Unsupervised Learning of Consensus Maximization for 3D Vision Problems

Thomas Probst, Danda Pani Paudel, Ajad Chhatkuli, Luc Van Gool

Computer Vision Laboratory, ETH Zurich, Switzerland

1. Supplementary Material

1.1. 2D-2D Matching

In addition to the operating points in Fig. 8 in the paper, we provide the ROC curves of the Dino [5] and KITTI [2] datasets in Fig. 1. Especially on the KITTI dataset one can see the benefits of unsupervised finetuning on the real data, as there is a clear gap to the synthtetic pretraining, altough it is the easiest of the three datasets according to the other methods.

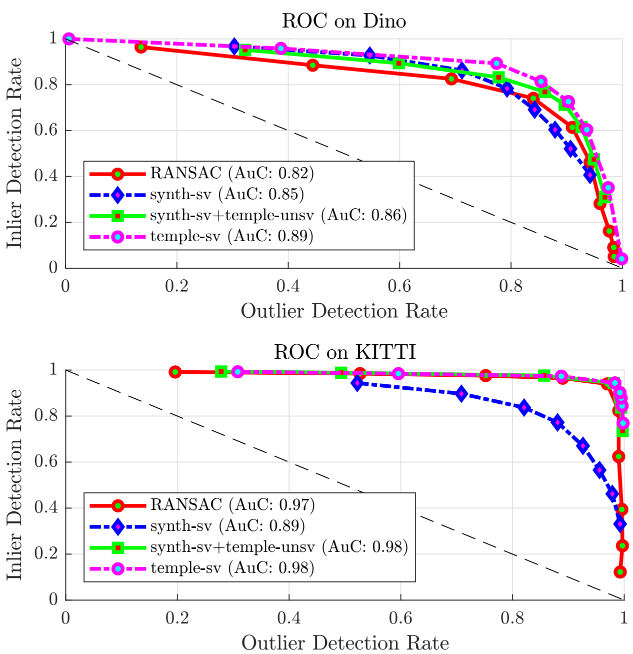

Figure 1. **2D-2D Fundamental matrix estimation.** ROC curve on Middlebury Dino (top) and KITTI (bottom).

1.2. Non-rigid matching

We present some qualitative results of our experiments on the FAUST dataset, using our network trained for non-rigid shape matching. In order to obtain challenging matches, we perform partial-to-full matching using matches from KM [3] and DFM [4]. In Fig. 2 and Fig. 3, each row shows the original matching between the reference and the target model, for different combinations of inter/intra subjects and matching methods KM and DFM, respectively. Our method is able to successfully remove outliers in the intra-subject cases (1st rows), at the expense of some inliers. For inter-subject case, our method can handle the matches provided by KM quite well (2nd row in Fig. 2), whereas DFM matches contain are very challenging outlier patterns (2nd row in Fig. 3), and only obvious sparse outliers can be removed.

References

[1] F. Bogo, J. Romero, M. Loper, and M. J. Black. FAUST: Dataset and evaluation for 3D mesh registration. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, Piscataway, NJ, USA, 2014. IEEE. 2, 3

[2] A. Geiger, P. Lenz, and R. Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012. 1

[3] Z. Lähner, M. Vestner, A. Boyarski, O. Litany, R. Slossberg, T. Remez, E. Rodolà, A. M. Bronstein, M. M. Bronstein, R. Kimmel, and D. Cremers. Efficient deformable shape correspondence via kernel matching. In *3DV*, 2017. 1, 2

[4] O. Litany, T. Remez, E. Rodola, A. M. Bronstein, and M. M. Bronstein. Deep functional maps: Structured prediction for dense shape correspondence. In *ICCV*, 2017. 1, 3

[5] S. M. Seitz, B. Curless, J. Diebel, D. Scharstein, and R. Szeliski. A comparison and evaluation of multi-view stereo reconstruction algorithms. *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, 1:519–528, 2006. 1

Figure 2. **Non-rigid matching on FAUST [1] using KM [3].** Each row shows the original matching between the reference and the target model, followed by the matches filtered by our network with color-coded correspondences. The top row shows an example for intra-subject, the bottom row inter-subject matching.

