

ViP-DeepLab: Learning Visual Perception with Depth-aware Video Panoptic Segmentation

Siyuan Qiao^{1*} Yukun Zhu² Hartwig Adam² Alan Yuille¹ Liang-Chieh Chen²

¹Johns Hopkins University ²Google Research

Abstract

In this paper, we present ViP-DeepLab, a unified model attempting to tackle the long-standing and challenging inverse projection problem in vision, which we model as restoring the point clouds from perspective image sequences while providing each point with instance-level semantic interpretations. Solving this problem requires the vision models to predict the spatial location, semantic class, and temporally consistent instance label for each 3D point. ViP-DeepLab approaches it by jointly performing monocular depth estimation and video panoptic segmentation. We name this joint task as *Depth-aware Video Panoptic Segmentation*, and propose a new evaluation metric along with two derived datasets for it, which will be made available to the public. On the individual sub-tasks, ViP-DeepLab also achieves state-of-the-art results, outperforming previous methods by 5.1% VPQ on Cityscapes-VPS, ranking 1st on the KITTI monocular depth estimation benchmark, and 1st on KITTI MOTS pedestrian. The datasets and the evaluation codes are made publicly available¹.

1. Introduction

The inverse projection problem, one of the most fundamental problems in vision, refers to the ambiguous mapping from the retinal images to the sources of retinal stimulation. Such a mapping requires retrieving all the visual information about the 3D environment using the limited signals contained in the 2D images [60, 62]. Humans are able to easily establish this mapping by identifying objects, determining their sizes, and reconstructing the 3D scene layout, etc. To endow machines with similar abilities to visually perceive the 3D world, we aim to develop a model to tackle the inverse projection problem.

As a step towards solving the inverse projection, the problem is simplified as restoring the 3D point clouds with semantic understandings from the perspective image se-




Figure 1: Projecting 3D points to the image plane results in 2D images. We study the inverse projection problem: how to restore the 3D points from 2D image sequences while providing temporally consistent instance-level semantic interpretations for the 3D points.

quences, which calls for vision models to predict the spatial location, semantic class, and temporally consistent instance label for each 3D point. Fig. 1 shows an example of the inverse projection problem we study in this paper. This simplified problem can be formulated as Depth-aware Video Panoptic Segmentation (DVPS) that contains two sub-tasks: (i) monocular depth estimation [68], which is used to estimate the spatial position of each 3D point that is projected to the image plane, and (ii) video panoptic segmentation [42], which associates the 3D points with temporally consistent instance-level semantic predictions.

For the new task DVPS, we present two derived datasets accompanied by a new evaluation metric named Depth-aware Video Panoptic Quality (DVPQ). DVPS datasets are hard to collect, as they need special depth sensors and a huge amount of labeling efforts. Existing datasets usually lack some annotations or are not in the format for DVPS. Our solution is to augment and convert existing datasets for DVPS, producing two new datasets, Cityscapes-DVPS and SemKITTI-DVPS. Cityscapes-DVPS is derived from Cityscapes-VPS [42] by adding depth annotations from Cityscapes dataset [18], while SemKITTI-DVPS is derived from SemanticKITTI [6] by projecting its annotated 3D point clouds to the image plane. Additionally, the proposed metric DVPQ includes the metrics for depth estimation and

*Work done while an intern at Google.

¹<https://github.com/joe-siyuan-qiao/ViP-DeepLab>

video panoptic segmentation, requiring a vision model to simultaneously tackle the two sub-tasks. To this end, we present ViP-DeepLab, a unified model that jointly performs video panoptic segmentation and monocular depth estimation for each pixel on the image plane. In the following, we introduce how ViP-DeepLab tackles the two sub-tasks.

The first sub-task of DVPS is video panoptic segmentation [42]. Panoptic segmentation [43] unifies semantic segmentation [36] and instance segmentation [34] by assigning every pixel a semantic label and an instance ID. It has been recently extended to the video domain, resulting in video panoptic segmentation [42], which further demands each instance to have the same instance ID throughout the video sequence. This poses additional challenges to panoptic segmentation as the model is now expected to be able to track objects in addition to detecting and segmenting them. Current approach VPSNet [42] adds a tracking head to learn the correspondence between the instances from different frames based on their regional feature similarity. By contrast, our ViP-DeepLab takes a different approach to tracking objects. Specifically, motivated by our finding that video panoptic segmentation can be modeled as concatenated image panoptic segmentation, we extend Panoptic-DeepLab [17] to perform center regression for *two* consecutive frames with respect to *only* the object centers that appear in the *first* frame. During inference, this offset prediction allows ViP-DeepLab to group all the pixels in the two frames to the same object that appears in the first frame. New instances emerge if they are not grouped to the previously detected instances. This inference process continues for every two consecutive frames (with one overlapping frame) in a video sequence, stitching panoptic predictions together to form predictions with temporally consistent instance IDs. Based on this simple design, our ViP-DeepLab outperforms VPSNet [42] by a large margin of 5.1% VPQ, setting a new record on the Cityscapes-VPS dataset [42]. Additionally, Multi-Object Tracking and Segmentation (MOTS) [77] is a similar task to video panoptic segmentation, but only segments and tracks two classes: pedestrians and cars. We therefore also apply our ViP-DeepLab to MOTS. As a result, ViP-DeepLab outperforms the current state-of-the-art PointTrack [92] by 7.2% and 2.5% sMOTSA on pedestrians and cars, respectively, and ranks 1st on the leaderboard for KITTI MOTS pedestrian.

The second sub-task of DVPS is monocular depth estimation, which is challenging for both computers [68] and humans [38]. The state-of-the-art methods are mostly based on deep networks trained in a fully-supervised way [20, 21, 22, 26]. Following the same direction, our ViP-DeepLab appends another depth prediction head on top of Panoptic-DeepLab [17]. Without using any additional *depth* training data, such a simple approach outperforms all the published and unpublished works on the KITTI benchmark [30].

Specifically, it outperforms DORN [26] by 0.97 SIlog, and even outperforms MPSD that uses extra planet-scale depth data [2], breaking the long-standing record on the challenging KITTI depth estimation [74]. Notably, the differences between top-performing methods are all around 0.1 SIlog, while our method significantly outperforms them.

To summarize, our contributions are listed as follows.

- We propose a new task Depth-aware Video Panoptic Segmentation (DVPS), as a step towards solving the inverse projection problem by formulating it as joint video panoptic segmentation [42] and monocular depth estimation [68].
- We present two DVPS datasets along with an evaluation metric Depth-aware Video Panoptic Quality (DVPQ). To facilitate future research, the datasets and the evaluation codes are made publicly available.
- We develop ViP-DeepLab, a unified model for DVPS. On the individual sub-tasks, ViP-DeepLab ranks 1st on Cityscapes-VPS [42], KITTI-MOTS pedestrian [77], and KITTI monocular depth estimation [30].

2. Related Work

Panoptic Segmentation Recent methods for image panoptic segmentation can be grouped into two types: top-down (proposal-based) methods and bottom-up (box-free) methods. Top-down methods employ a two-stage approach which generates object proposals followed by outputting panoptic predictions based on regional computations [16, 45, 50, 51, 52, 65, 71, 85, 88]. For example, Panoptic FPN [43] incorporates a semantic segmentation head into Mask R-CNN [35]. Porzi *et al.* [63] proposes a novel segmentation head to integrate FPN [53] features by a lightweight DeepLab-like module [13]. Bottom-up panoptic segmentation methods group pixels to form instances on top of semantic segmentation prediction [79, 80, 95]. For example, SSAP [28] uses pixel-pair affinity pyramid [54] and a cascaded graph partition module [41] to generate instances from coarse to fine. BBFNet [9] uses Hough-voting [5, 48] and Watershed transform [4, 76] to generate instance segmentation predictions. Panoptic-DeepLab [17] employs class-agnostic instance center regression [40, 59, 73] on top of semantic segmentation outputs from DeepLab [12, 14].

Object Tracking One of the major tasks in video panoptic segmentation is object tracking. Many trackers use tracking-by-detection, which divides the task into two sub-tasks where an object detector (*e.g.* [25, 66]) finds all objects and then an algorithm associates them [8, 23, 46, 64, 69, 70, 72, 84, 91, 105]. Another design is transforming object detectors to object trackers which detect and track objects at the same time [7, 24, 61, 82, 101, 102]. For example, CenterTrack [103] extends CenterNet [104] to predict offsets from the object center to its center in the previous frame. STEM-Seg [3] proposes to group all instance pixels in a video clip

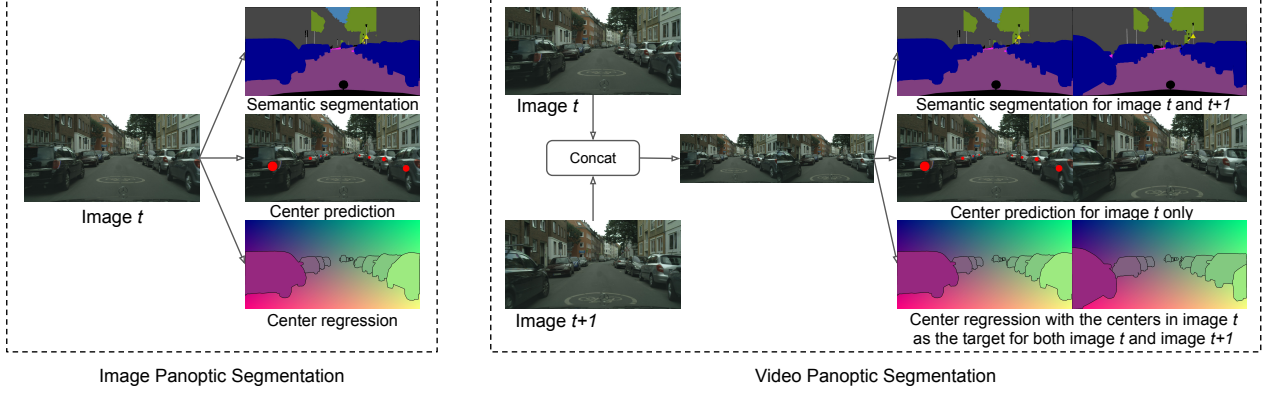


Figure 2: Comparing image panoptic segmentation and video panoptic segmentation. Our method is based on the finding that video panoptic segmentation can be modeled as concatenated image panoptic segmentation. Center regression is an offset map from each pixel to its object center. Here we draw the predicted centers instead of the offsets for clearer visualization.

by learning a spatio-temporal embedding. By contrast, our ViP-DeepLab implicitly performs object tracking by clustering all instance pixels in two consecutive video frames. Additionally, our method simply uses center regression and achieves better results on MOTs [77].

