

Video Generation Models are General-Purpose Vision Learners

Letian Wang^{1,2}, Chuhan Zhang¹, Rishabh Kabra^{1,3}, Jasper Uijlings¹, Steven Waslander², Andrew Zisserman^{1,4}, Joao Carreira¹, Kaiming He^{1,5}, Misha Andriluka¹, Eduard Gabriel Bazavan¹, Andrei Zanfir¹ and Cristian Sminchisescu^{1,6,*}

¹Google DeepMind, ²University of Toronto, ³University College London, ⁴University of Oxford, ⁵MIT, ⁶Lund University

Driven by next-token prediction, NLP shifted from task-specific models into powerful generalist foundation models. What, then, is the equivalent catalyst needed to achieve a general-purpose model in computer vision? In this paper, we contend that large-scale text-to-video generation serves as a strong pre-training paradigm for computer vision, providing the necessary spatiotemporal priors, vision-language alignment, and scalability required for general visual intelligence. We introduce GenCepcion, which leverages a pre-trained video generative diffusion backbone to define a feed-forward perception model, capable of performing various vision tasks steered by text instructions. Empirical results demonstrate that GenCepcion achieves state-of-the-art performance across a diverse suite of tasks, including depth, surface normal, and camera pose estimation, expression-referring segmentation, and 3D keypoint prediction, often matching or surpassing specialized models (e.g. DepthAnything3, SAM3, D4RT, VGGT- Ω , Sapiens, David, Genmo, and Lotus-2). Furthermore, the video generative pretrained backbone outperforms alternative pretraining paradigms (e.g., V-JEPA, and Video MAE) under comparable settings. Importantly, GenCepcion exhibits preliminary data and model scaling properties along with exceptional data efficiency where it achieves comparable performance with leading models like D4RT and VGGT- Ω with $7\times$ to $500\times$ less training data. Finally, GenCepcion also exhibits intriguing emergent behaviors: a model trained exclusively on synthetic human videos generalizes to real-world footage and out-of-distribution object categories (e.g., animals and robots). These findings suggest that video generation is not merely a synthesis tool, but a foundational path toward generalist vision intelligence for the physical world.

Project page: <https://genception.github.io>

1. Introduction

Natural language processing (NLP) evolved from an era of specialized models—where separate models were developed for each language task (e.g. translation, summarization)—to a singular, unified foundation model paradigm. This paradigm, supported by large-scale next-token prediction pre-training followed by task-aligned post-training, has effectively collapsed thousands of disparate linguistic challenges into a single generalist intelligence, ultimately unlocking emerging behaviors such as chain-of-thought and in-context learning.

In stark contrast, computer vision is still lingering in its "specialized model" stage. While recent years have yielded powerful foundation models such as Segment Anything series [10, 36, 53] for localization or Depth Anything series [41, 76, 77] for geometry, these remain fundamentally task-specific models that need customized architectures for each task. We have mastered specialized perception, but we have yet to realize a unified vision foundation model—a unified task-agnostic architecture that mirrors LLMs from general pre-training to versatile, emerging intelligence. To this end, we posit that the quest for a generalist vision model is essentially a search for a universal pre-training objective—a visual analog to next-token prediction. Such a pre-training paradigm should satisfy three core imperatives:

- 1) *Spatio-Temporal Evolution*: The world is a 4D continuum. The pre-training objective should force the model to internalize the 4D temporal causality and physics of a world in motion.
- 2) *Vision-Language Alignment*: To inherit the instruction-following capabilities and common-sense knowledge of language models, visual features should be natively aligned with linguistic semantics during both pre-training and post-training stages.
- 3) *Scaled up*: The paradigm should be scaled in both data and compute, ultimately enabling the emergence of vision intelligence.

* Work done while at Google Deepmind.




Figure 1 | **Methodology** (Left): GenCception treats a video generative diffusion model as a *pre-training* base to capture rich spatio-temporal world priors and native vision-language alignment at scale. During multi-task *post-training*, the model is adapted to feed-forward model fine-tuned on predominantly synthetic data to handle diverse perception tasks. GenCception shows strong performance with intriguing *Emerging Behaviors*, enabling seamless sim-to-real transfer and generalization to out-of-distribution object categories. **Paradigm Shift** (Right): This highlights a paradigm shift from specialized task-specific computer vision models, to fully unified generalist vision models.

In this paper, we demonstrate that large-scale text-to-video generation serves as this pivotal pre-training paradigm, uniquely satisfying all three requirements, compared to prior vision pretraining methods (e.g. contrastive learning, masked autoencoder). First, the objective of generating high-fidelity video sequences forces a model to internalize the spatio-temporal world priors, including 3D geometry, object permanence, and physical interactions. Second, as these models are natively conditioned on text, they provide the inherent vision-language alignment. Finally, modern text-to-video generation models have been scaled using massive datasets and compute due to their relatively low annotation cost and high commercial value, providing a direct path toward emergent intelligence.

We introduce GenCception, a generalist video perception model that instantiate this paradigm. As illustrated in Figure 1, we treat the pre-trained video diffusion backbone as our "base model", leveraging the rich priors learned during generative pre-training. We then apply the post-training phase, where we fine-tune the model across diverse visual tasks primarily with synthetic data. This enables GenCception to execute heterogeneous perception tasks within a single task-agnostic architecture, ranging from pixel-level segmentation to more complex 3D pose estimation and tracking. Specifically, by repurposing iterative diffusion into a feed-forward architecture, GenCception dynamically infers the intended modality in a single forward pass, steered by the text instruction. Notably, GenCception promotes a paradigm shift from specialized vision models to unified generalist intelligence that adopt unified backbone, head, and loss function across different vision tasks. Under this paradigm, task specifications are shifted from architectural modifications to data format designs, which enables scaling up with an ever-growing spectrum of visual tasks.

As shown in Figure 2, our empirical results demonstrate that GenCception achieves state-of-the-art performance across multiple perception tasks, performing on par with or even surpassing specialized models designed for single modalities, including DepthAnything V3 [41], SegmentAnything V3 (SAM3) [10], D4RT [82], VGGT- Ω [65], Sapiens [33], David [56], and Genmo [38]. Under comparable model size and same fine-tuning data, the video generative pretrained backbone is also found to outperforms the largest available variants of alternative pretraining paradigms (e.g., V-JEPA [4], VideoMAE V2 [61]). Importantly, GenCception exhibits preliminary scaling properties where the performance improves with more data and larger model size, along with exceptional data efficiency where it achieves comparable performance with leading models like D4RT and VGGT- Ω with 7×




Figure 2 | **SOTA Generalist Capability** (Left): GenCepcion achieves universally competitive performance on a wide range of vision tasks, matching or outperforming state-of-the-art models dedicated to individual tasks (e.g. DepthAnything3 [41], SAM3 [10], D4RT [82], VGGT-Ω [65], Sapiens [33], David [56], Genmo [38], Lotus-2 [22]). Our specialist denotes a model trained on each task individually, whereas the generalist represents a single model trained jointly across multiple tasks. **Data Efficiency in Finetuning** (Right): Validated on depth estimation, the video generative pretrained backbone (i) outperforms the largest available variants alternative pretraining paradigms (e.g., V-JEPA, and VideoMAE V2) under the same finetuning data. (ii) exhibits preliminary scaling properties, where the performance improves with more data and large model size; (iii) shows exceptional data efficiency, achieving comparable performance with leading models like D4RT [82] and VGGT-Ω [65] with 7× to 500× less training data.

to 500× less training data. Finally, this paradigm triggers emergent behaviors, including seamless sim-to-real transfer and out-of-distribution generalization to unseen object categories, as shown in Figure 10b and Figure 10a. This evidence suggests that video generation is not merely a synthesis tool, but the foundational path toward generalist vision intelligence for the physical world.

