Single Image Shadow Detection using Multiple Cues in a Supermodular MRF

Tomás F. Yago Vicente
tyagovicente@cs.stonybrook.edu
Chen-Ping Yu
cheyu@cs.stonybrook.edu
Dimitris Samaras
samaras@cs.stonybrook.edu

Image Analysis Lab
Computer Science Dept.
Stony Brook University
Stony Brook, NY, USA

In this work we present a complete methodology for single image shadow detection based on the learned appearance of shadows. Our main contributions are:

- A novel single region classifier with a multikernel model specifically tailored for shadow region classification with SVM that outperforms the state of the art.
- A novel boundary classifier for shadow boundaries cast over surfaces with the same material, and two improved paired regions classifiers.
- We pose shadow detection as an MRF binary labelling problem that combines our accurate single region classifier with both boundary cues and illumination relationships between pairs of regions of the same material obtained from precise classifiers.

We propose a single region shadow classifier based on a multikernel SVM. Our multikernel model is a linear combination of χ^2 and Earth Mover's Distance(EMD)[5] kernels that operate on texture and color histograms disjointly. This single region classifier already outperforms the more complex state of art methods, without performing MRF/CRF optimization.

The local appearance of a single region is often ambiguous. Even for a human observer it can be hard to discern if a region is in shadow or not, without considering its context. Hence, it is sensible to look beyond the boundaries of a single region to decide its shadow label [1] [6]. In contrast to previous work we strive to use such contextual information sparingly. For MRF optimization reasons we prefer that most of the work is handled by the single region classifier (unary MRF potentials), with sparse pairwise connections that smooth the label changes across regions. We build on the work of [1] to propose our own improved pairwise classifiers but constrained to adjacent regions: for pairs of regions sharing the same material and same illumination condition, and for same material pairs viewed under different illumination (first lit, second in shadow). We also propose a shadow boundary classifier. Since shadow boundaries often overlap with reflectance changes confounding the effects of the illumination change, our classifier focuses on boundaries of shadows cast over surfaces with the same underlying material.

We integrate our single region classifier, our pairwise classifiers, and our boundary classifier using an MRF. Confident positive predictions of the pairwise and boundary classifiers are used to define the pairwise potentials and the graph topology of the MRF. The unary potentials are defined based on the single region classifier. We want to minimize the following functional:

$$E(\mathbf{x}) = \sum_{i \in \mathbf{R}} \phi(x_i) + \sum_{i,j \in \Omega^s} \psi_s(x_i,x_j) + \sum_{i,j \in \Omega^d} \psi_d(x_i,x_j) + \sum_{i,j \in \beta} \psi_b(x_i,x_j)$$

where Ω^s, Ω^d denote pairs of regions of the same material confidently predicted to have same and different illumination respectively, and β is the set of confidently predicted shadow boundaries. The key idea is to combine a very strong region classifier with a set of precise boundary and paired region classifiers that define the smoothing in the MRF. In order to achieve sparsity and increase reliability we favor precision over recall in the outcome of the context based classifiers. The same material/same illumination relationships induce submodular potentials (ψ_s). Potentials due to the boundary classifier (ψ_b) are supermodular. The directionality of the same material different/illumination classifier introduces asymmetric potentials (ψ_d). We minimize the MRF energy using QPBO [3] [4], the asymmetric potentials led to a minor modification of the initial reparameterization stage of QPBO.

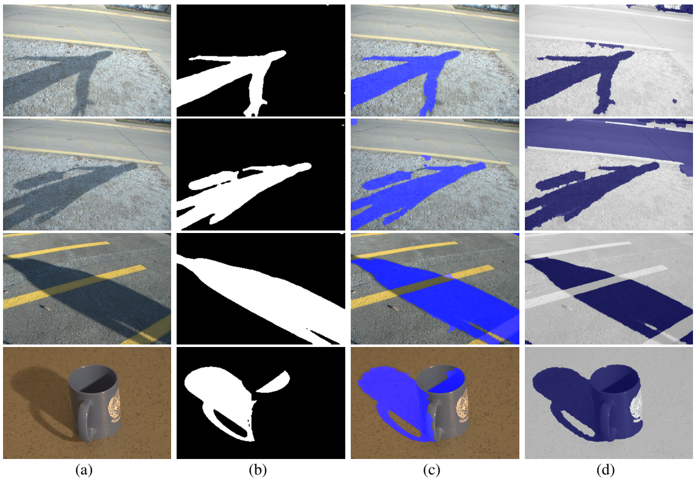

Figure 1: Comparison of shadow detection results. a) Input image. b) Ground truth shadow pixel mask. c) Our results overlaid in ligh blue d). Guo *et al.* [2] results in dark blue.

	Acc.	BER	Sha.	Non
Guo *et al.* [2]	0.891	0.166	0.716	0.952
single region	0.916	0.124	0.795	0.957
MRF ψ_b	0.926	0.101	0.840	0.958
MRF $\psi_s \psi_d$	0.937	0.087	0.862	**0.963**
MRF	**0.941**	**0.078**	**0.880**	0.962

Table 1: Shadow detection quantitative results.

Experimental results on the dataset collected by [1] show that our method clearly outperforms the state of the art. Our single region classifier reduces the balance error rate (BER) by a 25% with respect to [2], correctly detecting 11% more shadow pixels. Moreover, with our MRF optimization we further reduce the BER to a 7.8%. That is a 53% decrease compared to the state of the art. Our total pixel accuracy is 94.1% (88% on shadow pixels).

[1] Ruiqi Guo, Qieyun Dai, and D. Hoiem. Single-image shadow detection and removal using paired regions. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 2033 –2040, june 2011. doi: 10.1109/CVPR.2011.5995725.

[2] Ruiqi Guo, Qieyun Dai, and Derek Hoiem. Paired regions for shadow detection and removal. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 99(PrePrints):1, 2012. ISSN 0162-8828.

[3] Vladimir Kolmogorov and Carsten Rother. Minimizing nonsubmodular functions with graph cuts - a review. Technical report, TPAMI, 2007.

[4] C. Rother, V. Kolmogorov, V. Lempitsky, and M. Szummer. Optimizing binary mrfs via extended roof duality. In *Computer Vision and Pattern Recognition, 2007. CVPR '07. IEEE Conference on*, 2007.

[5] Yossi Rubner, Carlo Tomasi, and Leonidas J. Guibas. A metric for distributions with applications to image databases. pages 59–66, 1998.

[6] Jiejie Zhu, K.G.G. Samuel, S.Z. Masood, and M.F. Tappen. Learning to recognize shadows in monochromatic natural images. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 223 –230, june 2010.