

# Unified Video Dense Prediction from Disjoint Data

Yihong Sun<sup>1,2,\*</sup>, Seoung Wug Oh<sup>1</sup>, Jiahui Huang<sup>1</sup>,  
Bharath Hariharan<sup>2</sup>, and Joon-Young Lee<sup>1</sup>

<sup>1</sup>Adobe Research      <sup>2</sup>Cornell University

**Abstract.** Scene understanding requires simultaneous prediction about geometry, appearance, and semantics. However, existing task-specific annotations are fragmented across incompatible, domain-specific datasets. Current unified systems circumvent this by restricting training to fully co-annotated data, or by incurring the large computational cost of pseudo-labeling. To mitigate this, we introduce **UniD**, a unified video model that jointly predicts eight dense scene properties—depth, surface normals, semantic segmentation, boundaries, human parts, albedo, shading, and materials—all learned from disjoint, domain-specific datasets. We propose a simple yet effective distillation step in which per-task experts supervise a unified backbone through lightweight task projectors, eliminating the need for annotation overlap or pseudo-labeling. Our key insight is that the strong visual priors of a pretrained diffusion model are sufficient to bridge the domain gaps introduced by disjoint training sources, enabling robust generalization to scene-task combinations never seen during training. UniD achieves competitive performance against per-task specialists and multi-task baselines, with strong generalization to out-of-distribution scenarios and enhanced temporal and cross-task consistency. Code and video results are available at <https://unid-video.github.io>.


## 1 Introduction

To understand a scene, a vision system must answer many questions at once: how far away are the visible objects? What materials are they made of? And how does light interact with their surfaces? These questions correspond to distinct dense prediction tasks—depth, material segmentation, intrinsic decomposition—each of which has spawned its own specialized model, learning framework, and datasets. However, the real world is a combination of all of these properties and ever-changing: an embodied agent navigating or manipulating a scene would need all of these cues simultaneously and continuously.

This has motivated recent development of *unified* dense prediction models [13, 43, 77, 79–81, 87] that output multi-task predictions jointly. However, almost all existing methods assume that all tasks are co-annotated on the same images—an assumption that is unrealistic with how task-specific data is collected in practice. Geometric annotations such as depth and normals can be acquired at scale via stereo or RGB-D cameras, but are typically captured in

---

\* Work done during an internship at Adobe Research.



**Fig. 1:** Unified Video Dense Predictions by UniD.

environments with limited semantic diversity, *e.g.* architectural landscapes. Semantic annotations require manual labeling over a diverse set of scenes, yielding rich category diversity but rarely paired with geometric ground truth. Finally, intrinsic decompositions [5, 51] are prohibitively expensive to measure in the real world and largely confined to synthetic domains. Driven by these fundamental differences, forcing these fragmented datasets into a single co-annotated training set is either infeasible or requires large-scale pseudo-labeling that is unscalable: each new task requires re-annotating the entire dataset.

On top of this, extending into the video domain would require a substantial amount of annotated videos. While feasible for geometry (*e.g.* RGB-D video captures), this becomes prohibitively expensive for semantics and intrinsics tasks. As a result, existing video-capable methods are either confined to the self-driving domain [30] or focus mainly on low-level geometric estimations [3].

These constraints raise a natural question: *can we learn unified dense prediction for both images and videos from disjoint, domain-specific datasets?* We present **UniD**, a **unified** video model learned from **disjoint**, domain-specific data that jointly estimates eight dense scene properties—depth, surface normals, semantic segmentation, boundaries, human parts, albedo, shading, and materials (Figure 1). In order to bridge the domain gaps arising from disjoint training sources, we first learn a task-specific latent space with strong generalization to unseen examples. We achieve this generalization by leveraging the visual priors of a diffusion model that is pretrained on internet-scale data [57]. Then, we train a unified model to reconstruct all task-specific latent embeddings for any input, via a lightweight per-task latent projector that maps a shared backbone representation into each task’s latent space. By computing the task-specific embeddings on-the-fly, the unified training requires no co-annotated data or pseudo-labels.

UniD offers several key advantages. First, it achieves enhanced generalization and temporal stability. By jointly leveraging disparate sources across diverse distributions—such as synthetic and real-world, indoor and outdoor, and static images alongside dynamic videos—the model avoids overfitting and demonstrates significant robustness to out-of-distribution (OOD) scenarios. Crucially, tasks without video supervision inherit temporal coherence from other video-centric tasks, ensuring stable and consistent predictions over time.

Second, the unified representation provided by a shared backbone ensures both cross-task consistency and high inference efficiency. By predicting a single

unified feature manifold, our model enforces strong structural alignment across tasks. This leads to spatially congruent predictions across different tasks—such as consistent boundary delineations—which are often misaligned in independent task-specific models. Furthermore, this shared architecture significantly reduces computational costs, as the backbone features are computed only once to support multiple task-specific projectors, incurring minimal per-task overhead.

To demonstrate these advantages, we conducted extensive experiments across all tasks, where UniD achieves competitive performance compared to the existing task-specific models. Furthermore, UniD demonstrates superior OOD generalization, temporal stability, and cross-task consistency over existing unified models and a multi-task baseline utilizing a frozen DINOv3-H backbone, proving its effectiveness in complex, real-world video dense prediction.

## 2 Related Work

*Multi-task dense prediction.* Unifying dense prediction tasks within a single model improves efficiency and encourages cross-task complementarity [85]. Naturally, many existing works [42, 43, 46, 70, 71, 76, 79–81, 83] acquire multi-task capabilities by requiring every training image to be annotated for every task. To relax this constraint, recent works [40, 82] also investigated learning with partial per-image annotations to improve label efficiency. However, it still assumes that each image contains a subset of labeled tasks, drawn from the same co-annotation distribution as before. In addition, few-shot generalists [35, 36, 68] reduce per-task annotation requirements through meta-learning and episodic training. However, they still rely on shared meta-datasets across tasks and do not take full advantage of large-scale task-specific datasets.

In parallel, several works adopt pseudo-labeling to enforce large-scale annotation co-occurrence [1, 2, 30, 73, 87]. While effective, pseudo-labeling is costly, and the entire labeling pipeline must be re-run when a new task or dataset is added. Notably, 4M-21 [1] further requires discrete tokenization of all task outputs, and its masked modeling objective couples all tasks into a uniform token prediction problem, precluding task-specific optimization objectives. In contrast, UniD trains on disjoint, task-specific datasets while allowing flexible task-specific learning objectives, with no annotation co-occurrence required.

*Diffusion models for dense prediction.* Diffusion models [21, 27, 57, 63] have shown impressive image generation capabilities. Their rich visual priors learned from large-scale pretraining have enabled many applications, including conditional generation [28, 50, 84, 86], editing [10, 33, 53], and video generation [8, 26, 29, 38].

Following this, recent methods [25, 34, 64, 69, 78, 90] propose to leverage the generative priors to learn dense predictions (*e.g.* depth and surface normals) that generalize beyond the limited training data. In addition to these task-specialists, direct multi-task training with diffusion models have been proposed with either fully-annotated data [77], pseudo-labeling [87], or synthetic-only supervision [13]. Notably, DICEPTION [87] unifies multi-tasks by mapping all into RGB space

with conditioned task prompts (*e.g.* “image2depth”). However, the RGB unification cannot naively represent tasks with high-dimensional outputs (*e.g.* semantic segmentation), and text-conditioned generation requires  $K$  separate forward passes to obtain all  $K$  predictions, precluding simultaneous multi-task inference.

UniD instead leverages diffusion models to learn a task-specific latent space, enabling per-pixel, task-specific optimization objectives, while leveraging the generative priors of the diffusion backbone to bridge domain gaps from disjoint training sources. Architecturally, rather than iterative denoising, UniD adopts a single-step diffusion with a shared backbone and lightweight per-task latent projectors, producing all task predictions simultaneously in a single forward pass.

*Video dense prediction.* Dense prediction has been studied extensively in the video domain, including video depth estimation [15, 59, 72], surface normal estimation [7], and semantic segmentation [31, 49, 60]. Naturally, extending dense prediction to the video domain introduces additional annotation scarcity challenges. This limitation results in existing unified video prediction methods being restricted to the self-driving domain [30], primarily oriented toward video generation [73], or confined to low-level geometric tasks [3]. In contrast, UniD is not restricted to any domain, handles high-level semantic and intrinsic tasks beyond low-level geometry, and achieves temporal consistency through a lightweight parameter-free module without requiring task-specific video annotations.