Monocular Depth Estimation Monocular depth estimation predicts depth from a single image. It can be learned in a supervised way [10, 21, 26, 27, 47, 49, 68, 90, 96], by reconstructing images in the stereo setting [29, 31, 32, 44, 87], from videos [57, 78, 97], in relative order [15], or by joint learning with other vision tasks [75, 81, 89]. ViP-DeepLab models monocular depth estimation as a dense regression problem, and we train it in a fully-supervised manner.

3. ViP-DeepLab

In this section, we present ViP-DeepLab, which extends Panoptic-DeepLab [17] to jointly perform video panoptic segmentation [42] and monocular depth estimation [68].

3.1. Video Panoptic Segmentation

Rethinking Image and Video Panoptic Segmentation In the task of video panoptic segmentation, each instance is represented by a tube on the image plane and the time axis when the frames are stacked up. Given a clip $I^{t:t+k}$ with time window k , true positive (TP) is defined by $TP = \{(u, \hat{u}) \in U \times \hat{U} : \text{IoU}(u, \hat{u}) > 0.5\}$ where U and \hat{U} are the set of the ground-truth and predicted tubes, respectively. False positives (FP) and false negatives (FN) are defined accordingly. After accumulating the TP_c , FP_c , and FN_c on all the clips with window size k and class c , the evaluation metric Video Panoptic Quality (VPQ) [42] is defined by

$$\text{VPQ}^k = \frac{1}{N_{\text{classes}}} \sum_c \frac{\sum_{(u, \hat{u}) \in TP_c} \text{IoU}(u, \hat{u})}{|TP_c| + \frac{1}{2}|FP_c| + \frac{1}{2}|FN_c|} \quad (1)$$

PQ [43] is thus equal to VPQ^1 (i.e., $k = 1$).

Our method is based on the connection between PQ and VPQ. For an image sequence I_t ($t = 1, \dots, T$), let P_t denote the panoptic prediction and Q_t be the ground-truth panoptic segmentation. As VPQ^k accumulates the PQ-related statistics from P_t and Q_t within a window of size k , we have

$$\text{VPQ}^k \left(\left[P_t, Q_t \right]_{t=1}^T \right) = \text{PQ} \left(\left[\parallel_{i=t}^{t+k-1} P_i, \parallel_{i=t}^{t+k-1} Q_i \right]_{t=1}^{T-k+1} \right) \quad (2)$$

where $\parallel_{i=t}^{t+k-1} P_i$ denotes the horizontal concatenation of P_i from t to $t+k-1$, and $\left[P_t, Q_t \right]_{t=1}^T$ denotes a list of pairs of (P_t, Q_t) from 1 to T as the function input.

Equ. (2) reveals an interesting finding that video panoptic segmentation could be formulated as image panoptic segmentation with the images concatenated. Such a finding motivates us to extend image panoptic segmentation models to video panoptic segmentation with extra modifications.

From Image to Video Panoptic Segmentation Panoptic-DeepLab [17] approaches the problem of image panoptic segmentation by solving three sub-tasks: (1) semantic predictions for both ‘thing’ and ‘stuff’ classes, (2) center prediction for each instance of ‘thing’ classes, and (3) center regression for each pixel of objects. Fig. 2 shows an example of the tasks on the left. During inference, object centers with high confidence scores are kept, and each ‘thing’ pixel is associated with the closest object center to form object instances. Combining this ‘thing’ prediction and the ‘stuff’ prediction from semantic segmentation, Panoptic-DeepLab [17] generates the final panoptic prediction.

Our method extends Panoptic-DeepLab [17] to perform video panoptic segmentation. As the right part of Fig. 2 shows, it also breaks down the task of video panoptic segmentation into three sub-tasks: semantic segmentation, center prediction, and center regression. During inference, our method takes image t and $t+1$ concatenated horizontally

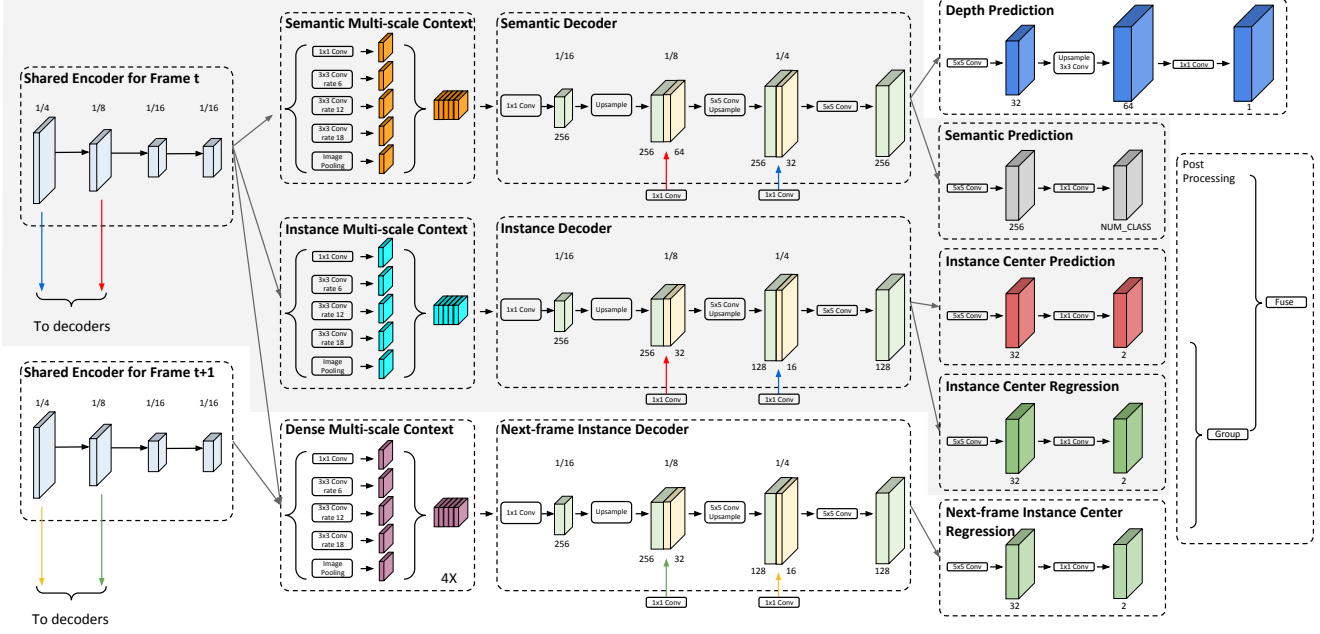


Figure 3: ViP-DeepLab extends Panoptic-DeepLab [17] (the gray part) by adding a depth prediction head to perform monocular depth estimation and a next-frame instance branch which regresses to the object centers in frame t for frame $t + 1$.

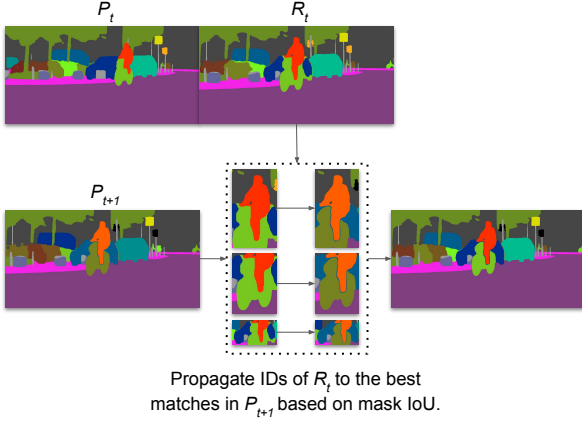


Figure 4: Visualization of stitching video panoptic predictions. It propagates IDs based on mask IoU between region pairs. ViP-DeepLab is capable of tracking objects with large movements, *e.g.*, the cyclist in the image. Panoptic prediction of R_t is of high quality, which is why a simple IoU-based stitching method works well in practice.

as input, and only predict the centers in image t . The center regression for both t and $t + 1$ will regress to the object centers in image t . By doing so, our method detects the objects in the first frame, and finds all the pixels belonging to them in the first and the second frames. Objects that appear only in the second frame are ignored here and will emerge again when the model works on the next image pair (*i.e.*, $(t + 1, t + 2)$). Our method models video panoptic segmen-

tation as concatenated image panoptic segmentation, highly consistent with the definition of the metric VPQ.

