blocks. We retain the original intermediate feature dimensionality of 1024 without feature compression. Additional details are provided in Appendix.

Evaluation Metrics. For depth prediction, we report the mean Absolute Relative Error (AbsRel) and depth accuracy (δ_1). Following prior future prediction works [6, 23], we evaluate on Cityscapes under short-term (3 frames, 187.5ms)

Table 1: Depth forecasting on KITTI. We evaluate representative world-model baselines on future scene depth prediction, reporting AbsRel ($\downarrow$) and δ_1 ($\uparrow$) for short- and mid-term horizons.

Method	2011_09_26_drive_0002_sync				2011_10_03_drive_0047_sync				Mean (Dataset Avg.)			
	AbsRel ($\downarrow$)		δ_1 ($\uparrow$)		AbsRel ($\downarrow$)		δ_1 ($\uparrow$)		AbsRel ($\downarrow$)		δ_1 ($\uparrow$)	
	Short	Mid	Short	Mid	Short	Mid	Short	Mid	Short	Mid	Short	Mid
Cosmos-4B [1]	0.138	0.148	84.0	81.4	0.155	0.156	77.5	75.2	0.155	0.165	77.1	74.1
Cosmos-12B [1]	0.131	0.139	84.5	79.7	0.149	0.176	77.6	75.9	0.154	0.185	77.5	71.5
Copy Last (VGGT [49])	0.112	0.125	89.6	82.6	0.140	0.153	84.4	79.6	0.137	0.144	83.7	78.5
Aether [43]	0.074	0.102	94.2	88.0	0.113	0.111	87.4	85.2	0.096	0.124	87.9	84.7
DINO-Foresight [23]	0.063	0.105	96.5	89.9	0.076	0.107	90.9	86.7	0.082	0.115	91.6	85.1
VGGT-World	0.050	0.078	98.3	93.2	0.064	0.101	94.9	92.7	0.065	0.098	94.0	89.4




Fig. 3: The short- and mid-term depth forecasting results on KITTI. “Upper Bound” denotes the depth obtained by feeding the real future images into VGGT. “DINO-Foresight” predicts depth in the scale of Depth Anything [58], leading to a different visualization appearance compared to others.

and mid-term (9 frames, 562.5ms) forecasting settings, and on KITTI under short-term (2 frames, 200 ms) and mid-term (6 frames, 600 ms) settings. To assess 3D point cloud quality, following the evaluation protocol of [19], we first align the predicted point clouds to the ground truth using Umeyama similarity alignment [46], then sample 20k points from both sets via Farthest Point Sampling (FPS) [38], and compute standard geometric metrics including Accuracy, Completeness, and Chamfer Distance (CD) [12].

Baselines. We compare our method with three categories of state-of-the-art world models: (1) *pixel-space 2D generation models*, including Cosmos 4B/12B [1] and VISTA [15] (we use its fine-tuned results following [23]); (2) *latent-space 2D generation models*, represented by DINO-Foresight [23]; and (3) *geometry-aware generation models*, such as Aether [43], WVD [60], and Gen3R [19]. Note that we also include a COPY-LAST baseline, where the depth predicted from the ground-truth RGB of the previous frame.

4.1 Depth Forecasting

Results on KITTI. Table 1 and Fig. 3 compare VGGT-World with representative world models on depth forecasting. *Quantitatively*, our method consistently achieves the best performance across all horizons, attaining a mean δ_1 of 94.0%/89.4% and reducing AbsRel to 0.065/0.098, corresponding to relative improvements of 21%/15% on the short- and mid-term horizons respectively. Notably, the pixel-space Cosmos (4B/12B) [1] underperforms even the naive COPY-LAST baseline [49], indicating that scaling video generation without 3D priors prioritizes photometric realism over geometric consistency. While Aether [43] incorporates geometry supervision, it still underperforms due to pixel-space representation. Additionally, although the latent-space DINO-Foresight [23] achieves competitive accuracy, its performance is largely driven by fine-tuning on ground-truth depth. In contrast, our method, trained in the geometry latent space, achieves superior performance without any fine-tuning on ground-truth depth. *Qualitatively*, our model produces more faithful depth maps shown in Fig. 3, mitigating distant-region errors and geometric distortions observed in Cosmos [1] and Aether [43], while preserving sharp structural details and temporal-geometric consistency, demonstrating stronger geometry modeling capability.

