This CVPR paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Unified Dense Prediction of Video Diffusion

Lehan Yang¹ Lu Qi² Xiangtai Li³ Sheng Li¹ Varun Jampani⁴ Ming-Hsuan Yang² ¹University of Virginia ²University of California, Merced

³Nanyang Technological University ⁴Stability AI

Abstract

We present a unified network for simultaneously generating videos and their corresponding entity segmentation and depth maps from text prompts. We utilize colormap to represent entity masks and depth maps, tightly integrating dense prediction with RGB video generation. Introducing dense prediction information improves video generation's consistency and motion smoothness without increasing computational costs. Incorporating learnable task embeddings brings multiple dense prediction tasks into a single model, enhancing flexibility and further boosting performance. We further propose a large-scale dense prediction video dataset Panda-Dense, addressing the issue that existing datasets do not concurrently contain captions, videos, segmentation, or depth maps. Comprehensive experiments demonstrate the high efficiency of our method, surpassing the state-of-the-art in terms of video quality, consistency, and motion smoothness.

1. Introduction

Video generation has witnessed substantial progress through diffusion models [3, 4, 15, 16, 46, 52, 64] and autoregressive models [22, 47, 48, 54]. The main goal is to create coherent video clips from text prompts, facilitating diverse applications like video editing [2, 11] and production [12, 57]. These video generation techniques have a wide range of use cases, ranging from creative content generation to immersive multimedia experiences.

While existing foundational models can generate coherent video clips, they often face consistency issues, such as visual misalignment between frames and discrepancies with text prompts, which hinder some downstream editing tasks. Several approaches [10, 58, 63] have focused on effective module structures, such as 3D VAEs and MM-DiT blocks, for various vision tasks. However, these models still struggle to learn robust representations due to the complex-



Figure 1. Visualization results of multi-task UDPDiff model on image generation and dense prediction. Our model can generate the video and the corresponding dense estimation. We incorporate two tasks in one multi-task model, including video entity segmentation and video depth estimation. Both segmentation and depth map have been encoded into RGB format as a video sequence, using Pixelplanes.

ity of the video generation task. The work of REPA [60] suggests that diffusion models can achieve faster and improved performance by aligning the representation with self-supervised methods. Nevertheless, these representations remain implicit without explicit semantic or geometry reasoning. An effective training signal should interpret generated results and integrate seamlessly with diffusion models. This insight motivates incorporating explicit signals into video diffusion training to enhance consistency, smoothness, and realism. Inspired by the unified representations in Emu Edit [44] and UniGS [40], we utilize two dense prediction signals: entity segmentation and depth estimation, both of which can be encoded as a unified representation, as shown in Fig. 1. Entity segmentation ensures realistic object shapes and adherence to motion laws, while depth estimation enhances the model's perception of depth, form, and positional differences, enabling more structured

Correspondence to Lu Qi and Xiangtai Li.

and context-aware generation.

The first issue we tackle in this work is the lack of joint video, segmentation, and depth datasets. For this, we curate a video subset from the public Panda-70M [6] by a series of data filtering such as video quality assessment and motion score. To address the scarcity of data for video entity segmentation, we have developed a new pipeline that initiates with the EntitySeg [30] to segment the initial frame, ensuring consistent granularity, and then employs the SAM2 [42] video segmentation framework to extend the segmentation through all the video frames. To obtain the depth data, we utilize the DepthCrafter [18], achieving consistent video depth estimation. Additionally, we have enhanced the existing captions using the Video-LLaVA [27], providing more detailed and comprehensive descriptions, and improving the alignment from text to video.

The second challenge we address is designing a unified representation and learning architecture for jointly predicting different signals of RGB video and their corresponding frame segmentation and depth maps. Compared to the use of colormaps [40] in image-level diffusion, to the use of colormaps for video-level diffusion presents additional complexities as it is necessary to adapt to multiple frames and consider the moving entities. Video-level representation is more challenging in two aspects, including the lack of large-scale training data wrapped with high-quality segmentation or depth annotations, as well as the lack of any unified architectural designs for multi-task training.

In this work, we introduce a Pixelplanes method to represent entity segmentation maps effectively. To better represent moving objects with varying locations and dense entities, we have randomized different colors to represent different entities. This approach also eliminates the potential ambiguity of predefined location-to-color mappings, allowing for less biased training.

We call our technique UDPDiff. Experimental results demonstrate that our method, benefiting from the assistance of segmentation and depth estimation, possesses better video synthesis quality and enhanced consistency, whether in the foreground or background. Additionally, our motion is more natural and smooth. The consistency and motion smoothness enhancements will provide stronger support and potential for downstream tasks such as video editing. The main contributions of this work are as follows:

- We introduce UDPDiff, the first adopted latent colormap representation in video generation and dense prediction, leveraging unified segmentation and depth estimation simultaneously with video generation.
- We introduce a new unified representation Pixelplanes designed for video segmentation, removes the ambiguity of location-to-color mapping, and supports more dense entities. It also unifies the encoding of the depth estimation map, featuring a multi-task unified

representation.

- We propose a new large-scale video dataset equipped with detailed captions, video segmentation, and depth estimation, powering the training of large-scale video diffusion dense prediction.
- Extensive experiments have demonstrated improvements in video quality, consistency, and motion smoothness in video generation from UDPDiff, with joint training. Our work can inspire the use of foundation models for dense prediction tasks in the video itself and enhance downstream tasks' capabilities.