2. Related Work

2.1. Perception Foundation Models.

Perception foundation models have recently attracted significant attention, with frameworks such as Segment Anything series [10, 36, 53] and Depth Anything series [41, 75, 76] leveraging massive datasets to achieve robust performance across diverse visual scenarios. Parallel efforts have sought to develop unified models capable of handling multiple tasks simultaneously. However, these works operate primarily in the image domain [2, 47, 56], lacking an inherent understanding of temporal dynamics. Furthermore, multitasking in these models is often implemented at a rigid architectural level—relying on task-specific encoder, decoder, or losses [56]—which lacks the flexibility required for truly general-purpose systems. Even among models that attempt to process video and flexibly handle varied tasks [26, 86], they have not yet reached the level of efficacy seen in specialized, task-specific models when evaluated on a wide range of video tasks. This stands in stark contrast to Large Language Models (LLMs), which have demonstrated that a single unified architecture can flexibly master a vast array of tasks with human-level or even superhuman proficiency. A truly unified video perception model that matches the general-purpose proficiency of LLMs has yet to be realized.

2.2. Visual Representation Learning

At the heart of this challenge lies the long-standing task of representation learning: learning rich and robust visual representations to enable a wide range of vision tasks. In the literature, unsupervised and self-supervised learning have become core paradigms for representation learning, as they scale efficiently with large amounts of unlabeled data. Among representative approaches, masked autoencoders [23, 49] (MAE), inspired by masked language modeling [8, 14], learn to reconstruct

missing image regions and have established a foundational framework for visual pretraining, where the learned representations can be effectively finetuned to various downstream vision tasks. Later works like VideoMAE [61] and RVM [87] have extended this to the video domain. In parallel, another prominent self-supervised paradigm leverages feature-level self-distillation, as exemplified by the DINO series [12, 48, 59], where a student network is trained to match the feature distributions of a dynamically updated teacher network across different global and local views of the same image. However, these vision-only methods inherently lack explicit multimodal vision-language alignment, lacking the semantic grounding required to link visual patterns to high-level concepts. Another prominent paradigm, contrastive learning, aims to align semantically similar samples while separating dissimilar ones. By performing contrastive learning across vision and language modalities, models such as CLIP [51] and SigLip [81] learn powerful multimodal representations that unlock a wide range of tasks such as open-vocabulary classification and segmentation. Despite the vibrancy of this field, multimodal representation learning at the video level remains a crucial problem, presenting numerous challenges. Chief among these is the severe limitation in model scaling; while language models routinely scale to hundreds of billions of parameters, prominent video representation learners remain orders of magnitude smaller. This bottleneck stems from, among other reasons, the prohibitive computational cost of training on dense video data, alongside the added complexity of modeling intricate temporal dynamics.

2.3. Re-purposing Diffusion Models.

Recently, an emerging direction is to leverage the rich features learned in pre-trained diffusion models, a path which our work follows. Early work in this area focused mainly on single-image prediction. *Marigold* [30] demonstrated how a pre-trained image diffusion model like Stable Diffusion [55] can be repurposed through fine-tuning on synthetic data to act as a monocular depth estimator. [?] improves the efficiency of prior diffusion-based depth estimators [17, 30] by aligning the time step with the noise level, showing that a simple end-to-end fine-tuning approach can create a single-step, deterministic model that is both fast and highly accurate. *GenPercept* [71] conducts a systematic study and identifies key components such as feature injection, decoding mechanisms, and training objectives, that are critical for successfully adapting pre-trained diffusion models to dense perception tasks. *Diception* [84] demonstrates strong vision generalists capable of addressing multiple tasks, steering perception tasks using text prompts.

An inherent limitation of these single-image models, however, is their inability to produce temporally consistent predictions when applied to video sequences. To address this, *BufferAnytime* [37] proposes to adapt the existing image diffusion model to the video domain by incorporating a temporal layer for depth and normal estimation problems. Going further, various works have explored directly employing native video diffusion models to address the temporal consistency issues. *DepthCrafter* [25] and *NormalCrafter* [6] incorporate alignment objectives with large-scale Vision Foundation Models, such as CLIP [51] or DINO [48], to facilitate training. [74] explores training with scalable synthetic data to overcome the scarcity of video ground truth data. *Geo4D* [28] adapts video diffusion models for 4D scene reconstruction. *DiffusionRenderer* [40] leverage video diffusion priors to tackle the dual problems of inverse rendering and forward rendering. *ReferEverything* [3] repurpose video diffusion models for expression-referring open-world segmentation. [1] provides related evidence that video diffusion models encode reusable visual priors, demonstrating few-shot LoRA adaptation to diverse tasks via input-output video transition. However, these advanced methods primarily concentrate on a single, often dense, vision task. We present a unified architecture for both dense and sparse video perception, leveraging the rich prior knowledge and instruction-following strengths of text-to-video models to steer diverse tasks via text prompts. Additionally, we replace slow iterative sampling with an efficient, direct feed-forward architecture, with superior inference efficiency and performance that matches or exceeds SOTA specialists. This work advances our previous human-centric model, THFM [66], into a unified general-purpose model, with broader task and benchmark evaluations, and more exhaustive analysis of model behaviors.

Notably, our work shares close conceptual ties with two concurrent works, while differing in both architecture and evaluation. First, both our work and [69] share the hypothesis that large-scale video




Figure 3 | **Capabilities Illustrations.** GenCception is a versatile video perception model with SOTA performance on a multitude of output modalities including surface normal estimation, depth estimation, foreground segmentation, expression-referring open-world segmentation, dense human pose estimation, 2D/3D keypoints estimation, and camera pose estimation. Our approach has been trained on human-centric synthetic videos, yet it generalizes to real videos both for people as well as other categories such as animals and anthropomorphic characters.

generative models learn powerful visual priors. If [69] demonstrates the existence of these priors, we provide the architecture to efficiently harness them. Specifically, while [69] investigates a training-free prompting approach via multi-step video generation with a focus primarily on qualitative validation, we introduce a dedicated post-training strategy and conduct rigorous quantitative evaluation across standardized benchmarks. Second, while *Vision Banana* [18] operates within the 2D image space, GenCception extends the paradigm to the native video domain to naturally capture temporal consistency. Additionally, instead of focusing on multi-step generation, our method adopts an efficient feed-forward architecture with improved inference efficiency and strong performance across a broader task spectrum. See Figure 3, Figure 6, Figure 7, and Figure 8 for more demonstrations of capabilities.

3. GenCception

3.1. Methodology

The core premise of GenCception is that large-scale video generation models are not merely synthesis engines, but powerful, universal visual representation learners. Inspired by the trajectory of Large Language Models (LLMs)—where generative pre-training (next-token prediction) inherently builds rich features that can be harnessed for diverse downstream tasks—we reformulate video perception as a post-training endeavor built atop a pretrained video generative model. Our methodology is driven by three foundational principles:

- 1) *Multimodal Generative Pre-training as Representation Learning:* We utilize a pre-trained text-to-video generative diffusion model as our universal backbone. The objective of generating high-fidelity video conditional on text inherently forces the model to internalize robust spatiotemporal priors, 3D geometry, and physical laws, while simultaneously fostering a deep, aligned understanding of visual-language concepts. To fully preserve these foundational representations, one core design principle is to maximize compliance with the original pre-training regime with minimal modifications.
- 2) *Task-Agnostic Post-Training in Unified Architecture:* To transition from a specialized generator to a generalist perceiver, we treat diverse vision tasks as unified sequence-to-sequence mapping problems, rather than engineering task-specific architectures or relying on specialized encoders, decoders, and loss functions. By formatting various target modalities into a shared representational space (e.g., standard RGB space for dense tasks), we allow a single, unified architecture to master multiple tasks guided seamlessly by text instructions. Under this paradigm, integrating new visual capabilities no longer necessitates altering the architecture or training recipes;




Figure 4 | Architecture overview of GenCception, a simple yet powerful architecture adapted from text-to-video diffusion models. Given an input video and a text prompt specifying the desired output, our unified model, trained majorly on synthetic data, is capable of performing a wide range of dense and sparse perception tasks, with a single forward-pass of the model. The dense vision tasks are unified in the RGB ambient space where supervision can be applied in latent space efficiently, and the sparse vision tasks are realized by adding learnable tokens as additional inputs to the diffusion transformer (DiT).

instead, task specifications are directly reflected in the data representation—a data-driven approach that enables scaling up efficiently across a continually expanding range of visual tasks.