### 3 Method

**Task Formulation.** We consider the problem of dense scene understanding for videos, where a desired system must simultaneously predict  $K$  tasks from an input video stream. Formally, given an input video  $v = \{\mathbf{x}_1, \dots, \mathbf{x}_T\}$  with  $T$  input frames  $\mathbf{x}_i \in \mathbb{R}^{H \times W \times 3}$ , the goal is to produce dense prediction maps  $\{\mathbf{y}_1^k, \dots, \mathbf{y}_T^k\}_{k=1}^K$ , where  $\mathbf{y}_i^k \in \mathbb{R}^{H \times W \times C^k}$  is  $i^{\text{th}}$ -frame prediction for task  $k$ . Depending on the task, the output dimension  $C^k$  can range from 1 for scalar depth maps to the number of classes for semantic segmentation. Furthermore, each task  $k$  is associated with annotated data  $\mathbf{D}^k$  that may be completely disjoint.

We first learn task-specific embedding spaces that generalize beyond each task’s limited training domain and modality (Section 3.1), and then distill all task embeddings into a single unified video model (Section 3.2).

#### 3.1 Learning Task-Specific Embeddings

Unifying a diverse set of tasks often requires learning from disjoint data, as different tasks often require different means of label collection (*e.g.* indoor RGB-D scans for geometric tasks versus manual annotations of internet images for semantic tasks). To bridge these domain gaps, we leverage the rich generative priors of diffusion models pretrained on internet-scale data [57].



**Dense prediction with diffusion models.** We build on a pretrained latent diffusion model (LDM) [57] with a VAE encoder  $\mathcal{E}$ , decoder  $\mathcal{D}$ , and a U-Net backbone  $\mathcal{F}_\theta$ . Given an input image  $\mathbf{x}_i \in \mathbb{R}^{H \times W \times 3}$ , the encoder maps it to a compact latent code  $\mathbf{z}_i = \mathcal{E}(\mathbf{x}_i) \in \mathbb{R}^{h \times w \times d}$ , while the decoder attempts to reconstruct the image  $\mathbf{x}_i$  from  $\mathbf{z}_i$  as  $\mathcal{D}(\mathbf{z}_i) \in \mathbb{R}^{H \times W \times 3}$ .

A key design choice is how to interface the network with the task output space. We adopt a *pixel decoding* strategy: the U-Net predicts a task latent  $\mathbf{l}_i = \mathcal{F}_\theta(\mathbf{z}_i) \in \mathbb{R}^{h \times w \times d}$ , which is decoded by the pretrained VAE decoder  $\mathcal{D}$  up to its final convolutional layer, which we denote as  $\mathcal{D}'(\mathbf{l}_i) \in \mathbb{R}^{H \times W \times d'}$ . We then project  $\mathcal{D}'(\mathbf{l}_i)$  to the target dimension  $C^k$  of task  $k$  via a lightweight task-specific pixel projector  $\mathcal{P}$ .

For each task  $k$ , the final prediction  $\hat{\mathbf{y}}$  is computed as follows:

$$\mathbf{l}_i^k = \mathcal{F}_\theta^k(\mathcal{E}(\mathbf{x}_i)) \in \mathbb{R}^{h \times w \times d}, \quad \hat{\mathbf{y}}_i^k = \mathcal{P}^k(\mathcal{D}'(\mathbf{l}_i^k)) \in \mathbb{R}^{H \times W \times C^k} \quad (1)$$

where the task latent  $\mathbf{l}_i^k$  is decoded by  $\mathcal{P}(\mathcal{D}'(\cdot))$  into  $\hat{\mathbf{y}}_i^k \in \mathbb{R}^{H \times W \times C^k}$ .

The pixel decoding strategy (i) supports any  $C^k$ -dimensional outputs of arbitrary task, (ii) enables task-specific optimizations on the pixel predictions, and (iii) implicitly learns the projection between the latent and prediction space.

**Video predictions with temporal connections.** Extending unified dense prediction to the video domain remains challenging due to annotation asymmetry across tasks (RGB-D videos for geometric tasks is more accessible than manual annotation for semantic or intrinsic tasks). Therefore, a temporal mechanism should remain compatible for tasks with no temporal annotations.

We adopt extended self-attention [41], which extends every self-attention layer in the U-Net backbone  $\mathcal{F}_\theta$  with features from past frames. At frame  $t$ , the queries attend jointly to the keys and values from both the current frame and a set of  $M$  historical frames  $\mathcal{M}_t \subseteq \{1, \dots, t-1\}$  stored in a memory bank:


$$\text{Attn}(\mathbf{Q}_t, [\mathbf{K}_s]_{s \in \mathcal{M}_t \cup \{t\}}, [\mathbf{V}_s]_{s \in \mathcal{M}_t \cup \{t\}}), \quad (2)$$

where  $[\cdot]$  denotes concatenation. Intuitively, even without training on videos, adjacent frames would share rich correspondences [64] that can encourage temporal smoothing as queries from frame  $t$  attend to historical keys and values.

Setting  $M = 0$  naturally reverts to single-frame predictions, making it compatible when training without temporal annotations. The memory queue length  $M$  and the enqueue/dequeue frequency can be chosen at inference, allowing flexibility to the degree of temporal smoothing and memory usage without retraining.

Now, an input video  $v = \{\mathbf{x}_1, \dots, \mathbf{x}_T\}$  is projected to a task-specific embedding  $\mathbf{l} = [\mathcal{F}_\theta(\mathbf{z}_1), \dots, \mathcal{F}_\theta(\mathbf{z}_T)] \in \mathbb{R}^{T \times h \times w \times d}$  via an  $\mathcal{F}_\theta$  processing each  $\mathbf{z}_i = \mathcal{E}(\mathbf{x}_i)$  with temporal connections, yielding a temporally-aware representation.

**Training task specialists.** As shown in Figure 2(a), we train a dedicated specialist  $(\mathcal{F}_\theta^k, \mathcal{P}^k)$  for each task  $k$  to learn a task-specific embedding that generalizes



**Fig. 2:** Overview of UniD. **(a) Task specialist training.** For each task  $k$ , a specialist backbone  $\mathcal{F}_\theta^k$  and pixel projector  $\mathcal{P}^k$  are finetuned on their task-specific dataset  $\mathbf{D}^k$  with loss  $\mathcal{L}^k$ . Note that both the encoder  $\mathcal{E}$  and decoder  $\mathcal{D}'$  (in gray) are frozen during training. **(b) Unified model training.** The unified backbone  $\mathcal{G}_\phi$  is then trained along with  $K$  latent projectors  $\{\Psi^k\}_{k=1}^K$  to reconstruct the latents  $\{\mathbf{l}_i^k\}$  predicted by the frozen specialist  $\{\mathcal{F}_\theta^k\}_{k=1}^K$  (in gray). **(c) Streaming video inference.** During inference,  $\mathcal{G}_\phi$  processes an input video via Extended Self-Attention (ESA), then the unified representation  $\mathbf{u}_i$  is decoded through each  $\Psi^k$ ,  $\mathcal{D}'$ , and  $\mathcal{P}^k$  to simultaneously produce temporally consistent predictions across all  $K$  tasks.

across disjoint, domain-specific data. In line with the pixel-decoding framework, we also adopt a *single-step* diffusion: rather than iterative denoising, the U-Net performs one forward pass to directly predict the latent  $\mathbf{l}_i^k$ . The single-step diffusion not only reduces latency during inference, but also improves training efficiency by simplifying the original denoising objective [25].

Formally, for each task  $k$ , we obtain the task-specialist  $(\mathcal{F}_\theta^k, \mathcal{P}^k)$  by minimizing the task-specific objective  $\mathcal{L}^k$  over the task-specific dataset  $\mathbf{D}^k$ :

$$(\mathcal{F}_\theta^k, \mathcal{P}^k) = \arg \min_{\mathcal{F}_\theta^k, \mathcal{P}^k} \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathbf{D}^k} \left[ \mathcal{L}^k \left( \mathcal{P}^k \left( \mathcal{D}' \left( \mathcal{F}_\theta^k (\mathcal{E}(\mathbf{x})) \right) \right), \mathbf{y} \right) \right] \quad (3)$$

Due to its large-scale pretraining and generative priors, the learned  $\mathcal{F}_\theta^k$  can bridge disjoint datasets when embedding unseen data during unified training.

### 3.2 Unified Video Dense Prediction

**Learning unified model with latent projectors.** Given  $K$  task specialists  $\{(\mathcal{F}_\theta^k, \mathcal{P}^k)\}_{k=1}^K$  in hand, one option would be to pseudo-label a large shared dataset with each specialist and train on the computed labels. While feasible, computing/storing  $K$  full-resolution label maps  $\{\mathbb{R}^{H \times W \times C^k}\}_{k=1}^K$  per frame is time-consuming and memory-intensive, with both costs growing linearly with  $K$ —all before unified training even begins.

Instead, we distill task embeddings directly in the latent space. This allows us to avoid pseudo-labeling entirely, skip gradient propagation through the decoder  $\mathcal{D}$ , and unify the learning objective for each task as latent reconstruction.