2. Related Work

Diffusion Model for Generation. Advances in latent diffusion models [43] have contributed significantly to generation fields like text-to-image [33, 34, 43] and video generation [15, 16, 46, 52]. Building on these pre-trained base models, techniques such as ControlNet [61] and LoRA [17] enable more practical generation and editing applications [3, 4, 7, 23, 59, 64] by incorporating conditional signals, such as segmentation layouts [40, 53], depth/motion information [24, 26], or reference images/videos [7, 45]. Unlike those works, we focus on developing a unified model that simultaneously generates visual content and dense predictions, allowing both tasks to enhance each other.

Diffusion Model for Dense Prediction. Dense prediction tasks like segmentation and depth estimation are fundamental in computer vision. While existing methods [5, 14, 28, 29, 35, 36, 39, 62] mainly rely on discriminative models, diffusion-based approaches [9, 20, 25, 32, 40, 51, 55] have demonstrated state-of-the-art results. For instance, Dataset Diffusion [31] synthesis pixel-level semantic segmentation, and InstructDiffusion [13] performs multi-task segmentation tasks. SemFlow [51] and UniGS [40] achieve comparable performance in semantic and entity segmentation [37, 38]. Latent Diffusion Segmentation [50] creates a new mask VAE for panoptic segmentation, while Marigold [20] surpasses prior arts in monocular depth estimation. In contrast to prior work focusing only on the accuracy of a single dense prediction task, we enhance both generation and downstream editing by jointly training multiple tasks.

Unified Representation. The unified representation [8, 12, 40, 47–49, 54] of generation and dense prediction is an emerging topic. Existing work generally follows two main pipelines, including auto-regressive-based and diffusion-based methods. Auto-regressive models [8, 12, 47–49, 54], such as the Emu series [47, 48, 54], predict image, video, or text tokens sequentially. In contrast, diffusion-based methods [20, 40, 51] treat dense prediction as a colormap that can be directly diffused. Our work focuses on the diffusion process, leveraging state-of-the-art performance demonstrated



Figure 2. **Overview of the Panda-Dense pipeline and the UDPDiff framework. Left:** For segmentation, we use the first frame's results from EntitySeg as a prompt for SAM2, which then performs video segmentation across the entire sequence. For depth estimation, DepthCrafter is used to generate video depth maps. For long prompts, Video LLaVA is used for captioning. **Right:** Similar to CogVideoX, our method UDPDiff denoises the feature sequence in the latent space, encoding and decoding the latent using a 3D VAE. Video generation and dense prediction share a similar paradigm, using the same VAE for encoding and decoding through a unified representation. Task embeddings are applied to the time step embeddings, enabling more powerful differentiation of various tasks under a multi-task joint training model.

in video generation. Unlike existing diffusion-based approaches that primarily target image-level generation, our study is the first to introduce a unified diffusion model for video-level generation and dense prediction.

3. Method

Based on CogVideoX [58], we propose the UDPDiff framework, which aims to denoise images and dense predictions simultaneously based on text prompts. On the right side of Fig. 2, we showcase the pipeline of UDPDiff that allows the integration of multiple tasks within the same model. It guides the model in completing different generation tasks by inputting various task IDs through a learned task embedding.

3.1. Dense Prediction Dataset

There exist no readily available video datasets with entity segmentation and depth annotations. We propose a large-scale captioned dense prediction video dataset Panda-Dense, as shown in the left part of Fig. 2, including video entity segmentation and depth maps. Based on the Panda-70M dataset, we randomly sampled a subset of approximately 300K samples, using a 13B Video-LLaVA for recaptioning. We retain the original brief captions from Panda-70M and the new detailed captions.

Entity Segmentation. We leverage the image entity segmentation model of EntitySeg CropFormer [30] coupled with the powerful SAM2 [42] video segmentation to construct the video segmentation dataset. We need to use the two models as the SAM model [21] has inconsistent granularity issues, where the point grid initialization method can

lead to overly fine or coarse segmentation maps. In particular, we extract the first frame for each video and use the EntitySeg CropFormer [30] model for entity segmentation. We then use this segmentation mask as the prompt input for SAM2 to complete the propagation throughout the entire video frames. We also post-process these frames to remove small holes and the empty area. Our segmentation dataset exhibits diversity in the number of entities, which can provide a rich learning signal to generate entities in both simple and complex scenes.

Depth Estimation. Video depth estimation has traditional methods based on single images, such as Depth Anything V2 [56], which can accurately estimate the depth of each frame. However, linking these frames does not ensure consistency between frames, resulting in a series of flickers and noise, which is detrimental to training video diffusion. Another way is video-based diffusion methods, such as DepthCrafter [18], which can jointly process multiple video frames and generate video depth maps with high precision and consistency. We use DepthCrafter to perform inference on a subset of our dataset and map the single-channel depth map to RGB space using a spectral-styled distance-based colormap projection.

3.2. Encoding Pixelplanes

Review of Location Aware Colormap. UniGS [40] colormap uses a pre-defined color grid with fixed size and colors. A location-based projection Ψ from the entity centroid coordinates to the colors is used during encoding. This converts entity-level binary segmentation masks $M \in \{0, 1\}^{n \times h \times w}$ into an RGB colormap $M_c \in [0, 255]^{h \times w \times 3}$: $M_c = \Psi(M)$.