- 3) *Feed-Forward Reformulation*: While diffusion models typically rely on slow, iterative sampling processes for synthesizing diverse outputs, perception tasks demand highly accurate and efficient predictions. Therefore, a critical design decision in GenCception is the transformation of the multi-step generative backbone into a single-step, feed-forward perception model.

Guided by these principles, we instantiate GenCception, a model capable of various video perception tasks. As illustrated in Figure 4, we repurpose text-to-video generation models to enable both dense vision tasks (normal estimation, depth estimation, foreground segmentation, expression referring open-world segmentation, dense human semantic estimation, raymap estimation) and sparse vision tasks (2D and 3D keypoint prediction) within a unified architecture. Given the input RGB video and the text prompt specifying the target modality, GenCception is able to produce target video in a single forward pass (Sec 3.2), with a unified representation for various perception tasks (Sec 3.3); A scalable data synthesis strategy is proposed to enable efficient low-cost collection of diverse high-quality data (Sec 3.4). We then discuss our training recipe in Sec 3.5.

3.2. From Diffusion Pre-training to Perception Finetuning

Preliminary. Our model is based on text-to-video diffusion model, with its core components being a VAE encoder-decoder pair, a text encoder, and a transformer-based latent diffusion model (DiT). The original DiT model takes a latent token initialized with Gaussian noise and text prompt as inputs and conducts a number of denoising steps, gradually transforming it into latent tokens, which can be decoded to the generated video. The index of the denoising step is provided to DiT as additional input to modulate the predictions with respect to amount of input noise. Let us denote the video token inputs to DiT at step t as x_t , with x_0 being the noise-free latents of the original video. Depending on the particular formulation, DiT module is trained to predict the noise component ϵ as in [24, 60], the velocity $v = \epsilon - x_0$ between the noise and original input as in the "Rectified Flow" variant [16, 44, 45], or the target x_0 directly [39].

Feedforward Perception Formulation. In this work, we repurpose the multi-step diffusion transformer into a feed-forward prediction model. To that end, instead of taking noisy video latents as inputs, we directly feed the clean latent of the input video as the inputs to DiT, and conduct only a single forward-pass of the model. To signify that the input is noise-free, the conditioning timestep is fixed to $t = 0$, corresponding to the termination of the generative Rectified Flow process. Since the

DiT is trained under the Rectified Flow objective to predict the velocity $v = \epsilon - x_0$, we thus negate the raw output of the DiT prior to passing it to the loss functions or decoder. The negated output $-v = x_0 - \epsilon$ aligns with the latent of the target video more closely, which, empirically accelerates model convergence and improves performance.

An intuitive way to view this formulation is that we recast the DiT as a powerful feature extractor. Crucially, by leveraging Rectified Flow rather than traditional noise prediction, conditioning on $t = 0$, and negating the velocity output, we strive to extract the most informative and richest features from the pretrained backbone. Furthermore, we harvest features directly from the final layer rather than intermediate ones; this ensures they natively align with the subsequent decoder, without any architectural modifications to the base model.

3.3. Unified Task Representation

Unify Dense Tasks in RGB Space. Unlike prior practices that often adopt separate backbones, decoders, and loss functions for different tasks, our framework promotes a unified approach, where all tasks share the same architecture and weights for both backbone and decoder, and the target task is simply modulated via the text prompt. To achieve this, the outputs for all dense perception tasks are represented in the 3-channel RGB space within the range $[0, 1]$. Specifically, the three RGB channels are replicated for single-dimensional outputs like depth and segmentation, whereas they represent distinct dimensions for three-dimensional tasks including surface normals and DensePose estimation. Importantly, this representation is also applicable to higher-dimensional vision modalities, many of which can be elegantly adapted to fit within this 3-channel constraint. For instance, to estimate camera poses which are traditionally represented as matrices, we employ a pixel-space raymap [28, 83, 85] that natively aligns with the video generative format (see Figure 5). To fit its 6-channel format within our 3-channel constraint, we spatially partition the frame—allocating ray origins to the central region and ray directions to the periphery. This simple yet effective layout preserves the integrity of our single-decoder framework while fully leveraging pre-trained visual priors. More importantly, this design philosophy scales naturally across the spectrum of vision tasks. Analogous to reformatting non-text data into text strings for LLMs, we advocate for projecting fundamentally visual tasks directly into the continuous pixel space—the native domain where pre-trained visual priors are best leveraged.




Figure 5 | The "Rothko" Raymap as an example of adapting high-dimensional modal data into standard 3 RGB channels. This representation effectively compresses the camera's multi-channel ray data by assembling rotation and translation components into a single three-channel map.

Enable Sparse Tasks with Learnable Tokens. While RGB is versatile, downstream tasks like robotic control often require directly manipulatable, structured outputs such as explicit 2D/3D coordinates. To enable these sparse predictions, we introduce a lightweight token-based extension. We append T learnable tokens—one for each video frame—to the video latents. After passing through the model, an MLP decodes each token to predict a K -dimensional target per frame. To maintain compliance with the base DiT model, we apply its native 3D RoPE to the additional tokens. Specifically, because the spatial positions of these additional tokens are unknown, we make them learnable. For temporal positioning, since the video frame count T is compressed to a latent length T' , we down-scale the frame indices of the T tokens using position interpolation [13] to stay within the temporal bounds seen during pre-training. Empirically, this additional-query approach outperforms methods that rely on adding extra attention layers.



Figure 6 | Demonstration of GenCeption’s depth and surface normal estimation capabilities.

3.4. Scalable Synthetic Data Generation

Training of our model requires high-quality and diverse video data with ground-truth for depth, normal, segmentation mask, dense pose, 2D keypoints, 3D keypoints, and camera poses. Since most real-world datasets may contain only a subset of the modalities in a limited scale, we resort to synthetic data as a scalable approach to address the data scarcity problem.

Specifically, we design a synthetic data generation workflow to produce high-quality human-centric video data, covering diverse human entities and motions. We use 800 RenderPeople assets [54], and animate them with 200 motions from the CMU motion capture dataset [11]. We use various 3D full scenes or HDRI backdrops similar to [5] for enhanced background variability, and vary focal lengths, camera positioning and trajectory for enriched camera conditions. In total, we generate a synthetic

dataset of 7,500 videos with various identities, motions and backgrounds. To support dense vision tasks, video normals, depth and segmentation masks are generated with separate render passes via Blender [7]. We recover the human joint positions from the rigged RenderPeople [54] assets and use them as ground truth for our 2D and 3D keypoint regression task. Each of the generated videos is trimmed to the target number of frames in the video model. We apply a offline data preprocessing stage, where we precompute and cache both the input RGB video and target-modality video into video latents, and the text conditioning prompt into text embeddings.

