Learning to Synthesize Motion Blur
Supplement

Tim Brooks Jonathan T. Barron
Google Research

1. Additional Results

Because our synthetic dataset contains a validation set, we report performance of our model and its ablations in Table 1. We do not report the performance of our baseline techniques, as their performance on this synthetic data is unlikely to be meaningful when compared to our real test dataset, and also because some of our baselines needed to be run by the respective authors of each paper whom we did not wish to burden by requesting they process 15000 images in addition to our test set. In the table we see that the relative ordering of our model with respect to its ablations is consistent with their ordering in our test-set, though absolute performance is consistently higher.

Algorithm	PSNR	SSIM
Ours (direct pred.)	35.371	0.9854
Ours (kernel pred.)	36.762	0.9873
Ours (uniform weight)	37.217	0.9866
Our Model	37.673	0.9881

Table 1. Performance of our model and its ablations on the validation set of our synthetic dataset.

See Figures 1-4 for additional results on our real dataset, in which we compare our model against a set of ablations as well as a set of optical flow and video frame interpolation methods that could also be used to synthesize motion blurred images.

References

[1] Huaizu Jiang, Deqing Sun, Varun Jampani, Ming-Hsuan Yang, Erik G. Learned-Miller, and Jan Kautz. Super slomo: High quality estimation of multiple intermediate frames for video interpolation. *CVPR*, 2018.

[2] Simon Niklaus, Long Mai, and Feng Liu. Video frame interpolation via adaptive separable convolution. *ICCV*, 2017.

[3] Jerome Revaud, Philippe Weinzaepfel, Zaid Harchaoui, and Cordelia Schmid. EpicFlow: Edge-Preserving Interpolation of Correspondences for Optical Flow. *CVPR*, 2015.

[4] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. *CVPR*, 2018.