Fig. 3 shows the architecture of our method. In order to perform the inference described above, we take image t and $t + 1$ as the input during training, and use the features of image t to predict semantic segmentation, object centers and center offsets for image t . In addition to that, we add a *next-frame instance branch* which predicts the center offsets for the pixels in image $t + 1$ with respect to the centers in image t . The backbone features of image t and $t + 1$ are concatenated along the feature axis before the *next-frame instance branch*. As their backbone features are separated before concatenation, the *next-frame instance branch* needs a large receptive field to perform long-range center regression. To address this, we use four ASPP modules in the branch, the output of which are densely-connected [39, 94] to dramatically increase the receptive field. We name this densely-connected module as Cascade-ASPP. Its architecture details are shown in Appendix. Finally, the decoder in the *next-frame instance branch* uses the backbone features of image $t + 1$ while the other branches use those of image t , as indicated by the colored arrows in the figure.

Stitching Video Panoptic Predictions Our method outputs panoptic predictions with temporally consistent IDs for two consecutive frames. To generate predictions for the entire sequence, we need to stitch the panoptic predictions. Fig 4 shows an example of our stitching method. For each image pair t and $t + 1$, we split the panoptic prediction of

the concatenated input in the middle, and use P_t to denote the left prediction, and R_t to denote the right one. By doing so, P_t becomes the panoptic prediction of image t , and R_t becomes the panoptic prediction of image $t + 1$ with instance IDs that are consistent with those of P_t . The goal of stitching is to propagate IDs from R_t to P_{t+1} so that each object in P_t and P_{t+1} will have the same ID.

The ID propagation is based on mask IoU between region pairs. For each region pair in R_t and P_{t+1} , if they have the same class, and both find each other to have the largest mask IoU, then we propagate the ID between them. Objects that do not receive IDs will become new instances. A formal algorithm can be found in Appendix.

3.2. Monocular Depth Estimation

We model monocular depth estimation as a dense regression problem [22], where each pixel will have an estimated depth. As shown in Fig. 3, we add a depth prediction head on top of the decoded features of the semantic branch (*i.e.*, Semantic Decoder), which upsamples the features by 2x and generates logits f_d for depth regression:

$$\text{Depth} = \text{MaxDepth} \times \text{Sigmoid}(f_d) \quad (3)$$

MaxDepth controls the range of the predicted depth, which is set to 88 for the range (about 0 to 80m) of KITTI [74].

Many metrics have been proposed to evaluate the quality of monocular depth prediction [30]. Among them, scale invariant logarithmic error [22] and relative squared error [30] are popular ones, which could also be directly optimized as training loss functions. We therefore combine them to train our depth prediction. Specifically, let d and \hat{d} denote the ground-truth and the predicted depth, respectively. Our loss function for depth estimation is then defined by

$$\mathcal{L}_{\text{depth}}(d, \hat{d}) = \frac{1}{n} \sum_i \left(\log d_i - \log \hat{d}_i \right)^2 - \frac{1}{n^2} \left(\sum_i \log d_i - \log \hat{d}_i \right)^2 + \left(\frac{1}{n} \sum_i \left(\frac{d_i - \hat{d}_i}{d_i} \right)^2 \right)^{0.5} \quad (4)$$

3.3. Depth-aware Video Panoptic Segmentation

Motivated by solving the inverse projection problem, we introduce a challenging task, Depth-aware Video Panoptic Segmentation (DVPS), unifying the problems of monocular depth estimation and video panoptic segmentation. In the task of DVPS, images are densely annotated with a tuple (c, id, d) for each labeled pixel, where c , id and d denote its semantic class, instance ID and depth. The model is expected to also generate a tuple $(\hat{c}, \hat{id}, \hat{d})$ for each pixel.

To evaluate methods for DVPS, we propose a metric called Depth-aware Video Panoptic Quality (DVPQ), which extends VPQ by additionally considering the depth prediction with the inlier metric. Specifically, let P and Q be the

prediction and ground-truth, respectively. We use P_i^c , P_i^{id} and P_i^d to denote the predictions of example i on the semantic class, instance ID, and depth. The notations also apply to Q . Let k be the window size (as in Equ. (2)) and λ be the depth threshold. Then, $\text{DVPQ}_\lambda^k(P, Q)$ is defined by

$$\text{PQ} \left(\left[\left\|_{i=t}^{t+k-1} (\hat{P}_i^c, P_i^{id}), \left\|_{i=t}^{t+k-1} (Q_i^c, Q_i^{id}) \right\|_{t=1}^{T-k+1} \right] \right) \quad (5)$$

where $\hat{P}_i^c = P_i^c$ for pixels that have *absolute relative* depth errors under λ (*i.e.*, $|P_i^d - Q_i^d| \leq \lambda Q_i^d$), and will be assigned a void label otherwise. In other words, \hat{P}_i^c filters out pixels that have large absolute relative depth errors. As a result, the metrics VPQ [42] (also image PQ [43]) and depth inlier metric (*i.e.*, $\max(P_i^d/Q_i^d, Q_i^d/P_i^d) = \delta < \text{threshold}$) [22] can be *approximately* viewed as special cases for DVPQ.

Following [42], we evaluate DVPQ_λ^k for four different values of k (depending on the dataset) and three values of $\lambda = \{0.1, 0.25, 0.5\}$. Those values of λ approximately correspond to the depth inlier metric $\delta < 1.1$, $\delta < 1.25$, and $\delta < 1.5$, respectively. They are harder than the thresholds 1.25, 1.25^2 and 1.25^3 that are commonly used in depth evaluation. We choose harder thresholds as many methods are able to get $> 99\%$ on the previous metrics [26, 47]. Larger k and smaller λ correspond to a higher accuracy requirement for a long-term consistency of joint video panoptic segmentation and depth estimation. The final number DVPQ is obtained by averaging over all values of k and λ .

4. Datasets

To evaluate on the new task, Depth-aware Video Panoptic Segmentation, we create two new datasets, Cityscapes-DVPS and SemKITTI-DVPS. Fig. 5 shows two examples, one for each dataset. The details are elaborated below.

4.1. Cityscapes-DVPS

The original Cityscapes [18] only contains image-level panoptic annotations. Recently, Kim *et al.* introduce a video panoptic segmentation dataset Cityscapes-VPS [42] by further annotating 6 frames out of each 30-frame video sequence (with a gap of 5 frames between each annotation), resulting in totally 3,000 annotated frames where the training, validation, and test sets have 2,400, 300, and 300 frames, respectively. In the dataset, there are 19 semantic classes, including 8 ‘thing’ and 11 ‘stuff’ classes.

Even though Cityscapes-VPS contains video panoptic annotations, the depth annotations are missing. We find that the depth annotations could be converted from the disparity maps via stereo images, provided by the original Cityscapes dataset [18]. However, the quality of the pre-computed disparity maps is not satisfactory. To improve it, we select several modern disparity estimation methods [33, 37, 99, 100]




Figure 5: Dataset examples of Cityscapes-DVPS (top) and SemKITTI-DVPS (bottom). From left to right: input image, video panoptic segmentation annotation, and depth map. Regions are black if they are not covered by the velodyne data or they are removed by the data preprocessing step including disparity consistency check and non-foreground suppression.




Figure 6: Top: Removing occluded but falsely visible points highlighted in red by disparity consistency check. Bottom: Removing the invading background points in pink for the thin object colored yellow by non-foreground suppression.

and follow the process similar to [18]. Nevertheless, to discourage reproducing the depth generation process (so that one may game the benchmark), we do not disclose the details (*e.g.*, the exact employed disparity method). The depth annotations are made publicly available.

4.2. SemKITTI-DVPS

SemanticKITTI dataset [6] is based on the odometry dataset of the KITTI Vision benchmark [30]. The dataset splits 22 sequences in to 11 training sequences and 11 test sequences. The training sequence 08 is used for validation. This dataset includes 8 ‘thing’ and 11 ‘stuff’ classes.

SemanticKITTI dataset provides perspective images and panoptic-labeled 3D point clouds (*i.e.*, semantic class and instance ID are annotated). To convert it for our use, we project the 3D point clouds into the image plane. However, there are two challenges when converting the dataset, as presented in Fig. 6. The first problem is that some point clouds are not visible to the camera but are recorded and labeled. For example, the first row of Fig. 6 shows that some regions behind the car become visible in the converted depth map due to the alignment of different sensors. To address this issue, we follow Uhrig *et al.* [74] and use the same disparity methods for Cityscapes-DVPS to remove the

sampled points that exhibit large relative errors, which are highlighted in red in the right figure. We refer to this processing as disparity consistency check. The second problem is that the regions of thin objects (*e.g.*, poles) are usually invaded by the far-away background point cloud after projection. To alleviate this problem, for a small image patch, the projected background points are removed if there exists at least one foreground point that is closer to the camera. We refer to this processing as non-foreground suppression. In practice, we use a small 7×7 image patch. Doing so leaves clear boundaries for thin objects so they can be identified without confusion as shown in the second row of Fig. 6.

5. Experiments

In this section, we first present our major results on the new task Depth-aware Video Panoptic Segmentation. Then, we show our method applied to three sub-tasks, including video panoptic segmentation [42], monocular depth estimation [30], and multi-object tracking and segmentation [77].

