Multiframe Scene Flow with Piecewise Rigid Motion
Supplementary Material

Vladislav Golyanik[1,2,3] Kihwan Kim[1] Robert Maier[1,4] Matthias Nießner[4,5] Didier Stricker[2,3] Jan Kautz[1]
[1]NVIDIA [2]University of Kaiserslautern [3]DFKI [4]Technical University of Munich [5]Stanford University

This document is intended to complement the draft presented at 3DV 2017. Note that all shown results are generated by the method and formulation described in the main paper. We describe further aspects of Multiframe Scene Flow (MSF) including:

- segment propagation in the multiframe setting,
- more qualitative results supporting the analysis in Sec. 4.3 (SLIC and regular segmentation as an alternative to the Felzenszwalb segmentation),
- more results on the SINTEL data set [2].

Additionally, we provide a supplementary video with results on the SINTEL data set.

A. Segment propagation (multiframe setting)

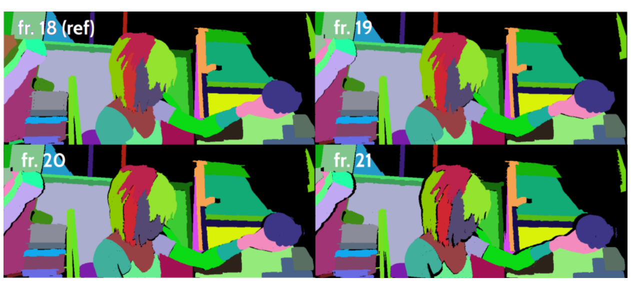

Figure 1: Segmentation transfer from the reference frame to three other frames in *alley1*. In this experiment, scene flow was computed in the global manner (every frame was optimized jointly with every other). In all frames, the same colours denote same segments, and the shapes are obtained by projection of the initial segment to the respective frames.

B. SLIC Segmentation

Sometimes, Felzenszwalb segmentation does not generate sufficiently accurate segmentation of a reference frame. In this case, other initialization methods can be used. One of the alternative algorithms is the SLIC superpixels approach [1]. Fig. 2 shows examples of the reference frame segmentations and segmentation transfer with the Felzenszwalb segmentation [3] and SLIC segmentation.

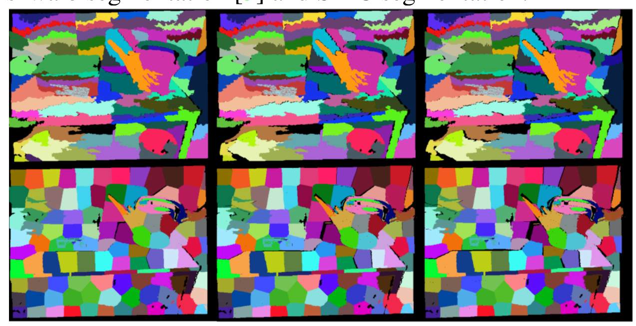

Figure 2: Segmentation transfer on the Bonn watering can sequence [4] with Felzenszwalb segmentation [3] (top row) and SLIC segmentation [1] (bottom row).

C. Additional visualisations

Fig. 3 contains an enlarged teaser from the main matter. Some excerpts from the *alley1*, *bandage1*, *sleeping1*, *sleeping2*, *shaman2*, *shaman3*, *mountain1* sequences are given in Fig. 4–6 (see also the video).

References

[1] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Susstrunk. Slic superpixels compared to state-of-the-art superpixel methods. *Transactions on Pattern Analysis and Machine Intelligence (T-PAMI)*, 34(11):2274–2282, 2012. 1

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

[3] P. F. Felzenszwalb and D. P. Huttenlocher. Efficient graph-based image segmentation. *International Journal of Computer Vision (IJCV)*, 59(2):167–181, 2004. 1

[4] J. Stueckler and S. Behnke. Efficient dense rigid-body motion segmentation and estimation in rgb-d video. *International Journal of Computer Vision (IJCV)*, 2015. 1

Figure 3: An enlarged version of the teaser picture.

Figure 4: Additional visualisations of the *alley1* sequence. For every block: reference RGB frame – segmentation of the reference frame (top row), our result, projection of the scene flow into the image plane (no depth thresholding) – projection of the segmentation into the current frame (second row), our result, projection of the segmentation into the current frame (depth therholded to two meters) – ground truth optical flow (bottom row).