Results on Cityscapes. On Cityscapes, the main challenge is not only short-term frame extrapolation, but maintaining metrically stable depth under continuous ego-motion and large perspective changes. This is reflected by the larger gap between short- and mid-term results across all methods. As shown in Table 2, our method achieves the best performance on both metrics across short- and mid-term horizons, improving AbsRel and δ_1 by up to 32%. Mid-term forecasting is inherently more challenging, as small structural errors in earlier predictions can accumulate and lead to larger depth deviations. More qualitative results are provided in the Appendix for reference.

Table 2: Depth forecasting on Cityscapes with AbsRel and depth accuracy δ_1 .

Method	AbsRel ($\downarrow$)		δ_1 ($\uparrow$)	
	Short	Mid	Short	Mid
Cosmos-4B [1]	0.189	0.247	76.3	70.5
Cosmos-12B [1]	0.181	0.241	78.2	72.4
COPY-LAST (VGGT [49])	0.154	0.212	84.1	77.8
VISTA _{ft} [15]	0.124	0.153	86.4	82.8
DINO-Foresight [23]	0.114	0.136	88.6	85.4
VGGT-World w/o Forcing	0.079	0.131	93.7	86.8
VGGT-World	0.078	0.126	93.8	87.3
VGGT-World ($\lambda = 0.1$)	0.084	0.142	93.1	85.0
VGGT-World ($\lambda = 0.3$)	0.089	0.148	92.6	84.9
VGGT-World ($\lambda = 0.7$)	0.109	0.165	90.1	81.4

4.2 Point Map Forecasting

Table 3 provides point cloud forecasting results on TartanAir, demonstrating that VGGT-WORLD achieves superior geometry accuracy by predicting scene directly within geometry-centric latent spaces. Compared to pixel-space generative model GEN3R which models dynamics in pixel space, our method significantly outperforms it on all metrics in the pose/text-free setting, while remaining competitive even when it is aided by explicit pose conditions. Furthermore, the

Table 3: Geometry generation on TartanAir with different conditioning (text, pose) under 1-view and 2-view settings.

Method	Condition		1-view			2-views		
	Text-free	Pose-free	Acc.↓	Comp.↓	CD↓	Acc.↓	Comp.↓	CD↓
AETHER [43]	✓	✗	3.1547	4.5366	3.8457	2.7745	3.4420	3.1082
WVD [60]	✓	✗	4.3944	3.0660	3.7302	4.3794	2.5268	3.4531
GEN3R [19]	✗	✗	3.0250	2.5367	2.7809	2.2825	1.6643	1.9734
GEN3R [19]	✓	✗	3.4579	4.7384	4.0982	2.6029	3.1154	2.8591
GEN3R [19]	✓	✓	3.6427	7.8528	5.7478	2.7441	6.1423	4.4432
VGGT-World w/o Forcing	✓	✓	1.2200	5.0435	3.0818	1.2200	5.0435	3.0818
VGGT-World	✓	✓	1.1604	4.5297	2.8451	1.1604	4.5297	2.8451

**Fig. 4: Point cloud forecasting on TartanAir.** Gen3R produces structurally disorganized geometry on walls and rooftops, whereas our method preserves coherent structure and yields more accurate predictions.

performance gain of the full model over the w/o Forcing variant validates that our Trajectory-Consistent Flow Forcing effectively mitigates exposure bias during rollout. These results confirm that by repurposing spatiotemporal features rather than reconstructing pixels, VGGT-WORLD ensures superior geometric fidelity at a fraction of the computational cost of existing generative baselines. Qualitative results in Fig. 4 further illustrate the superior geometric fidelity of VGGT-WORLD. Without pose/text conditioning, GEN3R exhibits significant structural disorganization and blurry distortions, failing to maintain scene integrity. In contrast, our method produces well-defined geometry that closely adheres to the ground-truth topology across both indoor and outdoor scenes.

