

# Filter Flow: Supplemental Material

Steven M. Seitz

University of Washington

Simon Baker

Microsoft Research

We include larger images and a number of additional results obtained using Filter Flow [5].

## 1 Focus

In Figure 1 we include a larger version of our linear depth-from-defocus results (Figure 4 in [5]), along with the corresponding result from [6] and from [2]. In Figure 2 we include similar results on another dataset from [6], along with corresponding result from [6] ([2] does not include results on this sequence.)

In Figure 3 we include a larger version of our depth-from-defocus and motion results (Figure 5 in [5]).

In Figure 4 we include results to illustrate robustness in the presence of blur.

## 2 Intensity Stereo

In Figure 5 we include larger images of our results on the Rocks1 images from [3]. We include similar results on the Cloth4 sequence in Figure 6 and on the Wood1 sequence in Figure 7.

## 3 Optical Flow


In Tables 1 and 2 we include our results on the Middlebury Flow benchmark [1]. We include results both with and without the kernel filters. Table 1 contains the end-point error measure. Table 2 contains the interpolation intensity error. See <http://vision.middlebury.edu/flow/> for the other measures and a full set of images of our results with the kernel filters.

In Figure 8 we include our results on the challenging (because it has a large search range) Urban sequence, which are amongst the best to date. In Figure 9 we include our results on interpolation results on the Backyard sequence. Our algorithm is one of the few (if any) to render the orange ball reasonably well.


In Figure 10 we include a comparison of our results with and without a kernel filter to explain the relative performance of the algorithms in Table 1.




Input (one of two)



Kernel Filter Width<sup>2</sup> (Computed)



Median Filtered Results from [6]




Results from [2]


Figure 1: Defocus result on an image pair from [6] using our linear MRF-based formulation, showing estimated kernel width ( $W^2$ ) for each pixel (white small, black large.) Light values correspond to sharpening, dark to blur. We also include results from [6] and [2].



Input (one of two)



Kernel Filter Width<sup>2</sup> (Computed)



Raw, Unfiltered Results from [6]

Figure 2: Defocus result on an image pair from [6] using our linear MRF-based formulation, showing estimated kernel width ( $W^2$ ) for each pixel (white small, black large.) Light values correspond to sharpening, dark to blur. We also include the corresponding results from [6].




(a) Far Focused Image



(b) Near Focused Image



(c) The Affine Alignment



(d) Flow Color Coding



(e) The Symmetric Kernel Filters



(f) The Filter Width<sup>2</sup>

Figure 3: Defocus and Motion Estimation. Two images (a–b) at different focus settings yields an affine flow field which is roughly a zoom (c) and symmetric per-pixel kernels (e). The filter width<sup>2</sup> (f) visually correlates to scene depth. (d) Contains the color coding of the flow, from [1].



Left Image



Blurred Right Image



Without Kernel Filter



With Kernel Filter


Figure 4: Robustness Results in the Presence of Blur. (a) The left Venus image from [4]. (b) The right Venus image, blurred with a Gaussian,  $\sigma = 1.4$  pixels. (c) Filter Flow results with no kernel filter. (d) Filter Flow results with a kernel filter to model the blur.



(a) input image



(b) input, different illumination



(c) true depth




(d) true shading change



(e) disparity with illum.



(f) estimated shading change



(g) disparity without illum.


Figure 5: Stereo with Illumination Changes–Rocks1. (d) shows the shading changes (ratio of intensities) between images (a) and (b), as a result of moving the light source. Without illumination compensation, stereo performs poorly (g). Adding an MRF over shading changes greatly improves results (e), and the estimated shading (f) closely match (a regularized version of) the ground truth (d).




(a) input image



(b) input, different illumination



(c) true depth



(d) true shading change



(e) disparity with illum.



(f) estimated shading change



(g) disparity without illum.


Figure 6: Stereo with Illumination Changes–Cloth4. (d) shows the shading changes (ratio of intensities) between images (a) and (b), as a result of moving the light source. Without illumination compensation, stereo performs poorly (g). Adding an MRF over shading changes greatly improves results (e), and the estimated shading (f) closely match (a regularized version of) the ground truth (d).




(a) input image



(b) input, different illumination




(c) true depth



(d) true shading change



(e) disparity with illum.



(f) estimated shading change



(g) disparity without illum.

Figure 7: Stereo with Illumination Changes–Wood1. (d) shows the shading changes (ratio of intensities) between images (a) and (b), as a result of moving the light source. Without illumination compensation, stereo performs poorly (g). Adding an MRF over shading changes greatly improves results (e), and the estimated shading (f) closely match (a regularized version of) the ground truth (d).



Table 1: Quantitative Result on Middlebury Flow Benchmark: End-Point Error in Pixels.

	Army	Mequon	Schefflera	Wooden	Grove	Urban	Yosemite	Teddy
No Kernel Filter	0.20	0.94	0.93	1.23	1.13	0.71	0.22	1.16
With Kernel Filter	0.17	0.43	0.75	0.70	1.13	0.57	0.22	0.96

Table 2: Quantitative Result on Middlebury Flow Benchmark: Interpolation Error in Grey-Levels.

