

Supplementary Materials for Neural Video Depth Stabilizer

Yiran Wang¹ Min Shi¹ Jiaqi Li¹ Zihao Huang¹

Zhiguo Cao¹ Jianming Zhang² Ke Xian^{3*} Guosheng Lin³

¹School of Artificial Intelligence and Automation, Huazhong University of Science and Technology

²Adobe Research ³S-Lab, Nanyang Technological University

{wangyiran,min_shi,lijiaqi_mail,zihaohuang,zgcao}@hust.edu.cn

jianmzha@adobe.com, {ke.xian,gslin}@ntu.edu.sg

<https://github.com/RaymondWang987/NVDS>

This supplementary contains the following contents:

- More details on the VDW dataset.
- More implementation details for NVDS.
- More details on experimental settings.
- More quantitative and qualitative results.

We also elaborate a demo consisting of many video visualizations and the illustration of our framework.

1. More details on the VDW Dataset

1.1. Dataset Construction

Data Acquisition and Pre-processing. Here we add more details on data acquisition and pre-processing (Sec. 4, line 466, main paper). Having obtained the raw videos, we use FFmpeg [4] and PySceneDetect [13] to split all the videos into 104,582 sequences. We manually check and remove the duplicated, chaotic, and blur scenes. Videos that are wrongly split by the scene detect tools are also removed. Finally, we reserve 32,405 videos with more than six million frames for disparity annotation.

Disparity Annotation. In Sec. 4, line 474 of the main paper, we mentioned that the disparity ground truth is obtained via sky segmentation and optical flow estimation. Here we specify the details. Compared with common practice [15, 20], we introduce a few engineering improvements to make the disparity maps more accurate. As the sky is considered to be infinitely far, pixels in the sky regions should be segmented and set to the minimum value in the disparity maps. We find that using a single segmentation model [1, 8] like prior arts [15, 20] causes errors and noises in the sky regions. Hence, we generate the sky masks in a model ensemble manner. Each frame along with its horizontally flipped copy are fed into two state-of-the-art semantic segmentation models SegFormer [21] and

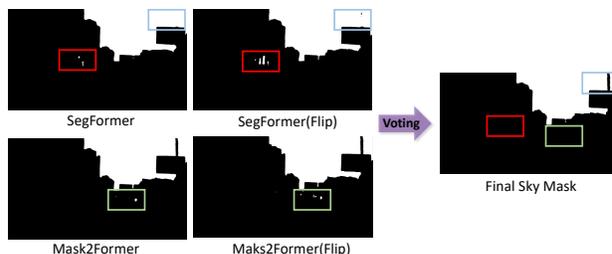


Figure 1: **Model ensemble strategy for sky segmentation on VDW dataset.** White area represents sky regions. Errors and noises in the rectangles are removed by model ensemble and voting, which improves the quality of the ground truth.

Mask2Former [3], which yields four sky masks in total. A pixel is considered as the sky when it is positive in more than two predicted sky masks. Besides, we also fill the connected regions with less than 50 pixels to further remove the noisy holes in the sky masks. Such ensemble strategy can improve the quality of the ground truth as shown in Fig. 1, and consequently improves the performance of the trained models, especially on skylines as shown in Fig. 6 and 8.

Following the practice of previous single-image depth datasets [15, 20], we adopt a state-of-the-art optical flow model GMFlow [22] to generate the ground truth disparity of the left- and right-eye views. The estimated optical flow is bidirectional. We perform a consistency check between the optical flow pairs to obtain the valid masks for training. We adopt the adaptive consistency threshold for each pixel as [11]. The ground truth of each video is normalized by its minimum and maximum disparity. Then, the disparity value is discretized into 65,535 intervals. Fig. 4 shows more examples of our VDW dataset.