4.3 Further Analysis

Impact of Forcing Curriculum. Bottom part of Table 2 evaluates the impact of the mixing ratio λ , where our VGGT-WORLD adopts a linear scheduling strategy that gradually increases λ as training progresses. Although the model without forcing achieves strong short-term performance, it degrades during rollout. The proposed dynamic scheduling alleviates this issue by progressively en-




Fig. 5: Long-term horizon forecasting on TartanAir.

Table 4: Short-term horizon foresight on TartanAir.

Metric	VGGT [49]	VGGT-World
Acc. ↓	0.8672	0.8764
Comp. ↓	1.2159	1.0646
CD ↓	1.0416	0.9705

hancing robustness to recursive conditioning while preserving the learning of scene dynamics. In contrast, static settings cause the model to learn a fixed and biased perturbation distribution; moreover, a static large mixing rate (*e.g.*, $\lambda = 0.7$) can induce catastrophic forgetting of the Stage-1 distribution. Overall, the curriculum-based design balances stable structural learning with long-horizon forecasting stability.

Can future prediction improve present-time perception? We further ask whether VGGT-WORLD can provide a better geometric estimation for the *current* moment than directly applying VGGT to the current input alone. The short-horizon comparison in Table 4 shows that this is largely the case. Although the accuracy metric changes only marginally, VGGT-WORLD achieves clearly better completeness and a lower overall CD value, indicating a more balanced and globally consistent reconstruction. This result suggests that predictive world modeling does not merely support future imagination but can also feed back to strengthen present-time perception. For embodied agents, such a capability is particularly appealing: instead of treating perception as a purely instantaneous process, the agent can leverage short-term foresight to refine its current geometric belief, leading to more reliable downstream planning and control.

Camera Parameter Forecasting. Beyond depth and point cloud forecasting, we further evaluate whether the predicted next-frame VGGT features preserve camera motion cues. We predict camera parameters from the generated VGGT features and compare them against ground truth trajectories on KITTI and TartanAir. After Umeyama alignment, the predicted trajectories exhibit relatively consistent global structure with the ground truth in Fig. 6. Quantitative analysis of ATE/RTE/RRE errors is reported in the Appendix. These results suggest that VGGT-World’s latent dynamics preserve both scene structure and camera evolution without explicit pose supervision, achieving reasonable consistency without the prohibitive computational costs of appearance-driven world models.

How far into the future should we forecast? We first ask whether forecasting a longer future horizon can further benefit geometry estimation when




Fig. 6: Qualitative study of camera parameter prediction on KITTI and TartanAir.

the predicted observations are aggregated with the current frame and passed to VGGT. As shown in Fig. 5, extending the foresight horizon initially improves the reconstructed geometry, suggesting that moderate future context provides complementary structural cues beyond the instantaneous observation. However, this benefit is not monotonic. The CD reaches its best value around frame 7, indicating a practical sweet spot for horizon length. Beyond this point, performance begins to degrade, and the overall geometry exhibits distortion.

Efficiency Analysis. We evaluate the computational efficiency of VGGT-WORLD against representative video-generation paradigms, including COSMOS-12B [1] and GEN3R [19]. Since they operate under varying input and rollout horizons, we adopt a normalized latency metric (seconds per generated frame) to ensure a fair, scale-invariant

comparison. This efficiency comes not only from forecasting in the compact VGGT latent space, but also from our frozen-backbone design. Instead of retraining a large video-generation or geometry backbone from scratch, VGGT-WORLD reuses the pretrained VGGT backbone as a fixed geometric prior and only optimizes the lightweight

temporal forecasting modules. As illustrated in Fig. 7(b), this greatly reduces the trainable parameter footprint and avoids the costly full-model retraining required by large-scale video generation world models. As shown in Fig. 7a, COSMOS-12B requires 9.50 s/frame to synthesize future frames on KITTI. In contrast, our approach achieves an end-to-end inference time of 1.9 s/frame. This represents a $5\times$ reduction in per-frame latency compared to COSMOS and a $3.6\times$ speedup over GEN3R (6.90s/frame).

