Supplementary Material: Directional Priors for Multi-Frame Optical Flow

Daniel Maurer
maurer@vis.uni-stuttgart.de
Michael Stoll
stoll@vis.uni-stuttgart.de
Andrés Bruhn
bruhn@vis.uni-stuttgart.de

Institute for Visualization and
Interactive Systems
University of Stuttgart, Germany

1 Results for the MPI Sintel Benchmark

Besides the KITTI benchmarks [7, 8] we also evaluated the quality of our method on the MPI Sintel benchmark [1]. The results for the training sequences can be found in Tab. 1. As one can see, also for this benchmark our method obtains an improvement compared to the two-frame baseline. The results for the test sequences can be found Tab. 2. Here, in case of the clean render pass our methods achieves Rank 3. Thereby, it clearly outperforms all two-frame pipeline approaches from the literature. In contrast, in case of the final render path, our method achieves Rank 18. In this context, however, one should note that we only used the training data of the clean render pass to adjust the parameters for our approach.

Method	#	◁	R	⇝	CPM Matches [4] Sintel AEE	DF Matches [5] Sintel AEE
EpicFlow refinement [2]	2	–	–	–	2.00	1.94
OIR refinement [4]	2	–	–	–	1.99	1.91
our method	3	✓	–	–	1.99	1.91
our method	3	✓	✓	–	1.97	**1.89**
our method	3	✓	✓	✓	1.94	**1.89**

#: number of frames, ◁: trajectorial filtering, R: multi-frame refinement, ⇝: directional regularizer

Table 1: Results for the the MPI-Sintel benchmark (clean render pass) [1] training sequences in terms of the average enpoint error (AEE).

2 Visual Analysis

In addition to the exemplary result for the KITTI 2015 benchmark in the main paper (Fig. 5), we further depict an exemplary result for the KITTI 2012 benchmark. For the sake

of completeness, both figures can be found in Fig. 1 and Fig. 2. Apart from the forward and backward flow they show the multiplicative and additive illumination coefficients for both directions as well as the selection map of the order-adaptive spatial regularization. Regarding the illumination changes, the KITTI 2012 result shows significant changes in forward direction, while the KITTI 2015 result shows almost constant illumination. One can also see illumination changes at specular reflections and occlusions. On the other hand, regarding the automatic selection of the regularization order, the KITTI 2012 result used second-order regularization for the entire scene, while the KITTI 2015 result used first-order regularization for the background and second-order regularization for the moving cars.

MPI Sintel clean	**all**	**matched**	**unmatched**	**MPI Sintel final**	**all**	**matched**	**unmatched**
MR-Flow[2]	2.527	0.954	15.365	PWC-Net	5.042	2.445	26.221
FlowFields+	3.102	0.820	21.718	DCFlow	5.119	2.283	28.228
our approach (DF)	**3.103**	**0.881**	**21.227**	FlowFieldsCNN	5.363	2.303	30.313
CPM2	3.253	0.980	21.812	MR-Flow[2]	5.376	2.818	26.235
MirrorFlow	3.316	1.338	19.470	S2F-IF	5.417	2.549	28.795
DF+OIR	3.331	0.942	22.817	InterpoNet_ff	5.535	2.372	31.296
S2F-IF	3.500	0.988	23.986	RicFlow	5.620	2.765	28.907
SPM-BPv2	3.515	1.020	23.865	InterpoNet_cpm	5.627	2.594	30.344
DCFlow	3.537	1.103	23.394	ProbFlowFields	5.696	2.545	31.371
RicFlow	3.550	1.264	22.220	FlowFields+	5.707	2.684	30.356
—	—	—	—	DF+OIR	5.862	2.864	30.303
CPM-Flow	3.557	1.189	22.889	CPM-Flow	5.960	2.990	30.177
—	—	—	—	**our approach (DF)**	**6.014**	**2.922**	**31.224**
DiscreteFlow	3.567	1.108	23.626	DiscreteFlow	6.077	2.937	31.685
EpicFlow	4.115	1.360	26.595	EpicFlow	6.285	3.060	32.564

Table 2: Results of the MPI Sintel [1] benchmark test sets. Top 10 non-anonymous methods and methods related to our approach.