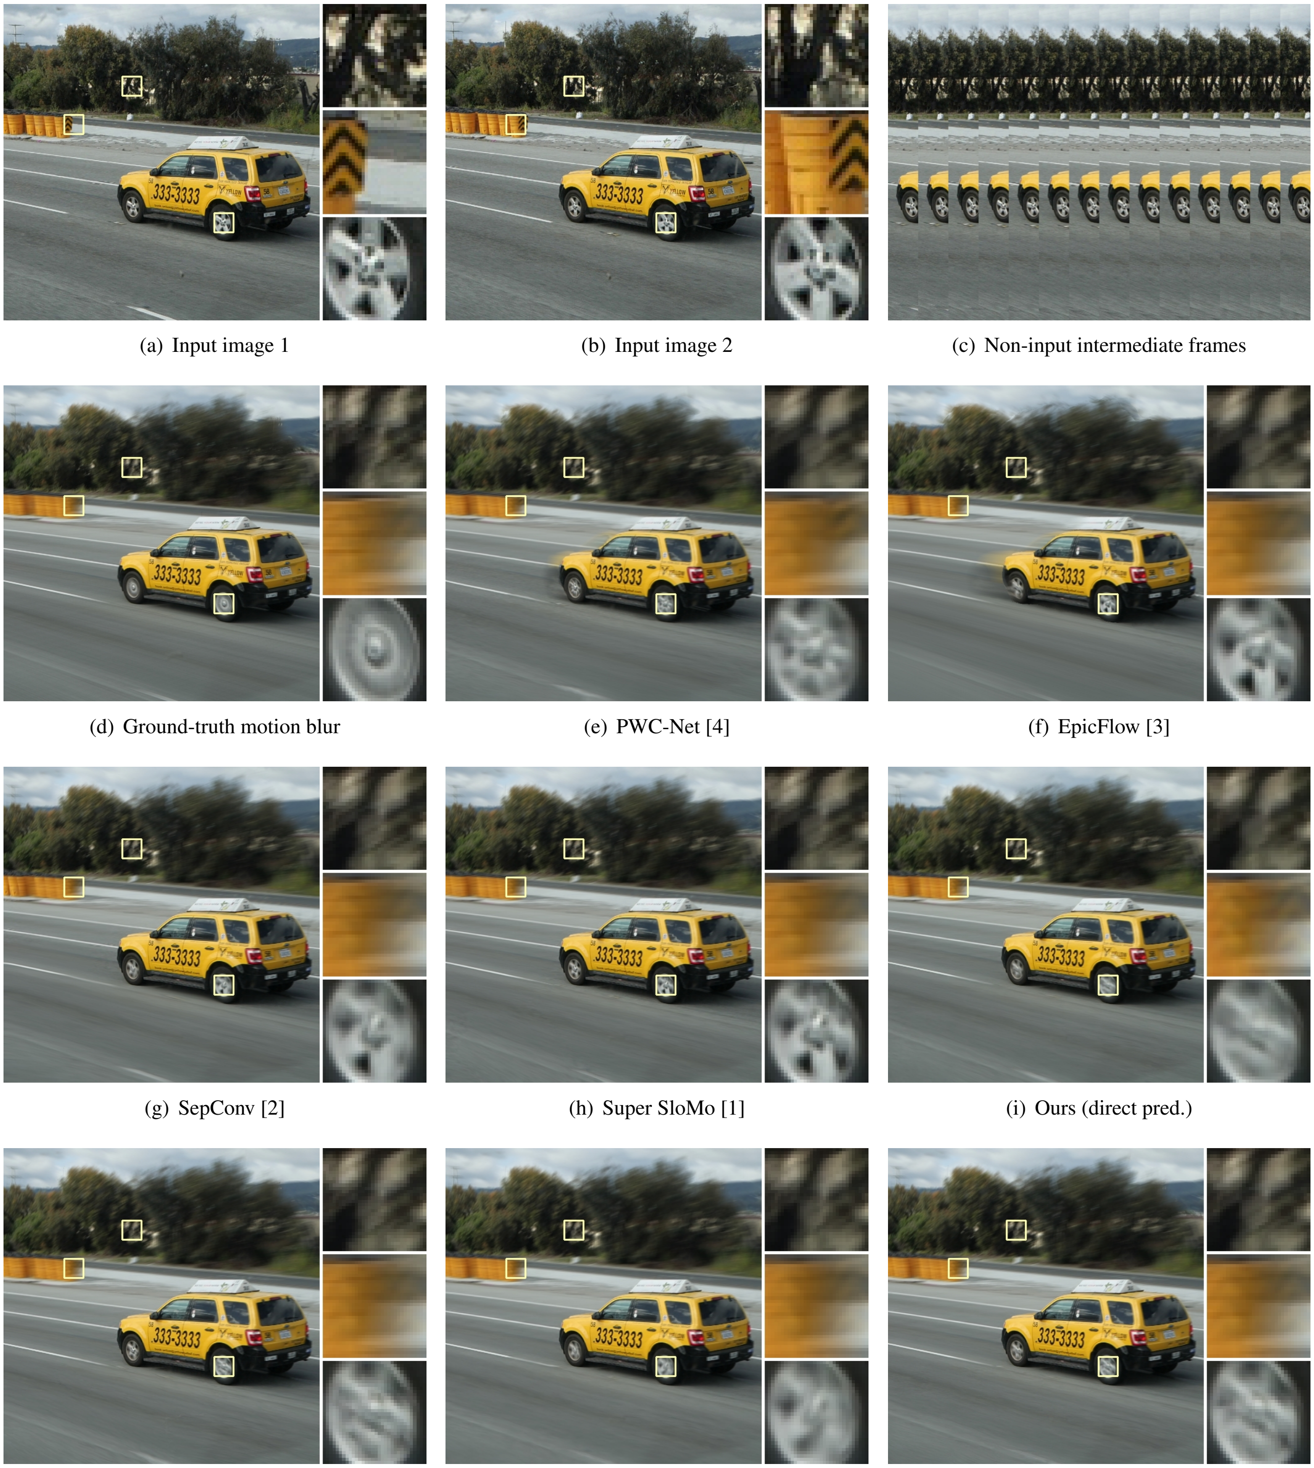

(a) Input image 1 (b) Input image 2 (c) Non-input intermediate frames

(d) Ground-truth motion blur (e) PWC-Net [4] (f) EpicFlow [3]

(g) SepConv [2] (h) Super SloMo [1] (i) Ours (direct pred.)