5.1. Depth-aware Video Panoptic Segmentation

Tab. 1 shows our results on Depth-aware Video Panoptic Segmentation. We evaluate our method on the datasets Cityscapes-DVPS and SemKITTI-DVPS so that the research community can compare their methods with it. The evaluation is based on our proposed $DVPQ_{\lambda}^k$ metric (Equ. (5)), where λ is the threshold of relative depth error, and k denotes the length of the short video clip used in evaluation. The training time is about 1 day using 32 TPUs.

Following [42], we set $k = \{1, 2, 3, 4\}$ out of the total 6 frames per video sequence for Cityscapes-DVPS. By contrast, we set $k = \{1, 5, 10, 20\}$ for SemKITTI-DVPS which contains much longer video sequences, and we aim to evaluate a longer temporal consistency. We study the drops of $DVPQ_{\lambda}^k$ as the number of frames k increases, where smaller performance drops indicate higher temporal consistency. Interestingly, as the number of frames k increases, the performance drops on SemKITTI-DVPS are smaller than that on Cityscapes-DVPS. For example, $DVPQ_{0.5}^1 - DVPQ_{0.5}^2$

DVPQ $_{\lambda}^k$ on Cityscapes-DVPS	k = 1			k = 2			k = 3			k = 4			Average		
$\lambda = 0.50$	68.7	61.4	74.0	61.7	48.5	71.3	58.4	42.1	70.2	56.3	38.0	69.5	61.3	47.5	71.2
$\lambda = 0.25$	66.5	60.4	71.0	59.5	47.6	68.2	56.2	41.3	67.1	54.2	37.3	66.5	59.1	46.7	68.2
$\lambda = 0.10$	50.5	45.8	53.9	45.6	36.9	51.9	42.6	31.7	50.6	40.8	28.4	49.8	44.9	35.7	51.5
Average	61.9	55.9	66.3	55.6	44.3	63.8	52.4	38.4	62.6	50.4	34.6	61.9	55.1	43.3	63.6

DVPQ $_{\lambda}^k$ on SemKITTI-DVPS	k = 1			k = 5			k = 10			k = 20			Average		
$\lambda = 0.50$	54.7	46.4	60.6	51.5	41.0	59.1	50.1	38.5	58.5	49.2	36.9	58.2	51.4	40.7	59.1
$\lambda = 0.25$	52.0	44.8	57.3	48.8	39.4	55.7	47.4	37.0	55.1	46.6	35.6	54.7	48.7	39.2	55.7
$\lambda = 0.10$	40.0	34.7	43.8	37.1	30.3	42.0	35.8	28.3	41.2	34.5	26.5	40.4	36.8	30.0	41.9
Average	48.9	42.0	53.9	45.8	36.9	52.3	44.4	34.6	51.6	43.4	33.0	51.1	45.6	36.6	52.2

Table 1: ViP-DeepLab performance for the task of *Depth-aware Video Panoptic Segmentation* (DVPS) evaluated on Cityscapes-DVPS and SemKITTI-DVPS. Each cell shows DVPQ $_{\lambda}^k$ | DVPQ $_{\lambda}^k$ -Thing | DVPQ $_{\lambda}^k$ -Stuff where λ is the threshold of relative depth error, and k is the number of frames. Smaller λ and larger k correspond to a higher accuracy requirement.

on Cityscapes-DVPS is 7%, while DVPQ $_{0.5}^1$ - DVPQ $_{0.5}^5$ on SemKITTI-DVPS is 3.2%. We speculate that this is because the annotation frame rate is higher on SemKITTI-DVPS (*cf.* only every 5th frame is annotated on Cityscapes-DVPS), making our ViP-DeepLab’s offsets prediction easier for the following frames, despite the evaluation clip length k is larger. At last, we use the mean of DVPQ $_{\lambda}^k$ with different λ and k as the final performance score. Fig. 7 visualizes the predictions of our method on the validation set of Cityscapes-DVPS (top) and SemKITTI-DVPS (bottom), where the second column shows P_t and R_t defined in prediction stitching. Although the training samples of SemKITTI-DVPS are sparse points, our method is able to predict smooth and sharp predictions, as the points are evenly distributed in the regions covered by the velodyne data. Please see the supplementary material for more visualization results. After experimenting on DVPS, we compare ViP-DeepLab with the previous state-of-the-arts on the sub-tasks to showcase its strong performance.

5.2. Video Panoptic Segmentation

The first sub-task of DVPS is Video Panoptic Segmentation (VPS). We conduct experiments on Cityscapes-VPS following the setting of [42] with the depth head removed. Tab. 2 shows our major results on their validation set (top) and the test set where the annotations are missing (bottom). As the table shows, our method outperforms VPSNet [42] by 5.6% VPQ on the validation set and 5.1% VPQ on the test set. STQ [83] comparisons are shown in Appendix.

Tab. 3 shows the ablation study on Cityscapes-VPS. The baseline is our method with backbone WR-41 [11, 86, 98] pre-trained on ImageNet [67]. Next, ‘MV’ initializes the model with a checkpoint pretrained on Mapillary Vistas [58]. ‘CS’ uses a model further pretrained on Cityscapes videos with pseudo labels [11] on the *train* sequence. Both ‘MV’ and ‘CS’ only involve image panoptic segmentation pretraining. Hence, they mainly improves image PQ

Val set	VPSNet [42]			ViP-DeepLab		
k = 1	65.0	59.0	69.4	70.4	63.2	75.7
k = 2	57.6	45.1	66.7	63.6	50.7	73.0
k = 3	54.4	39.2	65.6	60.1	44.0	71.9
k = 4	52.8	35.8	65.3	58.1	40.2	71.2
VPQ	57.5	44.8	66.7	63.1	49.5	73.0

Test set	VPSNet [42]			ViP-DeepLab		
k = 1	64.2	59.0	67.7	68.9	61.6	73.5
k = 2	57.9	46.5	65.1	62.9	51.0	70.5
k = 3	54.8	41.1	63.4	59.9	46.0	68.8
k = 4	52.6	36.5	62.9	58.2	42.1	68.4
VPQ	57.4	45.8	64.8	62.5	50.2	70.3

Table 2: VPQ on Cityscapes-VPS. Each cell shows VPQ k | VPQ k -Thing | VPQ k -Stuff. VPQ is averaged over $k = \{1, 2, 3, 4\}$. $k = \{0, 5, 10, 15\}$ in [42] correspond to $k = \{1, 2, 3, 4\}$ in this paper as we use different notations.

Method	k = 1	k = 2	k = 3	k = 4	VPQ
Baseline	65.7	58.9	55.8	53.6	58.5
+ MV	66.7	59.3	56.1	54.1	59.0
+ CS	67.9	60.4	56.8	54.7	59.9
+ DenseContext	68.2	61.3	58.2	56.1	60.9
+ AutoAug [19]	68.6	61.6	58.6	56.3	61.3
+ RFP [65]	69.2	62.3	59.2	57.0	61.9
+ TTA	70.3	63.2	59.9	57.5	62.7
+ SSL	70.4	63.6	60.1	58.1	63.1

Table 3: Ablation Study on Cityscapes-VPS.

(*i.e.* $k = 1$) but increases the gaps between VPQ k (*e.g.*, VPQ 1 - VPQ 2), showing the temporal consistency benefits less from the pretrained models. Then, ‘DenseContext’ increases the number of the context modules (from 1 to 4) for the next-frame instance branch, which narrows down the gaps between VPQ k . ‘AutoAug’ uses AutoAugment [19] to augment the data. ‘RFP’ adds Recursive Feature Pyramid (RFP) [65] to enhance the backbone. ‘TTA’ stands for test-time augmentation, which includes multi-scale inference at




Figure 7: Prediction visualizations on Cityscapes-DVPS (top) and SemKITTI-DVPS (bottom). From left to right: input image, temporally consistent panoptic segmentation prediction, monocular depth prediction, and point cloud visualization.

Method	Rank	SILog	sqRel	absRel	iRMSE
DORN [26]	10	11.77	2.23	8.78	12.98
BTS [47]	9	11.67	2.21	9.04	12.23
BANet [1]	8	11.61	2.29	9.38	12.23
MPSD [2]	2	11.12	2.07	8.99	11.56
ViP-DeepLab	1	10.80	2.19	8.94	11.77

Table 4: KITTI Depth Prediction Leaderboard. Ranking includes published and unpublished methods.

scales 0.5:1.75:0.25 and horizontal flipping. In ‘SSL’, we follow Naive-Student [11] to generate temporally consistent pseudo labels on the unlabeled *train* sequence in Cityscapes videos [18], which adds more training samples for temporal consistency, as demonstrated by +0.1% on VPQ¹ and +0.6% on VPQ⁴.

5.3. Monocular Depth Estimation

The second sub-task of DVPS is monocular depth estimation. We test our method on the KITTI depth benchmark [74]. Tab. 4 shows the results on the leaderboard. Our model is pretrained on Mapillary Vistas [58] and Cityscapes videos with pseudo labels [11] (*i.e.*, the same pretrained checkpoint we used in the previous experiments). Then the model is fine-tuned with the training and validation set provided by KITTI depth benchmark [74]. However, the model is slightly different from the previous ones in the following aspects. It does not use RFP [65]. In TTA, it only has horizontal flipping. We use ± 5 degrees of random rotation during training, which improves SILog by 0.27. The previous models use a decoder with stride 8 and 4. Here, we find it useful to further exploit decoder stride 2, which improves SILog by 0.17. After the above changes, our method achieves the best results on KITTI depth benchmark.