Invalid Sample Filtering. Having obtained the annotations, we further filter the videos that are not qualified for our dataset. According to optical flow and valid masks, samples with the following three conditions are removed:

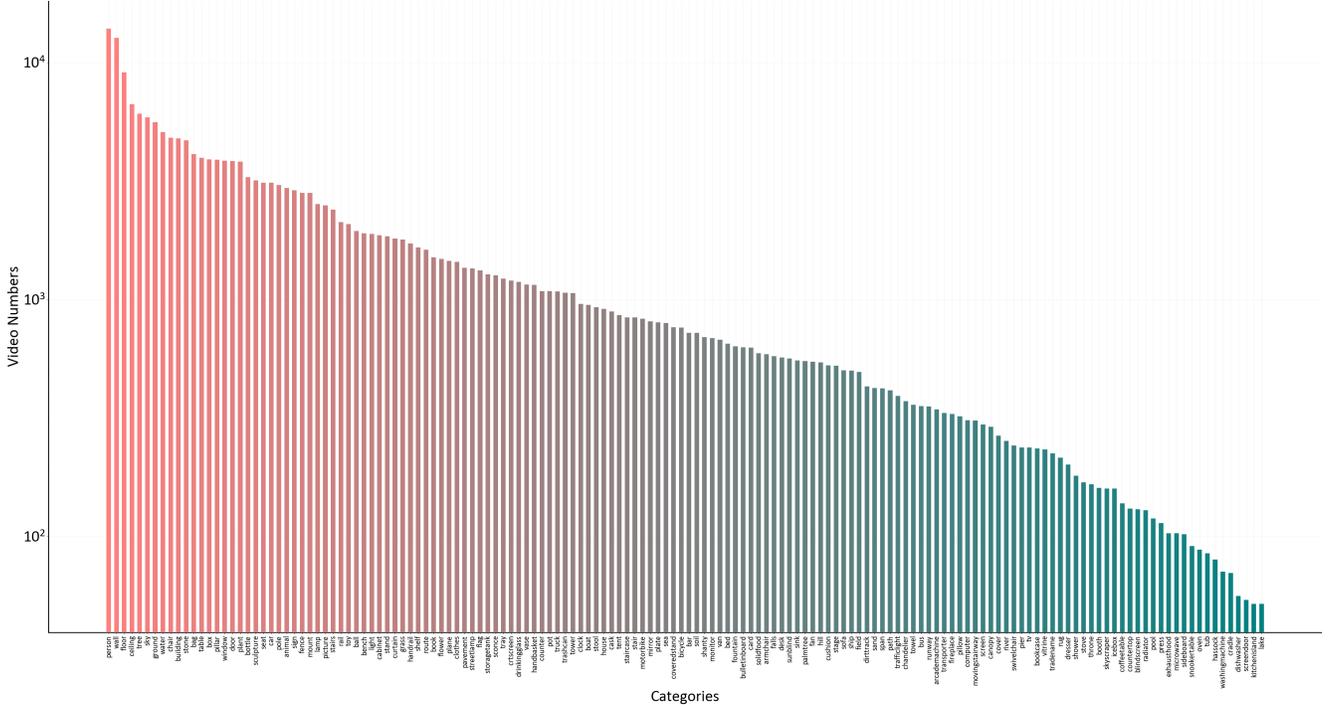


Figure 3: The statistics of the 150 semantic categories in VDW dataset.

scale and shift:

$$t(D) = \text{median}(D), s(D) = \frac{1}{M} \sum_{i=1}^M |D_i - t(D_i)|, \quad (1)$$

where M denotes the number of valid pixels. The prediction and the ground truth are aligned to zero translation and unit scale as follows:

$$\tilde{D} = \frac{D - t(D)}{s(D)}, \tilde{D}^* = \frac{D^* - t(D^*)}{s(D^*)}. \quad (2)$$

Then the affinity invariant loss can be formulated as:

$$\mathcal{L}_{af} = \frac{1}{M} \sum_{i=1}^M |\tilde{D} - \tilde{D}^*|. \quad (3)$$

Besides, we also adopt the multi-scale gradient matching loss [15], which can improve smoothness of homogeneous regions and sharpness of discontinuities in the disparity maps. The gradient matching loss is formulated as:

$$\mathcal{L}_{grad} = \frac{1}{M} \sum_{k=1}^K \sum_{i=1}^M (|\nabla_x R_i^k| + |\nabla_y R_i^k|), \quad (4)$$

where $R_i = \tilde{D}_i - \tilde{D}_i^*$, and R^k denotes the difference between the disparity maps at scale $k = 1, 2, 3, \dots, K$ (the resolution is halved at each level). Following [14], we set

