

Shadow Optimization from Structured Deep Edge Detection

Li Shen Teck Wee Chua Karianto Leman
Institute for Infocomm Research

Abstract

Local structures of shadow boundaries as well as complex interactions of image regions remain largely unexploited by previous shadow detection approaches. In this paper, we present a novel learning-based framework for shadow region recovery from a single image. We exploit local structures of shadow edges by using a structured CNN learning framework. We show that using structured label information in classification can improve local consistency over pixel labels and avoid spurious labelling. We further propose and formulate shadow/bright measure to model complex interactions among image regions. The shadow and bright measures of each patch are computed from the shadow edges detected by the proposed CNN. Using the global interaction constraints on patches, we formulate a least-square optimization problem for shadow recovery that can be solved efficiently. Our shadow recovery method achieves state-of-the-art results on major shadow benchmark databases collected under various conditions.

1. Introduction

Shadow detection has long been considered a crucial component of scene interpretation. Shadows in an image provide useful information about the scenes: the object shapes [18], the relative positions and 3D structures of the scene [2, 1], the camera parameters and geo-location cues [11], and the characteristics of the light sources [22, 14, 19]. However, they can also cause great difficulties to many computer vision algorithms, such as background subtraction, segmentation, tracking and object recognition. Despite its importance and long tradition, shadow detection remains an extremely challenging problem, particularly from a single image. The main difficulty is due to the complex interactions of geometry, albedo, and illumination in nature scenes. Since shadows correspond to a variety of visual phenomena, finding a unified approach to shadow detection is difficult.

Motivated by this observation, recent papers [10, 25, 15, 8, 6, 12] have explored the use of learning techniques for shadow detection. These approaches take an image patch

and compute the likelihood that the centre pixel contains a shadow edge. Such classifiers are limited by their locality since they treat each pixel independently. Optionally, the independent shadow predictions could then be combined using global reasoning by using CRF/GBP/MRF.

Shadow edges in a local patch are actually highly interdependent, and exhibit common forms of local structures: straight lines, corners, curves, parallel lines; while structures such as T-junctions or Y-junctions are highly unlikely to be on shadow boundaries [8, 15].

In this paper we propose a novel learning-based framework for shadow detection from a single image. We exploit the local structures of shadow edges by using a *structured* Convolutional Neural Networks (CNN) framework. A CNN learning framework is designed to capture the local structure information of shadow edges and automatically learn the most relevant features. We formulate the problem of shadow edge detection as predicting local shadow edge structures given input image patches. In contrast to unary classification, we take structured labelling information of the label neighbourhood into account. We show that using the structured label information in the classification can improve the local consistency and avoid spurious labelling.

We also propose a novel global shadow optimization framework. In previous learning approaches, CRF/GBP/MRF algorithm is usually employed to enforce local consistency over neighbouring labels [25, 15, 12, 6] and non-local constraints of region pairs with the same materials [6]. The size of the label images and presence of loops in such a algorithm make it intractable to compute the expectation values. Moreover, the memory requirement for loading all the training data is large and parameter updating requires tremendous computing resources. Here, we introduce novel shadow and bright measures to model non-local interactions based on the spatial layout of image regions. For each image patch, shadow and bright measures are computed according to its connectivities to all the shadow and bright boundaries in the image, respectively. The shadow/bright boundaries are extracted from the shadow edges detected by the proposed CNN. Using these shadow and bright measures, we formulate a least-square optimization problem for shadow recovery to solve for the

shadow map (locations). Our optimization framework combines the non-local cues of region interactions in a straightforward and efficient manner where all constraints are in linear form.

Experimental results on the major shadow benchmark databases demonstrate the effectiveness of the proposed technique.

1.1. Related Work

Early works for detecting shadows were motivated by physical models of illumination and colour [9, 21]. Finlayson et al. [5, 4] located shadows by using a 1-d shadow free image computed from a single image. Their methods can only work under the assumption of approximately Planckian lighting, and high-quality input images are needed.

To adapt to environment changes, statistical learning based approaches [20, 17, 7, 10] have been developed to learn shadow model parameters at each pixel from a video sequence. Recently, some data-driven learning approaches have been developed for single-image shadow detection. Lalonde et al. [15] detected cast shadow edges on the ground with a classifier trained on local features. [25] proposed a similar approach for monochromatic images. Every pixel is classified as either being inside a shadow or not. The per-pixel outputs are inherently noisy with poor contour continuity. To overcome this, the predicted posteriors are usually fed to a CRF formulation which defines the pairwise smoothness constraints across neighbouring pixels. Guo et al. [6] modelled long-range interactions using non-local cues of region pairs. They then incorporated the non-local pairwise constraints into a graph-cut optimization. Yago et al. [24] integrated region, boundary and paired regions classifiers using a MRF. All these learning approaches employ hand-crafted features as input.