5.4. Multi-Object Tracking and Segmentation

Finally, we evaluate our method on the KITTI MOTs benchmark [77]. Tab. 5 shows the leaderboard results. Different from the previous experiments, this benchmark only

Method	Pedestrians			Cars		
	Rank	sMOTSA	MOTSA	Rank	sMOTSA	MOTSA
TrackR-CNN [77]	20	47.3	66.1	19	67.0	79.6
MOTSFusion [55]	13	58.7	72.9	12	75.0	84.1
PointTrack [92]	11	61.5	76.5	5	78.5	90.9
ReMOTS [93]	6	66.0	81.3	9	75.9	86.7
ViP-DeepLab		67.7	83.4		80.6	90.3
ViP-DeepLab + KF	1	68.7	84.5	3	81.0	90.7

Table 5: KITTI MOTs Leaderboard. Ranking includes published and unpublished methods.

tracks pedestrians and cars. Adopting the same strategy as we used for Cityscapes-VPS, ViP-DeepLab outperforms all the published methods and achieves 67.7% and 80.6% sMOTSA for pedestrians and cars, respectively. To further improve our results, we use Kalman filter (KF) [82] to re-localize missing objects that are occluded or detection failures. This mechanism improves the sMOTSA by 1.0% and 0.4% for pedestrians and cars, respectively.

6. Conclusion

In this paper, we propose a new challenging task Depth-aware Video Panoptic Segmentation, which combines monocular depth estimation and video panoptic segmentation, as a step towards solving the inverse projection problem in vision. For this task, we propose Depth-aware Video Panoptic Quality as the evaluation metric along with two derived datasets. We present ViP-DeepLab as a strong baseline for this task. Additionally, our ViP-DeepLab also achieves state-of-the-art performances on several sub-tasks, including monocular depth estimation, video panoptic segmentation, and multi-object tracking and segmentation.

Acknowledgments This work is supported by Google Faculty Research Award. We thank Maxwell Collins for the feedback, and the support from Google Mobile Vision team. We thank Mark Weber for the help with STQ [83], and the authors of [42] for the support on their benchmark. Siyuan Qiao thanks Wanyu Huang for the support.

References

- [1] Shubhra Aich, Jean Marie Uwabeza Vianney, Md Amirul Islam, Mannat Kaur, and Bingbing Liu. Bidirectional attention network for monocular depth estimation. *arXiv preprint arXiv:2009.00743*, 2020. 8
- [2] Manuel López Antequera, Pau Gargallo, Markus Hofinger, and Samuel Rota. Mapillary planet-scale depth dataset. In *Proceedings of the European conference on computer vision (ECCV)*, 2020. 2, 8
- [3] Ali Athar, Sabarinath Mahadevan, Aljoša Ošep, Laura Leal-Taixé, and Bastian Leibe. Stem-seg: Spatio-temporal embeddings for instance segmentation in videos. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 2
- [4] Min Bai and Raquel Urtasun. Deep watershed transform for instance segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 2
- [5] Dana H Ballard. Generalizing the hough transform to detect arbitrary shapes. *Pattern recognition*, 13(2), 1981. 2
- [6] Jens Behley, Martin Garbade, Andres Milioto, Jan Quenzel, Sven Behnke, Cyrill Stachniss, and Jurgen Gall. Semantickitti: A dataset for semantic scene understanding of lidar sequences. In *Proceedings of the IEEE International Conference on Computer Vision*, 2019. 1, 6
- [7] Philipp Bergmann, Tim Meinhardt, and Laura Leal-Taixe. Tracking without bells and whistles. In *Proceedings of the IEEE international conference on computer vision*, 2019. 2
- [8] Alex Bewley, Zongyuan Ge, Lionel Ott, Fabio Ramos, and Ben Upcroft. Simple online and realtime tracking. In *2016 IEEE International Conference on Image Processing (ICIP)*. IEEE, 2016. 2
- [9] Ujwal Bonde, Pablo F Alcantarilla, and Stefan Leutenegger. Towards bounding-box free panoptic segmentation. *arXiv preprint arXiv:2002.07705*, 2020. 2
- [10] Yuanzhouhan Cao, Zifeng Wu, and Chunhua Shen. Estimating depth from monocular images as classification using deep fully convolutional residual networks. *IEEE Transactions on Circuits and Systems for Video Technology*, 28(11), 2017. 3
- [11] Liang-Chieh Chen, Raphael Gontijo Lopes, Bowen Cheng, Maxwell D Collins, Ekin D Cubuk, Barret Zoph, Hartwig Adam, and Jonathon Shlens. Naive-student: Leveraging semi-supervised learning in video sequences for urban scene segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, 2020. 7, 8
- [12] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Semantic image segmentation with deep convolutional nets and fully connected crfs. In *Proceedings of the International Conference on Learning Representations*, 2015. 2
- [13] 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. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017. 2
- [14] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 2
- [15] Weifeng Chen, Zhao Fu, Dawei Yang, and Jia Deng. Single-image depth perception in the wild. In *Advances in neural information processing systems*, 2016. 3
- [16] Yifeng Chen, Guangchen Lin, Songyuan Li, Omar Bourahla, Yiming Wu, Fangfang Wang, Junyi Feng, Mingliang Xu, and Xi Li. Banet: Bidirectional aggregation network with occlusion handling for panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3793–3802, 2020. 2
- [17] Bowen Cheng, Maxwell D Collins, Yukun Zhu, Ting Liu, Thomas S Huang, Hartwig Adam, and Liang-Chieh Chen. Panoptic-deeplab: A simple, strong, and fast baseline for bottom-up panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020. 2, 3, 4
- [18] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016. 1, 5, 6, 8
- [19] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation policies from data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019. 7
- [20] Raul Diaz and Amit Marathe. Soft labels for ordinal regression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4738–4747, 2019. 2, 14
- [21] David Eigen and Rob Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *Proceedings of the IEEE international conference on computer vision*, 2015. 2, 3
- [22] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *Advances in neural information processing systems*, 2014. 2, 5
- [23] Kuan Fang, Yu Xiang, Xiaocheng Li, and Silvio Savarese. Recurrent autoregressive networks for online multi-object tracking. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 2018. 2
- [24] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Detect to track and track to detect. In *Proceedings of the IEEE International Conference on Computer Vision*, 2017. 2
- [25] Pedro F Felzenszwalb, Ross B Girshick, David McAllester, and Deva Ramanan. Object detection with discriminatively trained part-based models. *IEEE transactions on pattern analysis and machine intelligence*, 32(9), 2009. 2
- [26] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep ordinal regression

- network for monocular depth estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 2, 3, 5, 8
- [27] Yukang Gan, Xiangyu Xu, Wenxiu Sun, and Liang Lin. Monocular depth estimation with affinity, vertical pooling, and label enhancement. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 3
- [28] Naiyu Gao, Yanhu Shan, Yupei Wang, Xin Zhao, Yinan Yu, Ming Yang, and Kaiqi Huang. Ssap: Single-shot instance segmentation with affinity pyramid. In *Proceedings of the IEEE International Conference on Computer Vision*, 2019. 2
- [29] Ravi Garg, Vijay Kumar Bg, Gustavo Carneiro, and Ian Reid. Unsupervised cnn for single view depth estimation: Geometry to the rescue. In *European conference on computer vision*. Springer, 2016. 3
- [30] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2012. 2, 5, 6
- [31] Clément Godard, Oisín Mac Aodha, and Gabriel J Brostow. Unsupervised monocular depth estimation with left-right consistency. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 3
- [32] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J Brostow. Digging into self-supervised monocular depth estimation. In *Proceedings of the IEEE international conference on computer vision*, 2019. 3
- [33] Xiaodong Gu, Zhiwen Fan, Siyu Zhu, Zuozhuo Dai, Feitong Tan, and Ping Tan. Cascade cost volume for high-resolution multi-view stereo and stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2495–2504, 2020. 5
- [34] Bharath Hariharan, Pablo Arbeláez, Ross Girshick, and Jitendra Malik. Simultaneous detection and segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2014. 2
- [35] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, 2017. 2
- [36] Xuming He, Richard S Zemel, and Miguel Á Carreira-Perpiñán. Multiscale conditional random fields for image labeling. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2004. 2
- [37] Heiko Hirschmüller. Stereo processing by semiglobal matching and mutual information. *IEEE Transactions on pattern analysis and machine intelligence*, 30(2):328–341, 2007. 5
- [38] Ian P Howard. *Perceiving in depth, volume 1: Basic mechanisms*. Oxford University Press, 2012. 2
- [39] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017. 4
- [40] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018. 2
- [41] Margret Keuper, Evgeny Levinkov, Nicolas Bonneel, Guillaume Lavoué, Thomas Brox, and Bjorn Andres. Efficient decomposition of image and mesh graphs by lifted multi-cuts. In *Proceedings of the IEEE International Conference on Computer Vision*, 2015. 2
- [42] Dahun Kim, Sanghyun Woo, Joon-Young Lee, and In So Kweon. Video panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9859–9868, 2020. 1, 2, 3, 5, 6, 7, 8, 14
- [43] Alexander Kirillov, Kaiming He, Ross Girshick, Carsten Rother, and Piotr Dollár. Panoptic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9404–9413, 2019. 2, 3, 5
- [44] Yevhen Kuznetsov, Jorg Stuckler, and Bastian Leibe. Semi-supervised deep learning for monocular depth map prediction. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017. 3
- [45] Justin Lazarow, Kwonjoon Lee, Kunyu Shi, and Zhuowen Tu. Learning instance occlusion for panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10720–10729, 2020. 2
- [46] Laura Leal-Taixé, Cristian Canton-Ferrer, and Konrad Schindler. Learning by tracking: Siamese cnn for robust target association. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2016. 2
- [47] Jin Han Lee, Myung-Kyu Han, Dong Wook Ko, and Il Hong Suh. From big to small: Multi-scale local planar guidance for monocular depth estimation. *arXiv preprint arXiv:1907.10326*, 2019. 3, 5, 8
- [48] Bastian Leibe, Ales Leonardis, and Bernt Schiele. Combined object categorization and segmentation with an implicit shape model. In *Workshop on statistical learning in computer vision, ECCV*, 2004. 2
- [49] Bo Li, Chunhua Shen, Yuchao Dai, Anton Van Den Hengel, and Mingyi He. Depth and surface normal estimation from monocular images using regression on deep features and hierarchical crfs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1119–1127, 2015. 3
- [50] Jie Li, Allan Raventos, Arjun Bhargava, Takaaki Tagawa, and Adrien Gaidon. Learning to fuse things and stuff. *arXiv preprint arXiv:1812.01192*, 2018. 2
- [51] Qizhu Li, Xiaojuan Qi, and Philip HS Torr. Unifying training and inference for panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020. 2
- [52] Yanwei Li, Xinze Chen, Zheng Zhu, Lingxi Xie, Guan Huang, Dalong Du, and Xingang Wang. Attention-guided unified network for panoptic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 2
- [53] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017. 2