However, a manually designed color grid may not necessarily conform to the optimal distribution for model learning. Moreover, this color grid is fixed in size, which can lead to issues of different entities sharing the same color in densely populated multi-entity scenes. Furthermore, entities are likely in motion in video-dense prediction, making location-aware methods effective only in the first frame and creating ambiguities in subsequent frames.

Pixelplanes. To address these issues with the above colormap representation, we designed a new colormap that reduces ambiguity and is adaptable for dense multi-object segmentation. It is suitable for both video entity segmentation and video depth estimation. We introduce the Pixelplanes as the unified representation. During the segmentation encoding, we use a random color generation function: random : [0, 255], which produces a random integer in the range [0,255] representing a color channel value.

For all entity-level binary segmentation masks $M \in \{0,1\}^{n \times h \times w}$, we randomly sample each of the RGB channels separately to obtain: $M_c = (r_n, g_n, b_n)$. In this way, the color of each entity is independent and random, and we ensure that no duplicate colors appear on different entities, thereby enhancing the entity-discrimination capability of this representation. Freed from the constraints of location coordinates, entities' motion will not experience the color ambiguities caused by changes in location.

For depth maps, we use a spectral style value projection Υ to map depth $D \in [0,1]^{h \times w}$ onto RGB spectral depth map $D_c \in [0,255]^{h \times w \times 3}$: $D_c = \Upsilon(D)$. Compared to directly duplicating the depth channel into a three-channel RGB image, this color projection method adapts better to the real image distribution of VAE. It conforms more closely to the color representation of the diffusion prior. Single video data lacks explicit motion actions and hierarchical relationships, segmentation and depth estimation can abstract these motions, reducing implicitness, thereby expecting to achieve a representation that conforms more closely to the real-world physical laws of motion.

3.3. Joint Generation Pipeline

Overview of CogVideoX. Our model is built upon CogVideoX 5B, which utilizes a 3D VAE for encoding the video sequence to a latent. In contrast to the 2D VAE used in Stable Diffusion, the 3D VAE can compress and encode latent temporally, offering better coherence and compressibility. Given an RGB video $\mathcal{X} \in \mathcal{R}^{T \times H \times W \times 3}$, where T, H and W are the frame number, height, and width of a video. The video is encoded by the 3D VAE into latent z. Compared to the UNet structure of diffusion models, Sora [1] used a scalable transformer for long video generation. The transformer architecture scales more efficiently and has greater model capacity, making it more suitable for temporally involved video generation. Just as

with image diffusion, as a Markov chain with forward and backward processes, the diffusion model gradually removes noise from a noisy video to obtain the final generated video.

Furthermore, video generation typically requires more detailed text prompts than image generation. The CLIP text encoder used in stable diffusion might not suffice for encoding long and complex prompts. Therefore, this model employs the more powerful T5 [41] model as its text encoder.

UDPDiff Framework. We input corresponding and more abstract video entity segmentation and depth estimation to enhance the guidance of dense prediction tasks on video generation tasks. This improves the explicitness of the input, enabling the model to understand scenes and motion better, thereby generating more realistic motion and consistent videos. As shown in the right part of Fig. 2, we expanded the input layer and output layer, doubling the video input and output channels, and performed channel-wise concatenation of the video latent with one type of dense prediction latent as the input. The advantage of this approach is that it allows the guidance from dense prediction to affect all layers, thereby enhancing consistency.

Specifically, our input has four components: the latent code of the noised video z_t^v , the latent code of the unified dense colormap z_t^c , the text prompt c_t , and the task id d. Under the condition of the text prompt, we compute the task embedding from the given d, which serves as a joint condition. We use a transformer f_{θ} to denoise z_0^v and z_0^c to z_t^v and z_t^c . During inference, z_0^v and z_0^c are encoded from random Gaussian noise. And through the channel-wise concatenation operation, the two groups of latents for video and dense prediction, z_t^v and z_t^c , are concatenated into z, and at the t-th step represented as z_t , resulting in a total of 32-channel latent: $z_t = \text{CONCAT}(z_t^v, z_t^c)$.

3.4. Learned Task Embedding

To adapt multiple tasks using a single model, the intuitive method is to distinguish tasks using text prompts. However, the text prompt condition based on semantics is an implicit task condition, which may not explicitly and correctly coordinate the target task sampling. Thus, this causes task semantic ambiguity and potentially misleads the guidance for dense prediction. Therefore, similar to the multi-task differentiation method used in Emu edit [44], we have effectively adapted it from image-level multi-task learning to video-level multi-task learning. We employ a learnable task embedding e_{θ}^{d} , which takes the task id d as input and adds its output to the time step embedding $e_{\theta}^{t}(t)$ to obtain t_{d} :

$$t_d = e_\theta^d(d) + e_\theta^t(t) \tag{1}$$

During training, we jointly optimize this embedding layer while optimizing the generative diffusion model. The objective can be described as:

$$\mathcal{L}_{\text{train}} = \frac{1}{2} ||f_{\theta}(z_t, t_d, c_t) - \epsilon||^2.$$
(2)

During inference, in addition to the context prompt guiding video generation, we input a task ID into the task embedding layer, with options for segmentation or depth estimation. The network will generate outputs according to the specified task ID.