For an unlabeled video  $\mathbf{x} = \{\mathbf{x}_1, \dots, \mathbf{x}_T\}$ , we first compute the target task latents  $\mathbf{l}_i^k$  with temporal connections on-the-fly via each frozen specialist backbone  $\mathcal{F}_\theta^k$ , where  $\mathbf{l}_i^k = \mathcal{F}_\theta^k(\mathbf{z}_i) \in \mathbb{R}^{h \times w \times d}$ ,  $\mathbf{z}_i = \mathcal{E}(\mathbf{x}_i)$ ,  $i = 1, \dots, T$ . Then, we formulate the unified model as a U-Net backbone  $\mathcal{G}_\phi$ , with same initialization and temporal memory as  $\mathcal{F}_\theta^k$ , and employ  $K$  lightweight task-specific latent projectors  $\{\Psi^k\}_{k=1}^K$  to project the unified representation to the task-specific embeddings. Given an input latent  $\mathbf{z}_i$ , the unified model predicts all task latents simultaneously:

$$\hat{\mathbf{l}}_i^k = \Psi^k(\mathcal{G}_\phi(\mathbf{z}_i)) \in \mathbb{R}^{h \times w \times d}, \quad k = 1, \dots, K. \quad (4)$$

As shown in Figure 2(b), the unified model  $(\mathcal{G}_\phi, \{\Psi^k\}_{k=1}^K)$  is learned by minimizing the latent reconstruction objective  $\mathcal{L}^{\text{rec}}$  as follows:

$$(\mathcal{G}_\phi, \{\Psi^k\}) = \arg \min_{\mathcal{G}_\phi, \{\Psi^k\}} \mathbb{E}_{v \sim \mathbf{D}} \sum_{k=1}^K \mathcal{L}^{\text{rec}}(\hat{\mathbf{l}}^k, \mathbf{l}^k) \quad (5)$$

where  $\hat{\mathbf{l}}^k$  and  $\mathbf{l}^k$  are both computed by running  $\Psi^k(\mathcal{G}_\phi(\cdot))$  and  $\mathcal{F}_\theta^k(\cdot)$  with temporal connections on the latent code  $\mathbf{z}_i$  from  $i = 1$  to  $i = T$ .

**Temporal latent stabilization.** If  $\mathcal{L}^{\text{rec}}$  only contains per-frame reconstruction losses, the unified backbone  $\mathcal{G}_\phi$  does not receive sufficient temporal regularization during unified training. To address this, we extend temporal gradient matching [15], originally proposed for video depth estimation, to a general latent formulation. Given predicted latents  $\hat{\mathbf{l}}^k \in \mathbb{R}^{t \times h \times w \times d}$  and target latents  $\mathbf{l}^k \in \mathbb{R}^{t \times h \times w \times d}$ , we enforce that the temporal gradient of the prediction matches that of the target at every spatial location:

$$\mathcal{L}_{\text{TGM}}(\hat{\mathbf{l}}, \mathbf{l}) = \left\| \left[ (\hat{\mathbf{l}}_{i+1} - \hat{\mathbf{l}}_i) - (\mathbf{l}_{i+1} - \mathbf{l}_i) \right] \cdot \mathcal{M}[\|\mathbf{l}_{i+1} - \mathbf{l}_i\|_2 < \varepsilon] \right\|_2 \quad (6)$$

where the mask  $\mathcal{M}[\cdot]$  suppresses learning at fast-changing regions where temporal regularization is not reliable. The final unified training objective combines per-frame  $\ell_1$  latent reconstruction with the temporal gradient matching loss:  $\mathcal{L}^{\text{rec}}(\hat{\mathbf{l}}^k, \mathbf{l}^k) = \left\| \hat{\mathbf{l}}_i^k - \mathbf{l}_i^k \right\|_1 + \lambda \mathcal{L}_{\text{TGM}}(\hat{\mathbf{l}}^k, \mathbf{l}^k)$ , where  $\lambda$  balances the two terms.

**Inference and extensibility.** At inference, the full-resolution predictions for all tasks with input  $\mathbf{x} \in \mathbb{R}^{T \times H \times W \times 3}$  are computed as follows:

$$\mathbf{u}_i = \mathcal{G}_\phi(\mathcal{E}(\mathbf{x}_i)), \quad \hat{\mathbf{y}}^k = \mathcal{P}^k(\mathcal{D}'(\Psi^k(\mathbf{u}_i))) \in \mathbb{R}^{H \times W \times C^k}, \quad i = 1, \dots, T \quad (7)$$

This framework is naturally extensible: adding a new task  $K+1$  requires only training a new specialist  $(\mathcal{F}_\theta^{K+1}, \mathcal{P}^{K+1})$  on its own labeled data and introducing a new latent projector  $\Psi^{K+1}$ , with no architectural change to the unified backbone  $\mathcal{G}_\phi$  and no existing data to be re-labeled.

### 3.3 Implementation Details

*Task specialists.* We use Stable Diffusion v2 [57] as our pretrained LDM backbone. Each of  $K = 8$  specialists finetunes the U-Net backbone along with the task-specific pixel projector  $\mathcal{P}^k$  (a single  $3 \times 3$  convolution) to convergence with a learning rate of  $3 \times 10^{-5}$  using the AdamW optimizer [45] with decay  $10^{-2}$ .

*Unified model.* The unified backbone  $\mathcal{G}_\phi$  is initialized from the same pretrained U-Net. Each latent projector  $\Psi^k$  is a DPT head that fuses 3 intermediate representations of  $\mathcal{G}_\phi$  with the final output feature at a fusion dimension of 256. The unified backbone  $\mathcal{G}_\phi$  and latent projectors  $\Psi^k$  are trained with a union of all labeled data for all tasks, sampled with equal weight between image and video data. The full unified model is trained for 50k iterations with batch size of 32 and learning rate  $3 \times 10^{-5}$ . Please refer to the Appendix for details regarding each task’s training datasets, learning objective, and other implementation details.

## 4 Experiments

**Datasets and Metrics.** The training data, evaluation data, task-specific output dimension  $C^k$  and metrics are detailed in Table 1.

In terms of evaluation, we select at least one out-of-domain dataset (unseen during training) to evaluate the generalization of UniD. For depth and surface normal, we follow the evaluation protocols of Marigold [34] and DSINE [4], respectively. For albedo and shading, we report Weighted Human Disagreement Rate (WHDR) [6] and Precision at 70% Recall ( $P@0.7$ ) [39], respectively. For instance boundary, we report optimal-dataset-scale F-measure (odsF) [79]. For segmentation tasks, we report meanIoU and class-frequency-weighted IoU (fwIoU).

To measure temporal consistency, we extend Video Consistency (mVC-t) [49] originally proposed for video semantic segmentation to also measure video normal, albedo, shading, and instance boundary. Please refer to the appendix for details regarding evaluation benchmarks and metrics.

**Baselines.** We compare UniD against per-task specialists and unified multi-task baselines. Among the unified baselines, we include 4M-21-XL [1] and DICEPTION [87], both requiring large-scale pseudo-labeling. Notably, DICEPTION [87] requires task-specific captions as input, requiring  $K$  forward passes to obtain all unified predictions. We further include StableMTL [13], which is trained only on synthetic data and likewise predicts tasks sequentially during inference. As our most directly comparable baseline, we extend a frozen DINOv3-H [62] backbone with task-specific DPT heads trained on exactly the same datasets with the same loss functions for 20k iterations. For a fair comparison, the DINOv3-H backbone is chosen to have the same capacity as our unified backbone  $\mathcal{G}_\phi$ . DINOv3-H simulates the alternative design of freezing a strong image representation and learning per-task decoders from disjoint sources, without any task-specific joint training or temporal modules.




Fig. 3: Qualitative Results of UniD compared with DINOv3-H.<sup>1</sup>

#### 4.1 Per-Task Performance

*Out-of-distribution generalization.* The main claim of UniD is that the generative priors of a pretrained diffusion backbone are sufficient to bridge the domain gaps introduced by disjoint training sources, enabling generalization to scene-task combinations never seen with annotations. To evaluate this, in-distribution benchmarks are shown in gray in Tables 4 and 5, and all remaining columns are out-of-distribution evaluations where labeled training data is not seen.

Across all geometric and intrinsic tasks, UniD ( $\mathcal{G}_\phi$ ) consistently outperforms DINOv3-H on OOD benchmarks, despite identical training datasets. For

<sup>1</sup> Visual examples taken from SA-V [55] and internet videos [54].