(j) Ours (uniform weight) (k) Ours (kernel pred.) (l) Our Model

Figure 1. Results for one scene from our test dataset. The ground truth image (d) is the sum of the input images (a) & (b) and of the frames between those two images (c). We programmatically select the three non-overlapping 32×32 sub-images with maximal variance across all frames in (c) and present crops of those regions, rendered with nearest-neighbor interpolation and sorted by their y-coordinates. We compare our model (l) against four baselines (e)-(h), and three ablations (i)-(k). Note that all techniques are unable to accurately blur the spinning wheel, which violates our model's and optical flow's assumption of linear motion.

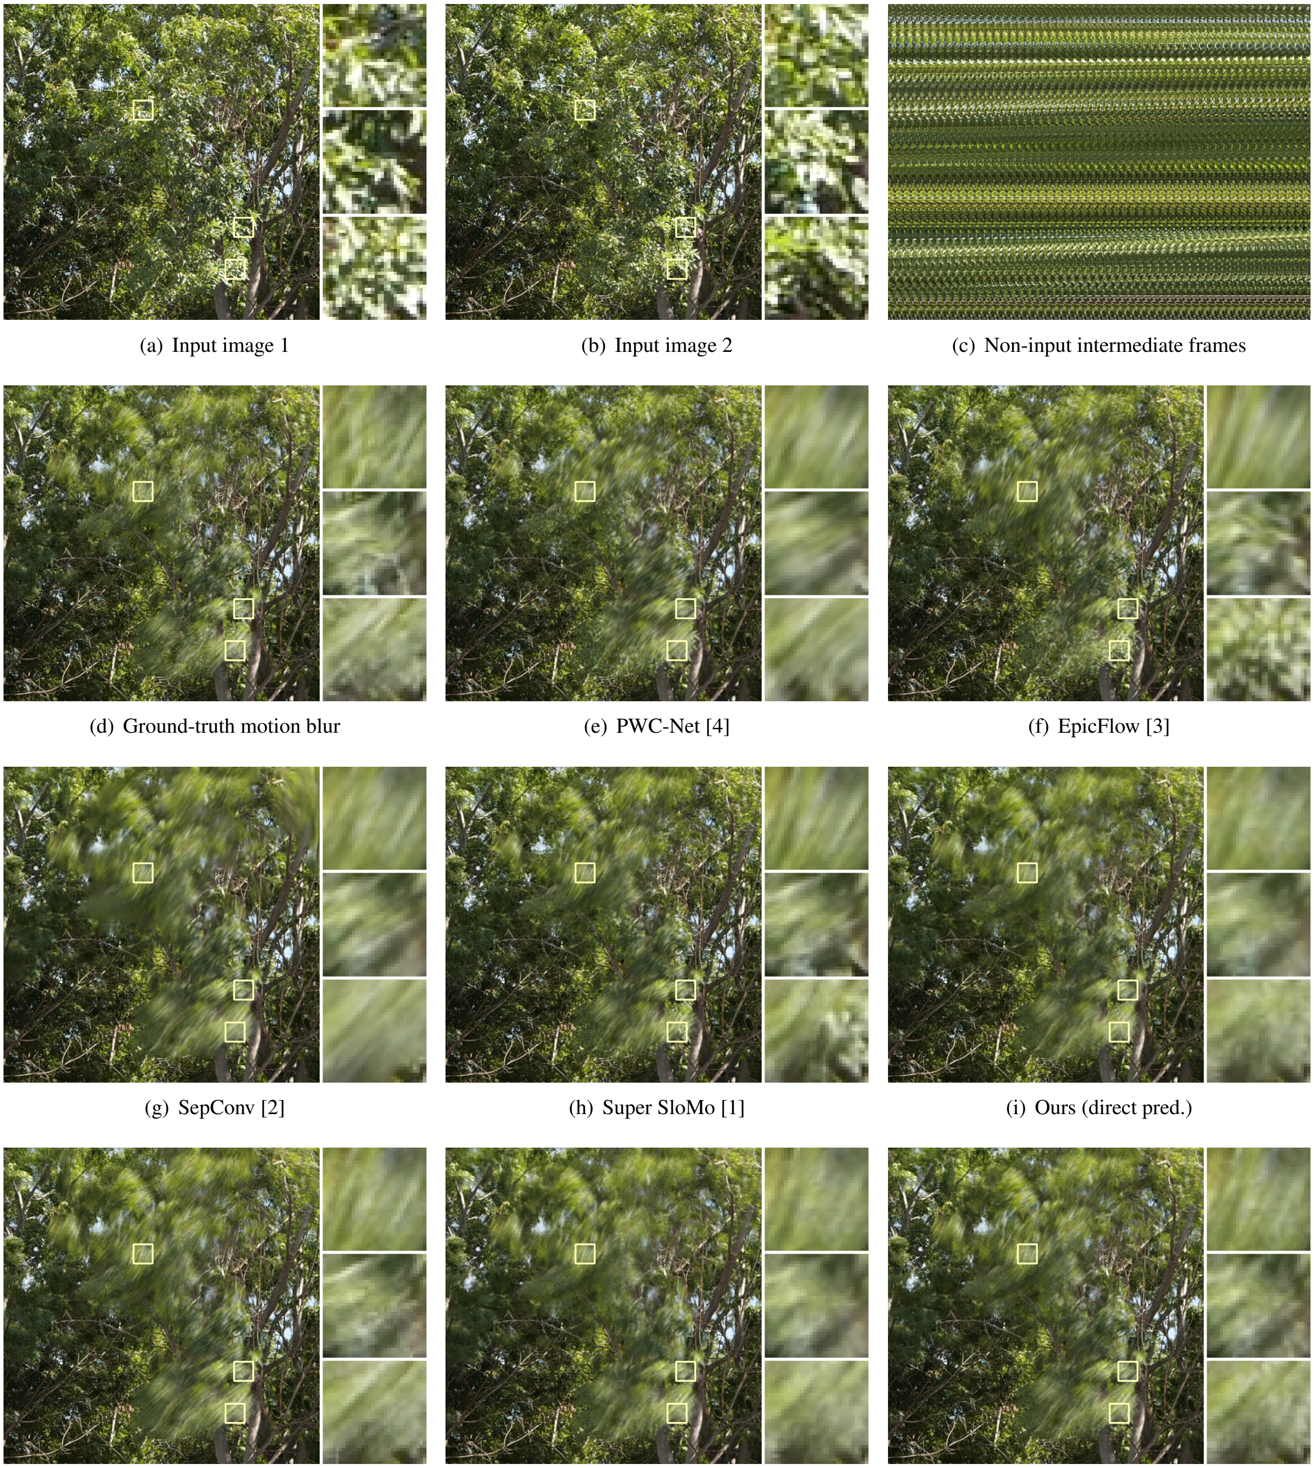

(a) Input image 1 (b) Input image 2 (c) Non-input intermediate frames

(d) Ground-truth motion blur (e) PWC-Net [4] (f) EpicFlow [3]

(g) SepConv [2] (h) Super SloMo [1] (i) Ours (direct pred.)

(j) Ours (uniform weight) (k) Ours (kernel pred.) (l) Our Model

Figure 2. Additional results in the same format as Figure 1.

