Supplementary Material

Hierarchically-Constrained Optical Flow

Ryan Kennedy and Camillo J. Taylor
Department of Computer and Information Science
University of Pennsylvania
{kenry,cjtaylor}@cis.upenn.edu

1. Effect of segmentation error

Because our graphical model is set up based on a hierarchical image segmentation, it is sensitive to segmentation errors. In order to get a sense of how this affects the results, we perform the following experiment. For the training set of MPI-Sintel, rather than using the image itself to generate a segmentation, we run the segmentation algorithm on the colorized flow image of the groundtruth optical flow. The resulting segmentation should divide the image into regions based explicitly on the true motion field, rather than on color differences in the original image.

Quantitative results are given in Table 1, and several example images are shown in Figure 1. By using an "oracle" segmentation of the flow image, we obtain a 15.138% improvement on the training dataset. This indicates that segmentation error is a significant source of error in our approach, and subsequent research might get improved results by focusing on this issue. Observe that if we obtained the same percentage improvement on the test set, the resulting endpoint error would be 7.446, putting us much closer to the current state-of-the-art.

Table 1: Effect of segmentation error on MPI-Sintel. Rather than segmenting the image, we apply the segmentation algorithm to the colorized groundtruth flow image on the training dataset. This results in an improved segmentation. Otherwise, all parameters are kept constant.

	EPE	s0-10	s10-40	s40+
HCOF+multi	5.265	2.800	5.928	15.892
HCOF+multi (oracle)	4.468	2.225	5.076	13.609
% improvement	15.1%	20.5%	14.4%	14.4%

2. Effect of sub-pixel localization

Although our MRF is discrete by nature, we have described a method for obtaining sub-pixel localization at the last step of backtracing by fitting a quadratic function separately for the x and y offsets. The effect of this is shown in Figure 2. Without sub-pixel localization, the results are very blocky and a grid-like pattern is clearly visible. Although our sub-pixel localization does not completely eliminate this pattern, the results are nonetheless significantly smoother while there is negligible additional computation time.

(a) Without sub-pixel localization (b) With sub-pixel localization (c) Groundtruth flow

Figure 2: Effect of sub-pixel localization on a portion of the RubberWhale image from the Middlebury [1] dataset.

3. Effect of approximations

Although our tree-structured MRF can be solved optimally in polynomial time, we outlined several approximations used in our actual implementation in order to speed up the computation and lessen the memory requirements. Still, it is possible to compute an exact solution at the cost of additional computation. We examine the effect of our approximations by evaluating our algorithm on the Middlebury [1] dataset without such approximations. In particular, we do not subsample the pixels when computing data cost matrices on the upward pass of the optimization and compute a full data cost matrix for each pixel on the downward pass when labeling each pixel.

A comparison between the exact solutions and the solutions obtained when using our proposed approximations is given in Figure 3 (in both cases, sub-pixel localization is also used). The approximate and exact solutions to the MRF are virtually indistinguishable and have nearly identical er-