3.5. Training Recipe

With the model architecture and data pipeline established, we now describe the training objectives employed in our framework. Conventionally, joint multi-task learning presents an inherent optimization challenge: distinct vision tasks demand specialized loss functions tailored to their unique modalities. In practice, this heterogeneity introduces substantial engineering overhead in manually balancing loss weights across different objectives, presenting a major bottleneck when scaling up to a larger variety of vision tasks, data, and larger model sizes. Instead, we adopt a fully unified loss design: our model is trained solely using a standard L_2 loss, mirroring how LLMs unify diverse language tasks under a single objective. This single loss function is applied directly in the latent space for dense tasks and in the output space for sparse tasks.

To achieve this unified optimization without compromising performance, we shift the responsibility of task-specific customization from the loss formulation to the data representation designs. Taking monocular depth estimation as a primary example—where scale ambiguity typically necessitates customized, scale-invariant losses—we resolve this constraint directly at the data level. Specifically, we normalize the depth map of each video using the median depth of the scene, thereby inherently eliminating scale ambiguity. To further align the normalized depth d into the standard $[0, 1]$ RGB range, similar to [18], we apply a nonlinear mapping function: $d' = \text{clip}(\alpha \log(d + 1), 0, 1)$, where α can dynamically adjust the model’s focus between near-field details and far-field structures. This data-level customization effectively eliminates the need for any task-specific objective functions, where balancing between different tasks only needs to be managed via the data mixture ratio, analogous to the prevailing practice in LLMs. We believe this data-centric harmonization principle can scale naturally across a wider spectrum of vision tasks, ensuring a streamlined and robust multi-task training paradigm.

4. Experiments

We comprehensively evaluate our model on various challenging real-world datasets across vision tasks: depth estimation, normal estimation, segmentation, 3D keypoint prediction, and camera pose estimation.

4.1. Implementation Details

Our approach is implemented based on an open-source and open-weights text-to-video diffusion model WAN 2.1 [63]. Following its original architecture, we train our model on video samples at a resolution of 480x832, comprising 81 frames captured at 24 frames per second (FPS). The model’s encoder downsamples the input dimensionality with a temporal factor of 4 and a spatial factor of 8 for both width and height. We train our model with a batch size of 64 on 256 v6e TPUs. We optimize the model using the Adam optimizer [35], with a learning rate of $5e^{-5}$ for 15000 training steps, implementing a linear warmup over the first 250 training steps. Crucially, for training stability, we apply both gradient clipping, constraining the gradient norm within a certain threshold, and gradient dropping, which discards any batch whose gradient norm exceeds a higher threshold. These techniques proved vital for stable training and high-quality performance. To enhance training data diversity, we augment our proposed synthetic data with a few publicly available synthetic datasets. Specifically, we incorporate TartanAir [67], Virtual KITTI [19], and MVS Synth [27] to provide additional depth annotations, while utilizing the camera trajectory data provided by TartanAir [67]. For expression-referring segmentation tasks, we leverage real-world datasets due to their immediate availability, integrating MeViS [15], Ref-COCO [80], and YouTube-VOS [72].




Figure 7 | Qualitative results of our approach on referring expression segmentation. Our model accurately recognizes objects, colors, spatial relationships, and motion. Furthermore, it demonstrates strong generalization by identifying unseen objects, effectively leveraging the knowledge acquired during text-to-video pretraining.




Figure 8 | **3D Pose Estimation**. Example 3D pose estimation results on challenging skiing (top) and snowboarding (bottom) video. We show every 10-th frame for a subsequence of the input video. Compared to the existing state-of-the-art 3D human pose estimation methods GenCepion takes a full video frame as input and does not require pre-processing steps such as person detection or 2d keypoint estimation.

4.2. Inference Cost

Our feedforward model eliminates the standard WAN’s slow 50-step diffusion. On a single v6e TPU (81 frames, 480x832), the 1.3B model takes 5.92s at 13.6 FPS (DiT 0.96s) and 15.3GB VRAM, while the 14B model takes 10.03s at 8.0 FPS (DiT 5.11s) and 42.8GB VRAM. Both share a 10GB text encoder and 0.25GB VAE, which can be offloaded to system memory to trade speed for VRAM.

Table 1 | Comparison to state-of-the-art specialized models. We show the performance for a common subset of tasks, including surface normal, depth estimation, camera pose estimation, foreground segmentation, expression-referring open-world segmentation, and 3D human keypoint estimation. "-" denotes that a given model cannot solve the task out-of-the-box, ~ denotes results are not obtained, * denotes results when trained with the same data as our model. Best results are **bold**, second best underlined. Note that on all tasks except for expression-referring segmentation, our method is trained purely on synthetic data.

Task	Benchmark	Normals		Depth				Cam. Pose			Foreground Seg.		Expression-Referring Seg.		3D Human
		Sintel [9]	Hi4D [79]	Sintel [9]	KITTI [21]	ETH3D [57]	Goliath [46]	Sintel [9]	RPE-T ↓	RPE-R ↓	V.Mat. [42]	P.Mat. [42]	Ref-DAVIS [34]	MEVIS [15]	EMDB [29]
Method	Backbone	mAE ↓		AbsRel ↓				ATE ↓			MSE ↓		J&F ↑		MPJPE ↓
David-Large [56]	VIT-L 0.34B	-	15.37	-	-	-	-	-	-	-	-	-	-	-	-
Sapiens [33]	VIT-L 2B	-	12.14	-	-	-	-	-	-	-	-	-	-	-	-
Marigold-E2E-FT [20]	SD V2 0.86B	33.5	~	-	-	-	-	-	-	-	-	-	-	-	-
NormalCrafter [6]	SVD 1.5B	30.7	~	-	-	-	-	-	-	-	-	-	-	-	-
Lotus-2 [22]	FLUX 12B	30.3	~	-	-	-	-	-	-	-	-	-	-	-	-
David-Large [56]	VIT-L 0.34B	-	-	-	-	-	0.012	-	-	-	-	-	-	-	-
Sapiens [33]	MAE 2B	-	-	-	-	-	0.033	-	-	-	-	-	-	-	-
MapAnything [32]	DinoV2 1.1B	-	-	0.306	0.096	~	-	0.202	0.089	2.383	-	-	-	-	-
VGGT [64]	DinoV2 1.1B	-	-	0.247	0.067	~	-	0.168	0.056	0.428	-	-	-	-	-
DepthAnything 3 [41]	DinoV2 1.15B	-	-	0.201	0.059	0.028	~	0.065	0.048	0.456	-	-	-	-	-
D4RT [82]	VideoMAE2 1B	-	-	0.148	0.051	~	-	0.065	0.024	0.126	-	-	-	-	-
VGGT-Q [65]	DinoV3 1B	-	-	<u>0.145</u>	0.041	0.016	~	<u>0.057</u>	0.059	<u>0.407</u>	-	-	-	-	-
MODNet [31]	MNetV2	-	-	-	-	-	-	-	-	-	0.0054	0.0304	-	-	-
RVM [43]	MNetV3	-	-	-	-	-	-	-	-	-	0.0010	-	-	-	-
David-Large [56]	VIT-L 0.34B	-	-	-	-	-	-	-	-	-	~	<u>0.0009</u>	-	-	-
ReferFormer [70]	Swin-L 0.19B	-	-	-	-	-	-	-	-	-	-	-	61.1	31.0	-
VISA [73]	Chat-Univ 13B	-	-	-	-	-	-	-	-	-	-	-	70.4	44.5	-
Vocap [62]	SAM2 0.4B	-	-	-	-	-	-	-	-	-	-	-	75.1	51.9	-
ReferEverything [3]	WAN 2.1 14B	-	-	-	-	-	-	-	-	-	-	-	75.0	60.3	-
SAM 3 [10]	PE 0.8B	-	-	-	-	-	-	-	-	-	-	-	41.3	38.3	-
SAM 3 + Gemini 3.5 Flash [10]	PE 0.8B	-	-	-	-	-	-	-	-	-	-	-	64.5	57.5	-
GVHMR [58]	VT	-	-	-	-	-	-	-	-	-	-	-	-	-	72.6
TRAM [68]	VTH	-	-	-	-	-	-	-	-	-	-	-	-	-	74.4
Genmo [38]	VTH	-	-	-	-	-	-	-	-	-	-	-	-	-	<u>73.0</u>
Diception - Generalist [84]	SD3 2B	37.9	20.4	0.500	0.145	0.177	~	-	-	-	-	-	-	-	-
Ours - Specialist - S	WAN 2.1 1.3B	32.7	~	0.167	0.060	0.056	~	<u>0.057</u>	<u>0.023</u>	0.720	~	~	69.7	59.9	~
Ours - Specialist - L	WAN 2.1 14B	<u>29.7</u>	10.99	0.130	<u>0.048</u>	0.037	0.008	0.050	0.018	0.464	<u>0.0027</u>	<u>0.0009</u>	76.4	70.0	71.8
Ours - Generalist - L	WAN 2.1 14B	29.3	<u>11.47</u>	0.156	<u>0.048</u>	0.044	<u>0.011</u>	0.062	0.018	0.485	0.0010	0.0008	<u>75.8</u>	<u>69.0</u>	~