Trainable Footprint. Beyond inference speed, the practicality of world models in resource-constrained settings depends heavily on the number of trainable parameters. In Fig. 7b, our approach optimizes only 0.43B parameters by keeping the 1.3B VGGT backbone frozen. In contrast, COSMOS-12B requires updating all 12.0B parameters, and GEN3R entails a heavy trainable DiT (2.7B) alongside a geometry adapter. By modeling dynamics directly in the geometry-centric latent space, VGGT-WORLD achieves a much smaller trainable footprint while keeping competitive capacity, enabling easier reproduction and faster iteration.




Fig. 7: Efficiency comparison with VideoGen world models. (a) Inference latency per generated frame. (b) Trainable vs. frozen parameter footprint.

5 Conclusion and Future Work

We presented VGGT-World, an early attempt at world modeling directly in a high-dimensional geometric latent space. Despite its simplicity, our approach already surpasses substantially heavier video-generation world models on both

forecasting quality and efficiency, suggesting that geometry-centric latent forecasting is a promising alternative to full-frame generation. Moreover, because VGGT-World requires no ground-truth depth or point-cloud supervision, it provides a flexible foundation for scaling geometry world models. Looking ahead, we can enable richer conditioning by incorporating camera pose and action information into the generative process; for embodied agents, action-conditioned forecasting may further support more controllable and interactive rollouts.

Acknowledgements

This work was partially supported by ARC DE240100105, DP240101814, DP230101196, BA24006, and ARC Industrial Transformation Research Hubs IH230100013.

References

1. Agarwal, N., Ali, A., Bala, M., Balaji, Y., Barker, E., Cai, T., Chattopadhyay, P., Chen, Y., Cui, Y., Ding, Y., Dworakowski, D., Fan, J., Fenzi, M., Ferroni, F., Fidler, S., Fox, D., Ge, S., Ge, Y., Gu, J., Gururani, S., He, E., Huang, J., Huffman, J.S., Jannaty, P., Jin, J., Kim, S.W., Klár, G., Lam, G., Lan, S., Leal-Taixé, L., Li, A., Li, Z., Lin, C., Lin, T., Ling, H., Liu, M., Liu, X., Luo, A., Ma, Q., Mao, H., Mo, K., Mousavian, A., Nah, S., Niverty, S., Page, D., Paschalidou, D., Patel, Z., Pavao, L., Ramezanali, M., Reda, F., Ren, X., Sabavat, V.R.N., Schmerling, E., Shi, S., Stefaniak, B., Tang, S., Tchapmi, L., Tredak, P., Tseng, W., Varghese, J., Wang, H., Wang, H., Wang, H., Wang, T., Wei, F., Wei, X., Wu, J.Z., Xu, J., Yang, W., Lin, Y., Zeng, X., Zeng, Y., Zhang, J., Zhang, Q., Zhang, Y., Zhao, Q., Zólkowski, A.: Cosmos world foundation model platform for physical AI. CoRR **abs/2501.03575** (2025) 10, 11, 14
2. Agarwal, S., Furukawa, Y., Snavely, N., Simon, I., Cxurless, B., Seitz, S.M., Szeliski, R.: Building rome in a day. Commun. ACM **54**(10), 105–112 (2011) 4
3. Baade, A., Chan, E.R., Sargent, K., Chen, C., Johnson, J., Adeli, E., Fei-Fei, L.: Latent forcing: Reordering the diffusion trajectory for pixel-space image generation. arXiv preprint arXiv:2602.11401 (2026) 7
4. Bai, X., He, G., Li, Z., Shechtman, E., Huang, X., Wu, Z.: Causality in video diffusers is separable from denoising. arXiv preprint arXiv:2602.10095 (2026) 7
5. Bai, Y., Fang, S., Yu, C., Wang, F., Huang, Q.: Geovideo: Introducing geometric regularization into video generation model. arXiv preprint arXiv:2512.03453 (2025) 2
6. Baldassarre, F., Szafraniec, M., Terver, B., Khalidov, V., Massa, F., LeCun, Y., Labatut, P., Seitzer, M., Bojanowski, P.: Back to the features: DINO as a foundation for video world models. CoRR **abs/2507.19468** (2025) 4, 9
7. Blattmann, A., Dockhorn, T., Kulal, S., Mendelevitch, D., Kilian, M., Lorenz, D., Levi, Y., English, Z., Voleti, V., Letts, A., et al.: Stable video diffusion: Scaling latent video diffusion models to large datasets. arXiv preprint arXiv:2311.15127 (2023) 4
8. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: 2021 IEEE/CVF