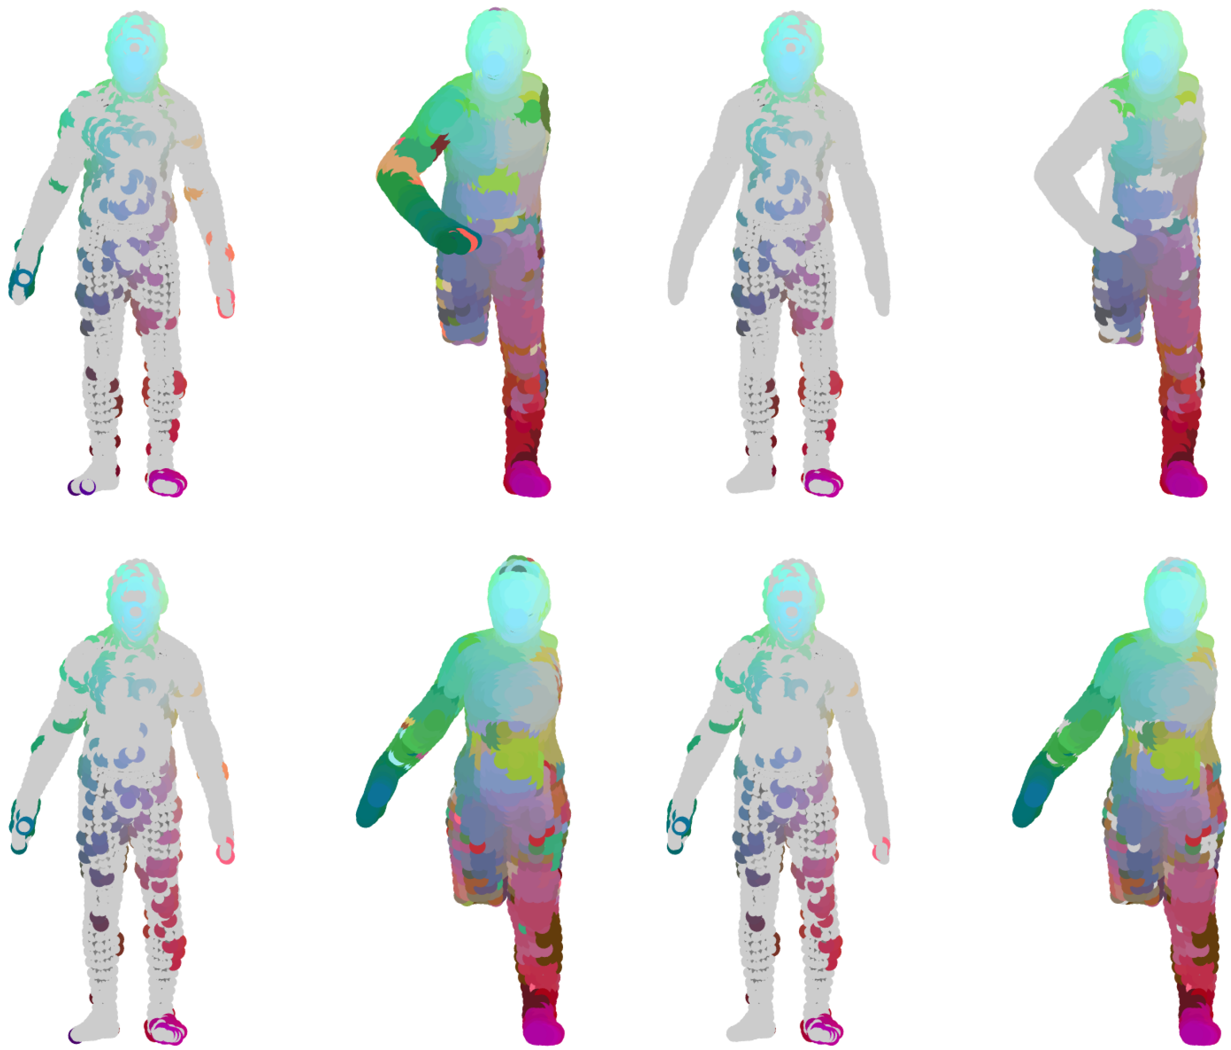

Figure 3. **Non-rigid matching on FAUST [1] using DFM [4].**. Each row shows the original matching between the reference and the target model, followed by the matches filtered by our network with color-coded correspondences. From top to bottom: The top row shows an example for intra-subject, the bottom row inter-subject matching, a failure case for our method.