$K = 4$ and set the weight β of \mathcal{L}_{grad} to 0.5. The spatial loss can be expressed as:

$$\mathcal{L}_s = \mathcal{L}_{af} + \beta \mathcal{L}_{grad}, \quad (5)$$

Temporal loss. In line 362 of the main paper, we mentioned that the temporal loss is masked with a visibility mask $O_{n \Rightarrow n-1}$ calculated from the warping discrepancy between frame F_n and the warped frame \hat{F}_{n-1} . This mask is obtained by:

$$O_{n \Rightarrow n-1} = \exp(-\gamma \|F_n - \hat{F}_{n-1}\|_2^2). \quad (6)$$

We set $\gamma = 50$ and use bilinear sampling layer for warping.

3. More Experimental Results

3.1. Depth Metrics

Here we specify the evaluation metrics for depth accuracy. we adopt commonly-applied depth evaluation metrics: Mean relative error (Rel) and accuracy with threshold t .

Mean relative error (Rel): $\frac{1}{M} \sum_{i=1}^M \frac{\|D_i - D_i^*\|_1}{D_i^*}$;

Accuracy with threshold t : Percentage of D_i such that $\max(\frac{D_i}{D_i^*}, \frac{D_i^*}{D_i}) = \delta < t \in [1.25, 1.25^2, 1.25^3]$, where M denotes pixel numbers, D_i and D_i^* are prediction and ground truth of pixel i .



Figure 4: More examples of our VDW dataset. Sky regions and invalid pixels are masked out.

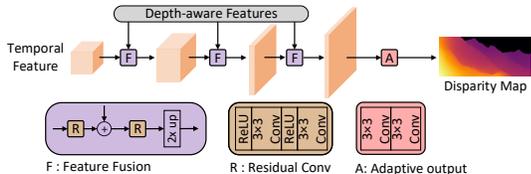


Figure 5: The architecture of decoder.

	DPT-L [14]	NeWCRFs [23]	Midas-v2 [15]	Stabilization Network
FLOPs (G)	1011.32	550.47	415.24	254.53
Params (M)	341.26	270.33	104.18	88.31

Table 3: Comparisons of FLOPs and model parameters.

Method	$\delta_1 \uparrow$	$\delta_2 \uparrow$	$\delta_3 \uparrow$	$Rel \downarrow$	$OPW \downarrow$
Midas [15]	0.644	0.853	0.928	0.347	0.647
DPT [14]	<u>0.724</u>	<u>0.890</u>	<u>0.950</u>	<u>0.266</u>	0.461
ST-CLSTM [24]	0.461	0.708	0.836	0.589	0.455
FMNet [19]	0.465	0.712	0.837	0.584	0.388
DeepV2D [17]	0.522	0.728	0.833	0.628	0.425
WSVD [18]	0.621	0.825	0.912	0.379	0.437
Robust-CVD [7]	0.658	0.855	0.928	0.334	0.251
Ours(Midas)	0.694	0.879	0.943	0.286	<u>0.164</u>
Ours(DPT)	0.731	0.895	0.952	0.259	0.138

Table 4: **Comparisons on VDW dataset.** The first 2 rows show the results of different single-image depth predictors. The next 5 rows contain video depth approaches. The last 2 rows consist of the results of our NVDS. Best performance is in boldface. Second best is underlined.

3.2. Model Efficiency

Here we evaluate the efficiency of the proposed Neural Video Depth Stabilizer (NVDS) in detail. Model parameters and FLOPs are reported in Table 3. The FLOPs are evaluated on a 384×384 video with four frames. The stabilization network of NVDS only introduces limited computation overhead compared with the off-the-shelf depth predictors.