	Army	Mequon	Urban	Teddy	Backyard	Basketball	Dumptruck	Evergreen
No Kernel Filter	1.84	3.00	4.17	5.71	10.1	5.82	7.68	7.09
With Kernel Filter	1.82	2.93	4.10	5.58	10.2	5.70	7.63	7.11

## 4 Higher Order Smoothness

In Figure 11 we include a larger version of Figure 2 in [5] which includes a comparison of local MRF smoothness terms (centroid smoothness, 2nd order smoothness, affine smoothness, and filter smoothness.)

## 5 Linear Global Affine

Figure 12 contains the ground-truth and computed flow fields for our results in Section 6.1.



Urban Result



Urban GT

Figure 8: Our results on the challenging Urban sequence from the Middlebury Flow benchmark [1] are amongst the best.



Backyard Result



Backyard GT


Figure 9: Our results on the Backyard sequence in the Middlebury Flow benchmark [1] are one of the few (if any) that render the orange ball anything close to correct.



(a) Mequon Image 10



(b) Mequon Image 11



(c) Mequon GT



(d) Mequon Without Kernel



(e) Mequon With Kernel

Figure 10: Results with/without a Kernel Filter on the Mequon image pair from [1]. The input images (a) and (b) contain a shadow to the left of the leftmost figure that moves across the cloth background. Comparing with the ground-truth flow (c), our estimate of the flow (d) is erroneous in that region without the kernel filter. With the addition of a  $1 \times 1$  kernel filter, our estimate of the flow (e) is much better.




(a) Centroid Smoothness



(b) 2nd Order Smoothness



(c) Affine Smoothness



(d) Filter Smoothness

Figure 11: Comparison of MRF smoothness terms. 2nd order smoothness produces smoother disparities compared to centroid smoothness, but over-smooths discontinuities. Affine smoothness yields both smooth disparities and relatively clean discontinuities. Filter smoothness is not quite as smooth as affine, but does best at discontinuities, due to the implicit truncation. These were computed using the filter flow stereo approach presented in Section 6.2.

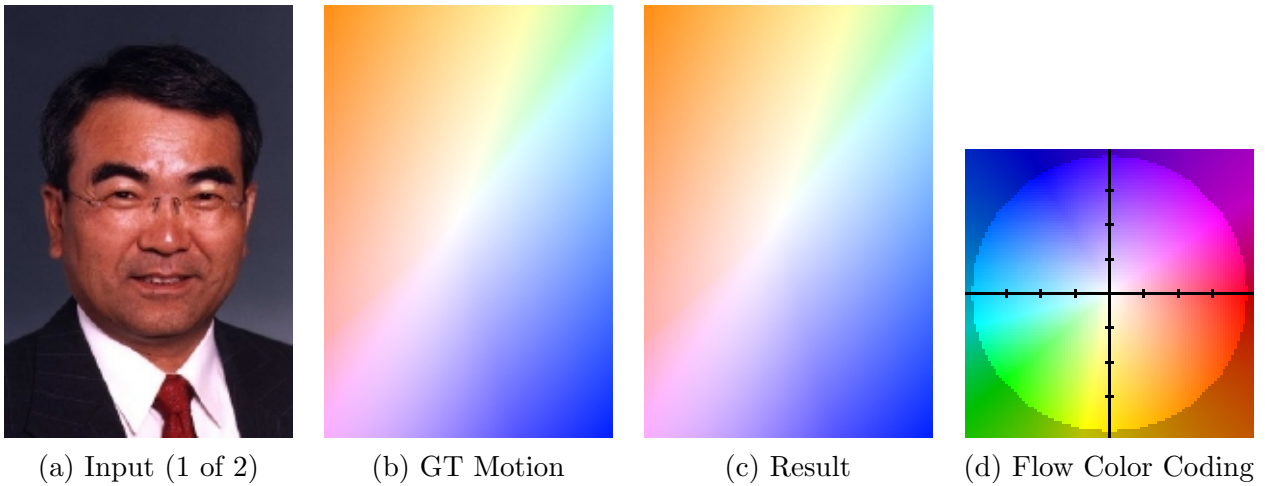


Figure 12: Linear affine alignment results. (a) The input image which was synthetically warped with an affine warp to generate the other input. (b) The ground-truth motion. (c) The motion computed by our linear algorithm. (d) The color coding of the flow, from [1].

## References

- [1] S. Baker, D. Scharstein, J. Lewis, S. Roth, M. Black, and R. Szeliski. A database and evaluation methodology for optical flow. In *Proc. ICCV*, 2007.
- [2] P. Favaro, S. Soatto, M. Burger, and S. Osher. Shape from defocus via diffusion. *PAMI*, 30(3):518–531, 2008.
- [3] H. Hirschmuller and D. Scharstein. Evaluation of stereo matching costs on images with radiometric differences. *PAMI*, 31(9):1582–1599, 2009.
- [4] D. Scharstein and R. Szeliski. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *IJCV*, 47(1-3):7–42, 2002.
- [5] S. Seitz and S. Baker. Filter Flow. In *Proc. Int. Conf. on Computer Vision*, 2009.
- [6] M. Watanabe and S. Nayar. Rational Filters for Passive Depth from Defocus. *IJCV*, 27(3):203–225, 1998.