**Table 1:** Summary of per-task output dimension  $C^k$ , training datasets (with frame counts), evaluation datasets, and metrics. Most training datasets are disjoint across task domains, with in-domain evaluation datasets shown in gray.

| Task              | $C^k$                   | Training Datasets                                                                                           | #Frames         | Evaluation Datasets                                             | Metrics                      |
|-------------------|-------------------------|-------------------------------------------------------------------------------------------------------------|-----------------|-----------------------------------------------------------------|------------------------------|
| Depth             | 1                       | <i>Image</i> : Hypersim [56], VKITTI2 [12]<br><i>Video</i> : PointOdyssey [88], Spring [47], TartanAir [67] | 80.8k<br>361.0k | NYUv2 [61], KITTI [23], ETH3D [58],<br>ScanNet [20], DIODE [66] | AbsRel↓, $\delta_1$ ↑        |
| Surface Normals   | 3                       | <i>Image</i> : Hypersim [56], VKITTI2 [12]                                                                  | 80.8k           | NYUv2 [61], ScanNet [20], iBims-1 [37], Sintel [11]             | Mean↓, 11.25° ↑              |
| Albedo            | 3                       | <i>Image</i> : Hypersim [56]                                                                                | 59.5k           | Hypersim [56], IIW [6]                                          | PSNR↑, WHDR <sub>15%</sub> ↓ |
| Shading           | 3                       | <i>Image</i> : Hypersim [56]                                                                                | 59.5k           | Hypersim [56], SAW [39]                                         | PSNR↑, P@0.7/AP↑             |
| Instance Boundary | 1                       | <i>Image</i> : COCO [44], Cityscapes [19]<br><i>Video</i> : VIPSeg [48]                                     | 121.3k<br>66.8k | COCO [44], YT-VIS [75], Mapillary [52]                          | odsF↑                        |
| Semantic Seg.     | 201/35/125<br>(C/CS/VS) | <i>Image</i> : COCO [44], Cityscapes [19]<br><i>Video</i> : VIPSeg [48]                                     | 121.3k<br>66.8k | Cityscapes [19], VIPSeg [48], ADE20K [89]                       | mIoU↑, fwIoU↑                |
| Material Seg.     | 57                      | <i>Image</i> : DMS [65]                                                                                     | 22.0k           | DMS [65], ADE20K [89]                                           | mIoU↑, fwIoU↑                |
| Human Part Seg.   | 16                      | <i>Image</i> : DensePose [24]                                                                               | 26.4k           | DensePose [24], Pascal-Human-Part [17]                          | mIoU↑, fwIoU↑                |

depth (Table 2), UniD ( $\mathcal{G}_\phi$ ) achieves the best AbsRel among all unified methods on NYUv2, ScanNet, and DIODE, while DINOv3-H underperforms by a consistent margin across all five datasets. For surface normals (Table 3), UniD ( $\mathcal{G}_\phi$ ) achieves state-of-the-art performance among unified methods on all but one OOD benchmark metric. For intrinsics and boundaries (Table 4), UniD ( $\mathcal{G}_\phi$ ) matches DINOv3-H in-domain, while substantially outperforming it on every OOD benchmark: improving albedo WHDR from 0.289 to 0.207, shading P@0.7 from 87.8 to 94.6, and boundary odsF from 55.0 to 57.2.

This trend of parity on in-domain and clear advantage out-of-domain verifies the efficacy of the diffusion prior in transferring robustly to unseen scenarios. To isolate the effect of diffusion prior, we apply color jittering to Hypersim val [56] and found the specialist UniD ( $\mathcal{F}_\theta^k$ ) to have substantially more stable predictions than DINOv3-H on depth (AbsRel deviation: 0.011 vs. 0.034) and surface normals (mean error deviation: 2.2° vs. 5.8°) — roughly 3× **more robust** despite identical training data. This suggests that discriminative features are prone to overfit to the appearance-geometry correlations in the training data, while generative pretraining on web-scale data encodes more appearance-invariant geometry. This is further evident in Figure 3, where DINOv3-H predictions degrade noticeably for unseen scenes while UniD ( $\mathcal{G}_\phi$ ) remains coherent across all tasks.

*Unified model versus task-specialist.* Across geometric and intrinsic tasks, UniD ( $\mathcal{G}_\phi$ ) achieves similar performance with per-task specialists  $\mathcal{F}_\theta^k$ . This demonstrates the efficacy of latent distillation that faithfully preserves specialist-level representation even after unifying all eight tasks into a single shared backbone. For surface normals, the unified model  $\mathcal{G}_\phi$  in fact *outperforms* its own specialist  $\mathcal{F}_\theta^N$  on NYUv2, ScanNet, and Sintel, a gain that can be attributed to cross-task regularization from depth, which shares the same geometric training domain and jointly reinforces surface geometry estimation within the unified backbone. Furthermore, for shading and instance boundaries (Table 4), UniD ( $\mathcal{G}_\phi$ ) matches the specialist  $\mathcal{F}_\theta^k$  on in-domain evaluations, while improving it marginally on out-of-domain evaluations.



**Table 2:** Monocular Depth Estimation Performance on NYUv2, KITTI, ETH3D, ScanNet, and DIODE. Best performance among unified methods is bolded.

| Method     | NYUv2                           |                     | KITTI       |                     | ETH3D       |                     | ScanNet     |                     | DIODE       |                     |             |
|------------|---------------------------------|---------------------|-------------|---------------------|-------------|---------------------|-------------|---------------------|-------------|---------------------|-------------|
|            | AbsRel ↓                        | $\delta_1 \uparrow$ | AbsRel ↓    | $\delta_1 \uparrow$ | AbsRel ↓    | $\delta_1 \uparrow$ | AbsRel ↓    | $\delta_1 \uparrow$ | AbsRel ↓    | $\delta_1 \uparrow$ |             |
| Specialist | DAv2 [74]                       | 0.043               | 98.1        | 0.076               | 94.7        | 0.127               | 88.2        | 0.043               | 98.1        | 0.260               | 75.9        |
|            | Depth Pro [9]                   | 0.042               | 97.7        | 0.055               | 97.4        | 0.043               | 97.4        | 0.041               | 97.8        | 0.217               | 76.4        |
|            | Marigold [34]                   | 0.055               | 96.4        | 0.099               | 91.6        | 0.065               | 95.9        | 0.064               | 95.2        | 0.308               | 77.3        |
|            | Genpercept [69]                 | 0.091               | 93.2        | 0.094               | 92.3        | 0.066               | 95.7        | 0.056               | 96.5        | 0.302               | 76.7        |
|            | GeoWizard [22]                  | 0.052               | 96.6        | 0.097               | 92.1        | 0.064               | 96.1        | 0.061               | 95.3        | 0.297               | 79.2        |
|            | Lotus [25]                      | 0.051               | 97.2        | 0.081               | 93.1        | 0.061               | 97.0        | 0.055               | 96.5        | 0.228               | 73.8        |
|            | UniD ( $\mathcal{F}_\theta^D$ ) | 0.059               | 96.0        | 0.102               | 90.9        | 0.070               | 95.4        | 0.069               | 95.0        | 0.299               | 77.4        |
| Unified    | 4M-21-XL [1]                    | 0.068               | 95.1        | 0.105               | 89.6        | <b>0.070</b>        | 95.3        | 0.065               | 95.5        | 0.331               | 73.4        |
|            | DICEPTION [87]                  | 0.060               | <b>96.3</b> | <b>0.085</b>        | <b>91.8</b> | <b>0.070</b>        | <b>96.6</b> | 0.071               | 94.8        | 0.304               | 72.2        |
|            | StableMTL [13]                  | 0.090               | 92.4        | 0.150               | 79.8        | 0.116               | 87.9        | 0.105               | 89.5        | 0.330               | 73.8        |
|            | DINOV3-H [62]                   | 0.092               | 92.9        | 0.121               | 86.5        | 0.087               | 93.7        | 0.089               | 93.6        | 0.314               | <b>77.5</b> |
|            | UniD ( $\mathcal{G}_\phi$ )     | <b>0.059</b>        | <b>96.3</b> | 0.110               | 89.5        | 0.073               | 95.5        | <b>0.063</b>        | <b>96.2</b> | <b>0.301</b>        | 77.3        |

**Table 3:** Surface Normal Estimation Performance on NYUv2, ScanNet, iBims-1, and Sintel. Best performance among unified methods is bolded.

| Method     |                                 | NYUv2       |             | ScanNet     |             | iBims-1     |             | Sintel      |             |
|------------|---------------------------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
|            |                                 | mean ↓      | 11.25°↑     | mean ↓      | 11.25°↑     | mean ↓      | 11.25°↑     | mean ↓      | 11.25°↑     |
| Specialist | OASIS [16]                      | 29.2        | 23.8        | 32.8        | 15.4        | 32.6        | 23.5        | 43.1        | 7.0         |
|            | Omnidata v2 [32]                | 17.2        | 55.5        | 16.2        | 60.2        | 18.2        | 63.9        | 40.5        | 14.7        |
|            | DSINE [4]                       | 16.4        | 59.6        | 16.2        | 61.0        | 17.1        | 67.4        | 34.9        | 21.5        |
|            | GeoWizard [22]                  | 18.9        | 50.7        | 17.4        | 53.8        | 19.3        | 63.0        | 40.3        | 12.3        |
|            | StableNormal [78]               | 18.6        | 53.5        | 17.1        | 57.4        | 18.2        | 65.0        | 36.7        | 14.1        |
|            | Lotus [25]                      | 16.2        | 59.8        | 14.7        | 64.0        | 17.1        | 66.4        | 32.3        | 22.4        |
|            | UniD ( $\mathcal{F}_\theta^N$ ) | 16.7        | 59.9        | 15.4        | 63.7        | 16.3        | 69.0        | 34.3        | 22.2        |
| Unified    | 4M-21-XL [1]                    | 18.4        | 49.8        | 17.4        | 48.9        | 18.6        | 61.4        | 41.6        | 12.4        |
|            | DICEPTION [87]                  | 19.7        | 51.4        | 19.3        | 54.5        | 17.8        | 66.2        | 37.2        | 19.6        |
|            | StableMTL [13]                  | 19.6        | 48.9        | 20.0        | 45.8        | 19.6        | 58.1        | 40.5        | 11.0        |
|            | DINov3-H [62]                   | 16.8        | 57.0        | 16.0        | 58.7        | 16.7        | 66.8        | <b>32.3</b> | 18.9        |
|            | UniD ( $\mathcal{G}_\phi$ )     | <b>16.2</b> | <b>59.9</b> | <b>14.6</b> | <b>65.1</b> | <b>16.6</b> | <b>67.4</b> | 33.0        | <b>21.4</b> |