3.3. More Quantitative Comparisons

In the main paper, only δ_1 , Rel , and OPW are reported. The additional results on the VDW and the Sintel [2] dataset are shown in Table 4 and Table 5. Besides, as CVD [10] and Zhang *et al.* [25] cannot produce results on 11 of 23 videos in Sintel [2] dataset, we additionally report the results on the other 12 videos in Table 6.

3.4. More Qualitative Results.

We show more visual comparisons in Fig. 6, 7, 8 and 9. Please refer to the supplementary video for video depth visualization results. We draw the scanline slice over time. Fewer zigzagging pattern means better consistency.

Method	$\delta_1 \uparrow$	$\delta_2 \uparrow$	$\delta_3 \uparrow$	$Rel \downarrow$	$OPW \downarrow$
Midas [15]	0.485	0.693	0.787	0.410	0.843
DPT [14]	0.597	<u>0.768</u>	<u>0.846</u>	<u>0.339</u>	0.612
ST-CLSTM [24]	0.351	0.571	0.706	0.517	0.585
FMNet [19]	0.357	0.579	0.712	0.513	0.521
DeepV2D [17]	0.486	0.674	0.760	0.526	0.534
WSVD [18]	0.501	0.709	0.804	0.439	0.577
CVD [10]	0.518	0.741	0.832	0.406	0.497
Robust-CVD [7]	0.521	0.727	0.833	0.422	0.475
Zhang <i>et al.</i> [25]	0.522	0.727	0.831	0.342	0.481
Ours(Midas)	0.532	0.731	0.833	0.374	<u>0.469</u>
Ours(DPT)	<u>0.591</u>	0.770	0.849	0.335	0.424

Table 5: **Comparisons on the Sintel dataset.** We only report CVD [10] and Zhang *et al.* [25] on the 12 videos with valid outputs, while other methods are on the 23 videos.

Method	$\delta_1 \uparrow$	$\delta_2 \uparrow$	$\delta_3 \uparrow$	$Rel \downarrow$	$OPW \downarrow$
Midas [15]	0.670	0.853	0.902	0.246	0.712
DPT [14]	0.747	0.874	0.917	0.196	0.671
ST-CLSTM [24]	0.477	0.711	0.827	0.366	0.547
FMNet [19]	0.492	0.728	0.825	0.363	0.516
DeepV2D [17]	0.509	0.735	0.827	0.384	0.575
CVD [10]	0.518	0.741	0.832	0.406	0.497
Zhang <i>et al.</i> [25]	0.522	0.727	0.831	0.342	0.481
WSVD [18]	0.621	0.822	0.891	0.305	0.581
Robust-CVD [7]	0.673	0.848	0.888	0.284	0.447
Ours(Midas)	0.700	0.866	0.918	0.226	<u>0.425</u>
Ours(DPT)	<u>0.741</u>	0.876	0.926	<u>0.205</u>	0.411

Table 6: **Comparisons on the 12 videos of Sintel [2] dataset.** We test the 12 videos that CVD [10] and Zhang *et al.* [25] can produce results for fair comparisons.

References

- [1] Samuel Rota Buló, Lorenzo Porzi, and Peter Kotschieder. In-place activated batchnorm for memory-optimized training of dnns. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5639–5647, 2018. 1
- [2] Daniel J Butler, Jonas Wulff, Garrett B Stanley, and Michael J Black. A naturalistic open source movie for optical flow evaluation. In *European Conference on Computer Vision (ECCV)*, pages 611–625. Springer, 2012. 5, 7
- [3] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1290–1299, 2022. 1, 2
- [4] FFmpeg developers. FFmpeg. <https://ffmpeg.org>. [Online; Accessed 2022]. 1
- [5] Simon Hadfield and Richard Bowden. Hollywood 3d: Recognizing actions in 3d natural scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern*

