

Distribution Fields with Adaptive Kernels for Large Displacement Image Alignment-Supplementary Material

Benjamin Mears
bmears@cs.umass.edu

Laura Sevilla Lara
lsevilla@cs.umass.edu

Erik Learned-Miller
elm@cs.umass.edu

School of Computer Science
UMass Amherst
Amherst, MA

1 Supplementary Material

The supplementary material provides our alignment code and the images we tested on in the paper along with additional images. See the README included with the supplementary material for the usage of the code.

See Figure 1 for example image regions we tested on. Note that the set of images (Takeo, Car, Knee, and Simon) we tested on in the paper includes the images used by Baker and Matthews [1] and Evangelidis and Psarakis [2]. See Figure 2 for the results of the parameter experiments and Figure 3 for the results of the convergence experiments. The details of each of these experiments are described in the paper and the figures presented here give results for an extended set of images.

For the parameter experiments, the "Search" method performs equal to or comparably to the best set of fixed parameters for all of the image regions. For the convergence experiments, the DF method is either the best performing algorithm or among the best performing algorithms on all the images except Image 8. Interestingly, the DF SIFT method performs very well on this image. A possible explanation for this is that the image region is relatively uniform which causes the DF method based on image intensity to get stuck in local minima while the DF SIFT method is based on gradients and the more sparse gradients are less likely to cause the method to get stuck in local minima.

References

- [1] S. Baker and I. Matthews. Lucas-Kanade 20 years on: A unifying framework. *IJCV*, 56(3):221–255, 2004.
- [2] G.D. Evangelidis and E.Z. Psarakis. Parametric image alignment using enhanced correlation coefficient maximization. *PAMI*, 2008.

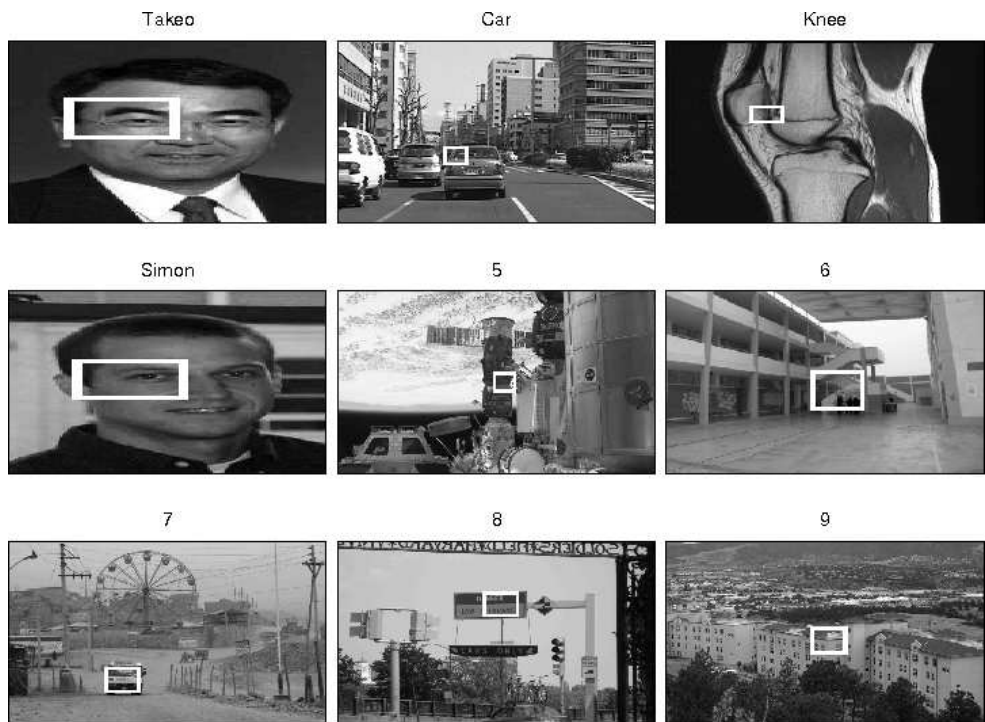


Figure 1: Example image regions tested on, including those used in the paper (Takeo, Car, Knee, and Simon).