- [54] Yiding Liu, Siyu Yang, Bin Li, Wengang Zhou, Jizheng Xu, Houqiang Li, and Yan Lu. Affinity derivation and graph merge for instance segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 2
- [55] Jonathon Luiten, Tobias Fischer, and Bastian Leibe. Track to reconstruct and reconstruct to track. *IEEE Robotics and Automation Letters*, 2020. 8, 13
- [56] Jonathon Luiten, Aljosa Osep, Patrick Dendorfer, Philip Torr, Andreas Geiger, Laura Leal-Taixé, and Bastian Leibe. Hota: A higher order metric for evaluating multi-object tracking. *International Journal of Computer Vision*, pages 1–31, 2020. 8, 14
- [57] Reza Mahjourian, Martin Wicke, and Anelia Angelova. Unsupervised learning of depth and ego-motion from monocular video using 3d geometric constraints. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 3
- [58] Gerhard Neuhold, Tobias Ollmann, Samuel Rota Buló, and Peter Kotschieder. The mapillary vistas dataset for semantic understanding of street scenes. In *Proceedings of the IEEE International Conference on Computer Vision*, 2017. 7, 8
- [59] Davy Neven, Bert De Brabandere, Marc Proesmans, and Luc Van Gool. Instance segmentation by jointly optimizing spatial embeddings and clustering bandwidth. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 2
- [60] Stephen E Palmer. *Vision science: Photons to phenomenology*. MIT press, 1999. 1
- [61] Jinlong Peng, Changan Wang, Fangbin Wan, Yang Wu, Yabiao Wang, Ying Tai, Chengjie Wang, Jilin Li, Feiyue Huang, and Yanwei Fu. Chained-tracker: Chaining paired attentive regression results for end-to-end joint multiple-object detection and tracking. In *European Conference on Computer Vision*, pages 145–161. Springer, 2020. 2
- [62] Zygmunt Pizlo. Perception viewed as an inverse problem. *Vision research*, 41(24), 2001. 1
- [63] Lorenzo Porzi, Samuel Rota Buló, Aleksander Colovic, and Peter Kotschieder. Seamless scene segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 2
- [64] Lorenzo Porzi, Markus Hofinger, Idoia Ruiz, Joan Serrat, Samuel Rota Buló, and Peter Kotschieder. Learning multi-object tracking and segmentation from automatic annotations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6846–6855, 2020. 2, 13
- [65] Siyuan Qiao, Liang-Chieh Chen, and Alan Yuille. Detectors: Detecting objects with recursive feature pyramid and switchable atrous convolution. *arXiv preprint arXiv:2006.02334*, 2020. 2, 7, 8
- [66] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015. 2
- [67] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. ImageNet large scale visual recognition challenge. *International journal of computer vision*, 2015. 7
- [68] Ashutosh Saxena, Sung H Chung, and Andrew Y Ng. Learning depth from single monocular images. In *Advances in neural information processing systems*, 2006. 1, 2, 3
- [69] Samuel Schulter, Paul Vernaza, Wongun Choi, and Manmohan Chandraker. Deep network flow for multi-object tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 2
- [70] Sarthak Sharma, Junaid Ahmed Ansari, J Krishna Murthy, and K Madhava Krishna. Beyond pixels: Leveraging geometry and shape cues for online multi-object tracking. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018. 2
- [71] Konstantin Sofiiuk, Olga Barinova, and Anton Konushin. Adaptis: Adaptive instance selection network. In *Proceedings of the IEEE International Conference on Computer Vision*, 2019. 2
- [72] Siyu Tang, Mykhaylo Andriluka, Bjoern Andres, and Bernt Schiele. Multiple people tracking by lifted multicut and person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 2
- [73] Jonas Uhrig, Eike Rehder, Björn Fröhlich, Uwe Franke, and Thomas Brox. Box2pix: Single-shot instance segmentation by assigning pixels to object boxes. In *2018 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2018. 2
- [74] Jonas Uhrig, Nick Schneider, Lukas Schneider, Uwe Franke, Thomas Brox, and Andreas Geiger. Sparsity invariant cnns. In *2017 international conference on 3D Vision (3DV)*. IEEE, 2017. 2, 5, 6, 8, 14
- [75] Simon Vandenhende, Stamatios Georgoulis, and Luc Van Gool. Mti-net: Multi-scale task interaction networks for multi-task learning. In *European Conference on Computer Vision*, 2020. 3
- [76] Luc Vincent and Pierre Soille. Watersheds in digital spaces: an efficient algorithm based on immersion simulations. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 1991. 2
- [77] Paul Voigtlaender, Michael Krause, Aljosa Osep, Jonathon Luiten, Berin Balachandar Gnana Sekar, Andreas Geiger, and Bastian Leibe. Mots: Multi-object tracking and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2019. 2, 3, 6, 8, 13
- [78] Chaoyang Wang, José Miguel Buenaposada, Rui Zhu, and Simon Lucey. Learning depth from monocular videos using direct methods. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 3
- [79] Haochen Wang, Ruotian Luo, Michael Maire, and Greg Shakhnarovich. Pixel consensus voting for panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020. 2
- [80] Huiyu Wang, Yukun Zhu, Bradley Green, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. Axial-deeplab: Stand-alone axial-attention for panoptic segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 2

- [81] Peng Wang, Xiaohui Shen, Zhe Lin, Scott Cohen, Brian Price, and Alan L Yuille. Towards unified depth and semantic prediction from a single image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015. 3
- [82] Zhongdao Wang, Liang Zheng, Yixuan Liu, and Shengjin Wang. Towards real-time multi-object tracking. In *Proceedings of the European conference on computer vision (ECCV)*, 2020. 2, 8
- [83] Mark Weber, Jun Xie, Maxwell Collins, Yukun Zhu, Paul Voigtlaender, Hartwig Adam, Bradley Green, Andreas Geiger, Bastian Leibe, Daniel Cremers, et al. Step: Segmenting and tracking every pixel. *arXiv preprint arXiv:2102.11859*, 2021. 7, 8, 14
- [84] Nicolai Wojke, Alex Bewley, and Dietrich Paulus. Simple online and realtime tracking with a deep association metric. In *2017 IEEE international conference on image processing (ICIP)*. IEEE, 2017. 2
- [85] Yangxin Wu, Gengwei Zhang, Yiming Gao, Xiajun Deng, Ke Gong, Xiaodan Liang, and Liang Lin. Bidirectional graph reasoning network for panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020. 2
- [86] Zifeng Wu, Chunhua Shen, and Anton Van Den Hengel. Wider or deeper: Revisiting the ResNet model for visual recognition. *Pattern Recognition*, 2019. 7
- [87] Junyuan Xie, Ross Girshick, and Ali Farhadi. Deep3d: Fully automatic 2d-to-3d video conversion with deep convolutional neural networks. In *European Conference on Computer Vision*. Springer, 2016. 3
- [88] Yuwen Xiong, Renjie Liao, Hengshuang Zhao, Rui Hu, Min Bai, Ersin Yumer, and Raquel Urtasun. Upsnet: A unified panoptic segmentation network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 2
- [89] Dan Xu, Wanli Ouyang, Xiaogang Wang, and Nicu Sebe. Pad-net: Multi-tasks guided prediction-and-distillation network for simultaneous depth estimation and scene parsing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 3
- [90] Dan Xu, Elisa Ricci, Wanli Ouyang, Xiaogang Wang, and Nicu Sebe. Multi-scale continuous crfs as sequential deep networks for monocular depth estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 3
- [91] Jiarui Xu, Yue Cao, Zheng Zhang, and Han Hu. Spatial-temporal relation networks for multi-object tracking. In *Proceedings of the IEEE International Conference on Computer Vision*, 2019. 2
- [92] Zhenbo Xu, Wei Zhang, Xiao Tan, Wei Yang, Huan Huang, Shilei Wen, Errui Ding, and Liusheng Huang. Segment as points for efficient online multi-object tracking and segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, 2020. 2, 8, 13
- [93] Fan Yang, Xin Chang, Chenyu Dang, Ziqiang Zheng, Sakriani Sakti, Satoshi Nakamura, and Yang Wu. Remots: Self-supervised refining multi-object tracking and segmentation. *arXiv preprint arXiv:2007.03200*, 2020. 8
- [94] Maoke Yang, Kun Yu, Chi Zhang, Zhiwei Li, and Kuiyuan Yang. Denseaspp for semantic segmentation in street scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3684–3692, 2018. 4
- [95] Tien-Ju Yang, Maxwell D Collins, Yukun Zhu, Jyh-Jing Hwang, Ting Liu, Xiao Zhang, Vivienne Sze, George Papandreou, and Liang-Chieh Chen. Deeperlab: Single-shot image parser. *arXiv preprint arXiv:1902.05093*, 2019. 2
- [96] Wei Yin, Yifan Liu, Chunhua Shen, and Youliang Yan. Enforcing geometric constraints of virtual normal for depth prediction. In *Proceedings of the IEEE International Conference on Computer Vision*, 2019. 3
- [97] Zhichao Yin and Jianping Shi. Geonet: Unsupervised learning of dense depth, optical flow and camera pose. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 3
- [98] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2016. 7
- [99] Feihu Zhang, Victor Prisacariu, Ruigang Yang, and Philip HS Torr. Ga-net: Guided aggregation net for end-to-end stereo matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 185–194, 2019. 5
- [100] Youmin Zhang, Yimin Chen, Xiao Bai, Suihanjin Yu, Kun Yu, Zhiwei Li, and Kuiyuan Yang. Adaptive unimodal cost volume filtering for deep stereo matching. In *AAAI*, pages 12926–12934, 2020. 5
- [101] Yifu Zhang, Chunyu Wang, Xinggang Wang, Wenjun Zeng, and Wenyu Liu. Fairmot: On the fairness of detection and re-identification in multiple object tracking. *arXiv preprint arXiv:2004.01888*, 2020. 2
- [102] Zheng Zhang, Dazhi Cheng, Xizhou Zhu, Stephen Lin, and Jifeng Dai. Integrated object detection and tracking with tracklet-conditioned detection. *arXiv preprint arXiv:1811.11167*, 2018. 2
- [103] Xingyi Zhou, Vladlen Koltun, and Philipp Krähenbühl. Tracking objects as points. *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 2
- [104] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *arXiv preprint arXiv:1904.07850*, 2019. 2
- [105] Ji Zhu, Hua Yang, Nian Liu, Minyoung Kim, Wenjun Zhang, and Ming-Hsuan Yang. Online multi-object tracking with dual matching attention networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 2