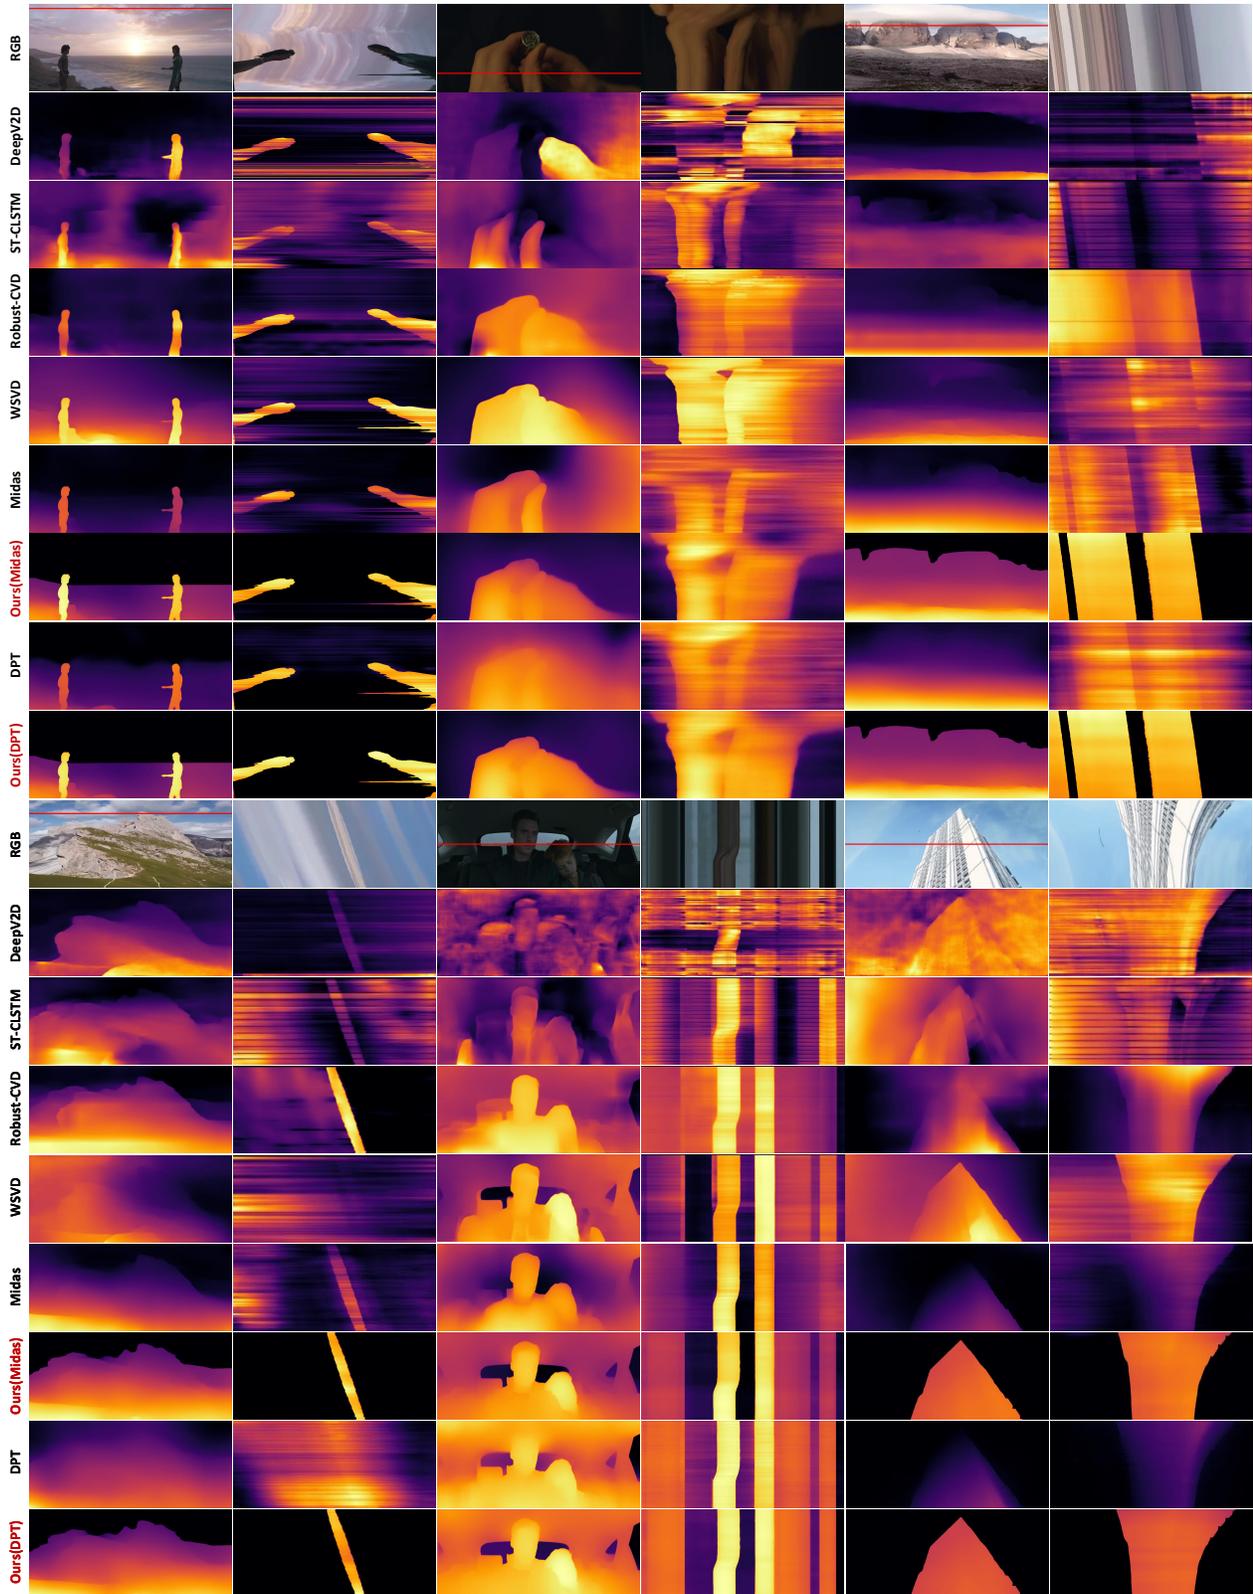


Figure 6: **More qualitative results on natural scenes.** The first image in each pair is the RGB frame, while the second is the scanline slice over time. Fewer zigzagging pattern means better consistency.

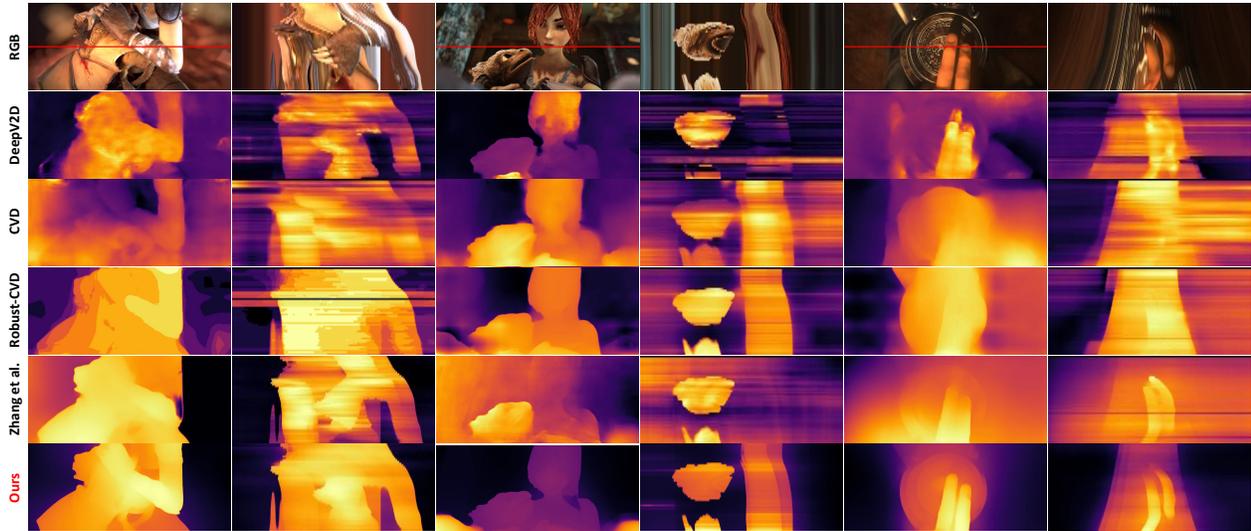


Figure 7: **Qualitative results on Sintel [2] dataset.** We compare the results of DeepV2D [17], CVD [10], Robust-CVD [7], and Zhang *et al.* [25]. Without relying on test-time training [7, 10, 25], we conduct zero-shot evaluations on Sintel [2] and achieve significantly better performance than those TTT-based methods [7, 10, 25].