978-1-4673-6964-0/15/$31.00 ©2015 IEEE

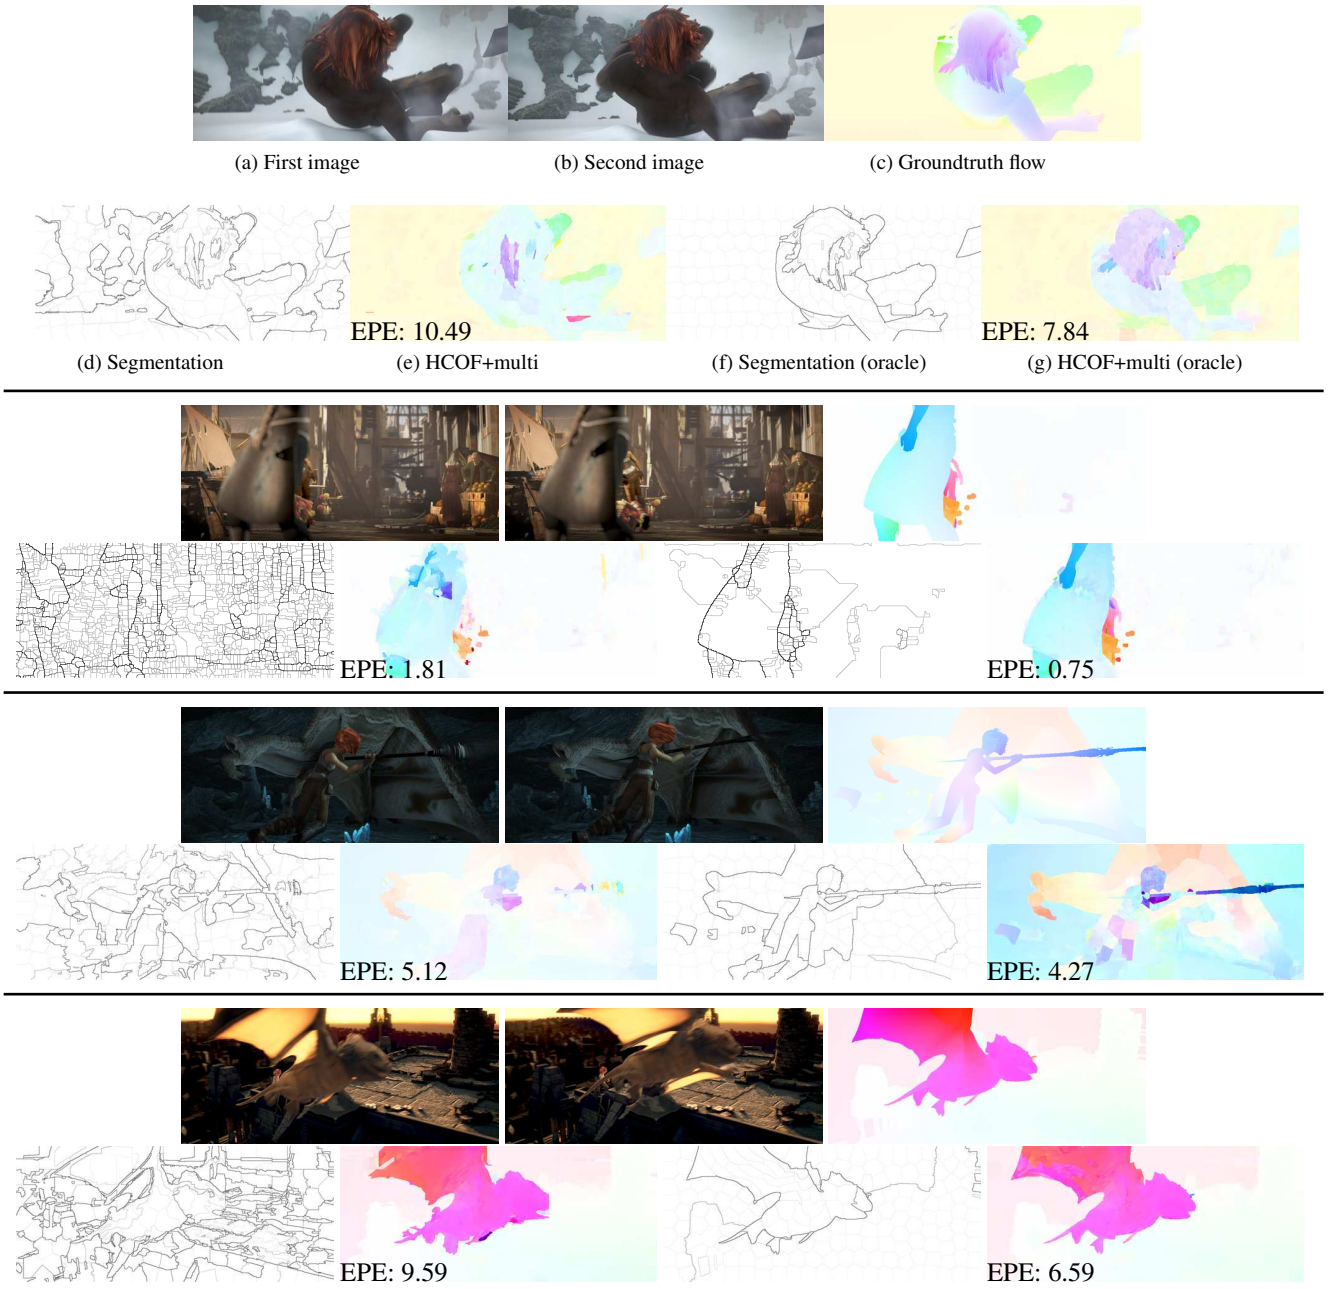

Figure 1: Example images showing the effect of segmentation error on optical flow results on MPI-Sintel. We show the optical flow results using a segmentation of the image, as well as when segmentation is performed on the colorized groundtruth flow image, denoted as the "oracle". The improved segmentation often results in improved flow estimates.

ror; the mean EPE for the approximate and exact solutions is 0.556 and 0.564, respectively. However, the computation time is very different: while the approximate solution is computed in about 5 minutes the exact solution takes nearly 45 minutes.

References

[1] S. Baker, D. Scharstein, J. Lewis, S. Roth, M. J. Black, and R. Szeliski. A database and evaluation methodology for optical flow. *IJCV*, 92(1):1–31, 2011. 1