- International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021. pp. 9630–9640. IEEE (2021) 4
9. Chen, B., Martí Monsó, D., Du, Y., Simchowitz, M., Tedrake, R., Sitzmann, V.: Diffusion forcing: Next-token prediction meets full-sequence diffusion. *Advances in Neural Information Processing Systems* **37**, 24081–24125 (2024) 8
 10. Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016. pp. 3213–3223. IEEE Computer Society (2016) 3, 9
 11. Deng, H., Pan, T., Diao, H., Luo, Z., Cui, Y., Lu, H., Shan, S., Qi, Y., Wang, X.: Autoregressive video generation without vector quantization. *arXiv preprint arXiv:2412.14169* (2024) 4
 12. Fan, H., Su, H., Guibas, L.J.: A point set generation network for 3d object reconstruction from a single image. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017. pp. 2463–2471. IEEE Computer Society (2017) 10
 13. Feiyang, J., Liu, L., Song, Z., Jia, C., Ye, H., Hao, X., Chen, L., et al.: Driveworld-vla: Unified latent-space world modeling with vision-language-action for autonomous driving. *arXiv preprint arXiv:2602.06521* (2026) 4
 14. Frahm, J., Georgel, P.F., Gallup, D., Johnson, T., Raguram, R., Wu, C., Jen, Y., Dunn, E., Clipp, B., Lazebnik, S.: Building rome on a cloudless day. In: Daniilidis, K., Maragos, P., Paragios, N. (eds.) *Computer Vision - ECCV 2010*, 11th European Conference on Computer Vision, Heraklion, Crete, Greece, September 5-11, 2010, Proceedings, Part IV. Lecture Notes in Computer Science, vol. 6314, pp. 368–381. Springer (2010) 4
 15. Gao, S., Yang, J., Chen, L., Chitta, K., Qiu, Y., Geiger, A., Zhang, J., Li, H.: Vista: A generalizable driving world model with high fidelity and versatile controllability. In: Globersons, A., Mackey, L., Belgrave, D., Fan, A., Paquet, U., Tomczak, J.M., Zhang, C. (eds.) *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024*, Vancouver, BC, Canada, December 10 - 15, 2024 (2024) 10, 11
 16. Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Vision meets robotics: The KITTI dataset. *Int. J. Robotics Res.* **32**(11), 1231–1237 (2013) 3, 9
 17. Ho, J., Salimans, T., Gritsenko, A., Chan, W., Norouzi, M., Fleet, D.J.: Video diffusion models. *Advances in neural information processing systems* **35**, 8633–8646 (2022) 4
 18. Hong, Y., Zhang, K., Gu, J., Bi, S., Zhou, Y., Liu, D., Liu, F., Sunkavalli, K., Bui, T., Tan, H.: LRM: large reconstruction model for single image to 3d. In: *The Twelfth International Conference on Learning Representations, ICLR 2024*, Vienna, Austria, May 7-11, 2024. OpenReview.net (2024) 1
 19. Huang, J., Yang, Y., Yang, B., Ma, L., Ma, Y., Liao, Y.: Gen3r: 3d scene generation meets feed-forward reconstruction. *CoRR* **abs/2601.04090** (2026) 2, 4, 9, 10, 12, 14
 20. Huang, N., Wei, X., Zheng, W., An, P., Lu, M., Zhan, W., Tomizuka, M., Keutzer, K., Zhang, S.: S^3 gaussian: Self-supervised street gaussians for autonomous driving. *CoRR* **abs/2405.20323** (2024) 1
 21. Huang, X., Li, Z., He, G., Zhou, M., Shechtman, E.: Self forcing: Bridging the train-test gap in autoregressive video diffusion. *arXiv preprint arXiv:2506.08009* (2025) 8