*Segmentation and lightweight fine-tuning.* In contrast, Table 5 reveals a systematic gap between UniD ( $\mathcal{G}_\phi$ ) and the per-task specialist for classification tasks. We find that reconstruction via  $\ell_1$  is insufficient to fully match the specialists learned under cross-entropy (CE) instead of regression-based objectives, and this degradation persists even when  $K=1$  (*i.e.* when only distilling a single segmentation specialist), indicating the limitation is intrinsic to the latent reconstruction objective. We trace this to a projector-backbone misalignment: CE-task specialist latents exhibit significantly higher spatial Laplacian norms (2.52 vs. 1.40) than regression-task latents, suggesting that  $\ell_1$  distillation may suppress the high-frequency structure that CE specialists encode. To address this, we perform a lightweight fine-tuning step that only updates the task-specific projectors ( $\Psi^k, \mathcal{P}^k$ ) while keeping  $\mathcal{G}_\phi$  entirely frozen. This targeted fine-tuning (reported as UniD ( $\mathcal{G}_\phi$ ) + FT) substantially closes the gap with the specialist, recovering



**Table 4:** Albedo, Shading, and Instance Boundary Estimation Performance. Best performance among unified methods is bolded.

| Method |                                 | Albedo          |                                  | Shading         |                      | Boundary        |                 |
|--------|---------------------------------|-----------------|----------------------------------|-----------------|----------------------|-----------------|-----------------|
|        |                                 | Hypersim        | IIW                              | Hypersim        | SAW                  | COCO            | Mapillary       |
|        |                                 | PSNR $\uparrow$ | WHDR <sub>15%</sub> $\downarrow$ | PSNR $\uparrow$ | P@0.7/AP% $\uparrow$ | odsF $\uparrow$ | odsF $\uparrow$ |
| Spec.  | ColorfulShading [14]            | 17.2            | 0.175                            | 14.7            | 81.3 / 79.4          | -               | -               |
|        | Mask2Former [18]                | -               | -                                | -               | -                    | 67.7            | 35.6            |
|        | UniD ( $\mathcal{F}_\theta^k$ ) | 17.8            | 0.201                            | 20.3            | 93.5 / 91.5          | 79.8            | 56.3            |
| Unif.  | DINOv3-H [62]                   | 17.8            | 0.289                            | 19.5            | 87.8 / 88.0          | 81.4            | 55.0            |
|        | StableMTL [13]                  | -               | <b>0.190</b>                     | -               | 82.7 / 63.7          | -               | -               |
|        | UniD ( $\mathcal{G}_\phi$ )     | 17.6            | 0.207                            | <b>20.2</b>     | <b>94.6 / 92.6</b>   | 80.8            | <b>57.2</b>     |

**Table 5:** Semantic, Material, and Human Part Segmentation Performance.

| Method                                | Semantic                         |              |              | Material     |              | Human Part   |                   |
|---------------------------------------|----------------------------------|--------------|--------------|--------------|--------------|--------------|-------------------|
|                                       | Cityscapes                       | VISeg        | ADE20K       | DMS          | ADE20K       | DensePose    | Pascal-Human-Part |
|                                       | mIoU / fwIoU                     | mIoU / fwIoU | mIoU / fwIoU | mIoU / fwIoU | mIoU / fwIoU | mIoU / fwIoU | mIoU / fwIoU      |
| Spec. UniD ( $\mathcal{F}_\theta^k$ ) | 64.1 / 85.3                      | 45.8 / 67.2  | 47.6 / 71.6  | 47.1 / 74.1  | 67.2 / 76.2  | 68.2 / 93.3  | 66.9 / 88.4       |
| Unif.                                 | StableMTL [13]                   | 55.8 / -     | -/-          | -/-          | -/-          | -/-          | -/-               |
|                                       | DINOV3-H [62]                    | 65.6 / 83.8  | 52.8 / 66.8  | 62.4 / 74.8  | 50.6 / 75.1  | 69.1 / 75.7  | 67.7 / 93.0       |
|                                       | UniD ( $\mathcal{G}_\phi$ )      | 37.8 / 79.7  | 46.1 / 67.1  | 35.7 / 66.2  | 40.6 / 72.1  | 44.1 / 81.4  | 64.3 / 92.9       |
|                                       | UniD ( $\mathcal{G}_\phi$ ) + FT | 58.4 / 83.3  | 47.0 / 67.0  | 44.8 / 68.9  | 47.3 / 73.6  | 66.0 / 76.0  | 66.1 / 92.9       |

from 37.8 to 58.4 in mIoU for Cityscapes and from 35.7 to 44.8 in mIoU for ADE20K. The rapid recovery from projector-only finetuning confirms that the backbone  $\mathcal{G}_\phi$  *does* learn semantically rich representations under  $\ell_1$  distillation, preserving the unified representation across all tasks.

## 4.2 Temporal Consistency

Table 6 evaluates temporal consistency across depth, normals, albedo, shading, boundary, and semantic segmentation on video benchmarks. UniD ( $\mathcal{G}_\phi$ ) outperforms all unified baselines across every task and, notably, is competitive with video-specific specialists trained for individual tasks. On depth, our streaming model outperforms DICEPTION and DINOV3-H, while marginally underperforming Video Depth Anything that processes video in a chunk-based fashion. For surface normals, UniD ( $\mathcal{G}_\phi$ ) outperforms DINOV3-H by a large margin, and even surpasses NormalCrafter [7] that utilizes a video diffusion backbone.

To verify that the advantage over unified baselines is not merely an artifact of the baselines lacking temporal module, we additionally equip DINOV3-H and DICEPTION with the same extended self-attention (ESA) for a fair comparison. We find that ESA yields marginal but inconsistent improvements for both baselines, while UniD ( $\mathcal{G}_\phi$ ) still outperforms all variants across all tasks. This demonstrates that temporal stability arises not from the parameter-free ESA module alone, but from its combination with unified training on video data, which together propagate temporal smoothness across all tasks simultaneously without requiring task-specific video annotations.

**Table 6:** Temporal Consistency performance on video prediction tasks. Best performance among unified methods is bolded. Extended Self-Attention [41] is additionally added to image models with 16-frame memory dequeued at every frame (+ESA[16,1]).

| Method  | Depth                           |              | Normal      |                    | Albedo             |                    | Shading     |             | Boundary | Semantic |
|---------|---------------------------------|--------------|-------------|--------------------|--------------------|--------------------|-------------|-------------|----------|----------|
|         | ScanNet                         |              | InteriorNet |                    | InteriorNet        |                    | InteriorNet |             | YT-VIS   | VIPSeg   |
|         | AbsRel ↓                        | $\delta_1$ ↑ | mVC-4/16 ↑  |                    | mVC-4/16 ↑         |                    | mVC-4/16 ↑  |             | mVC-4 ↑  | mVC-4 ↑  |
| Spec.   | Video Depth Anything [15]       | 0.063        | 96.8        | -/-                | -/-                | -/-                | -/-         | -           | -        | -        |
|         | NormalCrafter [7]               | -            | -           | 84.6 / 70.9        | -/-                | -/-                | -/-         | -           | -        | -        |
|         | ColorfulShading [14]            | -            | -           | - / -              | 78.9 / 55.1        | 73.5 / 41.5        | -           | -           | -        | -        |
|         | UniD ( $\mathcal{F}_\theta^k$ ) | 0.101        | 91.2        | 85.9 / 76.4        | 86.9 / 70.1        | 92.5 / 81.9        | 93.6        | 93.8        |          |          |
| Unified | DICEPTION [87]                  | 0.123        | 85.7        | 68.0 / 47.1        | -/-                | -/-                | -           | -           | -        | -        |
|         | DINOV3-H [62]                   | 0.120        | 86.9        | 78.2 / 62.2        | 77.8 / 53.4        | 90.2 / 77.6        | 92.0        | 87.8        |          |          |
|         | DICEPTION [87] + ESA[16,1]      | 0.118        | 85.9        | 66.0 / 48.6        | -/-                | -/-                | -           | -           | -        | -        |
|         | DINOV3-H [62] + ESA[16,1]       | 0.150        | 79.8        | 79.0 / 63.6        | 76.4 / 54.1        | 92.4 / 82.7        | 92.4        | 89.7        |          |          |
|         | UniD ( $\mathcal{G}_\phi$ )     | <b>0.103</b> | <b>91.0</b> | <b>84.7 / 75.8</b> | <b>88.8 / 73.7</b> | <b>94.5 / 85.2</b> | <b>93.7</b> | <b>90.3</b> |          |          |