Algorithm 1: Stitching Video Panoptic Predictions

Input : Panoptic prediction P_t for image t , and panoptic prediction R_t for image $t + 1$ when concatenated with image t .

Output: In-place stitched panoptic predictions P_t .

for $t = 1, 2, \dots, T - 1$ **do**

Increase all instance IDs of P_{t+1} by the maximum of the instance IDs of P_1 to P_t ;

// Find all overlapping region pairs \mathbb{C} between R_t and P_{t+1} with the same classes.

Let $\mathbb{S} = [(r, p) \mid r \in R_t \wedge p \in P_{t+1} \wedge r \cap p > 0]$;

Let $\mathbb{C} = [(r, p) \mid (r, p) \in \mathbb{S} \wedge \text{cls}(r) = \text{cls}(p)]$ where $\text{cls}(\cdot)$ denotes the class;

Sort $\mathbb{C} = [(r, p)]$ with respect to $\text{IoU}(r, p)$ in ascending order;

// For each region in R_t , \mathbb{M} stores the region in P_{t+1} having the largest IoU with it. \mathbb{N} stores the opposite direction.

Let \mathbb{M}, \mathbb{N} be empty dictionaries;

for $(r, p) \in \mathbb{C}$ **do**

$\mathbb{M}[r] = p$ and $\mathbb{N}[p] = r$;

for $r \rightarrow p \in \mathbb{M}$ **do**

// Propagate IDs from R_t to P_{t+1} and R_{t+1}

// if r and p map to each other in \mathbb{M} and \mathbb{N} .

if $r = \mathbb{N}[\mathbb{M}[r]]$ **then**

if $t < T - 1$ **then**

Assign the region in R_{t+1} that has the ID of p with the ID of r ;

Assign the region of p in P_{t+1} with the instance ID of r in R_t ;

A. Stitching Algorithm

Alg. 1 shows the details of the algorithm to stitch video panoptic predictions to form predictions with consistent IDs throughout the entire sequence. We split the panoptic prediction of the concatenated image pair t and $t + 1$ in the middle, and use P_t and R_t to denote the left and the right prediction, respectively. This makes P_t the panoptic prediction of image t , and R_t the panoptic prediction of image $t + 1$ with instance IDs that are consistent with those of P_t . The objective of the algorithm is to propagate IDs from R_t to P_{t+1} so that each object in P_t and P_{t+1} will have the same ID. The ID propagation is based on mask IoU between region pairs. For each region r in R_t , we find the region p in P_{t+1} that has the same class and the largest IoU with it. We use \mathbb{M} to store this mapping. Similarly, for each region p in P_{t+1} , we also find the region r in R_t that has the same class and the largest IoU with it. We use \mathbb{N} to store this mapping. If a region r of R_t and a region p of P_{t+1} are matched to each other (*i.e.* $\mathbb{M}(r) = p$ and $\mathbb{N}(p) = r$), then we propagate the ID from r to p .

Method	Pedestrians			Cars		
	sMOTSA	MOTSA	IDS	sMOTSA	MOTSA	IDS
TrackR-CNN [77]	46.8	65.1	78	76.2	87.8	93
MOTSNet [64]	54.6	69.3	-	78.1	87.2	-
MOTSFusion [55]	58.9	71.9	36	82.6	90.2	51
PointTrack [92]	62.4	77.3	19	85.5	94.9	22
ViP-DeepLab + KF	68.3	83.2	15	86.0	94.7	52

Table 6: Results on KITTI MOTS validation set.

$\mathcal{L}_{\text{depth}}$	weight	k = 1	k = 2	k = 3	k = 4	VPQ	absRel	DVPQ
0.1		68.9	61.9	58.8	56.5	61.5	9.51	51.3
1.0		69.0	62.0	58.7	56.5	61.6	7.21	55.1
10		67.8	61.1	57.5	55.5	60.5	6.54	54.3

Table 7: ViP-DeepLab trained with different training weights for $\mathcal{L}_{\text{depth}}$ on Cityscapes-DVPS.**B. More Experiments**

KITTI MOTS Validation Set We first evaluate ViP-DeepLab on the validation set of KITTI MOTS benchmark [77]. Tab. 6 shows the comparisons between ViP-DeepLab and the previous methods. We adopt the same strategy as we used for training models for KITTI MOTS Leaderboard except that the training data used in here does not include the validation set. As shown in the table, our method equipped with Kalman filter outperforms the previous methods by a large margin.

Effects of Depth Loss Weight Next, we study the effects of different training weights for the depth loss $\mathcal{L}_{\text{depth}}$. In the previous experiments on Cityscapes-DVPS and SemKITTI-DVPS, we use the depth loss defined by Equ. (4), which has a loss weight of 1.0. For the purpose of ablation study, we change the training weight from 1.0 to 10 and 0.1. The results are shown in Tab. 7. From the table we can see that as the $\mathcal{L}_{\text{depth}}$ weight increases, ViP-DeepLab performs better on the sub-task monocular depth estimation (*i.e.* absRel becomes lower), but worse on the sub-task video panoptic segmentation (*i.e.* VPQ becomes lower). This is consistent with our intuition that a larger $\mathcal{L}_{\text{depth}}$ weight makes the model focus more on the task of monocular depth estimation. The metric that matters most here is DVPQ, which unifies the metrics of both sub-tasks. In order to get a high DVPQ score, the predictions must be accurate on both tasks. Therefore, finding a balanced $\mathcal{L}_{\text{depth}}$ weight is critical to get a high DVPQ. As the table shows, setting $\mathcal{L}_{\text{depth}}$ weight to 1.0 achieves the best results among the three choices.

KITTI Depth Validation Set Finally, we show the performance of ViP-DeepLab on the official validation set of

Method	$\delta < 1.25 \uparrow$	$\delta < 1.25^2 \uparrow$	$\delta < 1.25^3 \uparrow$	absRel \downarrow	sqRel \downarrow	RMSE \downarrow	RMSElog \downarrow	SILog \downarrow
[20]	95.77	99.21	99.75	6.99	1.27	2.86	0.104	9.73
Ours	96.27	99.41	99.81	5.72	0.96	2.58	0.092	8.47

Table 8: Results on the official KITTI depth validation set. \uparrow : The higher the better. \downarrow : The lower the better.

Class	HOTA	DetA	AssA	DetRe	DetPr	AssRe	AssPr	LocA	sMOTSA
Cars	76.38	82.70	70.93	88.70	88.77	75.86	86.00	90.75	81.03
Pedestrians	64.31	70.69	59.48	75.71	81.77	67.52	74.92	84.40	68.76

Table 9: ViP-DeepLab performance on the KITTI MOTS test set for the new metrics.




Figure 8: The architecture of Cascade-ASPP, which is employed as *Dense Multi-scale Context* in the next-frame instance branch. It cascades four ASPP modules with the outputs densely connected.