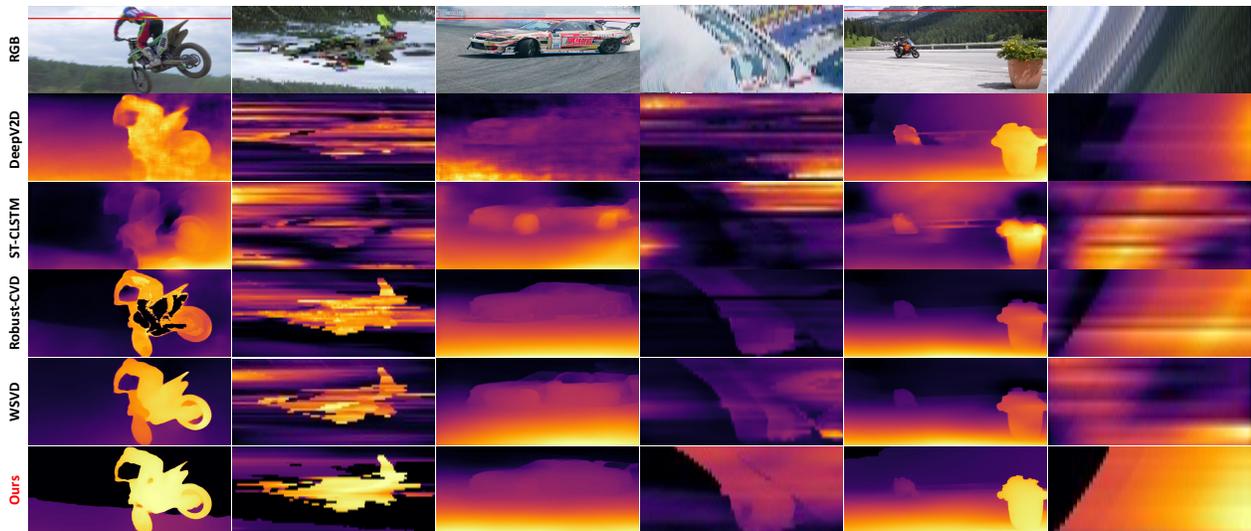


Figure 8: **Qualitative results on DAVIS [12] dataset.** We achieve better performance than prior arts on natural scenes.

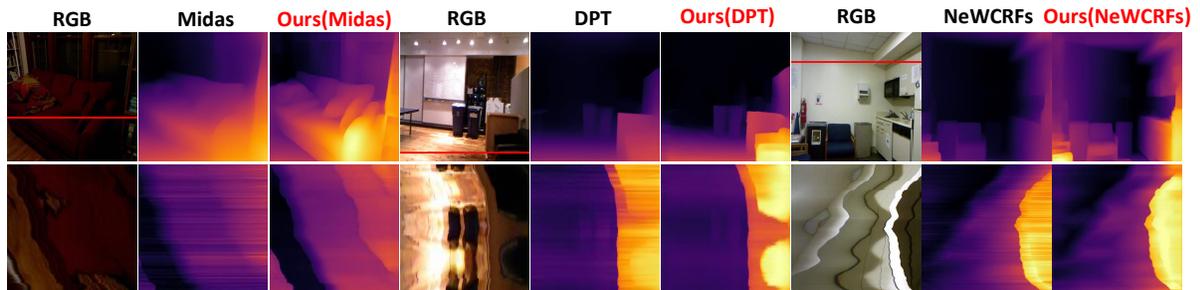


Figure 9: **Qualitative results on NYUDV2 [16] dataset.** We compare three different single-image depth predictors.