4. Experiments

In this section, we begin by evaluating the performance of our UDPDiff on a single auxiliary perception task for video generation, including segmentation and depth estimation. Then, we demonstrate the overall effects of multi-task training. We utilize FVD and three temporal metrics from VBench [19] to evaluate all video synthesis tasks to access the generation quality and consistency. For subject consistency, we calculate the similarity between frames using DINO features to measure whether the subject appearance remains consistent. For background consistency, we use CLIP features to calculate the similarity between frames, assessing whether the background scene remains consistent. For motion smoothness, we employ motion priors to evaluate the smoothness of the generated motions, checking whether the motion is fluid and conforms to the physical laws of the real world. In the ablation study, we examine the benefits of learned task embeddings and unified dense prediction encoding to show the effectiveness of our method.

4.1. Experimental Settings

For all experiments, we utilize our Panda-Dense dataset based on the Panda-70M [6], having about 300K samples in total. To enhance the model's prompt generalization capabilities, we employ 80% Video-LLaVA generated captions and 20% original captions from Panda-70M. We initialize our model using the CogVideoX 5B text-to-video model weights, resampling the original channel weights for the newly added channels. Our experiments are conducted on 8 NVIDIA A100-80G GPUs, with a batch size of 1 per GPU, fine-tuning for 40000 steps at a learning rate of 1e-5. We note that all experiments are fine-tuned by CogVideoX due to the limited computing sources. This training scheme ensures a fair comparison, we fine-tuned the original CogVideoX under equivalent conditions, allowing the experiments to yield consistent and reliable conclusions. During the evaluation, we randomly select 1024 samples from the previously unused Panda-70M dataset for evaluation using FVD and VBench.

4.2. Single-task Training

For video entity segmentation and depth estimation tasks, we individually train separate models for each, producing

Model	$SC(\uparrow)$	BC (†)	MS (†)
CogVideoX 5B	94.57	95.80	97.67
UDPDiff (segmentation)	<u>94.98</u>	<u>95.92</u>	98.62
UDPDiff (depth estimation)	95.91	96.24	<u>98.10</u>

Table 1. Evaluation of consistency and smoothness in single task models. SC stands for subject consistency, BC stands for background consistency, and MS stands for motion smoothness.

Model	SC (†)	BC (†)	MS (†)	$FVD(\downarrow)$
CogVideoX 5B	94.57	95.80	97.67	343.92
UDPDiff (seg)	<u>95.21</u>	95.69	98.24	<u>316.76</u>
UDPDiff (depth)	97.07	96.89	99.23	302.55

Table 2. Evaluation of consistency and smoothness in multitask models. SC stands for subject consistency, BC stands for background consistency, MS stands for motion smoothness, FVD stands for Frechet Video Distance.

outputs consisting of a video and one corresponding dense prediction output. As shown in Tab. 1, our method surpasses the original CogVideoX 5B in single-model performance for both the model with segmentation and the model with depth estimation, excelling in subject consistency, background consistency, and motion smoothness. Thus, it is clear that introducing guidance from segmentation or depth estimation alone can enhance consistency.

4.3. Multi-Task Training

We evaluate the multi-task model using the same architecture design, identical text prompts, and random seeds but with different task IDs as inputs. As shown in Table 2, when comparing FVD, our model's inference on depth estimation and segmentation significantly outperforms CogVideoX, demonstrating superior generation quality. Additionally, our model's inference on depth estimation outperforms CogVideoX regarding subject consistency, background consistency, and motion smoothness. Meanwhile, our model's inference on segmentation surpasses CogVideoX across most metrics, except background consistency, where the performance is comparable. In addition, the jointly trained multi-task model outperforms the singletask model, especially for depth estimation, This improvement can be attributed to the complementary signals from the two dense prediction tasks: segmentation enhances the realistic object shapes and motion adherence, while depth estimation strengthens the perception of depth, form, and positional differences.

Video Depth Estimation We also verify the effectiveness of the depth maps generated by our model. Since the output of our dense prediction is an RGB image, we conduct

Model	Data	$\delta_1 (\uparrow)$	RMSE (\downarrow)
Depth Anything V2	62.6 M	0.5808	0.0907
UDPDiff (multi-depth)	300 K	0.4176	0.1634

Table 3. Evaluation of depth estimation. δ_1 stands for the delta accuracy, and RMSE stands for the root mean square error.

a reverse process, converting the RGB image back into a depth map according to the original depth degree-to-color projection. Our model employs a text-to-video synthesis approach, but the generated dense predictions lack ground truth. Therefore, we input the videos produced by our model into DepthCrater to generate a pseudo-depth ground truth and evaluate Delta accuracy and RMSE on these 1024 samples. Table 3 compares the results using Depth Any-thing V2 with ViT-L as backbone trained on 595K precise synthetic images and 62M pseudo-labeled real images. Considering our model only trained on 300K videos, which is significantly smaller than Depth Anything V2, our model achieves an acceptable level of accuracy.