Method	STQ	AQ	SQ
VPSNet [42]	0.50	0.35	0.72
Ours	0.64	0.52	0.78

Table 10: STQ comparison on Cityscapes-VPS.

KITTI depth benchmark [74]. The validation set has 1,000 cropped images. Tab. 8 compares our method with previous methods that report their performances on it. Our method outperforms the previous methods by a large margin on all the metrics.

New Metrics on KITTI MOTS The KITTI MOTS benchmark changed their ranking metrics [56]. Tab. 9 reports the performance of ViP-DeepLab for the new metrics.

STQ on Cityscapes-VPS Table 10 also reports the STQ [83] performance comparison between our ViP-DeepLab and VPSNet [42].

C. Cascade-ASPP

Fig. 8 shows the architecture of Cascade-ASPP. It is used as the module *Dense Multi-scale Context* in the next-frame instance branch shown in Fig. 3. It cascades four ASPP

modules with their outputs densely connected. The motivation of Cascade-ASPP is to dramatically increase the receptive field of the next-frame instance branch. As demonstrated in Tab. 3, Cascade-ASPP (*i.e.* DenseContext) improves the performances of video panoptic segmentation on Cityscapes-VPS compared with the single ASPP variant.

D. More Visualizations

We show more prediction visualizations in Fig. 9, Fig. 10, Fig. 11, and Fig. 12. We choose four sequences from 50 validation sequences of Cityscapes-DVPS, and the results are shown in Fig. 9 and Fig. 10. As each sequence contains only 6 frames, the figures show all the frames of the four sequences. Here, the video panoptic predictions demonstrate the results after the stitching algorithm, so each instance has the same instance ID in all the frames. SemKITTI-DVPS results are shown in Fig. 11 and Fig. 12. We present the results on two 16-frame video clips from the validation sequence. From the visualizations we can see that ViP-DeepLab is capable of outputting accurate video panoptic predictions and high quality depth predictions.




Figure 9: Prediction visualizations on Cityscapes-DVPS. From left to right: input image, temporally consistent panoptic segmentation prediction, monocular depth prediction, and point cloud visualization.

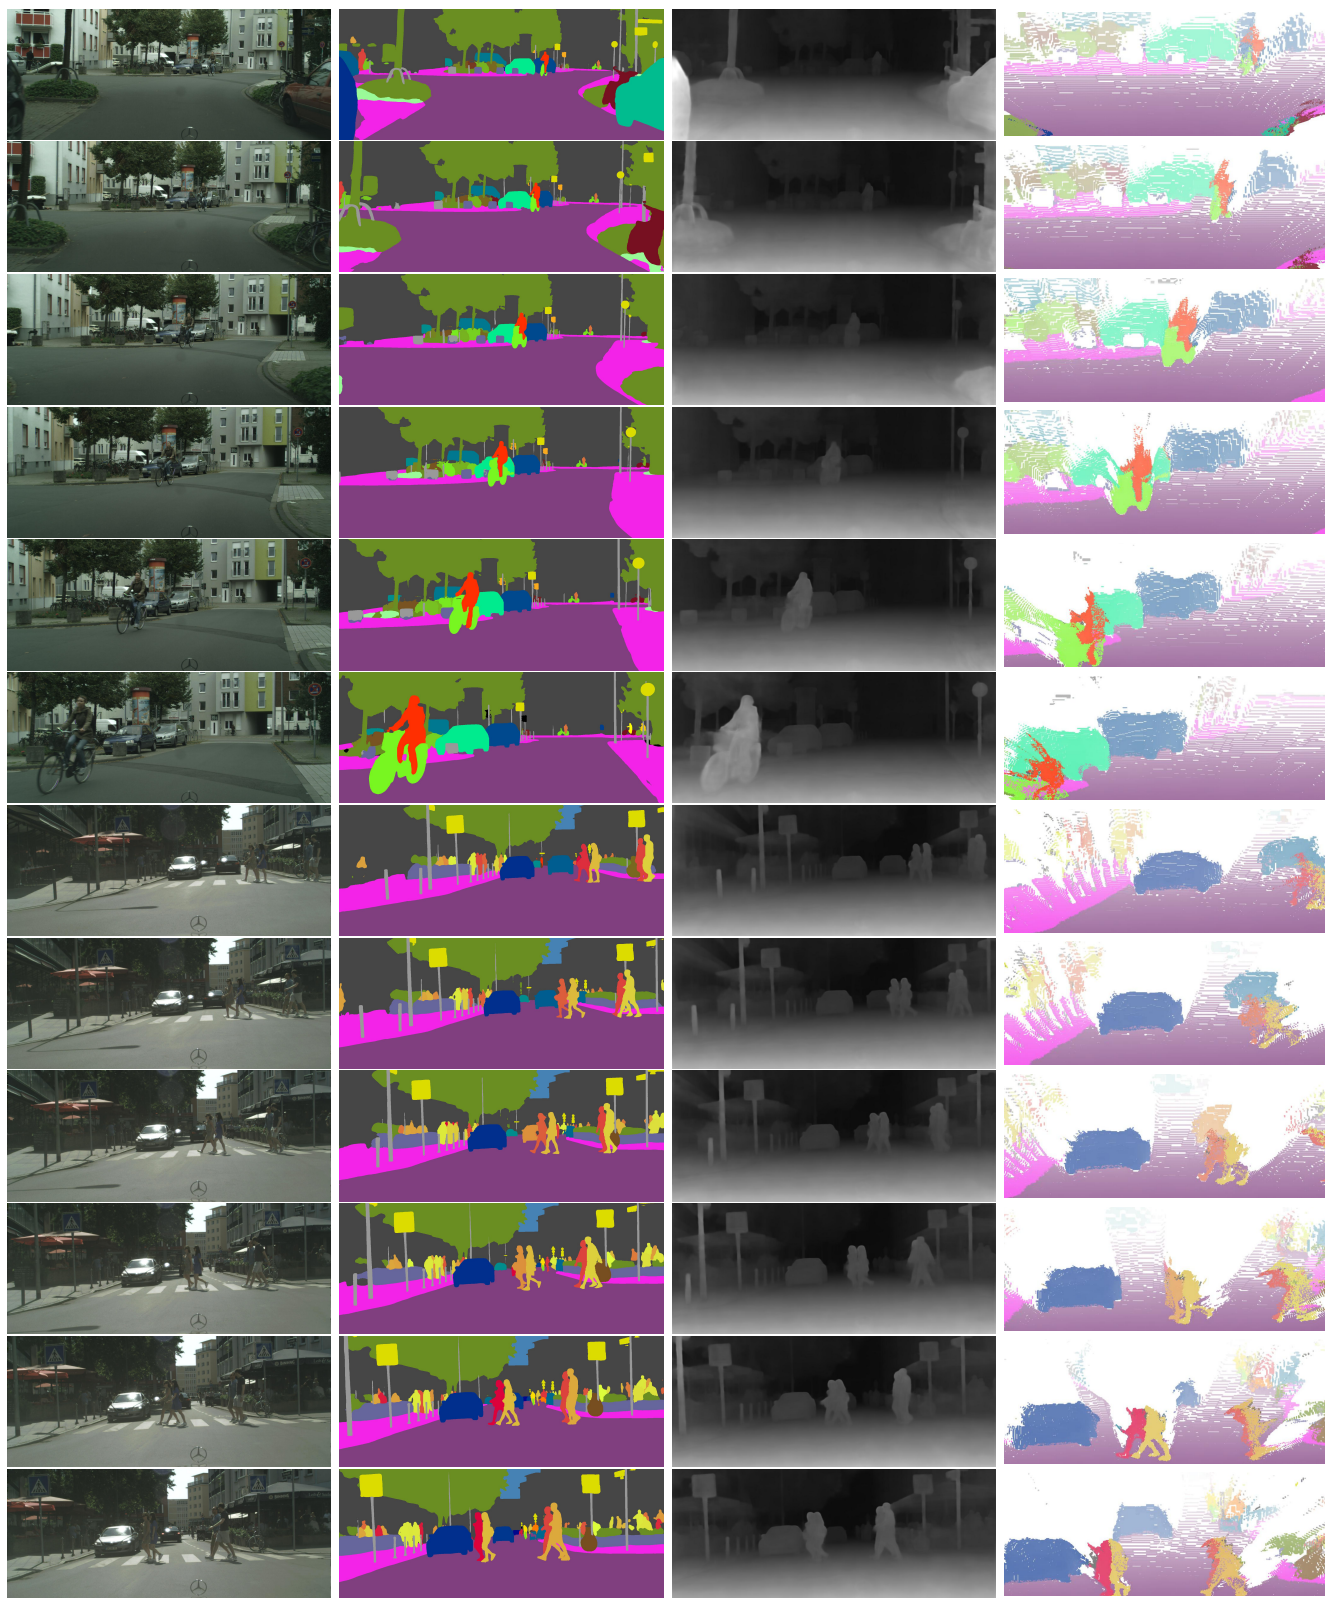


Figure 10: Prediction visualizations on Cityscapes-DVPS. From left to right: input image, temporally consistent panoptic segmentation prediction, monocular depth prediction, and point cloud visualization.




Figure 11: Prediction visualizations on SemKITTI-DVPS. From left to right: input image, temporally consistent panoptic segmentation prediction, monocular depth prediction, and point cloud visualization.



Figure 12: Prediction visualizations on SemKITTI-DVPS. From left to right: input image, temporally consistent panoptic segmentation prediction, monocular depth prediction, and point cloud visualization.