Table 2 | Validated on depth estimation, the generative backbone outperform other visual pretraining (V-JEPA and Video MAE v2), and show preliminary data and model scaling properties, where performance improves with data volume and model size Note that our method was trained exclusively on synthetic data.

Method	Model Size	Number of Training			Sintel		KITTI		ETH3D		Average	
		Datasets	Videos	Frames	AbsRel↓	$\delta_1\uparrow$	AbsRel↓	$\delta_1\uparrow$	AbsRel↓	$\delta_1\uparrow$	AbsRel↓	$\delta_1\uparrow$
Ours - V-JEPA - H	0.6B	1	7.5K	0.9M	0.422	37.3	0.226	47.5	0.196	71.8	0.281	52.2
Ours - VideoMAe V2 - H	0.6B	1	7.5K	0.9M	0.275	37.4	0.124	64.6	0.126	84.1	0.175	62.0
Ours - VideoMAe V2 - G	1B	1	7.5K	0.9M	0.260	41.4	0.105	70.2	0.099	89.3	0.154	66.9
Ours - WAN 2.1 - S	1.3B	1	7.5K	0.9M	0.201	72.7	0.099	89.2	0.068	95.6	0.122	85.8
Ours - WAN 2.1 - L	14B	1	7.5K	0.9M	0.181	76.8	0.060	97.1	0.039	98.2	0.093	90.7
DepthAnything V3 - G [41]	1.15B	22+	1.23M	~ 200M	0.201	72.1	0.059	96.6	<u>0.028</u>	<u>98.8</u>	0.096	89.1
D4RT [82]	1B	11+	~ 1M	~ 86M	0.148	80.3	0.055	97.9	0.045	97.0	0.082	91.7
VGGT - Q [65]	1B	33+	~ 3M	~ 600M	<u>0.145</u>	87.6	<u>0.041</u>	98.5	0.016	99.6	0.067	95.2
Ours - WAN 2.1 - S	1.3B	4	8.08K	1.23M	0.167	76.9	0.060	97.3	0.056	97.6	0.094	90.6
Ours - WAN 2.1 - L	14B	4	8.08K	1.23M	0.130	<u>84.5</u>	<u>0.048</u>	<u>98.3</u>	0.037	<u>98.8</u>	<u>0.071</u>	<u>93.8</u>

4.3. Evaluation Protocol

Evaluation Datasets. We evaluate our approach on a set of challenging benchmarks that encompass both video and image data. For surface normal estimation, we evaluate on Hi4D datasets [79] and SINTEL dataset [9]. For Hi4D, we adhere to the evaluation protocol established in Sapiens [33] and DAVID [56], selecting the same sequences from pairs 28, 32, and 37. These sequences feature six unique subjects captured by camera 4, resulting in a total of 1,195 frames for testing. For depth estimation, we evaluate on KITTI [21], SINTEL [9], ETH3D [57], and Goliath [46] dataset. For Goliath [46], we follow DAVID [56] to perform evaluation to ensure fair comparison, where we utilize the identical selection of 12 cameras, 16 frames, and 4 subjects as in DAVID [56], which are categorized into data splits of face, upper and full body, with approximately 2.2k frames in total. For soft foreground segmentation, we report results on VideoMatte [42], PhotoMatte 85 [42], and PPM-100 [31]. The original VideoMatte dataset does not provide an official split, thus we evaluate on the static split composite from [43]. For camera pose estimation, we evaluate on SINTEL dataset [9]. For 3D keypoint estimation, we evaluate on EMDB dataset [29]. For expression-referring open-world segmentation, we follow VoCap [62] to reports results on RefVOS-DAVIS [34] and MeViS [15].

Evaluation Metrics. To assess our model’s performance, we report standard metrics appropriate for each task. For surface normal estimation, we report the common metrics [33, 56] including mean and median angular error, alongside the percentage of pixels where the error is within the threshold $t \in \{11.25, 22.5, 30\}$. For depth estimation, we follow standard practice [33, 52, 76] by adopting the mean absolute value of the relative depth (AbsRel) and the root mean square error (RMSE). For soft foreground segmentation, we report mean squared error (MSE) as in prior work [42, 43, 56, 78]. For



Figure 9 | Effect of transferring increasing number of layers from the pre-trained text-to-video model.

expression referring open-world segmentation, we report J&F scores [50] (mean of IoU and contour accuracy) averaged on all text queries. For 3D keypoint estimation, we follow [38, 68] to report mean per joint position error (MPJPE). For camera pose estimation, we report average translation error (ATE), relative pose error - translation (RPE-T), and relative pose error - rotation (RPE-R).

4.4. Comparison to the state of the art

We compare our methods with the state-of-the-art models on various challenging benchmarks. Note that our model was trained without using any of the training sets provided alongside the evaluation benchmarks. Furthermore, our training relies entirely on synthetic data, with the sole exception of the expression-referring segmentation task which incorporates real-world data. As summarized in Table 1, despite its unified and simple architecture, our approach surpasses or delivers highly competitive performance against heavily tailored, task-specific state-of-the-art methods. For geometric understanding, our method exceeds specialized models like NormalCrafter [6] and Lotus-2 [22] in surface normal estimation, and outperforms or matches dedicated foundation models such as DepthAnything 3 [41], D4RT [82], and VGGT Ω [65] in depth estimation and camera pose estimation. In foreground segmentation, our approach achieves competitive results in soft foreground segmentation across both video and image datasets. On expression-referring segmentation, our work show stronger ability in understanding complex sentences compared to SegmentAnything3 [10]. Finally, for sparse tasks, it also delivers stronger 3D keypoint prediction performance compared to specialized models including Genmo [38] and TRAM [68]. Qualitative examples are show in Figure 6, Figure 7, and Figure 8.

4.5. Ablation Studies

Joint training vs. task-specific training: As demonstrated in Table 1, transitioning from a specialist to a unified generalist yields mixed behaviors across tasks and benchmarks. For dense prediction modalities, surface normal estimation shows diver trends on benchmarks, while depth and camera pose estimation generally undergo regression or remain intact (e.g., on KITTI). In contrast, foreground segmentation tasks consistently benefit from joint training, whereas expression-referring segmentation remains largely unchanged.