4.4. Qualitative Evaluation

We demonstrate the improvement of consistency in Fig. 3. In the videos generated by CogVideoX, inconsistencies are annotated with red bounding boxes to highlight the issues. For instance, the party hat disappears after several frames in the first example, and the watch's color changes in the second example. In contrast, the videos generated by our method exhibited greater consistency and smoother motion. At the same time, we demonstrate in Fig. 4 the higher quality of our video generation, where we can create more realistic motion trajectories for moving objects and motions that adhere more closely to the physical laws of the real world. Our method can generate clearer and sharper entities in dense scenes, such as those with multiple people or interactions among multiple individuals. This all benefits from incorporating segmentation and depth estimation tasks, which enhance the video generation model's capabilities in scene perception and entity interaction. In Fig. 5, we demonstrate the results of our method simultaneously generating video and corresponding dense predictions, including video entity segmentation and video depth estimation. Dense predictions are represented by the unified representation Pixelplanes, directly generated by the diffusion model and decoded using the same video 3D VAE. Our model can achieve high-quality segmentation and depth estimation without significantly increasing computational cost, bringing about greater efficiency.

4.5. Ablation Study

Model Inference Speed. Specifically, our model only increases the input and output channels and introduces an ad-

Model	Running time(s)
CogVideoX 5B	204.46
Ours (single task)	205.75
Ours (multi task)	205.81

Table 4. **Model inference speed comparison.** The inference speed of our single-task and multi-task models is compared with CogVideoX 5B, measured in seconds.

Method	SC (†)	BC (†)	MS (†)	$FVD(\downarrow)$
Text prompt	95.17	95.78	98.67	321.43
Task embedding	97.07	96.89	99.23	302.55

Table 5. **Ablation study on task embedding**, SC stands for subject consistency, BC stands for background consistency, MS stands for motion smoothness, FVD stands for Frechet Video Distance.

Method	SC (†)	BC (†)	MS (†)
Location-aware colormap	81.26	79.33	88.79
Pixelplanes	94.98	95.92	98.62

Table 6. Effectiveness of Pixelplanescompared to UniGS's location-aware colormap. SC stands for subject consistency, BC stands for background consistency, and MS stands for motion smoothness.

ditional task embedding without significantly increasing the number of parameters in the original network. Here, we compare the inference speed of the original CogVideoX 5B with our single-task model and the multi-task model with task embedding. On a single NVIDIA A100-80G, performing inference for 50 steps for a 49-frame video, all using the same text prompt, the original CogVideoX outputs videos, while our model outputs videos and dense predictions. As shown in Tab 4, our inference speeds are almost similar within the range of uncertainty. The added channels do not affect the inference speed, demonstrating the efficiency of our method.

Learned Task Embedding. We compare the video generation metrics of the multi-task model under the condition of inferring depth estimation, using both text prompts and task embeddings to distinguish between segmentation and depth estimation tasks. Tab. 5 shows the results under the task embedding method are improved in consistency, motion smoothness, and FVD compared to those using text prompts. This demonstrates the effectiveness of task embeddings in distinguishing different dense prediction tasks more effectively. It allows the model to understand the current task more explicitly and receive the correct guidance.

Unified Dense Prediction Encoding To demonstrate that our Pixelplanes is more suitable for video generation than the location-aware colormap proposed by UniGS, we follow



Figure 3. Consistency qualitative comparison between CogVideoX and UDPDiff. Six frames are evenly sampled from the generated video, with the horizontal axis representing time and the frame index gradually increasing. Inconsistent parts are annotated with red bounding boxes, including disappearances, color changes, and shape changes.



Figure 4. **Quality qualitative comparison between CogVideoX and UDPDiff**. Six frames are evenly sampled from the generated video, with the horizontal axis representing time and the frame index gradually increasing. Our advantages are reflected in clearer and sharper entities, more realistic motion, and better generation of dense scenes.

the colormap encoding method mentioned in UniGS, projecting the centroid of an entity in the initial frame onto an RGB color and populating it into subsequent frames. Tab. 6 shows the proposed unified model performs better than the single-task method using a random dense prediction encoding scheme, On the other hand, the colormap approach does not work as effectively, proving that our method eliminates the ambiguity of location changes and provides better guidance for multiple objects. **Multi Task Data Partition** Considering the varying difficulty and guidance strength levels for segmentation and depth estimation tasks, we employ different partitions for segmentation and depth estimation during the joint training of the multi-task model. We ablate partition ratios of 0.3, 0.5, and 0.7 for each dense data type. As shown in Tab 7, the best performance is observed when the data volumes for segmentation and depth estimation are equal. The guidance provided by segmentation and depth estimation can be con-



Figure 5. **Visualization of the video generation associated with dense prediction.** For each sample, we evenly sample four frames, with the left column of each sample representing the generated video and the right column representing the generated dense prediction. The left half of the samples is video entity segmentation, and the right half is video depth estimation. Our model can generate high-quality, dense predictions simultaneously with almost no increase in computational cost.

Partition	$SC(\uparrow)$	BC (\uparrow)	MS (†)
(0.3, 0.7)	95.89	96.54	98.16
(0.7, 0.3)	96.72	96.37	96.97
(0.5, 0.5)	97.07	96.89	99.23

Table 7. Multi task data modality partition for segmentation and depth estimation. In the partition column bracket, the first element indicates the data partition for segmentation, and the second indicates the data partition for depth estimation. SC stand for subject consistency, BC stand for background consistency, MS stand for motion smoothness.