**Table 7:** Cross-Task Consistency Evaluation. Best performance among unified methods is bolded.

| Method  | Depth↔Normal                    |             | Part↔Semantic |             | Albedo↔Shading |             |
|---------|---------------------------------|-------------|---------------|-------------|----------------|-------------|
|         | ScanNet                         |             | COCO          |             | HyperSim       |             |
|         | mean ↓                          | 11.25° ↑    | % inclusion ↑ |             | PSNR ↑         | LPIPS ↓     |
| Spec.   | ColorfulShading [14]            | -           | -             | -           | 20.8           | 0.27        |
|         | UniD ( $\mathcal{F}_\theta^k$ ) | 23.9        | 39.0          | 97.6        | 24.6           | 0.22        |
| Unified | 4M-21-XL [1]                    | 23.4        | 37.2          | -           | -              | -           |
|         | DICEPTION [87]                  | 27.3        | 27.4          | -           | -              | -           |
|         | DINOV3-H [62]                   | 36.0        | 14.1          | 98.1        | 23.9           | 0.26        |
|         | UniD ( $\mathcal{G}_\phi$ )     | <b>22.6</b> | <b>39.0</b>   | <b>98.3</b> | <b>24.3</b>    | <b>0.24</b> |

### 4.3 Cross-Task Consistency

Table 7 evaluates geometric, semantic, and photometric consistency between pairs of related task predictions.

To measure consistency between depth and normal, we compute error between surface normal derived from depth predictions and direct normal prediction. UniD ( $\mathcal{G}_\phi$ ) achieves the best consistency on ScanNet among all methods, outperforming DICEPTION and DINOV3-H via a wide margin, and surpassing the independent specialist ensemble UniD ( $\mathcal{F}_\theta^k$ ). For consistency between human part and semantic segmentation, we measure the % inclusion of predicted human part pixels within the human semantic category. Here, the unified model marginally outperforms DINOV3-H and the specialist ensemble, confirming that the shared backbone encourages semantic predictions to remain coherent with each other. Finally, to measure intrinsic consistency, we directly measure reconstruction performance of diffuse reconstruction of the predicted albedo and shading. Similarly, UniD ( $\mathcal{G}_\phi$ ) remains competitive with the specialist ensemble, while still outperforming DINOV3-H.

This suggests that the joint latent distillation objective—where a unified backbone is supervised by all task specialists simultaneously—implicitly encourages cross-task geometric and semantic coherence that independent specialists or frozen-backbone baselines cannot achieve.

**Table 8:** Per-component inference analysis on a single A100 GPU per-frame at  $432 \times 768$  resolution with varying  $M$  number of past frames in memory. UniD ( $\mathcal{F}_\theta^k$ ) denotes the specialist ensemble ( $K=8$  sequential U-Net passes); UniD ( $\mathcal{G}_\phi$ ) denotes the unified model (one shared U-Net pass with  $K$  DPT connectors  $\Psi^k$ ).

| Method                          | $M$ | $\mathcal{E}$ (ms) | U-Net (ms)       | $\Psi^k$ (ms)  | $\mathcal{D}$ (ms) | $\mathcal{P}^k$ (ms) | Total (s)                                | FPS         |
|---------------------------------|-----|--------------------|------------------|----------------|--------------------|----------------------|------------------------------------------|-------------|
| StableMTL [13]                  | 0   | -                  | -                | -              | -                  | -                    | $0.03 + 0.58 \times K$                   | 0.21        |
| 4M-21-XL [1]                    | 0   | -                  | -                | -              | -                  | -                    | $0.29 + 0.43 \times K$                   | 0.27        |
| DICEPTION [87]                  | 0   | -                  | -                | -              | -                  | -                    | $0.07 + 4.15 \times K$                   | 0.03        |
| DINov3-H [62]                   | 0   | -                  | -                | -              | -                  | -                    | <b><math>0.21 + 0.02 \times K</math></b> | <b>2.67</b> |
| UniD ( $\mathcal{F}_\theta^k$ ) | 0   | 35.9               | $58.6 \times K$  | -              | $42.3 \times K$    | $1.1 \times K$       | $0.04 + 0.10 \times K$                   | 1.17        |
| UniD ( $\mathcal{G}_\phi$ )     | 0   | 35.9               | 58.6             | $3.1 \times K$ | $42.3 \times K$    | $1.1 \times K$       | $0.09 + 0.05 \times K$                   | 2.14        |
| DICEPTION [87] + ESA            | 16  | -                  | -                | -              | -                  | -                    | $0.07 + 5.05 \times K$                   | 0.02        |
| DINov3-H [62] + ESA             | 16  | -                  | -                | -              | -                  | -                    | $1.30 + 0.02 \times K$                   | 0.69        |
| UniD ( $\mathcal{F}_\theta^k$ ) | 16  | 36.5               | $235.0 \times K$ | -              | $42.2 \times K$    | $1.1 \times K$       | $0.04 + 0.28 \times K$                   | 0.44        |
| UniD ( $\mathcal{G}_\phi$ )     | 16  | 36.5               | 235.0            | $3.0 \times K$ | $42.2 \times K$    | $1.1 \times K$       | <b><math>0.27 + 0.05 \times K</math></b> | <b>1.56</b> |

**Table 9:** Ablation Study on temporal consistency performance on video prediction tasks. If *Train w/ Video* is set to  $\times$ , all videos are treated as independent images.

| Train w/     | Inf.        | Depth               |                     | Normal              |  | Albedo              |  | Shading             |  | Boundary         | Semantic         |
|--------------|-------------|---------------------|---------------------|---------------------|--|---------------------|--|---------------------|--|------------------|------------------|
|              |             | ScanNet             |                     | InteriorNet         |  | InteriorNet         |  | InteriorNet         |  | YT-VIS           | VIPSeg           |
| Video        | Memory      | AbsRel $\downarrow$ | $\delta_1 \uparrow$ | mVC-4/16 $\uparrow$ |  | mVC-4/16 $\uparrow$ |  | mVC-4/16 $\uparrow$ |  | mVC-4 $\uparrow$ | mVC-4 $\uparrow$ |
| $\times$     | $\emptyset$ | 0.121               | 87.1                | 79.0 / 65.6         |  | 84.4 / 65.0         |  | 92.7 / 81.5         |  | 92.9             | 85.9             |
| $\times$     | [16, 1]     | 0.115               | 88.1                | 82.9 / 72.7         |  | 87.3 / 70.7         |  | 93.6 / 82.9         |  | 93.0             | 89.1             |
| $\checkmark$ | $\emptyset$ | 0.109               | 90.0                | 80.9 / 68.5         |  | 86.0 / 67.1         |  | 93.7 / 83.5         |  | 93.6             | 86.5             |
| $\checkmark$ | [16, 1]     | 0.103               | 91.0                | 84.7 / 75.8         |  | 88.8 / 73.7         |  | 94.5 / 85.2         |  | 93.7             | 90.3             |

#### 4.4 Efficiency Analysis & Ablation Studies

**Efficiency Analysis.** In Table 8, UniD ( $\mathcal{G}_\phi$ ) is efficient across all memory settings, remaining competitive with the discriminative DINov3-H [62] at  $M=0$  while substantially outperforming all baselines at  $M=16$ . StableMTL [13], DICEPTION [87], and 4M-21-XL [1] underperform due to their prohibitive per-task costs, whereas UniD ( $\mathcal{G}_\phi$ ) sustains 1.56 FPS even with a 16-frame memory bank.

The key to this efficiency is that nearly all computation is shared across tasks. Compared to the specialist ensemble UniD ( $\mathcal{F}_\theta^k$ ), which requires  $K$  sequential U-Net passes with memory attentions, UniD ( $\mathcal{G}_\phi$ ) processes a single *unified* historical feature for all tasks at once, yielding a  $1.8\times$  speedup at  $M=0$  that compounds to over  $3.5\times$  at  $M=16$  as temporal context grows.