Most notably, we empirically observe that joint training severely degrades 3D human keypoint estimation and compromises other dense tasks. We presume this stems from two architectural factors. First, the explicit token-based coordinate regression diverges from the continuous pixel-space pre-training of the generative video model. Second, the additional learnable tokens, trained from scratch, disrupt the native attention mechanics of pre-trained DiT layers, damaging their optimized features and requiring more data to converge. This phenomenon suggest a valuable insight for post-training: one should strive for minimal architectural modifications to the pre-trained backbone,




Figure 10 | **Generalization capabilities.** Left: Trained on synthetic data with one single object within each video, our method generalizes in zero-shot to real videos with multiple objects. Right: Our approach has been trained only on videos of a single class (humans), but generalizes to a variety of other classes of articulated objects.

or alternatively, fundamentally redesign the pre-training paradigm to natively accommodate highly diverse downstream tasks.

Importance of pretraining. We compare video generative pretraining backbone with other representation learning approaches. We use the largest available variants of V-JEPA [4] and VideoMae V2 [61] and fine-tune all on the same training set. As in Table 2, our diffusion-based pretraining (WAN 2.1) significantly outperforms other pretraining approaches. This serves as a suggestive empirical evidence that the generative diffusion objective itself—rather than just the dataset or model scale—is essential for extracting superior spatiotemporal priors. Furthermore, VideoMae and V-JEPA lack native support for a unified multi-task model, often requiring task-specific training, while our model can natively use text prompts to steer between tasks. Furthermore, as in Figure 9, training the randomly initialized DiT from scratch yields a nearly flat learning curve. Performance improves as more pre-trained layers are used. This confirms that pre-training is essential for downstream convergence and performance.

Preliminary scaling property and exceptional data efficiency. As in Figure 2 and Table 2, the generative backbone demonstrates promising preliminary scaling properties across varying model sizes and datasets, with performance generally improving as model size and data volume increase. The model also shows exceptional data efficiency where it achieves comparable performance with leading models like D4RT and VGGT- Ω with 7 $\times$ to 500 $\times$ less training data.

4.6. Emergent Behaviors

In addition to competitive performance that matches or surpasses state-of-the-art methods, our model also exhibits intriguing emergent behaviors beyond its intended training objectives. We believe these behaviors arise from the rich representations learned in large-scale text-to-video generation models, shedding light on the broader potential of this paradigm for future exploration.

Generalization from synthetic training data: Our approach is trained purely on synthetic videos, yet it generalizes remarkably well to real-world videos. Quantitatively, as shown in Table 1 and Table 2, our model establishes new state-of-the-art performance across all benchmarks. Qualitatively, we observe that the fine-grained details present in the output of our model are exceeding the quality of our training data generated with conventional computer graphics methods in Blender, as illustrated in Figure 3, Figure 6, Figure 7, and Figure 8. For example, see the whiskers on the cat in Figure 10b (first column) or the soft segmentation of the hair in Figure 3 (second row).

Generalization to multiple instances: trained on synthetic data with one single object within each video, our method generalizes in zero-shot to real videos with multiple objects as shown in Figure 10a.

Generalization to OOD classes: our approach is trained purely on synthetic videos of human entities, yet it generalizes well to real-world videos—both of humans and of other categories such as animals

and anthropomorphic characters—as shown in Figure 10b.

5. Conclusion

We have presented GenCepion, a unified framework that identifies large-scale video generation as a foundational pre-training paradigm for visual perception. By repurposing a pre-trained video diffusion backbone into a high-efficiency, feed-forward model, we demonstrate that the rich spatiotemporal priors and vision-language alignment required for synthesis at large scale can be directly translated into precise perceptual capabilities without the need for costly iterative sampling. Our experiments across diverse vision tasks confirm that GenCepion not only matches the performance of specialized state-of-the-art models but also exhibits significant emergent behaviors, provides strong evidence for the existence of a universal "world model" within generative video backbones. We believe this paradigm shift, from task-specific engineering to scalable generative pre-training, paves a promising path toward a truly unified and generalist vision intelligence for the physical world.

References

- [1] Acuaviva, P., Davtyan, A., Hassan, M., Stapf, S., Rahimi, A., Alahi, A., Favaro, P.: From generation to generalization: Emergent few-shot learning in video diffusion models. arXiv preprint arXiv:2506.07280 (2025) 4
- [2] Bachmann, R., Kar, O.F., Mizrahi, D., Garjani, A., Gao, M., Griffiths, D., Hu, J., Dehghan, A., Zamir, A.: 4m-21: An any-to-any vision model for tens of tasks and modalities. Advances in Neural Information Processing Systems 37, 61872–61911 (2024) 3
- [3] Bagchi, A., Bao, Z., Wang, Y.X., Tokmakov, P., Hebert, M.: Refereverything: Towards segmenting everything we can speak of in videos. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 23221–23231 (2025) 4, 11
- [4] Bardes, A., Garrido, Q., Ponce, J., Chen, X., Rabbat, M., LeCun, Y., Assran, M., Ballas, N.: Revisiting feature prediction for learning visual representations from video. arXiv preprint arXiv:2404.08471 (2024) 2, 13
- [5] Bazavan, E.G., Zanzir, A., Zanzir, M., Freeman, W.T., Sukthankar, R., Sminchisescu, C.: Hspace: Synthetic parametric humans animated in complex environments (2022), <https://arxiv.org/abs/2112.12867> 8
- [6] Bin, Y., Hu, W., Wang, H., Chen, X., Wang, B.: Normalcrafter: Learning temporally consistent normals from video diffusion priors. arXiv preprint arXiv:2504.11427 (2025) 4, 11, 12
- [7] Blender Online Community: Blender - a 3d modeling and rendering package (2025), <https://www.blender.org> 9
- [8] Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al.: Language models are few-shot learners. Advances in neural information processing systems 33, 1877–1901 (2020) 3
- [9] Butler, D.J., Wulff, J., Stanley, G.B., Black, M.J.: A naturalistic open source movie for optical flow evaluation. In: European conference on computer vision. pp. 611–625. Springer (2012) 11
- [10] Carion, N., Gustafson, L., Hu, Y.T., Debnath, S., Hu, R., Suris, D., Ryali, C., Alwala, K.V., Khedr, H., Huang, A., et al.: Sam 3: Segment anything with concepts. arXiv preprint arXiv:2511.16719 (2025) 1, 2, 3, 11, 12
- [11] Carnegie Mellon University: CMU Graphics Lab Motion Capture Database. <http://mocap.cs.cmu.edu/> (2001), database created with funding from NSF EIA-0196217 8

