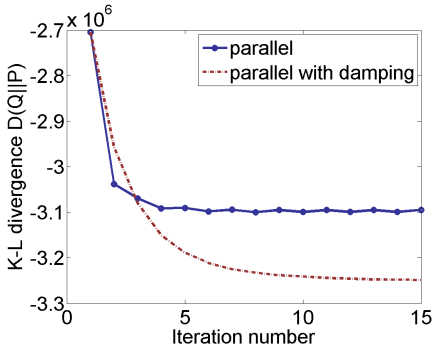

Figure 5. **Damping helps convergence.** The parallel mean field algorithm (blue circle) fails to converge for the fully connected layered model, while damping (red dot) helps the algorithm to converge to a better local minimum.

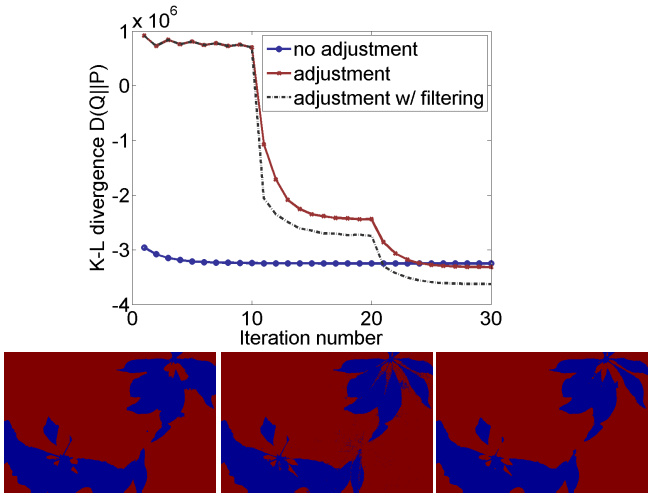

(a) no adjustment (b) adjustment (c) w/ filtering

Figure 6. **Adjustment and median filtering helps convergence.** They also lead to fewer speckles; see the image in Figure 7. (***better viewed on a computer screen***)

ally, the segmentation has fewer speckles with the adjusting and filtering.

4.1. Benchmark Sequences

Middlebury optical flow. The proposed method, **FC-2Layers-FF**, is ranked 11th on EPE and 7th on AAE in the public table at the time of writing (April 2013). Without the FlowFusion step, the algorithm still obtains reasonable results, as shown in Tables 1 and 2. We perform a bootstrap statistical significance test of the flow estimation results on the Middlebury training and test set for the algorithm with and without the FlowFusion step. The P-Values are 0.9602 and 0.8954, suggesting that the two have similar performance. For practical purposes, we can drop the computationally expensive FlowFusion step and still obtain acceptable results. Our dense two-layer **FC-2Layers-FF** model does not outperform the sparse multi-layer **Layers++** method; we

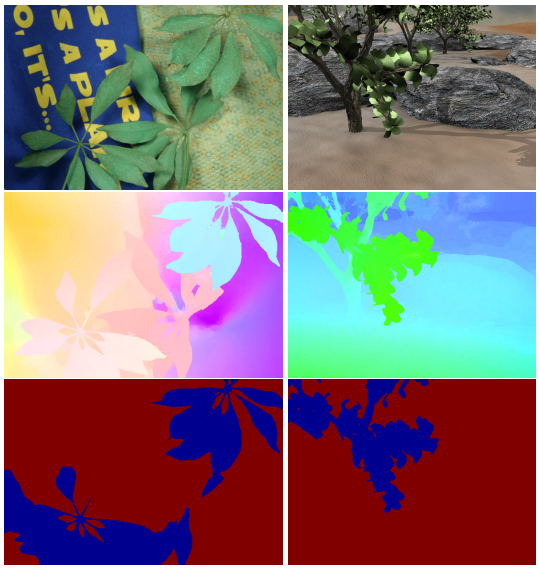

Figure 7. Example results on Middlebury. Top to bottom: First frame; Estimated flow; Segmentation.

expect future multi-layer formulations to further improve the performance of fully-connected models of layer support.

MIT layered segmentation. Figures 1 and 9 show some results on the MIT dataset. **FC-2Layers-FF** correctly segments the hand from the background. Compared with a local model, the fully-connected model can recover the background holes between the fingers. **FC-2Layers-FF** also recovers the fine structure of the bicycle in "Car3" (Figure 9). However, there are also failure cases that reveal the limitation of the fully-connected layered model. The reflections on the car and the woman's forehead are misleading color cues for the segmentation. As discussed in [15], long-range connections may propagate wrong information for small regions with different appearance from that of their true layer.

Historical sequences. We also apply **FC-2Layers-FF** to several other sequences, as shown in Figure 8. **FC-2Layers-FF** correctly segments the tree from the background in "flowergarden" and the person from the background in "8_org".

MPI Sintel. We apply **FC-2Layers-FF** to the MPI Sintel dataset [6] using the same parameters tuned on the Middlebury training set. As summarized in Table 3, **FC-2Layers-FF** performs better than **MDP-Flow2** on the more challenging final set. In the unmatched (occlusion) regions, **FC-2Layers-FF** is better on both sets than **MDP-Flow2**. As shown in Figure 10, **FC-2Layers-FF** captures the major occlusions in the scene and the segmentation is consistent with the scene structure. The estimated flow fields are visually close to the ground truth. Note that because the head in "shaman_2" has very different motion from the body in the four frames we used, it is reasonable that **FC-2Layers-FF** separates the head from the body.