22. Jia, X., Liu, Y., You, J., Xia, R., Hong, Y., Yan, J.: Drivevggt: Visual geometry transformer for autonomous driving. CoRR **abs/2511.22264** (2025) 1
23. Karypidis, E., Kakogeorgiou, I., Gidaris, S., Komodakis, N.: DINO-foresight: Looking into the future with DINO. In: The Thirty-ninth Annual Conference on Neural Information Processing Systems (2025) 4, 7, 9, 10, 11
24. Keetha, N.V., Müller, N., Schönberger, J., Porzi, L., Zhang, Y., Fischer, T., Knapitsch, A., Zauss, D., Weber, E., Antunes, N., Luiten, J., Lopez-Antequera, M., Bulò, S.R., Richardt, C., Ramanan, D., Scherer, S.A., Kotschieder, P.: Mapanything: Universal feed-forward metric 3d reconstruction. CoRR **abs/2509.13414** (2025) 1, 4
25. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. In: Bengio, Y., LeCun, Y. (eds.) 2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings (2014) 2
26. Labs, B.F., Batifol, S., Blattmann, A., Boesel, F., Consul, S., Diagne, C., Dockhorn, T., English, J., English, Z., Esser, P., Kulal, S., Lacey, K., Levi, Y., Li, C., Lorenz, D., Müller, J., Podell, D., Rombach, R., Saini, H., Sauer, A., Smith, L.: FLUX.1 kontext: Flow matching for in-context image generation and editing in latent space. CoRR **abs/2506.15742** (2025) 9
27. Leroy, V., Cabon, Y., Revaud, J.: Grounding image matching in 3d with mast3r. In: Leonardi, A., Ricci, E., Roth, S., Russakovsky, O., Sattler, T., Varol, G. (eds.) Computer Vision - ECCV 2024 - 18th European Conference, Milan, Italy, September 29-October 4, 2024, Proceedings, Part LXXII. Lecture Notes in Computer Science, vol. 15130, pp. 71–91. Springer (2024) 1
28. Li, F., Song, W., Zhao, H., Wang, J., Ding, P., Wang, D., Zeng, L., Li, H.: Spatial forcing: Implicit spatial representation alignment for vision-language-action model. arXiv preprint arXiv:2510.12276 (2025) 1
29. Li, T., He, K.: Back to basics: Let denoising generative models denoise. CoRR **abs/2511.13720** (2025) 7
30. Li, Y., Fan, L., He, J., Wang, Y., Chen, Y., Zhang, Z., Tan, T.: Enhancing end-to-end autonomous driving with latent world model. arXiv preprint arXiv:2406.08481 (2024) 4
31. Li, Y., Wang, Y., Liu, Y., He, J., Fan, L., Zhang, Z.: End-to-end driving with online trajectory evaluation via bev world model. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 27137–27146 (2025) 4
32. Liu, X., Gong, C., Liu, Q.: Flow straight and fast: Learning to generate and transfer data with rectified flow. In: The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023. OpenReview.net (2023) 6, 7
33. Liu, Y., Zhang, K., Li, Y., Yan, Z., Gao, C., Chen, R., Yuan, Z., Huang, Y., Sun, H., Gao, J., et al.: Sora: A review on background, technology, limitations, and opportunities of large vision models. arXiv preprint arXiv:2402.17177 (2024) 4
34. Lourakis, M.I.A., Argyros, A.A.: SBA: A software package for generic sparse bundle adjustment. ACM Trans. Math. Softw. **36**(1), 2:1–2:30 (2009) 4
35. Ma, B., Yang, J., Di, D., Zhang, X., Cui, J., Li, H., Xie, Y., Chen, W.: Metricanything: Scaling metric depth pretraining with noisy heterogeneous sources. arXiv preprint arXiv:2601.22054 (2026) 1, 4
36. NVIDIA: World simulation with video foundation models for physical AI. CoRR **abs/2511.00062** (2025) 2, 4