- [12] Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 9650–9660 (2021) 4
- [13] Chen, S., Wong, S., Chen, L., Tian, Y.: Extending context window of large language models via positional interpolation, 2023. URL <https://arxiv.org/abs/2306.15595> 7
- [14] Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. In: Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers). pp. 4171–4186 (2019) 3
- [15] Ding, H., Liu, C., He, S., Jiang, X., Loy, C.C.: Mevis: A large-scale benchmark for video segmentation with motion expressions. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 2694–2703 (2023) 9, 11
- [16] Esser, P., Kulal, S., Blattmann, A., Entezari, R., Müller, J., Saini, H., Levi, Y., Lorenz, D., Sauer, A., Boesel, F., et al.: Scaling rectified flow transformers for high-resolution image synthesis. In: Forty-first international conference on machine learning (2024) 6
- [17] Fu, X., Yin, W., Hu, M., Wang, K., Ma, Y., Tan, P., Shen, S., Lin, D., Long, X.: Geowizard: Unleashing the diffusion priors for 3d geometry estimation from a single image. In: European Conference on Computer Vision (ECCV) (2024) 4
- [18] Gabeur, V., Long, S., Peng, S., Voigtlaender, P., Sun, S., Bao, Y., Truong, K., Wang, Z., Zhou, W., Barron, J.T., et al.: Image generators are generalist vision learners. arXiv preprint arXiv:2604.20329 (2026) 5, 9
- [19] Gaidon, A., Wang, Q., Cabon, Y., Vig, E.: Virtual worlds as proxy for multi-object tracking analysis. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4340–4349 (2016) 9
- [20] Garcia, G.M., Zeid, K.A., Schmidt, C., De Geus, D., Hermans, A., Leibe, B.: Fine-tuning image-conditional diffusion models is easier than you think. In: Proceedings of the Winter Conference on Applications of Computer Vision. pp. 753–762 (2025) 11
- [21] Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Vision meets robotics: The kitti dataset. The international journal of robotics research 32(11), 1231–1237 (2013) 11
- [22] He, J., Li, H., Sheng, M., Chen, Y.C.: Lotus-2: Advancing geometric dense prediction with powerful image generative model. arXiv preprint arXiv:2512.01030 (2025) 3, 11, 12
- [23] He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 16000–16009 (2022) 3
- [24] Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models (2020), <https://arxiv.org/abs/2006.11239> 6
- [25] Hu, W., Gao, X., Li, X., Zhao, S., Cun, X., Zhang, Y., Quan, L., Shan, Y.: Depthcrafter: Generating consistent long depth sequences for open-world videos (2024), <https://arxiv.org/abs/2409.02095> 4
- [26] Huang, J., Zhang, Y., He, X., Gao, Y., Cen, Z., Xia, B., Zhou, Y., Tao, X., Wan, P., Jia, J.: Unityvideo: Unified multi-modal multi-task learning for enhancing world-aware video generation. arXiv preprint arXiv:2512.07831 (2025) 3
- [27] Huang, P.H., Matzen, K., Kopf, J., Ahuja, N., Huang, J.B.: Deepmvs: Learning multi-view stereopsis. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2821–2830 (2018) 9

- [28] Jiang, Z., Zheng, C., Laina, I., Larlus, D., Vedaldi, A.: Geo4d: Leveraging video generators for geometric 4d scene reconstruction. arXiv preprint arXiv:2504.07961 (2025) 4, 7
- [29] Kaufmann, M., Song, J., Guo, C., Shen, K., Jiang, T., Tang, C., Zárate, J.J., Hilliges, O.: EMDb: The Electromagnetic Database of Global 3D Human Pose and Shape in the Wild. In: International Conference on Computer Vision (ICCV) (2023) 11
- [30] Ke, B., Obukhov, A., Huang, S., Metzger, N., Dautt, R.C., Schindler, K.: Repurposing diffusion-based image generators for monocular depth estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9492–9502 (June 2024) 4
- [31] Ke, Z., Sun, J., Li, K., Yan, Q., Lau, R.W.: Modnet: Real-time trimap-free portrait matting via objective decomposition. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 36, pp. 1140–1147 (2022) 11
- [32] Keetha, N., Müller, N., Schönberger, J., Porzi, L., Zhang, Y., Fischer, T., Knapitsch, A., Zauss, D., Weber, E., Antunes, N., et al.: Mapanything: Universal feed-forward metric 3d reconstruction. arXiv preprint arXiv:2509.13414 (2025) 11
- [33] Khrodar, R., Bagautdinov, T., Martinez, J., Zhaoen, S., James, A., Selednik, P., Anderson, S., Saito, S.: Sapiens: Foundation for human vision models. In: Leonardis, A., Ricci, E., Roth, S., Russakovsky, O., Sattler, T., Varol, G. (eds.) Computer Vision – ECCV 2024. pp. 206–228. Springer Nature Switzerland, Cham (2025) 2, 3, 11
- [34] Khoreva, A., Rohrbach, A., Schiele, B.: Video object segmentation with language referring expressions. In: Asian conference on computer vision. pp. 123–141. Springer (2018) 11
- [35] Kingma, D.P.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014) 9
- [36] Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 4015–4026 (2023) 1, 3
- [37] Kuang, Z., Zhang, T., Zhang, K., Tan, H., Bi, S., Hu, Y., Xu, Z., Hasan, M., Wetzstein, G., Luan, F.: Buffer anytime: Zero-shot video depth and normal from image priors (2024), <https://arxiv.org/abs/2411.17249> 4
- [38] Li, J., Cao, J., Zhang, H., Rempe, D., Kautz, J., Iqbal, U., Yuan, Y.: Genmo: Generative models for human motion synthesis. arXiv preprint arXiv:2505.01425 (2025) 2, 3, 11, 12
- [39] Li, T., He, K.: Back to basics: Let denoising generative models denoise. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 36115–36125 (2026) 6
- [40] Liang, R., Gojcic, Z., Ling, H., Munkberg, J., Hasselgren, J., Lin, Z.H., Gao, J., Keller, A., Vijaykumar, N., Fidler, S., Wang, Z.: Diffusionrenderer: Neural inverse and forward rendering with video diffusion models (2025), <https://arxiv.org/abs/2501.18590> 4
- [41] Lin, H., Chen, S., Liew, J., Chen, D.Y., Li, Z., Shi, G., Feng, J., Kang, B.: Depth anything 3: Recovering the visual space from any views. arXiv preprint arXiv:2511.10647 (2025) 1, 2, 3, 11, 12
- [42] Lin, S., Ryabtsev, A., Sengupta, S., Curless, B.L., Seitz, S.M., Kemelmacher-Shlizerman, I.: Real-time high-resolution background matting. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8762–8771 (2021) 11
- [43] Lin, S., Yang, L., Saleemi, I., Sengupta, S.: Robust high-resolution video matting with temporal guidance. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 238–247 (2022) 11

- [44] Lipman, Y., Chen, R.T., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. arXiv preprint arXiv:2210.02747 (2022) 6
- [45] Liu, X., Gong, C., Liu, Q.: Flow straight and fast: Learning to generate and transfer data with rectified flow. arXiv preprint arXiv:2209.03003 (2022) 6
- [46] Martinez, J., Kim, E., Romero, J., Bagautdinov, T., Saito, S., Yu, S.I., Anderson, S., Zollhöfer, M., Wang, T.L., Bai, S., et al.: Codec avatar studio: Paired human captures for complete, driveable, and generalizable avatars. *Advances in Neural Information Processing Systems* **37**, 83008–83023 (2024) 11
- [47] Mizrahi, D., Bachmann, R., Kar, O., Yeo, T., Gao, M., Dehghan, A., Zamir, A.: 4m: Massively multimodal masked modeling. *Advances in Neural Information Processing Systems* **36**, 58363–58408 (2023) 3
- [48] Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al.: Dinov2: Learning robust visual features without supervision. arXiv preprint arXiv:2304.07193 (2023) 4
- [49] Pathak, D., Krahenbuhl, P., Donahue, J., Darrell, T., Efros, A.A.: Context encoders: Feature learning by inpainting. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 2536–2544 (2016) 3
- [50] Perazzi, F., Pont-Tuset, J., McWilliams, B., Van Gool, L., Gross, M., Sorkine-Hornung, A.: A benchmark dataset and evaluation methodology for video object segmentation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 724–732 (2016) 12
- [51] Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: *International conference on machine learning*. pp. 8748–8763. PmLR (2021) 4
- [52] Ranftl, R., Lasinger, K., Hafner, D., Schindler, K., Koltun, V.: Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE transactions on pattern analysis and machine intelligence* **44**(3), 1623–1637 (2020) 11
- [53] Ravi, N., Gabeur, V., Hu, Y.T., Hu, R., Ryali, C., Ma, T., Khedr, H., Rädle, R., Rolland, C., Gustafson, L., et al.: Sam 2: Segment anything in images and videos. arXiv preprint arXiv:2408.00714 (2024) 1, 3
- [54] Renderpeople: Renderpeople: 3d people for renderings., <https://renderpeople.com> 8, 9
- [55] Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)*. pp. 10684–10695 (2022) 4
- [56] Saleh, F., Aliakbarian, S., Hewitt, C., Petikam, L., Xiao-Xian, Criminisi, A., Cashman, T.J., Baltrušaitis, T.: DAViD: Data-efficient and accurate vision models from synthetic data (2025), <https://arxiv.org/abs/2507.15365> 2, 3, 11
- [57] Schops, T., Schonberger, J.L., Galliani, S., Sattler, T., Schindler, K., Pollefeys, M., Geiger, A.: A multi-view stereo benchmark with high-resolution images and multi-camera videos. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 3260–3269 (2017) 11
- [58] Shen, Z., Pi, H., Xia, Y., Cen, Z., Peng, S., Hu, Z., Bao, H., Hu, R., Zhou, X.: World-grounded human motion recovery via gravity-view coordinates. In: *SIGGRAPH Asia Conference Proceedings* (2024) 11