More recently, [12] proposed a deep learning framework to automatically learn the most relevant features for shadow detection. They showed that the CNN with learned features outperforms the current state-of-the-art with hand-crafted features. They trained a unary classifier where separated CNNs learned features for detecting shadows at boundaries and uniform regions, respectively. The per-pixel predictions are then fed to a CRF for enforcing local consistency. In contrast to unary classification, we predict the structure of shadow edges in a local patch. Our work is inspired by the recent works on learning structured labels for semantic image labelling [13] and edge detection [16, 3] in Random Forests. Our aim is to explore the structured learning in CNN for local shadow edge detection.

We are also inspired by the work on saliency estimation [26]. They propose a background measure based on the boundary connectivity prior that a salient region is less likely connected to the image boundary. Utilizing the

connectivity definition, we derive our shadow measures to model the region interactions based on the spatial layout of image regions.

2. Structured Deep Shadow Edge Detection

In this section, we present our structured CNN learning framework, then we explain how to apply it to shadow edge detection.

2.1. Structured Convolutional Neural Networks

Convolutional neural networks actually can have high dimensional and complex output which makes structured output prediction possible. We use $x \in X = \mathbb{R}^{d \times d \times K}$ to denote a color image patch, and $y \in Y = \mathbb{R}^{s \times s}$ to denote the target structured label, where d and s indicate the patch widths of x and the y , respectively. K is the number of channels. The structured prediction can be expressed as $C : X \rightarrow Y$ mapping from an input domain X to a structured output domain Y . Here, $C(\cdot)$ is the structured CNN.

Fig. 1 shows our network architecture with 7-layers. Our learning approach predicts a structured $y^{25 \times 1}$ label from a larger 28×28 image patch. The network consists of two alternating convolutional and max-pooling layers, followed by a fully connected layer and finally a logistic regression output layer with softmax nonlinear function. The first and second convolutional layers consist of six and twelve 5×5 kernels, respectively, with unit pixel stride. The pooling size is 3×3 with unit-pixel stride, while the fully connected layer has 64 hidden units.

One challenge for training CNN with structured labels is that structured output spaces are high dimensional, which causes long training duration. We show in the experiments that with the setting of $y^{5 \times 5}$ we can capture wide variety of local shadow structures sufficiently while still keeping the training complexity low.

2.2. Structured Learning Shadow Edges

We employ the proposed structured CNN for feature learning for shadow edge detection. It takes a 28×28 color image patch as an input, and outputs a $\{1, 0\}^{25 \times 1}$ vector which is corresponding to the 5×5 shadow probability map of the central patch.

Assume we have a set of images X with a corresponding set of binary images Y of shadow edges. The structured CNN operates on the patches at image edges: only the patches that contain image edges at its 5×5 central area are used. The input patches $\{x^{28 \times 28}\}$ are extracted from X . The corresponding groundtruth $\{y^{5 \times 5}\}$ are extracted from the binary images Y . We randomly sample $N (< 800)$ shadow and non-shadow edge patches, respectively, per image for training. Before feeding the extracted patches to the CNN, the data is zero-centered and normalized.




Figure 1: CNN architecture used for learning the structure of shadow edges. It takes a 28×28 color image patch as an input, and outputs a $\{0, 1\}^{25 \times 1}$ vector which is corresponding to the shadow edge structure of the 5×5 central patch.

Specifically, we first apply Canny edge detector to extract all shadows and non-shadows. Positive patches are obtained from the pixels on Canny edges that coincide with the groundtruth shadow edges. Likewise, negative (non-shadow) input patches are obtained from the edge pixels that do not overlap with the groundtruth shadow edges. In the experiments, as the groundtruth is hand-labelled by human, the actual shadow edge may not coincide with the groundtruth edges. We dilate the groundtruth edges before overlap with the Canny edges. For datasets with region-based groundtruth such as UCF and UIUC, we extract the blob boundaries as the groundtruth of shadow edges. As pointed out in [12] non-shadow pixels usually outnumber shadow pixels by approximately 6:1 ratio. We address this class imbalance problem by setting the number of positive samples as the upper bound of the number of negative samples to be sampled. To reduce the number of redundant samples, we only randomly sample the patches on 3×3 grid. During the training process, we use stochastic gradient descent.

Fig. 2 shows the shadow and non-shadow edges learned by the proposed structure CNN. We can see that our CNN does capture the local structures of shadow edges. Besides learning the classification of shadow and non-shadow edges, the proposed network also learns valid labelling transitions among adjacent pixels (i.e. interactions among the pixels in a local patch).