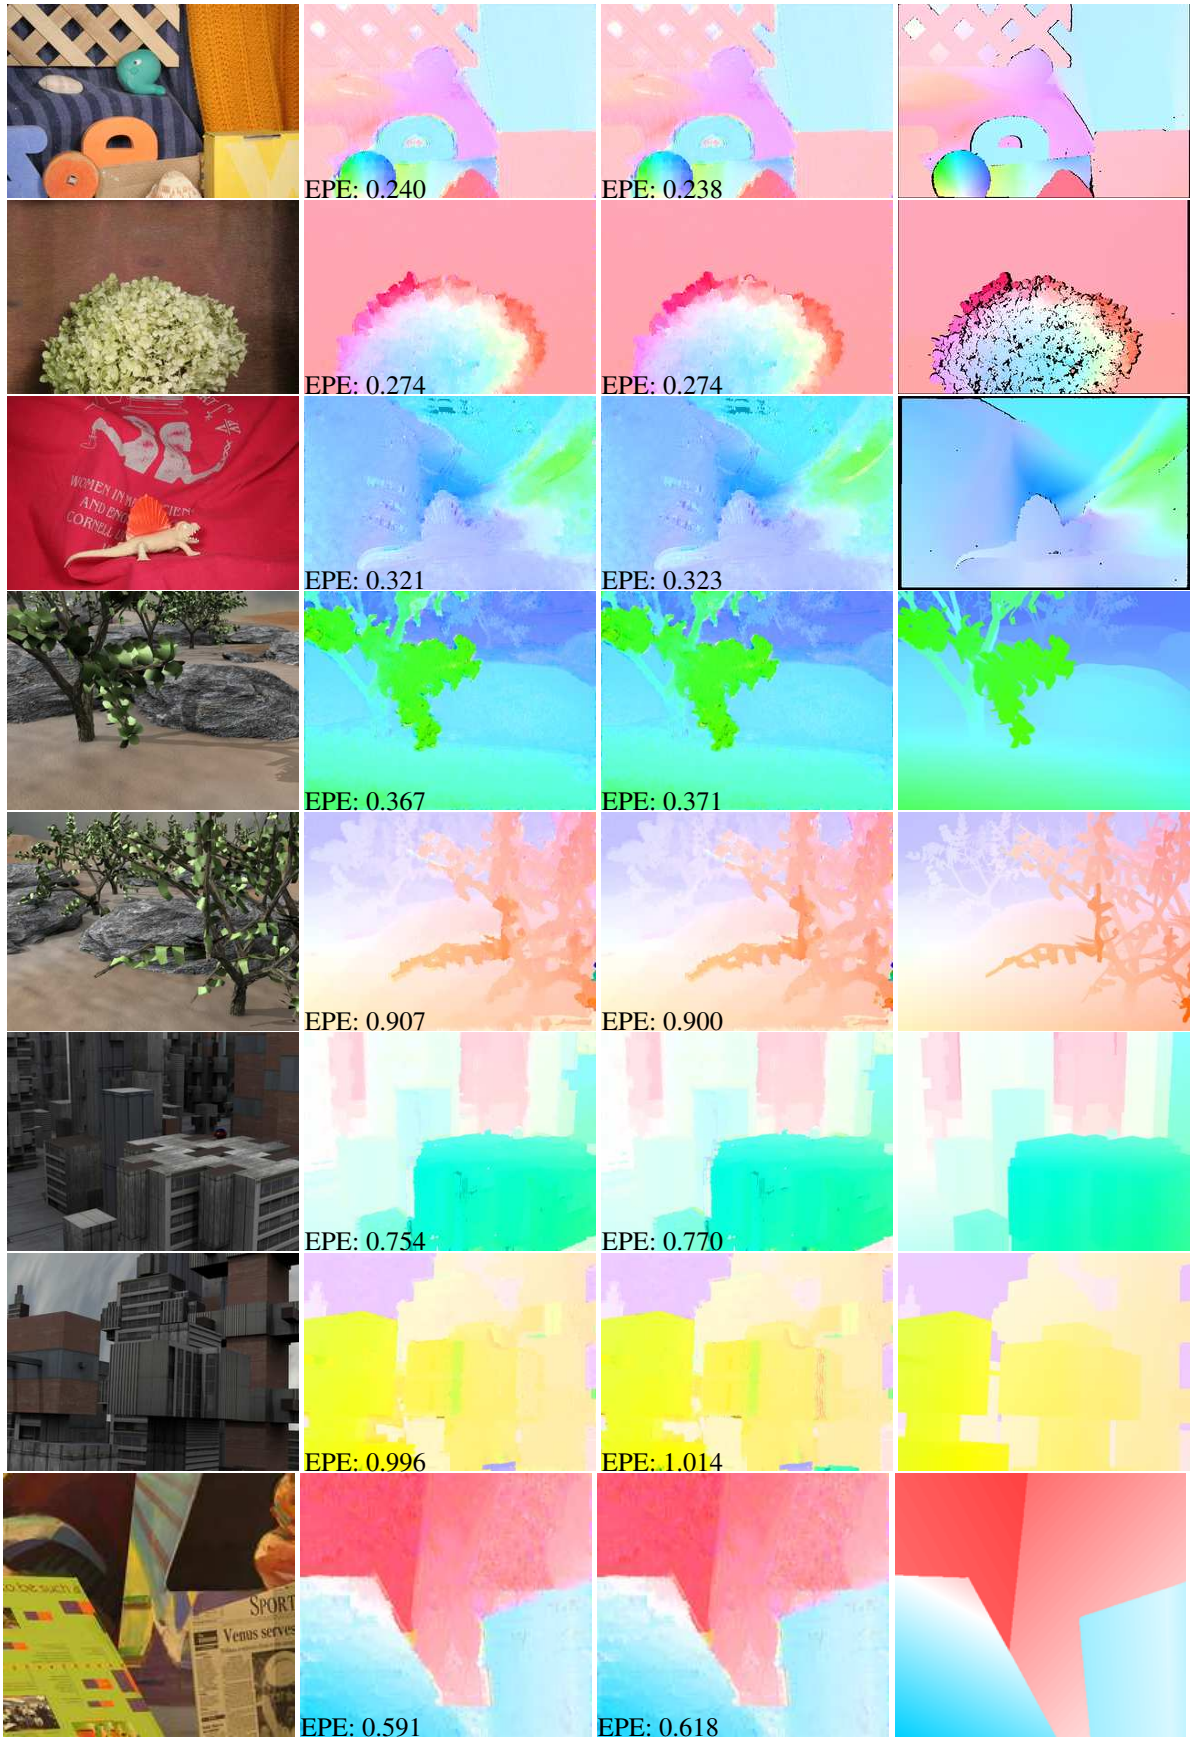

(a) First image (b) Approx. solution (c) Exact solution (d) Groundtruth flow

Figure 3: Results on the Middlebury training set. Colors for the estimated flow values are scaled based on the maximum offset of the groundtruth flow.