Method	$SC(\uparrow)$	BC (†)	MS (†)
Depth Anything V2	94.87	96.29	96.13
DepthCrafter	95.91	96.24	<u>98.10</u>

Table 8. **Depth Estimation Method for Panda-Dense Pipeline.** Different depth data creation method, evaluate our model's performance on each dataset. SC stands for subject consistency, BC stands for background consistency, MS stands for motion smoothness.

sidered to be of a comparable degree.

Dataset Pipeline For depth estimation, we compare the usage of Depth Anything V2 and DepthCrafter to generate two different depth datasets and train our single depth model on these datasets. After training on both datasets, as shown in Tab. 8, we evaluate the video generation quality and find that DepthCrafter outperformed Depth Anything V2. This is because DepthCrafter natively supports video depth estimation, whereas Depth Anything V2 performs estimation frame by frame in independent frames. Such an image-level approach loses inter-frame consistency, potentially introducing additional noise into the training.

Method	$SC(\uparrow)$	BC (†)	MS (†)
8×8	93.12	93.73	96.18
12×12	92.96	93.59	96.35
EntitySeg	94.98	95.92	98.62

Table 9. SAM 2 Init Prompt for Panda-Dense Pipeline. Point grid initialization, compared to EntitySeg initialization, evaluate our model's performance on each dataset. SC stands for subject consistency, BC stands for background consistency, MS stands for motion smoothness.

We employ a point grid as a prompt to intuitively automate SAM-2 video segmentation. As shown in Tab. 9, we compare 8×8 and 12×12 point grids of SAM-2 with our EntitySeg initialization prompt. The observation is that inconsistencies in segmentation granularity exist and can affect the quality of video generation.

5. Conclusion

This paper introduces a novel and effective improvement in video generation by incorporating dense prediction tasks such as segmentation and depth estimation to guide higherquality and consistent video generation. The key to our method is a redesigned color projection method that conforms to video encoding patterns. This method simultaneously generates the segmentation mask and depth map alongside the RGB video. Additionally, a large-scale captioned video dense prediction dataset is constructed with the proposed method. Extensive experiments prove the effectiveness of each of our modules, and the joint training of the multi-task model enhances efficiency and further improves consistency and quality.

6. Acknowledgement

Supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior/ Interior Business Center (DOI/IBC) contract number 140D0423C0074. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DOI/IBC, or the U.S. Government.

References

- Tim Brooks, Bill Peebles, Connor Holmes, Will DePue, Yufei Guo, Li Jing, David Schnurr, Joe Taylor, Troy Luhman, Eric Luhman, Clarence Ng, Ricky Wang, and Aditya Ramesh. Video generation models as world simulators. 2024. 4
- [2] Duygu Ceylan, Chun-Hao P Huang, and Niloy J Mitra. Pix2video: Video editing using image diffusion. In CVPR, 2023. 1
- [3] Haoxin Chen, Menghan Xia, Yingqing He, Yong Zhang, Xiaodong Cun, Shaoshu Yang, Jinbo Xing, Yaofang Liu, Qifeng Chen, Xintao Wang, et al. Videocrafter1: Open diffusion models for high-quality video generation. In *arXiv*, 2023. 1, 2
- [4] Haoxin Chen, Yong Zhang, Xiaodong Cun, Menghan Xia, Xintao Wang, Chao Weng, and Ying Shan. Videocrafter2: Overcoming data limitations for high-quality video diffusion models. In *CVPR*, 2024. 1, 2
- [5] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. In *TPAMI*, 2017. 2
- [6] Tsai-Shien Chen, Aliaksandr Siarohin, Willi Menapace, Ekaterina Deyneka, Hsiang-wei Chao, Byung Eun Jeon, Yuwei Fang, Hsin-Ying Lee, Jian Ren, Ming-Hsuan Yang, et al. Panda-70m: Captioning 70m videos with multiple cross-modality teachers. In CVPR, 2024. 2, 5
- [7] Xi Chen, Lianghua Huang, Yu Liu, Yujun Shen, Deli Zhao, and Hengshuang Zhao. Anydoor: Zero-shot object-level image customization. In CVPR, 2024. 2
- [8] Runpei Dong, Chunrui Han, Yuang Peng, Zekun Qi, Zheng Ge, Jinrong Yang, Liang Zhao, Jianjian Sun, Hongyu Zhou, Haoran Wei, et al. Dreamllm: Synergistic multimodal comprehension and creation. In *arXiv*, 2023. 2
- [9] Yiqun Duan, Xianda Guo, and Zheng Zhu. Diffusiondepth: Diffusion denoising approach for monocular depth estimation. In *arXiv*, 2023. 2
- [10] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *ICML*, 2024. 1
- [11] Ruoyu Feng, Wenming Weng, Yanhui Wang, Yuhui Yuan, Jianmin Bao, Chong Luo, Zhibo Chen, and Baining Guo.