Figure 1: Exemplary result for a sequence of the KITTI 2012 [2] benchmark. **First row:** Input frames f^{t-1}, f^t and f^{t+1}. **Second row:** Ground truth forward flow, bad pixel visualization and computed order selection map o (first-order: turquoise, second-order: brown). **Third row:** Computed forward flow $\mathbf{w}^t$, coefficients c_1^t and c_2^t (shifted and rescaled, zero maps to white). **Fourth row:** Computed backward flow $\mathbf{w}^{t-1}$, coefficients c_1^{t-1} and c_2^{t-1}.

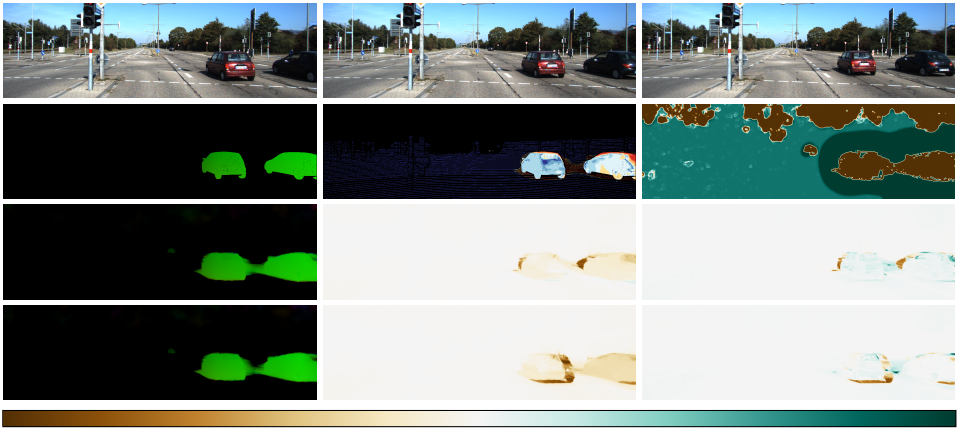

Figure 2: Exemplary result for a sequence of the KITTI 2015 [6] benchmark. **First row:** Input frames f^{t-1}, f^t and f^{t+1}. **Second row:** Ground truth forward flow, bad pixel visualization and computed order selection map o (first-order: turquoise, second-order: brown). **Third row:** Computed forward flow $\mathbf{w}^t$, coefficients c_1^t and c_2^t (shifted and rescaled, zero maps to white). **Fourth row:** Computed backward flow $\mathbf{w}^{t-1}$, coefficients c_1^{t-1} and c_2^{t-1}.

References

[1] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black. A naturalistic open source movie for optical flow evaluation. In *Proc. European Conference on Computer Vision*, pages 611–625, 2012.

[2] A. Geiger, P. Lenz, and R. Urtasun. Are we ready for autonomous driving? the KITTI vision benchmark suite. In *In Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 3354–3361, 2012.

[3] Y. Hu, R. Song, and Y. Li. Efficient coarse-to-fine PatchMatch for large displacement optical flow. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 5704–5712, 2016.

[4] D. Maurer, M. Stoll, and A. Bruhn. Order-adaptive and illumination-aware variational optical flow refinement. In *Proc. British Machine Vision Conference*, pages 1–13, 2017.

[5] B. Menze and A. Geiger. Discrete optimization for optical flow. In *Proc. German Conference on Pattern Recognition*, pages 16–28, 2015.

[6] M. Menze and A. Geiger. Object scene flow for autonomous vehicles. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 3061–3070, 2015.

[7] J. Revaud, P. Weinzaepfel, Z. Harchaoui, and C. Schmid. Epicflow: Edge-preserving interpolation of correspondences for optical flow. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 1164–1172, 2015.