37. Polyak, A., Zohar, A., Brown, A., Tjandra, A., Sinha, A., Lee, A., Vyas, A., Shi, B., Ma, C.Y., Chuang, C.Y., et al.: Movie gen: A cast of media foundation models. arXiv preprint arXiv:2410.13720 (2024) 2
38. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In: Guyon, I., von Luxburg, U., Bengio, S., Wallach, H.M., Fergus, R., Vishwanathan, S.V.N., Garnett, R. (eds.) *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*. pp. 5099–5108 (2017) 10
39. Schönberger, J.L., Frahm, J.: Structure-from-motion revisited. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*. pp. 4104–4113. IEEE Computer Society (2016) 4
40. Shi, M., Wang, H., Zheng, W., Yuan, Z., Wu, X., Wang, X., Wan, P., Zhou, J., Lu, J.: Latent diffusion model without variational autoencoder. *ICLR abs/2510.15301* (2026) 4
41. Siméoni, O., Vo, H.V., Seitzer, M., Baldassarre, F., Oquab, M., Jose, C., Khalidov, V., Szafraniec, M., Yi, S.E., Ramamonjisoa, M., Massa, F., Haziza, D., Wehrstedt, L., Wang, J., Darcet, T., Moutakanni, T., Sentana, L., Roberts, C., Vedaldi, A., Tolan, J., Brandt, J., Couprie, C., Mairal, J., Jégou, H., Labatut, P., Bojanowski, P.: Dinov3. *CoRR abs/2508.10104* (2025) 4
42. Snavely, N., Seitz, S.M., Szeliski, R.: Photo tourism: exploring photo collections in 3d. *ACM Trans. Graph.* **25**(3), 835–846 (2006) 4
43. Team, A., Zhu, H., Wang, Y., Zhou, J., Chang, W., Zhou, Y., Li, Z., Chen, J., Shen, C., Pang, J., He, T.: Aether: Geometric-aware unified world modeling. *CoRR abs/2503.18945* (2025) 10, 11, 12
44. Tian, K., Jiang, Y., Yuan, Z., Peng, B., Wang, L.: Visual autoregressive modeling: Scalable image generation via next-scale prediction. *Advances in neural information processing systems* **37**, 84839–84865 (2024) 4
45. Triggs, B., McLauchlan, P.F., Hartley, R.I., Fitzgibbon, A.W.: Bundle adjustment - A modern synthesis. In: Triggs, B., Zisserman, A., Szeliski, R. (eds.) *Vision Algorithms: Theory and Practice, International Workshop on Vision Algorithms, held during ICCV '99, Corfu, Greece, September 21-22, 1999, Proceedings. Lecture Notes in Computer Science*, vol. 1883, pp. 298–372. Springer (1999) 4
46. Umeyama, S.: Least-squares estimation of transformation parameters between two point patterns. *IEEE Trans. Pattern Anal. Mach. Intell.* **13**(4), 376–380 (1991) 10
47. Wan, T., Wang, A., Ai, B., Wen, B., Mao, C., Xie, C.W., Chen, D., Yu, F., Zhao, H., Yang, J., et al.: Wan: Open and advanced large-scale video generative models. arXiv preprint arXiv:2503.20314 (2025) 2, 4
48. Wan, Y., Liu, L., Zhou, J., Zhou, Z., Zhang, X., Zhang, D., Jiao, S., Hou, Q., Cheng, M.: Geoworld: Unlocking the potential of geometry models to facilitate high-fidelity 3d scene generation. *CoRR abs/2511.23191* (2025) 2, 4
49. Wang, J., Chen, M., Karaev, N., Vedaldi, A., Rupprecht, C., Novotny, D.: Vggt: Visual geometry grounded transformer. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 5294–5306 (2025) 1, 4, 10, 11, 13
50. Wang, S., Leroy, V., Cabon, Y., Chidlovskii, B., Revaud, J.: Dust3r: Geometric 3d vision made easy. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*. pp. 20697–20709. IEEE (2024) 1, 4
51. Wang, T., Mao, X., Zhu, C., Xu, R., Lyu, R., Li, P., Chen, X., Zhang, W., Chen, K., Xue, T., Liu, X., Lu, C., Lin, D., Pang, J.: Embodiedscan: A holistic multi-modal 3d perception suite towards embodied AI. In: *IEEE/CVF Conference on*

- Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024. pp. 19757–19767. IEEE (2024) 1
52. Wang, W., Zhu, D., Wang, X., Hu, Y., Qiu, Y., Wang, C., Hu, Y., Kapoor, A., Scherer, S.A.: Tartanair: A dataset to push the limits of visual SLAM. In: IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2020, Las Vegas, NV, USA, October 24, 2020 - January 24, 2021. pp. 4909–4916. IEEE (2020) 3, 9
 53. Wang, Y., Zhou, J., Zhu, H., Chang, W., Zhou, Y., Li, Z., Chen, J., Pang, J., Shen, C., He, T.: pi3: Permutation-equivariant visual geometry learning. arXiv preprint arXiv:2507.13347 (2025) 1, 4
 54. Wu, C.: Towards linear-time incremental structure from motion. In: 2013 International Conference on 3D Vision, 3DV 2013, Seattle, Washington, USA, June 29 - July 1, 2013. pp. 127–134. IEEE Computer Society (2013) 4
 55. Wu, H., Wu, D., He, T., Guo, J., Ye, Y., Duan, Y., Bian, J.: Geometry forcing: Marrying video diffusion and 3d representation for consistent world modeling. CoRR **abs/2507.07982** (2025) 1, 2, 4
 56. Wu, Y., Zheng, W., Zuo, S., Huang, Y., Zhou, J., Lu, J.: Embodiedocc: Embodied 3d occupancy prediction for vision-based online scene understanding. CoRR **abs/2412.04380** (2024) 1
 57. Yan, W., Zhang, Y., Abbeel, P., Srinivas, A.: Videogpt: Video generation using vq-vae and transformers. arXiv preprint arXiv:2104.10157 (2021) 4
 58. Yang, L., Kang, B., Huang, Z., Zhao, Z., Xu, X., Feng, J., Zhao, H.: Depth anything V2. In: Globersons, A., Mackey, L., Belgrave, D., Fan, A., Paquet, U., Tomczak, J.M., Zhang, C. (eds.) Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024 (2024) 10
 59. Yin, T., Zhang, Q., Zhang, R., Freeman, W.T., Durand, F., Shechtman, E., Huang, X.: From slow bidirectional to fast autoregressive video diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 22963–22974 (2025) 4
 60. Zhang, Q., Zhai, S., Martin, M.Á.B., Miao, K., Toshev, A., Susskind, J.M., Gu, J.: World-consistent video diffusion with explicit 3d modeling. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2025, Nashville, TN, USA, June 11-15, 2025. pp. 21685–21695. Computer Vision Foundation / IEEE (2025) 10, 12
 61. Zheng, Y., Yang, P., Xing, Z., Zhang, Q., Zheng, Y., Gao, Y., Li, P., Zhang, T., Xia, Z., Jia, P., et al.: World4drive: End-to-end autonomous driving via intention-aware physical latent world model. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 28632–28642 (2025) 4