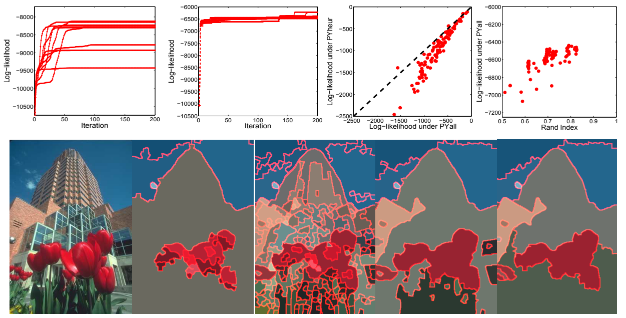

Figure 3. **Model and inference comparison**. *TOP (Left to right)* Log-likelihood (ll) trace plots of mean field runs, search runs, scatter plot comparing PYall and PYheur, scatter plot of ll vs Rand index. *BOTTOM (Left to right)* Test image, partitions with highest and lowest ll found by mean field, best and worst search partitions.

	BSDS300							LabelMe	
	Ncuts	MS	FH	gPb	PYheur	PYdist	PYall	gPb	PYall
PRI	0.73	0.77	0.77	0.80	0.60	0.69	0.76	0.74	0.73
segCover	0.40	0.48	0.53	0.58	0.45	0.50	0.54	0.54	0.55

Table 1. Quantitative performance of various algorithms on BSDS300 and LabelMe.

log likelihoods and rand indexes, again verifying that the partitions favored by our model are also favored by humans.

Comparison against competing methods. In this paper, our goal is not to produce one "optimal" segmentation but to provide a tractable handle on the posterior distribution over image partitions. Nevertheless, here we demonstrate that by summarizing the posterior with the MAP partition we produce results which are competitive with the state-of-the-art segmentation techniques. We compare against four popular segmentation techniques: Mean Shift (MS) [5], Felzenszwalb and Huttenocher's graph based segmentation (FH) [7], Normalized cuts [23] and gPb contour based segmentation [2][7]. In addition, we also compare against a version of our model which uses only distance cues for learning the covariance kernel (PYdist). Table 1 displays the quantitative numbers achieved on the BSDS300 test set. Figure 4 demonstrates qualitative differences amongst the methods. PYall is significantly better than both PYheur and PYdist. According to a Wilcoxon's signed rank test (at an 0.01 significance level) it is also significantly better than Ncuts and MS (on segCover metric, within noise on PRI), within noise of FH and statistically worse than gPb on the BSDS300 dataset.

Next, in order to test generalizability, we compare PYall against the top performing method on BSDS – gPb on the LabelMe dataset. The parameters for either method were tuned on BSDS and were not re-tuned to the LabelMe dataset. Table 1 displays the results. PYall and gPb are now statistically indistinguishable.

[7] All model parameters were tuned by performing a grid search on the training set. See supplement for more details.

Figure 5. **Diverse Segmentations.** Each row depicts multiple partitions for a given image. Partitions in the second column are the MAP estimates. Other partitions with significant probability masses are shown in the third and fourth columns.

Posterior Summary. Perhaps, a more accurate assessment of our model involves exploring the posterior distribution over partitions. In Figure 5 we summarize the posterior distributions, for a few randomly chosen test images, by presenting a set of high probability partitions discovered by our algorithm. It is worth noting that the set of multiple partitions produced by our method is richer than those produced by a single multi-resolution segmentation tree [2]. For instance, the partitions in the third and fourth columns of the first two rows of Figure 5 are mutually inconsistent with any one segmentation tree, but are nonetheless produced by our algorithm. More interesting ways of leveraging the distribution over partitions is an important direction of future work.

7. Discussion

Starting with a promising Bayesian nonparametric model of images partitions, we have developed substantially improved algorithms for learning from example human segmentations, and robustly inferring multiple plausible segmentations of novel images. By defining a consistent distribution on segmentations of varying resolution, this dependent PY process provides a promising building block for other high-level vision tasks.

References

[1] M. Andreetto, L. Zelnik-Manor, and P. Perona. Non-parametric probabilistic image segmentation. In *ICCV*, 2007.

[2] P. Arbelaez, M. Maire, C. Fowlkes, and J. Malik. From contours to regions: An empirical evaluation. *CVPR*, 0:2294–2301, 2009.

[3] R. Borsdorf, N. J. Higham, and M. Raydan. Computing a nearest correlation matrix with factor structure. *SIAM J. Matrix Analysis App.*, 31(5):2603–2622, 2010.

[4] C. Carson, S. Belongie, H. Greenspan, and J. Malik. Blobworld: Image segmentation using expectation-maximization and its application to image querying. *PAMI*, 24(8):1026–1038, Aug. 2002.