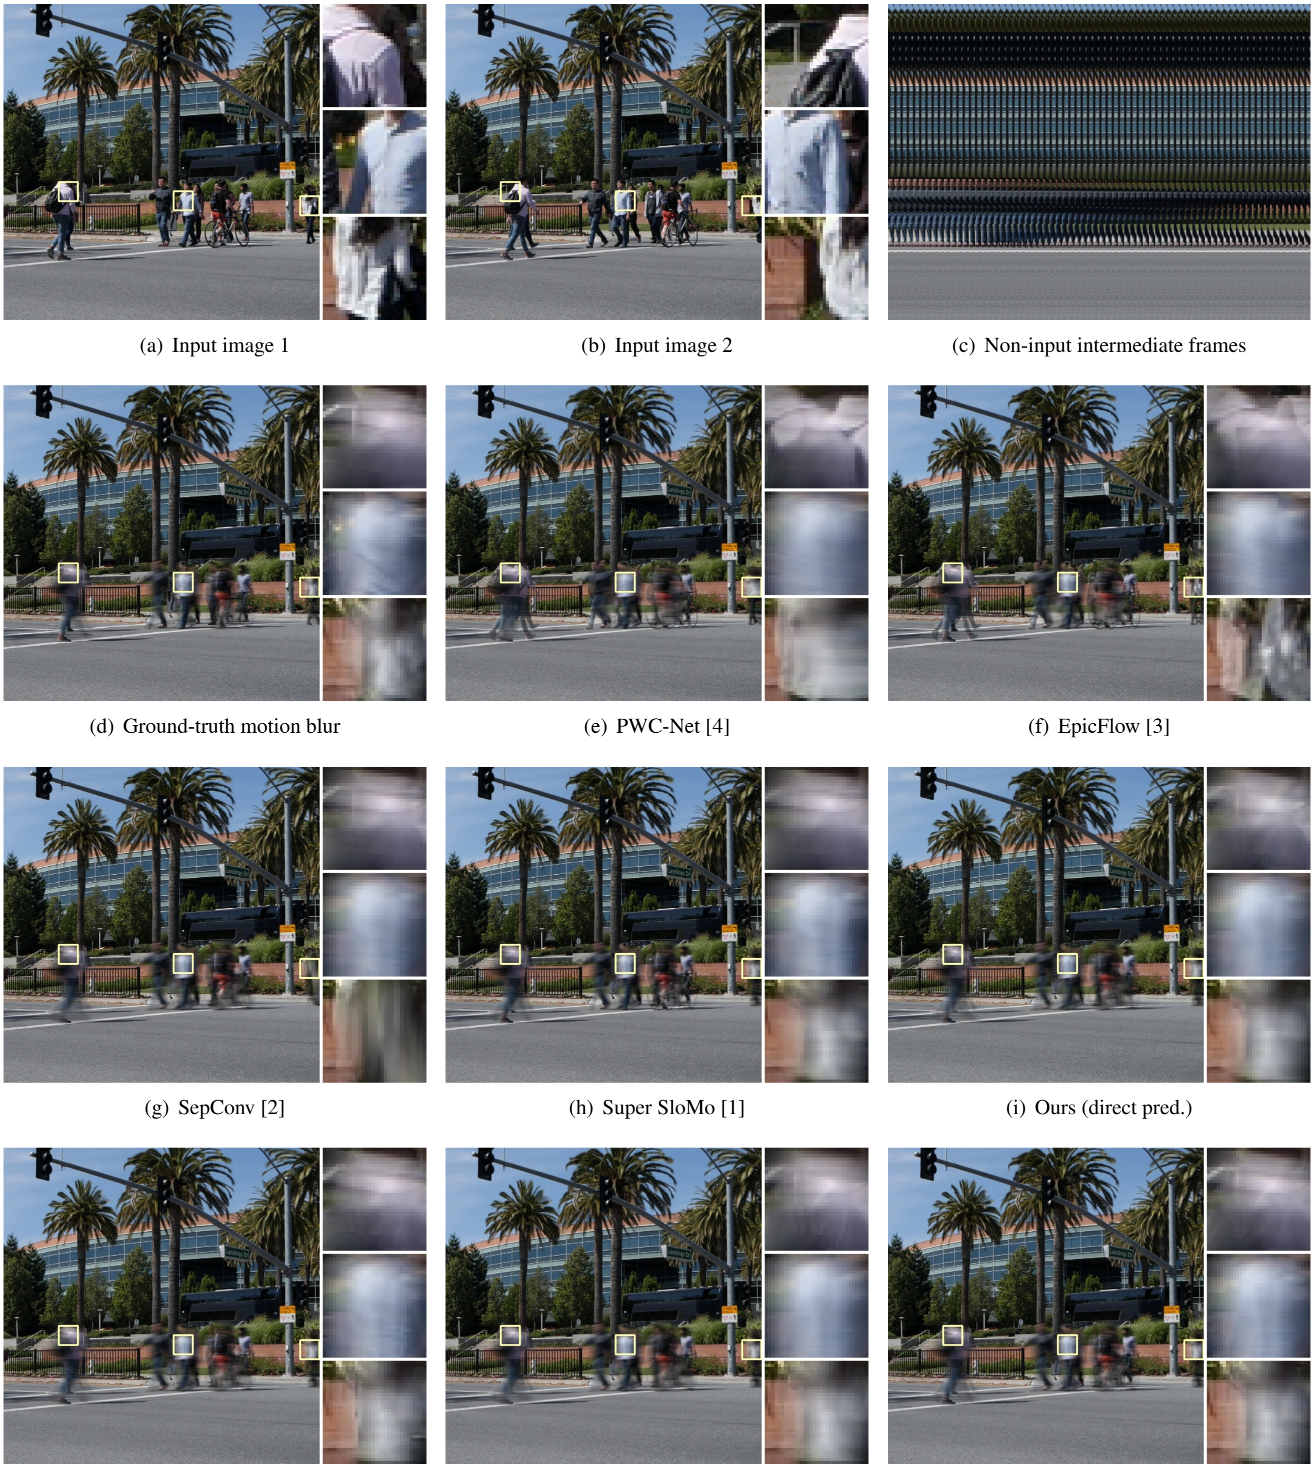

(a) Input image 1 (b) Input image 2 (c) Non-input intermediate frames

(d) Ground-truth motion blur (e) PWC-Net [4] (f) EpicFlow [3]

(g) SepConv [2] (h) Super SloMo [1] (i) Ours (direct pred.)

(j) Ours (uniform weight) (k) Ours (kernel pred.) (l) Our Model

Figure 3. Additional results in the same format as Figure 1.