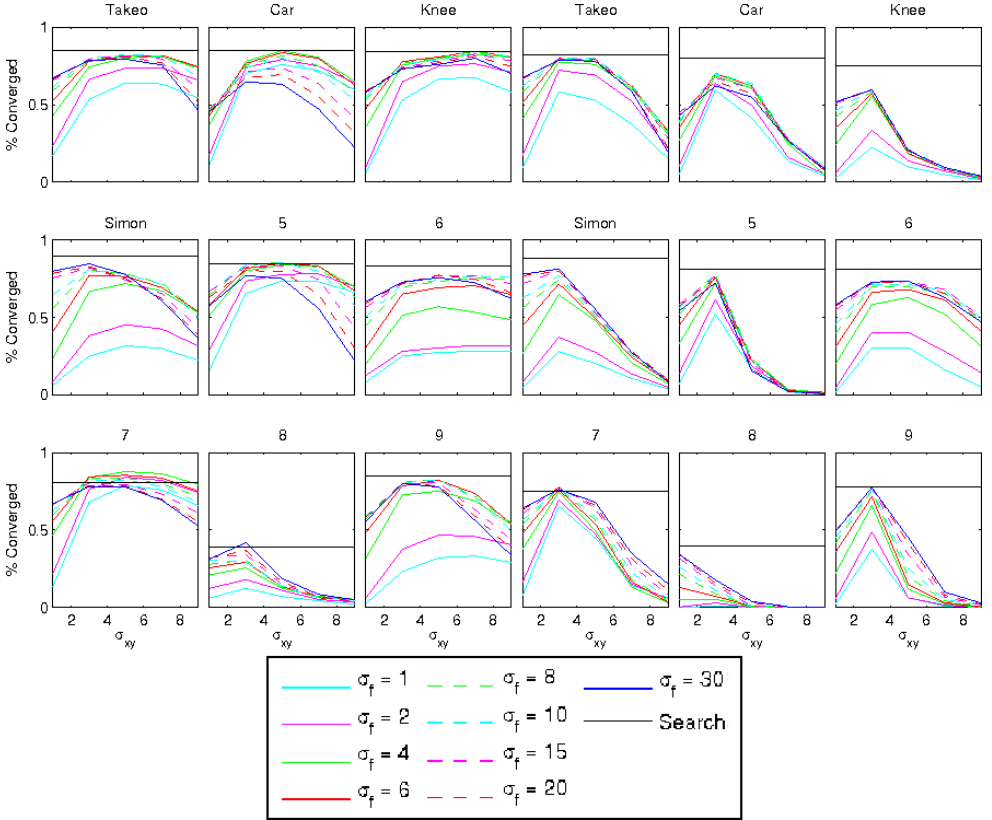


Figure 2: Results of experiments testing the use of fixed kernel parameters versus automatic selection. The first three columns show convergence rates for image pairs that differ by a geometric, but not a photometric, transformation and the last three columns show convergence rates for image pairs that differ by both a geometric and photometric transformation. “Search” refers to the automatic selection of kernel parameters.

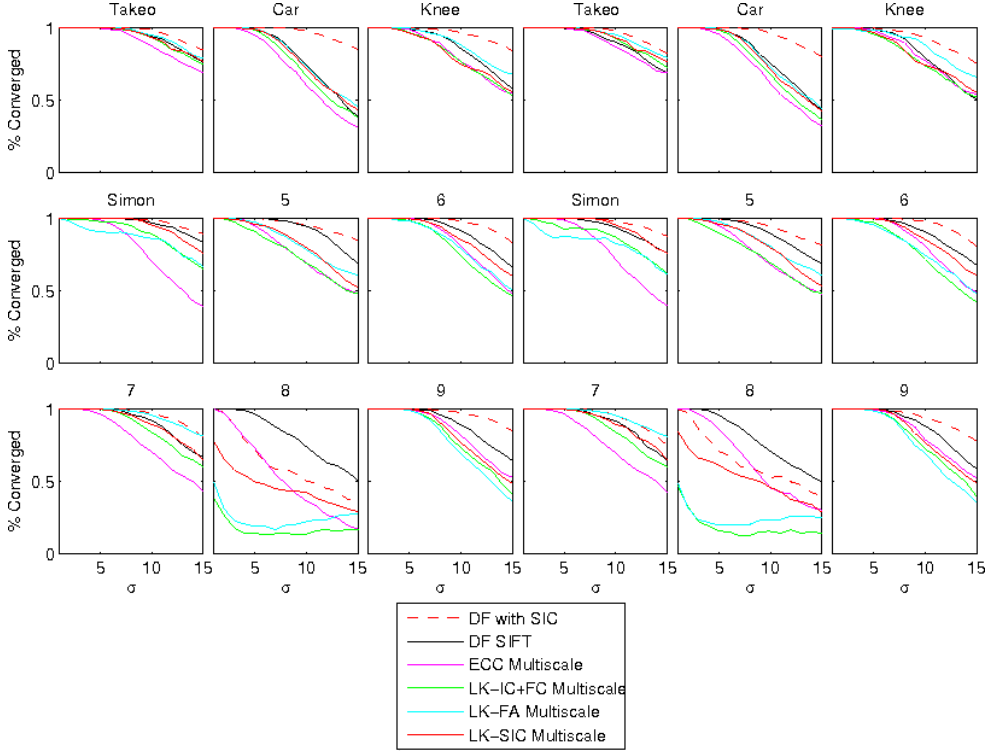


Figure 3: Results of experiments comparing convergence rates of different algorithms. The first three columns show convergence rates for image pairs that differ by a geometric, but not a photometric, transformation. The last three columns show convergence rates for image pairs that differ by both a geometric and photometric transformation, and that have Gaussian noise with standard deviation of 8 added to their pixel values.