**Latent versus per-pixel distillation.** A natural alternative to our proposed latent distillation would be to directly supervise the unified model with per-pixel predictions with either pre-computed pseudo-labels or on-the-fly. However, they are infeasible at scale due to the prohibitive nature of per-pixel backpropagation during unified training. Computing gradients through the decoder  $\mathcal{D}'$  per task in a training step requires holding  $K$  separate decoder activation graphs in GPU memory simultaneously, growing memory consumption linearly with  $K$  and rendering joint training intractable on standard hardware. Latent distilla-

tion sidesteps this bottleneck: latent targets  $\mathbf{l}_i^k = \mathcal{F}_\theta^k(\mathbf{z}_i) \in \mathbb{R}^{h \times w \times d}$  are cheap to compute, small to store, and avoid any computation with the decoder  $\mathcal{D}$ . As a result, with identical training setup, the peak memory usage per device can decrease from an estimated 650GB to 78GB, when switching to latent distillation.

**Impact of video training and inference-time memory.** Shown in Table 9, when all videos are treated as independent images during training, temporal consistency degrades across all benchmarks. This directly ablates the temporal gradient matching (TGM) loss, as static-image training removes all cross-frame supervision. This suggests that, with only depth and semantic segmentation containing video annotations, unified training propagates temporal consistency to all tasks, even those without task-specific video annotations.

Notably, this holds even with no inference memory ( $\emptyset$ ): video training alone yields consistently more stable representations across all tasks, before any inference-time memory module is introduced. Enabling extended self-attention memory bank at inference ([16, 1], a queue of 16 frames dequeued every frame) further improves temporal consistency across all tasks, regardless of whether video training was used. Interestingly, even without video training, inference-time memory provides non-trivial gains across tasks without video annotations, suggesting the memory mechanism can compensate for the absence of temporal training signal when we distill the task-specialist  $\mathcal{F}_\theta^k$  during unified training.

While video training drives temporal consistency, it is *not* responsible for our per-task accuracy advantage over the unified image baselines. Treating videos as independent images yields nearly identical per-task performance as shown in Appendix, confirming that the gains over baselines stem from the unified architecture and latent distillation rather than from video supervision.

## 5 Conclusion

Real-world data are inherently fragmented across incompatible domains. In this work, we present UniD, learning unified video dense prediction from these disjoint sources. It first learns task-specific embeddings on disjoint datasets, then distills them into a unified backbone via per-task latent projectors. To bridge these domain gaps, we leverage the diffusion priors from internet-scale pretraining to enable generalization to unseen scene-task combinations. UniD achieves competitive performance with per-task specialists and outperforms unified baselines across OOD scenarios with superior temporal and cross-task consistency.

**Limitations.** Latent reconstruction is less effective for classification tasks, requiring additional fine-tuning to recover segmentation performance; future work could explore objectives that better preserve classification representations.

**Acknowledgements.** This work was conducted in part during an internship at Adobe Research and supported in part by NSF (IIS-2144117). Yihong Sun is supported by an NSF Graduate Research Fellowship. We thank Yao-Chih Lee, Wei-Hong Li, Youming Deng, and Ashley P. Tsang for their valuable feedback.

## References

1. 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)
2. Bachmann, R., Mizrahi, D., Atanov, A., Zamir, A.: Multimaes: Multi-modal multi-task masked autoencoders. In: *European conference on computer vision*. pp. 348–367. Springer (2022)
3. Badki, A., Su, H., Wen, B., Gallo, O.: L4p: Towards unified low-level 4d vision perception. *arXiv preprint arXiv:2502.13078* (2025)
4. Bae, G., Davison, A.J.: Rethinking inductive biases for surface normal estimation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9535–9545 (2024)
5. Barron, J.T., Malik, J.: Shape, illumination, and reflectance from shading. *IEEE transactions on pattern analysis and machine intelligence* **37**(8), 1670–1687 (2014)
6. Bell, S., Bala, K., Snavely, N.: Intrinsic images in the wild. *ACM Trans. on Graphics (SIGGRAPH)* **33**(4) (2014)
7. Bin, Y., Hu, W., Wang, H., Chen, X., Wang, B.: Normalcrafter: Learning temporally consistent normals from video diffusion priors. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 8330–8339 (2025)
8. Blattmann, A., Dockhorn, T., Kulal, S., Mendelevitch, D., Kilian, M., Lorenz, D., Levi, Y., English, Z., Voleti, V., Letts, A., et al.: Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127* (2023)
9. Bochkovskii, A., Delaunoy, A., Germain, H., Santos, M., Zhou, Y., Richter, S.R., Koltun, V.: Depth pro: Sharp monocular metric depth in less than a second. *arXiv preprint arXiv:2410.02073* (2024)
10. Brooks, T., Holynski, A., Efros, A.A.: Instructpix2pix: Learning to follow image editing instructions. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 18392–18402 (2023)
11. 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)
12. Cabon, Y., Murray, N., Humenberger, M.: Virtual kitti 2. *arXiv preprint arXiv:2001.10773* (2020)
13. Cao, A.Q., Lopes, I., de Charette, R.: Stablemtl: Repurposing latent diffusion models for multi-task learning from partially annotated synthetic datasets. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2026)
14. Careaga, C., Aksoy, Y.: Colorful diffuse intrinsic image decomposition in the wild. *ACM Transactions on Graphics (TOG)* **43**(6), 1–12 (2024)
15. Chen, S., Guo, H., Zhu, S., Zhang, F., Huang, Z., Feng, J., Kang, B.: Video depth anything: Consistent depth estimation for super-long videos. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 22831–22840 (2025)
16. Chen, W., Qian, S., Fan, D., Kojima, N., Hamilton, M., Deng, J.: Oasis: A large-scale dataset for single image 3d in the wild. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 679–688 (2020)
17. Chen, X., Mottaghi, R., Liu, X., Fidler, S., Urtasun, R., Yuille, A.: Detect what you can: Detecting and representing objects using holistic models and body parts. In:

- Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1971–1978 (2014)
18. Cheng, B., Misra, I., Schwing, A.G., Kirillov, A., Girdhar, R.: Masked-attention mask transformer for universal image segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 1290–1299 (2022)
  19. Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3213–3223 (2016)
  20. Dai, A., Chang, A.X., Savva, M., Halber, M., Funkhouser, T., Nießner, M.: Scannet: Richly-annotated 3d reconstructions of indoor scenes. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5828–5839 (2017)
  21. 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)
  22. 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. pp. 241–258. Springer (2024)
  23. Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? the kitti vision benchmark suite. In: 2012 IEEE conference on computer vision and pattern recognition. pp. 3354–3361. IEEE (2012)
  24. Güler, R.A., Neverova, N., Kokkinos, I.: Densepose: Dense human pose estimation in the wild. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7297–7306 (2018)
  25. He, J., Li, H., Yin, W., Liang, Y., Li, L., Zhou, K., Zhang, H., Liu, B., Chen, Y.C.: Lotus: Diffusion-based visual foundation model for high-quality dense prediction. arXiv preprint arXiv:2409.18124 (2024)
  26. Ho, J., Chan, W., Saharia, C., Whang, J., Gao, R., Gritsenko, A., Kingma, D.P., Poole, B., Norouzi, M., Fleet, D.J., et al.: Imagen video: High definition video generation with diffusion models. arXiv preprint arXiv:2210.02303 (2022)
  27. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
  28. Ho, J., Salimans, T.: Classifier-free diffusion guidance. arXiv preprint arXiv:2207.12598 (2022)
  29. Ho, J., Salimans, T., Gritsenko, A., Chan, W., Norouzi, M., Fleet, D.J.: Video diffusion models. *Advances in neural information processing systems* **35**, 8633–8646 (2022)
  30. Huang, T.E., Liu, Y., Van Gool, L., Yu, F.: Video task decathlon: Unifying image and video tasks in autonomous driving. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8647–8657 (2023)
  31. Jain, S., Wang, X., Gonzalez, J.E.: Accel: A corrective fusion network for efficient semantic segmentation on video. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 8866–8875 (2019)
  32. Kar, O.F., Yeo, T., Atanov, A., Zamir, A.: 3d common corruptions and data augmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18963–18974 (2022)