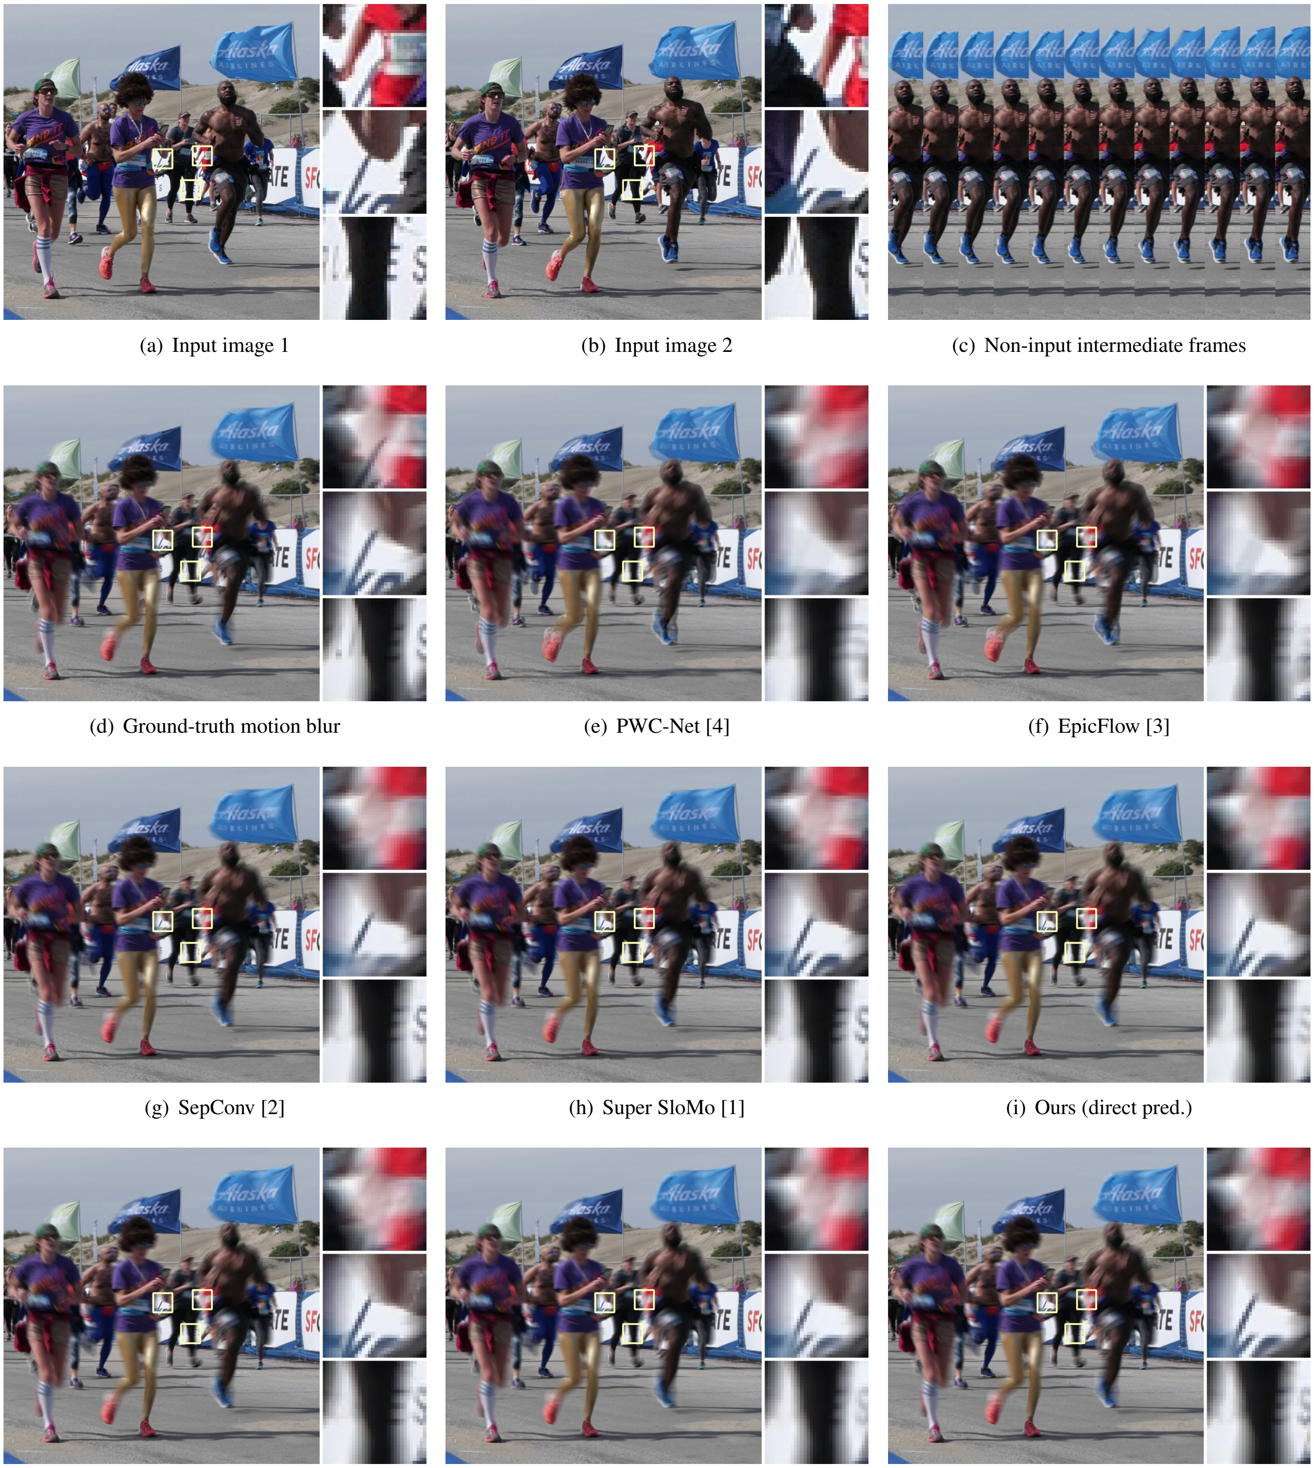

(a) Input image 1

(b) Input image 2

(c) Non-input intermediate frames

(d) Ground-truth motion blur

(e) PWC-Net [4]

(f) EpicFlow [3]

(g) SepConv [2]

(h) Super SloMo [1]

(i) Ours (direct pred.)

(j) Ours (uniform weight)

(k) Ours (kernel pred.)

(l) Our Model

Figure 4. Additional results in the same format as Figure 1.