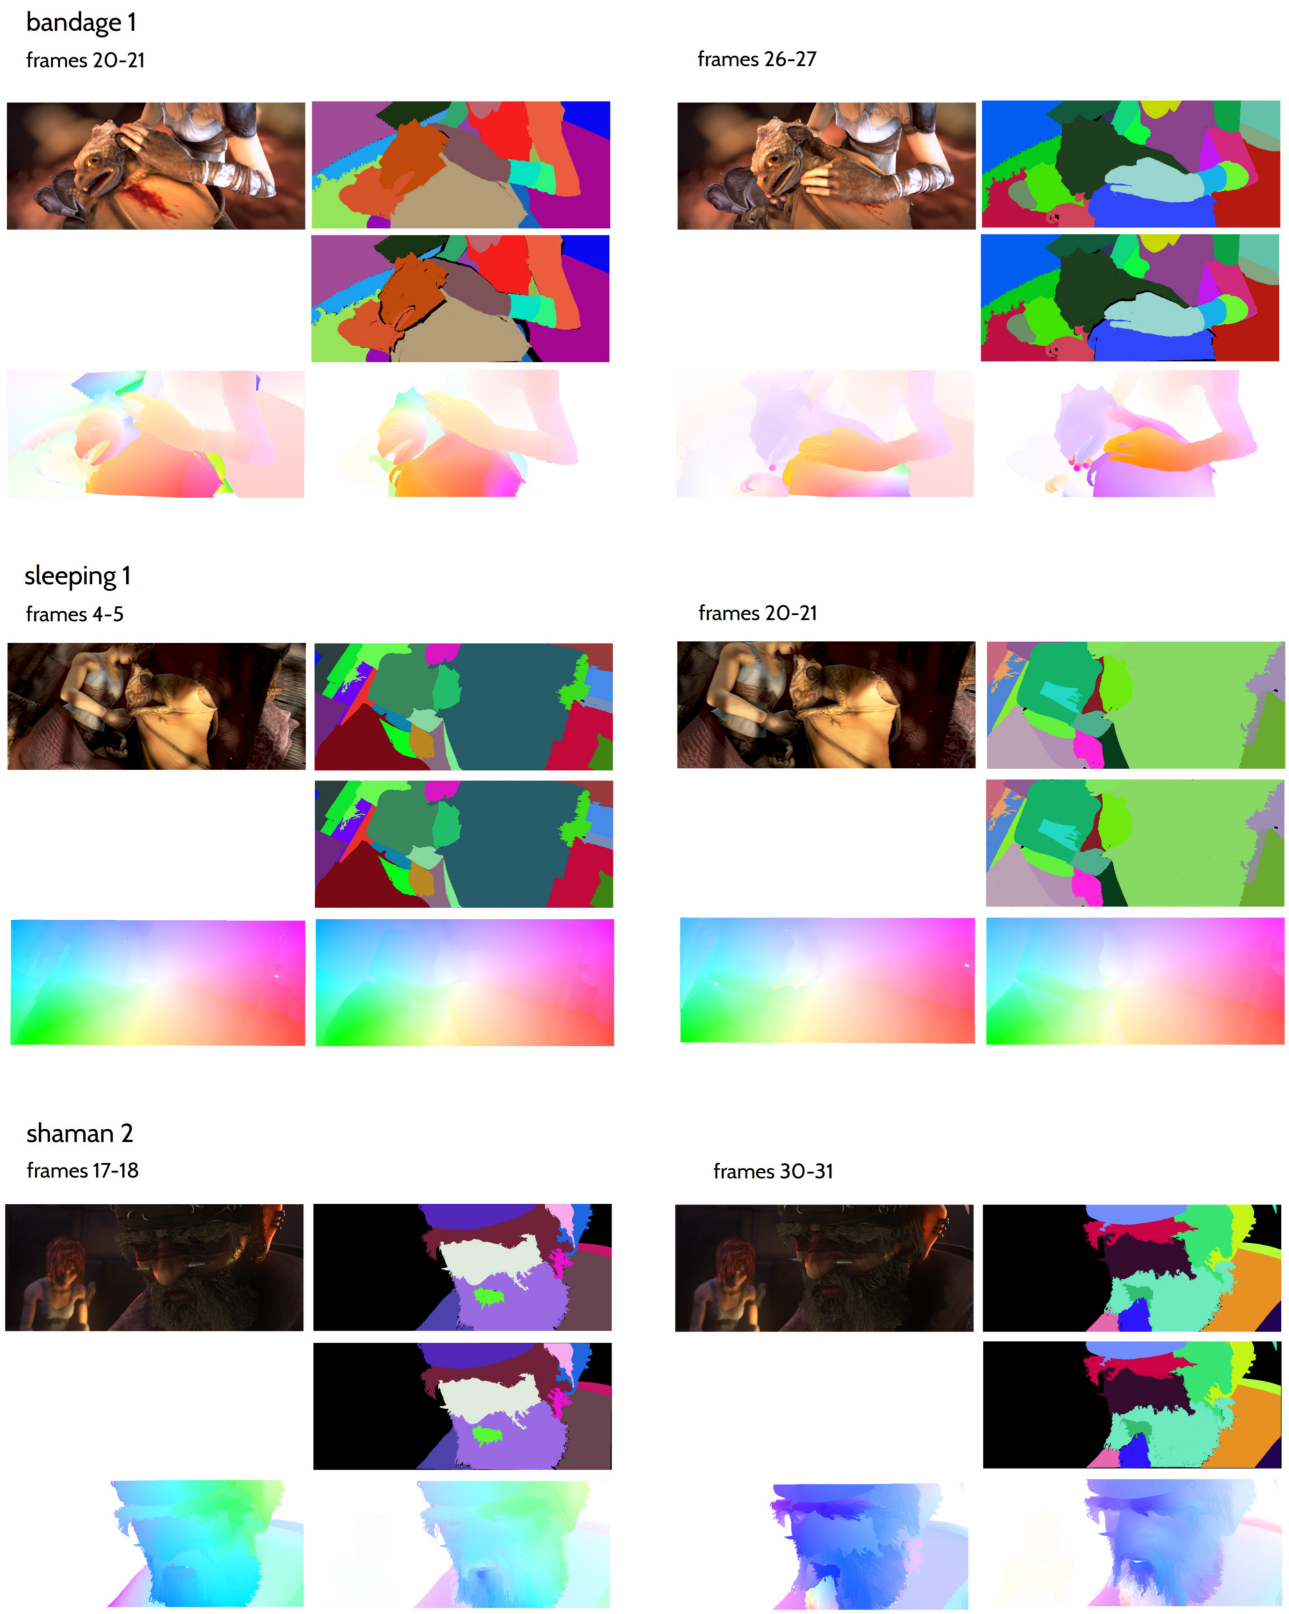

Figure 5: Additional visualisations of the *bandage1*, *sleeping1* and *shaman2* sequences. See Fig. 4 for the legend.