Ccedit: Creative and controllable video editing via diffusion models. In *CVPR*, 2024. 1

- [12] Yuying Ge, Sijie Zhao, Jinguo Zhu, Yixiao Ge, Kun Yi, Lin Song, Chen Li, Xiaohan Ding, and Ying Shan. Seed-x: Multimodal models with unified multi-granularity comprehension and generation. In *arXiv*, 2024. 1, 2
- [13] Zigang Geng, Binxin Yang, Tiankai Hang, Chen Li, Shuyang Gu, Ting Zhang, Jianmin Bao, Zheng Zhang, Houqiang Li, Han Hu, et al. Instructdiffusion: A generalist modeling interface for vision tasks. In CVPR, 2024. 2
- [14] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, 2017. 2
- [15] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition video generation with diffusion models. In *arXiv*, 2022. 1, 2
- [16] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. In *NeurIPS*, 2022. 1, 2
- [17] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. In *arXiv*, 2021. 2
- [18] Wenbo Hu, Xiangjun Gao, Xiaoyu Li, Sijie Zhao, Xiaodong Cun, Yong Zhang, Long Quan, and Ying Shan. Depthcrafter: Generating consistent long depth sequences for open-world videos. arXiv preprint arXiv:2409.02095, 2024. 2, 3
- [19] Ziqi Huang, Yinan He, Jiashuo Yu, Fan Zhang, Chenyang Si, Yuming Jiang, Yuanhan Zhang, Tianxing Wu, Qingyang Jin, Nattapol Chanpaisit, Yaohui Wang, Xinyuan Chen, Limin Wang, Dahua Lin, Yu Qiao, and Ziwei Liu. VBench: Comprehensive benchmark suite for video generative models. In *CVPR*, 2024. 5
- [20] Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. Repurposing diffusion-based image generators for monocular depth estimation. In CVPR, 2024. 2
- [21] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *ICCV*, 2023. 3
- [22] Dan Kondratyuk, Lijun Yu, Xiuye Gu, José Lezama, Jonathan Huang, Grant Schindler, Rachel Hornung, Vighnesh Birodkar, Jimmy Yan, Ming-Chang Chiu, et al. Videopoet: A large language model for zero-shot video generation. In *ICML*, 2024. 1
- [23] Nupur Kumari, Bingliang Zhang, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Multi-concept customization of text-to-image diffusion. In CVPR, 2023. 2
- [24] Ariel Lapid, Idan Achituve, Lior Bracha, and Ethan Fetaya. Gd-vdm: Generated depth for better diffusion-based video generation. In *arXiv preprint arXiv:2306.11173*, 2023. 2
- [25] Ziyi Li, Qinye Zhou, Xiaoyun Zhang, Ya Zhang, Yanfeng Wang, and Weidi Xie. Open-vocabulary object segmentation with diffusion models. In *CVPR*, pages 7667–7676, 2023. 2

- [26] Jingyun Liang, Yuchen Fan, Kai Zhang, Radu Timofte, Luc Van Gool, and Rakesh Ranjan. Movideo: Motion-aware video generation with diffusion model. In *ECCV*, 2025. 2
- [27] Bin Lin, Yang Ye, Bin Zhu, Jiaxi Cui, Munan Ning, Peng Jin, and Li Yuan. Video-llava: Learning united visual representation by alignment before projection. *arXiv preprint arXiv:2311.10122*, 2023. 2
- [28] Shu Liu, Lu Qi, Haifang Qin, Jianping Shi, and Jiaya Jia. Path aggregation network for instance segmentation. In *CVPR*, 2018. 2
- [29] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015. 2
- [30] Qi Lu, Jason Kuen, Shen Tiancheng, Gu Jiuxiang, Guo Weidong, Jia Jiaya, Lin Zhe, and Yang Ming-Hsuan. Highquality entity segmentation. In *ICCV*, 2023. 2, 3
- [31] Quang Nguyen, Truong Vu, Anh Tran, and Khoi Nguyen. Dataset diffusion: Diffusion-based synthetic data generation for pixel-level semantic segmentation. *NeurIPS*, 36, 2024. 2
- [32] Suraj Patni, Aradhye Agarwal, and Chetan Arora. Ecodepth: Effective conditioning of diffusion models for monocular depth estimation. In *CVPR*, 2024. 2
- [33] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *ICCV*, 2023. 2
- [34] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*, 2023. 2
- [35] Lu Qi, Li Jiang, Shu Liu, Xiaoyong Shen, and Jiaya Jia. Amodal instance segmentation with kins dataset. In CVPR, 2019. 2
- [36] Lu Qi, Yi Wang, Yukang Chen, Yingcong Chen, Xiangyu Zhang, Jian Sun, and Jiaya Jia. Pointins: Point-based instance segmentation. In *TPAMI*, 2021. 2
- [37] Lu Qi, Jason Kuen, Yi Wang, Jiuxiang Gu, Hengshuang Zhao, Zhe Lin, Philip Torr, and Jiaya Jia. Open-world entity segmentation. In *TPAMI*, 2022. 2
- [38] Lu Qi, Jason Kuen, Weidong Guo, Tiancheng Shen, Jiuxiang Gu, Wenbo Li, Jiaya Jia, Zhe Lin, and Ming-Hsuan Yang. Fine-grained entity segmentation. In *ICCV*, 2023. 2
- [39] Lu Qi, Yi-Wen Chen, Lehan Yang, Tiancheng Shen, Xiangtai Li, Weidong Guo, Yu Xu, and Ming-Hsuan Yang. Generalizable entity grounding via assistance of large language model. arXiv preprint arXiv:2402.02555, 2024. 2
- [40] Lu Qi, Lehan Yang, Weidong Guo, Yu Xu, Bo Du, Varun Jampani, and Ming-Hsuan Yang. Unigs: Unified representation for image generation and segmentation. In *CVPR*, 2024. 1, 2, 3
- [41] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *JMLR*, 21(140):1–67, 2020. 4
- [42] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, et al. Sam 2:

Segment anything in images and videos. *arXiv preprint* arXiv:2408.00714, 2024. 2, 3

- [43] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 2
- [44] Shelly Sheynin, Adam Polyak, Uriel Singer, Yuval Kirstain, Amit Zohar, Oron Ashual, Devi Parikh, and Yaniv Taigman. Emu edit: Precise image editing via recognition and generation tasks. In CVPR, 2024. 1, 4
- [45] Qingyu Shi, Lu Qi, Jianzong Wu, Jinbin Bai, Jingbo Wang, Yunhai Tong, Xiangtai Li, and Ming-Husang Yang. Relationbooth: Towards relation-aware customized object generation. In arXiv, 2024. 2
- [46] Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual, Oran Gafni, et al. Make-a-video: Text-to-video generation without text-video data. In *arXiv*, 2022. 1, 2
- [47] Quan Sun, Qiying Yu, Yufeng Cui, Fan Zhang, Xiaosong Zhang, Yueze Wang, Hongcheng Gao, Jingjing Liu, Tiejun Huang, and Xinlong Wang. Emu: Generative pretraining in multimodality. In *ICLR*, 2023. 1, 2
- [48] Quan Sun, Yufeng Cui, Xiaosong Zhang, Fan Zhang, Qiying Yu, Yueze Wang, Yongming Rao, Jingjing Liu, Tiejun Huang, and Xinlong Wang. Generative multimodal models are in-context learners. In CVPR, 2024. 1, 2
- [49] Changyao Tian, Xizhou Zhu, Yuwen Xiong, Weiyun Wang, Zhe Chen, Wenhai Wang, Yuntao Chen, Lewei Lu, Tong Lu, Jie Zhou, et al. Mm-interleaved: Interleaved image-text generative modeling via multi-modal feature synchronizer. In *arXiv*, 2024. 2
- [50] Wouter Van Gansbeke and Bert De Brabandere. A simple latent diffusion approach for panoptic segmentation and mask inpainting. *arXiv preprint arXiv:2401.10227*, 2024. 2
- [51] Chaoyang Wang, Xiangtai Li, Lu Qi, Henghui Ding, Yunhai Tong, and Ming-Hsuan Yang. Semflow: Binding semantic segmentation and image synthesis via rectified flow. In *NeurIPS*, 2024. 2
- [52] Jiuniu Wang, Hangjie Yuan, Dayou Chen, Yingya Zhang, Xiang Wang, and Shiwei Zhang. Modelscope text-to-video technical report. In *arXiv*, 2023. 1, 2
- [53] X Wang, Siming Fu, Qihan Huang, Wanggui He, and Hao Jiang. Ms-diffusion: Multi-subject zero-shot image personalization with layout guidance. In *arXiv*, 2024. 2
- [54] Xinlong Wang, Xiaosong Zhang, Zhengxiong Luo, Quan Sun, Yufeng Cui, Jinsheng Wang, Fan Zhang, Yueze Wang, Zhen Li, Qiying Yu, et al. Emu3: Next-token prediction is all you need. In *arXiv*, 2024. 1, 2
- [55] Kangyang Xie, Binbin Yang, Hao Chen, Meng Wang, Cheng Zou, Hui Xue, Ming Yang, and Chunhua Shen. Zippo: Zipping color and transparency distributions into a single diffusion model. In *arXiv*, 2024. 2
- [56] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything v2. arXiv preprint arXiv:2406.09414, 2024. 3
- [57] Shuai Yang, Yuying Ge, Yang Li, Yukang Chen, Yixiao Ge, Ying Shan, and Yingcong Chen. Seed-story: Multimodal long story generation with large language model. arXiv preprint arXiv:2407.08683, 2024. 1

- [58] Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, et al. Cogvideox: Text-to-video diffusion models with an expert transformer. In *arXiv*, 2024. 1, 3
- [59] Xuanyu Yi, Zike Wu, Qingshan Xu, Pan Zhou, Joo-Hwee Lim, and Hanwang Zhang. Diffusion time-step curriculum for one image to 3d generation. In *CVPR*, 2024. 2
- [60] Sihyun Yu, Sangkyung Kwak, Huiwon Jang, Jongheon Jeong, Jonathan Huang, Jinwoo Shin, and Saining Xie. Representation alignment for generation: Training diffusion transformers is easier than you think. In *arXiv*, 2024. 1
- [61] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *ICCV*, 2023. 2
- [62] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *CVPR*, 2017. 2
- [63] Sijie Zhao, Yong Zhang, Xiaodong Cun, Shaoshu Yang, Muyao Niu, Xiaoyu Li, Wenbo Hu, and Ying Shan. Cv-vae: A compatible video vae for latent generative video models. arXiv preprint arXiv:2405.20279, 2024. 1
- [64] Daquan Zhou, Weimin Wang, Hanshu Yan, Weiwei Lv, Yizhe Zhu, and Jiashi Feng. Magicvideo: Efficient video generation with latent diffusion models. In arXiv, 2022. 1, 2