- Recognition (CVPR)*, pages 3398–3405, 2013. 2
- [6] Qingqiu Huang, Yu Xiong, Anyi Rao, Jiase Wang, and Dahua Lin. Movienet: A holistic dataset for movie understanding. In *European Conference on Computer Vision (ECCV)*, pages 709–727. Springer, 2020. 2
- [7] Johannes Kopf, Xuejian Rong, and Jia-Bin Huang. Robust consistent video depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1611–1621, 2021. 5, 7
- [8] Guosheng Lin, Anton Milan, Chunhua Shen, and Ian Reid. Refinenet: Multi-path refinement networks for high-resolution semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1925–1934, 2017. 1, 2
- [9] 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/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2117–2125, 2017. 2
- [10] Xuan Luo, Jia-Bin Huang, Richard Szeliski, Kevin Matzen, and Johannes Kopf. Consistent video depth estimation. *ACM Transactions on Graphics (ToG)*, 39(4):71–1, 2020. 5, 7
- [11] Simon Meister, Junhwa Hur, and Stefan Roth. Unflow: Unsupervised learning of optical flow with a bidirectional census loss. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018. 1
- [12] Jordi Pont-Tuset, Federico Perazzi, Sergi Caelles, Pablo Arbeláez, Alexander Sorkine-Hornung, and Luc Van Gool. The 2017 davis challenge on video object segmentation. *arXiv preprint arXiv:1704.00675*, 2017. 7
- [13] PySceneDetect developers. PySceneDetect. <http://scenedetect.com>. [Online; Accessed 2022]. 1
- [14] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12179–12188, 2021. 2, 3, 5
- [15] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE transactions on pattern analysis and machine intelligence*, 44(03):1623–1637, 2020. 1, 2, 3, 5
- [16] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgb-d images. In *European Conference on Computer Vision (ECCV)*, pages 746–760. Springer, 2012. 7
- [17] Zachary Teed and Jia Deng. Deepv2d: Video to depth with differentiable structure from motion. In *International Conference on Learning Representations*, 2019. 5, 7
- [18] Chaoyang Wang, Simon Lucey, Federico Perazzi, and Oliver Wang. Web stereo video supervision for depth prediction from dynamic scenes. In *IEEE International Conference on 3D Vision*, pages 348–357. IEEE, 2019. 5
- [19] Yiran Wang, Zhiyu Pan, Xingyi Li, Zhiguo Cao, Ke Xian, and Jianming Zhang. Less is more: Consistent video depth estimation with masked frames modeling. In *Proceedings of the 30th ACM International Conference on Multimedia*, MM ’22, page 6347–6358, New York, NY, USA, 2022. Association for Computing Machinery. 5
- [20] Ke Xian, Chunhua Shen, Zhiguo Cao, Hao Lu, Yang Xiao, Ruibo Li, and Zhenbo Luo. Monocular relative depth perception with web stereo data supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 311–320, 2018. 1
- [21] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *Advances in neural information processing systems*, 34:12077–12090, 2021. 1, 2
- [22] Haoifei Xu, Jing Zhang, Jianfei Cai, Hamid Rezaatofghi, and Dacheng Tao. Gmflow: Learning optical flow via global matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8121–8130, 2022. 1
- [23] Weihao Yuan, Xiaodong Gu, Zuozhuo Dai, Siyu Zhu, and Ping Tan. Newcrfs: Neural window fully-connected crfs for monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3916–3925, 2022. 5
- [24] Haokui Zhang, Chunhua Shen, Ying Li, Yuanzhouhan Cao, Yu Liu, and Youliang Yan. Exploiting temporal consistency for real-time video depth estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1725–1734, 2019. 5
- [25] Zhoutong Zhang, Forrester Cole, Richard Tucker, William T Freeman, and Tali Dekel. Consistent depth of moving objects in video. *ACM Transactions on Graphics (TOG)*, 40(4):1–12, 2021. 5, 7
- [26] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 633–641, 2017. 2