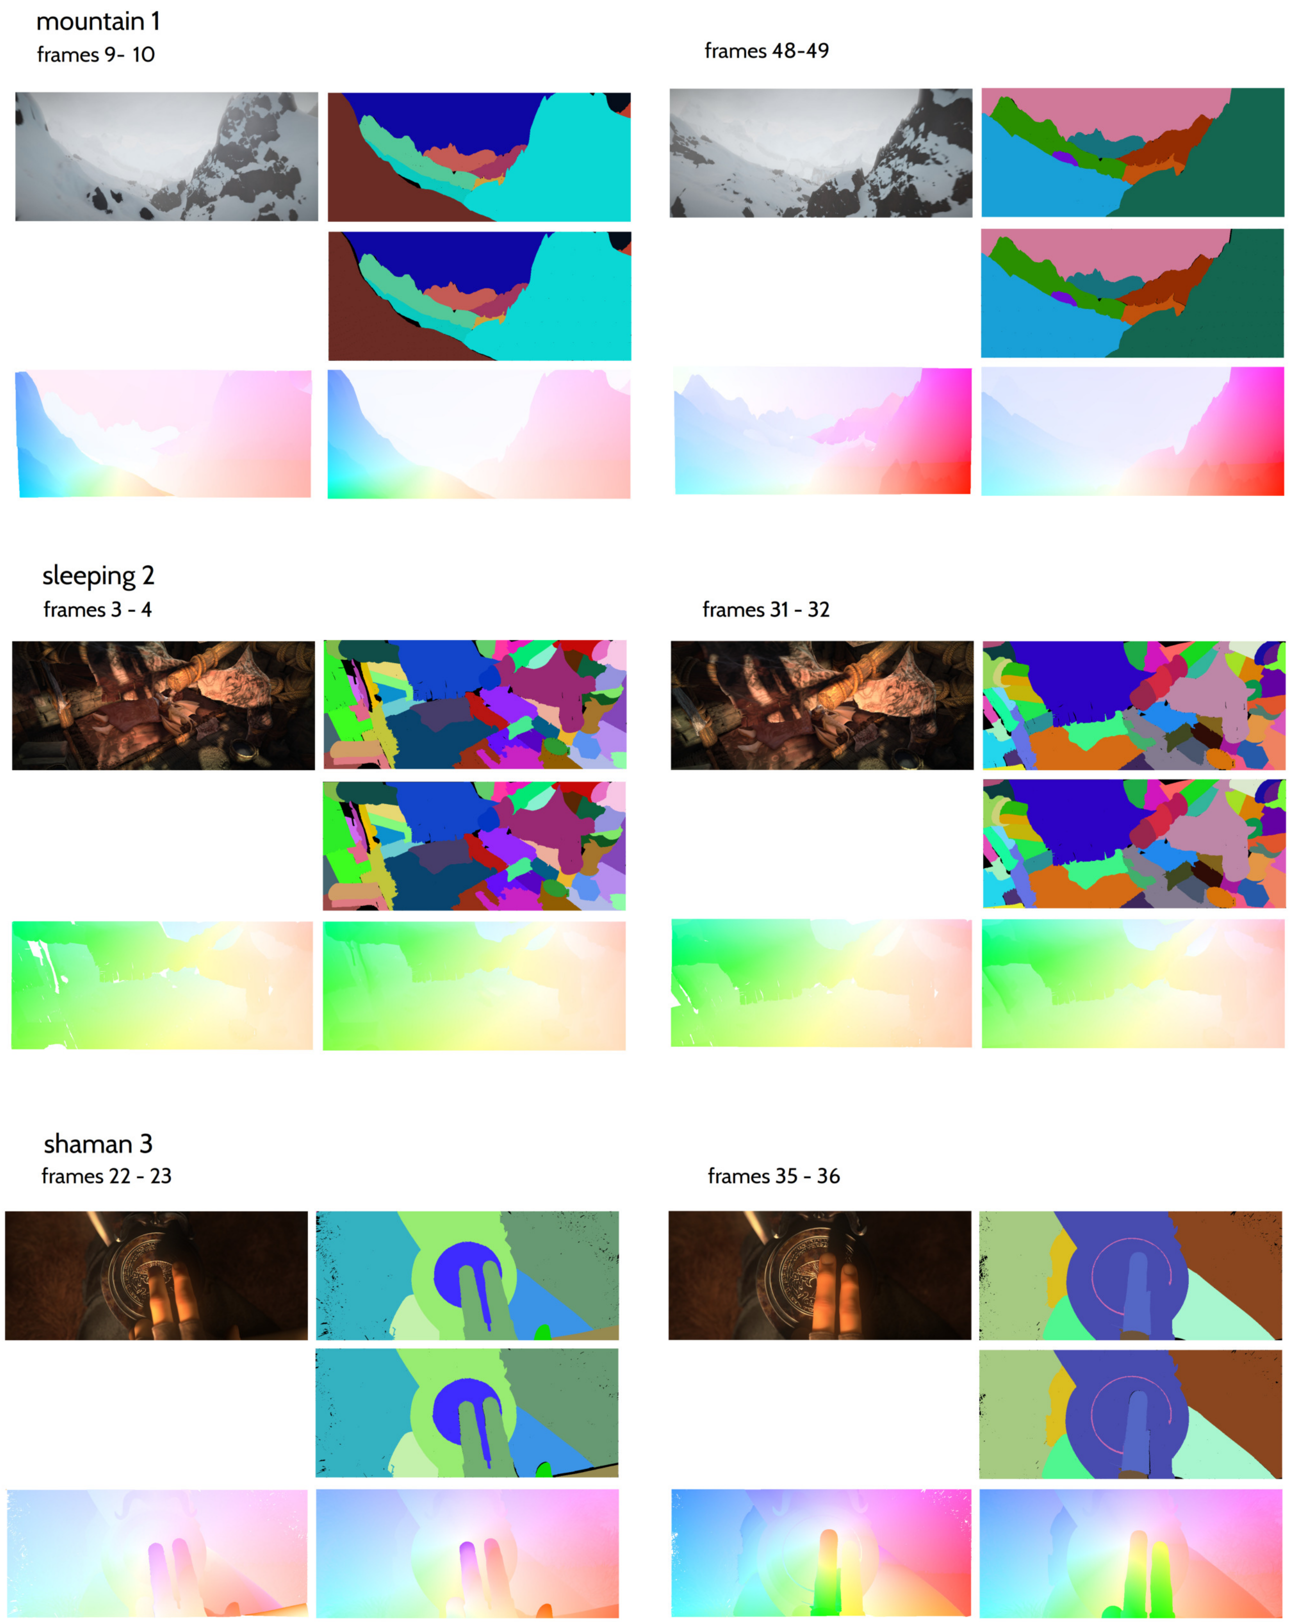

Figure 6: Additional visualisations of the *mountain1*, *sleeping2* and *shaman3* sequences. See Fig. 4 for the legend.