- [59] Siméoni, O., Vo, H.V., Seitzer, M., Baldassarre, F., Oquab, M., Jose, C., Khalidov, V., Szafraniec, M., Yi, S., Ramamonjisoa, M., et al.: Dinov3. arXiv preprint arXiv:2508.10104 (2025) 4
- [60] Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. In: International Conference on Learning Representations (2021) 6
- [61] Tong, Z., Song, Y., Wang, J., Wang, L.: Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *Advances in neural information processing systems* **35**, 10078–10093 (2022) 2, 4, 13
- [62] Uijlings, J., Zhou, X., Gu, X., Nagrani, A., Arnab, A., Fathi, A., Ross, D., Schmid, C.: Vocap: Video object captioning and segmentation from any prompt. arXiv preprint arXiv:2508.21809 (2025) 11
- [63] Wan, T.: Wan: Open and advanced large-scale video generative models. arXiv preprint arXiv:2503.20314 (2025) 9
- [64] 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) 11
- [65] Wang, J., Chen, M., Zhang, S., Karaev, N., Schönberger, J., Labatut, P., Bojanowski, P., Novotny, D., Vedaldi, A., Rupprecht, C.: Vggt-omega. arXiv preprint arXiv:2605.15195 (2026) 2, 3, 11, 12
- [66] Wang, L., Zanfir, A., Bazavan, E.G., Andriluka, M., Sminchisescu, C.: Thfm: A unified video foundation model for 4d human perception and beyond. arXiv preprint arXiv:2603.25892 (2026) 4
- [67] Wang, W., Zhu, D., Wang, X., Hu, Y., Qiu, Y., Wang, C., Hu, Y., Kapoor, A., Scherer, S.: Tartanair: A dataset to push the limits of visual slam. In: *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. pp. 4909–4916. IEEE (2020) 9
- [68] Wang, Y., Wang, Z., Liu, L., Daniilidis, K.: Tram: Global trajectory and motion of 3d humans from in-the-wild videos. In: *European Conference on Computer Vision*. pp. 467–487. Springer (2024) 11, 12
- [69] Wiedemer, T., Li, Y., Vicol, P., Gu, S.S., Matarese, N., Swersky, K., Kim, B., Jaini, P., Geirhos, R.: Video models are zero-shot learners and reasoners. arXiv preprint arXiv:2509.20328 (2025) 4, 5
- [70] Wu, J., Jiang, Y., Sun, P., Yuan, Z., Luo, P.: Language as queries for referring video object segmentation. In: *CVPR* (2022) 11
- [71] Xu, G., Ge, Y., Liu, M., Fan, C., Xie, K., Zhao, Z., Chen, H., Shen, C.: What matters when repurposing diffusion models for general dense perception tasks? arXiv preprint arXiv:2403.06090 (2024) 4
- [72] Xu, N., Yang, L., Fan, Y., Yang, J., Yue, D., Liang, Y., Price, B., Cohen, S., Huang, T.: Youtube-vos: Sequence-to-sequence video object segmentation. In: *Proceedings of the European conference on computer vision (ECCV)*. pp. 585–601 (2018) 9
- [73] Yan, C., Wang, H., Yan, S., Jiang, X., Hu, Y., Kang, G., Xie, W., Gavves, E.: Visa: Reasoning video object segmentation via large language models. In: *European Conference on Computer Vision*. pp. 98–115. Springer (2024) 11
- [74] Yang, H., Huang, D., Yin, W., Shen, C., Liu, H., He, X., Lin, B., Ouyang, W., He, T.: Depth any video with scalable synthetic data (2025), <https://arxiv.org/abs/2410.10815> 4

- [75] Yang, L., Kang, B., Huang, Z., Xu, X., Feng, J., Zhao, H.: Depth anything: Unleashing the power of large-scale unlabeled data. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10371–10381 (2024) [3](#)
- [76] Yang, L., Kang, B., Huang, Z., Zhao, Z., Xu, X., Feng, J., Zhao, H.: Depth anything v2. Advances in Neural Information Processing Systems **37**, 21875–21911 (2024) [1](#), [3](#), [11](#)
- [77] Yang, L., Kang, B., Huang, Z., Zhao, Z., Xu, X., Feng, J., Zhao, H.: Depth anything v2. Advances in Neural Information Processing Systems **37**, 21875–21911 (2025) [1](#)
- [78] Yang, P., Zhou, S., Zhao, J., Tao, Q., Loy, C.C.: Matanyone: Stable video matting with consistent memory propagation. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 7299–7308 (2025) [11](#)
- [79] Yin, Y., Guo, C., Kaufmann, M., Zarate, J., Song, J., Hilliges, O.: Hi4d: 4d instance segmentation of close human interaction. In: Computer Vision and Pattern Recognition (CVPR) (2023) [11](#)
- [80] Yu, L., Poirson, P., Yang, S., Berg, A.C., Berg, T.L.: Modeling context in referring expressions. In: European conference on computer vision. pp. 69–85. Springer (2016) [9](#)
- [81] Zhai, X., Mustafa, B., Kolesnikov, A., Beyer, L.: Sigmoid loss for language image pre-training. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 11975–11986 (2023) [4](#)
- [82] Zhang, C., Moing, G.L., Koppula, S., Rocco, I., Momeni, L., Xie, J., Sun, S., Sukthankar, R., Barral, J.K., Hadsell, R., et al.: Efficiently reconstructing dynamic scenes one d4rt at a time. arXiv preprint arXiv:2512.08924 (2025) [2](#), [3](#), [11](#), [12](#)
- [83] Zhang, J., Lin, A., Kumar, M., Yang, T.H., Ramanan, D., Tulsiani, S.: Cameras as rays: Pose estimation via ray diffusion. In: International Conference on Learning Representations. vol. 2024, pp. 23345–23366 (2024) [7](#)
- [84] Zhao, C., Sun, Y., Liu, M., Zheng, H., Zhu, M., Zhao, Z., Chen, H., He, T., Shen, C.: Diception: A generalist diffusion model for visual perceptual tasks. arXiv preprint arXiv:2502.17157 (2025) [4](#), [11](#)
- [85] Zhao, Q., Lin, A., Tan, J., Zhang, J.Y., Ramanan, D., Tulsiani, S.: Diffusionsfm: Predicting structure and motion via ray origin and endpoint diffusion. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 6317–6326 (2025) [7](#)
- [86] Zhu, X., Zhu, J., Li, H., Wu, X., Li, H., Wang, X., Dai, J.: Uni-perceiver: Pre-training unified architecture for generic perception for zero-shot and few-shot tasks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 16804–16815 (2022) [3](#)
- [87] Zoran, D., Parthasarathy, N., Yang, Y., Hudson, D.A., Carreira, J., Zisserman, A.: Recurrent video masked autoencoders. arXiv preprint arXiv:2512.13684 (2025) [4](#)