33. Kawar, B., Zada, S., Lang, O., Tov, O., Chang, H., Dekel, T., Mosseri, I., Irani, M.: Imagic: Text-based real image editing with diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 6007–6017 (2023)
34. Ke, B., Obukhov, A., Huang, S., Metzger, N., Daut, 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. pp. 9492–9502 (2024)
35. Kim, D., Cho, S., Kim, S., Luo, C., Hong, S.: Chameleon: A data-efficient generalist for dense visual prediction in the wild. In: European Conference on Computer Vision. pp. 422–441. Springer (2024)
36. Kim, D., Kim, J., Cho, S., Luo, C., Hong, S.: Universal few-shot learning of dense prediction tasks with visual token matching. arXiv preprint arXiv:2303.14969 (2023)
37. Koch, T., Liebel, L., Fraundorfer, F., Korner, M.: Evaluation of cnn-based single-image depth estimation methods. In: Proceedings of the European Conference on Computer Vision (ECCV) Workshops. pp. 0–0 (2018)
38. Kondratyuk, D., Yu, L., Gu, X., Lezama, J., Huang, J., Schindler, G., Hornung, R., Birodkar, V., Yan, J., Chiu, M.C., et al.: Videopoet: A large language model for zero-shot video generation. arXiv preprint arXiv:2312.14125 (2023)
39. Kovacs, B., Bell, S., Snavely, N., Bala, K.: Shading annotations in the wild. Computer Vision and Pattern Recognition (CVPR) (2017)
40. Li, W.H., Liu, X., Bilen, H.: Learning multiple dense prediction tasks from partially annotated data. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18879–18889 (2022)
41. Liang, F., Kodaira, A., Xu, C., Tomizuka, M., Keutzer, K., Marculescu, D.: Looking backward: Streaming video-to-video translation with feature banks. arXiv preprint arXiv:2405.15757 (2024)
42. Lin, B., Jiang, W., Chen, P., Liu, S., Chen, Y.C.: Mtmamba++: Enhancing multi-task dense scene understanding via mamba-based decoders. IEEE Transactions on Pattern Analysis and Machine Intelligence (2025)
43. Lin, B., Jiang, W., Chen, P., Zhang, Y., Liu, S., Chen, Y.C.: Mtmamba: Enhancing multi-task dense scene understanding by mamba-based decoders. In: European conference on computer vision. pp. 314–330. Springer (2024)
44. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: European conference on computer vision. pp. 740–755. Springer (2014)
45. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)
46. Lu, Y., Sirejiding, S., Ding, Y., Wang, C., Lu, H.: Prompt guided transformer for multi-task dense prediction. IEEE Transactions on Multimedia **26**, 6375–6385 (2024)
47. Mehl, L., Schmalfluss, J., Jahedi, A., Nalivayko, Y., Bruhn, A.: Spring: A high-resolution high-detail dataset and benchmark for scene flow, optical flow and stereo. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4981–4991 (2023)
48. Miao, J., Wang, X., Wu, Y., Li, W., Zhang, X., Wei, Y., Yang, Y.: Large-scale video panoptic segmentation in the wild: A benchmark. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21033–21043 (2022)



49. Miao, J., Wei, Y., Wu, Y., Liang, C., Li, G., Yang, Y.: Vspw: A large-scale dataset for video scene parsing in the wild. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4133–4143 (2021)
50. Mou, C., Wang, X., Xie, L., Wu, Y., Zhang, J., Qi, Z., Shan, Y.: T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. In: Proceedings of the AAAI conference on artificial intelligence. vol. 38, pp. 4296–4304 (2024)
51. Narihira, T., Maire, M., Yu, S.X.: Direct intrinsics: Learning albedo-shading decomposition by convolutional regression. In: Proceedings of the IEEE international conference on computer vision (2015)
52. Neuhold, G., Ollmann, T., Rota Bulò, S., Kotschieder, P.: The mapillary vistas dataset for semantic understanding of street scenes. In: Proceedings of the IEEE international conference on computer vision. pp. 4990–4999 (2017)
53. Nichol, A., Dhariwal, P., Ramesh, A., Shyam, P., Mishkin, P., McGrew, B., Sutskever, I., Chen, M.: Glide: Towards photorealistic image generation and editing with text-guided diffusion models. arXiv preprint arXiv:2112.10741 (2021)
54. Pexels Website: Pexels — free stock photos & videos you can use everywhere. <https://www.pexels.com/>
55. 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)
56. Roberts, M., Ramapuram, J., Ranjan, A., Kumar, A., Bautista, M.A., Paczan, N., Webb, R., Susskind, J.M.: Hypersim: A photorealistic synthetic dataset for holistic indoor scene understanding. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 10912–10922 (2021)
57. 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. pp. 10684–10695 (2022)
58. 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)
59. Shao, J., Yang, Y., Zhou, H., Zhang, Y., Shen, Y., Guizilini, V., Wang, Y., Poggi, M., Liao, Y.: Learning temporally consistent video depth from video diffusion priors. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 22841–22852 (2025)
60. Shin, I., Kim, D., Yu, Q., Xie, J., Kim, H.S., Green, B., Kweon, I.S., Yoon, K.J., Chen, L.C.: Video-kmax: A simple unified approach for online and near-online video panoptic segmentation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 229–239 (2024)
61. Silberman, N., Hoiem, D., Kohli, P., Fergus, R.: Indoor segmentation and support inference from rgb-d images. In: European conference on computer vision. pp. 746–760. Springer (2012)
62. 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)
63. Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. arXiv preprint arXiv:2011.13456 (2020)

64. Tang, L., Jia, M., Wang, Q., Phoo, C.P., Hariharan, B.: Emergent correspondence from image diffusion. *Advances in neural information processing systems* **36**, 1363–1389 (2023)
65. Upchurch, P., Niu, R.: A dense material segmentation dataset for indoor and outdoor scene parsing. In: *European conference on computer vision*. pp. 450–466. Springer (2022)
66. Vasiljevic, I., Kolkin, N., Zhang, S., Luo, R., Wang, H., Dai, F.Z., Daniele, A.F., Mostajabi, M., Basart, S., Walter, M.R., et al.: Diode: A dense indoor and outdoor depth dataset. *arXiv preprint arXiv:1908.00463* (2019)
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)
68. Xia, C., Jia, C., Dang, Z., Luo, M., Li, Z., Chang, X.: From ideal to real: Unified and data-efficient dense prediction for real-world scenarios. *arXiv preprint arXiv:2506.20279* (2025)
69. 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)
70. Xu, Y., Li, X., Yuan, H., Yang, Y., Zhang, L.: Multi-task learning with multi-query transformer for dense prediction. *IEEE Transactions on Circuits and Systems for Video Technology* **34**(2), 1228–1240 (2023)
71. Xu, Y., Yang, Y., Zhang, L.: Demt: Deformable mixer transformer for multi-task learning of dense prediction. In: *Proceedings of the AAAI conference on artificial intelligence*. vol. 37, pp. 3072–3080 (2023)
72. 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. *arXiv preprint arXiv:2410.10815* (2024)
73. Yang, L., Qi, L., Li, X., Li, S., Jampani, V., Yang, M.H.: Unified dense prediction of video diffusion. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 28963–28973 (2025)
74. 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)
75. Yang, L., Fan, Y., Xu, N.: Video instance segmentation. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 5188–5197 (2019)
76. Yang, Y., Jiang, P.T., Hou, Q., Zhang, H., Chen, J., Li, B.: Multi-task dense prediction via mixture of low-rank experts. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 27927–27937 (2024)
77. Yang, Y., Jiang, P.T., Hou, Q., Zhang, H., Chen, J., Li, B.: Multi-task dense predictions via unleashing the power of diffusion. In: *The Thirteenth International Conference on Learning Representations* (2025)
78. Ye, C., Qiu, L., Gu, X., Zuo, Q., Wu, Y., Dong, Z., Bo, L., Xiu, Y., Han, X.: Stablenormal: Reducing diffusion variance for stable and sharp normal. *ACM Transactions on Graphics (ToG)* **43**(6), 1–18 (2024)
79. Ye, H., Xu, D.: Inverted pyramid multi-task transformer for dense scene understanding. In: *European Conference on Computer Vision*. pp. 514–530. Springer (2022)
80. Ye, H., Xu, D.: Taskexpert: Dynamically assembling multi-task representations with memorial mixture-of-experts. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 21828–21837 (2023)

81. Ye, H., Xu, D.: Taskprompter: Spatial-channel multi-task prompting for dense scene understanding. In: The Eleventh International Conference on Learning Representations (2023)
82. Ye, H., Xu, D.: Diffusionmtl: Learning multi-task denoising diffusion model from partially annotated data. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 27960–27969 (2024)
83. Ye, H., Xu, D.: Invpt++: Inverted pyramid multi-task transformer for visual scene understanding. *IEEE transactions on pattern analysis and machine intelligence* **46**(12), 7493–7508 (2024)
84. Ye, H., Zhang, J., Liu, S., Han, X., Yang, W.: Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721* (2023)
85. Zamir, A.R., Sax, A., Shen, W., Guibas, L.J., Malik, J., Savarese, S.: Taskonomy: Disentangling task transfer learning. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3712–3722 (2018)
86. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 3836–3847 (2023)
87. 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)
88. Zheng, Y., Harley, A.W., Shen, B., Wetzstein, G., Guibas, L.J.: Pointodyssey: A large-scale synthetic dataset for long-term point tracking. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 19855–19865 (2023)
89. Zhou, B., Zhao, H., Puig, X., Fidler, S., Barriuso, A., Torralba, A.: Scene parsing through ade20k dataset. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 633–641 (2017)
90. Zhu, M., Liu, Y., Luo, Z., Jing, C., Chen, H., Xu, G., Wang, X., Shen, C.: Unleashing the potential of the diffusion model in few-shot semantic segmentation. *Advances in Neural Information Processing Systems* **37**, 42672–42695 (2024)