Our CNN was implemented in unoptimized Matlab code. The training took ~ 5 hours on 5×5 pixels structured output (~ 700 iteration) given input size of $500K+$ patches and consumed $\sim 6GB$ memory on Intel Quad-Core $2.4GHz$ PC.

2.3. Structured Labelling Shadow Edges

Given an input image, Canny edge detector is applied to find significant edges in the image. Then, we extracted windows along the image edges. The overlapping edge patches are then fed to the proposed CNN for labelling. The trained structured CNN differentiates between shadow and reflectance edges and predicts shadow edge structure of the




Figure 2: (a) Shadow and (b) non-shadow patches learned by the proposed structure CNN. Left: input 28×28 patches (red rectangle indicates the central 5×5 region). Center: groundtruth 5×5 patches. Right: output 5×5 patches.

central region. Our structured CNN achieves robust results by combining the neighbouring prediction results. Instead of independently assigning a class label to each pixel, our structured labels predict the interactions among the neighbouring pixels of a local patch. Each pixel collects class hypotheses from the structured labels predicted for itself and neighboring pixels. We employ a simple voting scheme to combine the multiple predictions at each pixel at the image edges.

Fig. 3 illustrates the advantage of shadow edge detection with structured output CNN. As can be seen, the proposed structured CNN can recover better local edge structures (local consistency), and avoid assigning implausible label transitions.

3. Shadow optimization

We first derive the local and global shadow/bright measures to model the interactions among the regions across




Figure 3: Structured shadow edge detection results. (a) input image. We compare the detection results using (b) 1x1 labelling CNN(previous method), (c) 5x5 structured labelling CNN. (d) Zoom-in of green and blue patches. Top to bottom:original, 1x1 and 5x5 outputs. We can see that 5x5 structured CNN is able to learn fine shadow details. 1x1 output has serious spurious noise.

the image. Then, we present our optimization framework to solve for the shadow map.

3.1. Global and local shadow(bright) probability

We observe that both shadow and bright regions have the following characteristic in their spatial layout: the shadow regions are much more connected to shadow boundaries, and the bright regions are more connected to bright boundaries. We define the dark side region of a shadow edge as the *shadow boundary*, while the bright side region as the *bright boundary*. In Fig. 5, we illustrate a tree and its shadow. The blue and pink regions are the shadow and bright boundaries, respectively. The grey region is clearly a shadow region as it significantly touches the shadow boundary, while the white region is clearly a bright region as it largely touches the bright boundary.

The geodesic distance between any two patches p, q in an image is defined as the accumulated edge weights along their shortest path on the graph:

$$d_{geo}(p, q) = \min_{p_1=p, p_2, \dots, p_n=q} \sum_{i=1}^{n-1} d_{app}(p_i, p_{i+1})$$

where $d_{app}(p, q)$ is the Euclidean distance between the average colours of the two patches. The normalized $D(p, q) \in (0, 1]$ is thus defined as $D(p, q) = \exp(-\frac{d_{geo}^2(p, q)}{2\sigma_{clr}^2})$ which characterizes how much patch p connects (or contributes) to patch q . $D(p, q) \approx 0$ when $d_{geo}^2(p, q) \gg \sigma_{clr}^2$.




Figure 5: An illustrative example. The orange lines are shadow edges. The blue and pink regions adjacent to the shadow edges are the shadow and bright boundaries, respectively.

Let $\mathbf{shd}$ be the set of shadow boundary patches, and $\mathbf{lit}$ be the set of bright boundary patches. Following the boundary connectivity introduced in [26], we formulate the shadow and bright boundary connectivities of a patch p as:

$$con_{shd}(p) = \frac{Len(p, \mathbf{shd})}{\sqrt{Area(p)}} \text{ and } con_{lit}(p) = \frac{Len(p, \mathbf{lit})}{\sqrt{Area(p)}}, \quad (1)$$

respectively. $Len(p, \mathbf{x}) = \sum_{q \in \mathbf{x}} D(p, q)$ is the connectivity of p to boundary $\mathbf{x}$. $Area(p) = \sum_{i=1}^N D(p, p_i)$ is the spanning area of p , where N is the number of patches in the image. Note that we set $d_{app}(p \in \mathbf{shd}, q \in \mathbf{lit}) = \inf$,




Figure 4: Global shadow optimization. (a) input images with superpixel boundaries overlaid. (b) shadow and bright boundary patches extracted from the local detected shadow edges. (c) local shadow(L) and bright(R) measures in Eq. 2; (d) global shadow(L) and bright(R) measures in Eq. 3. The predicts are propagated over the image. (e) optimized shadow maps by minimizing Eq. 4.

which implies that the bright and shadow boundary sets are not connected, and no path can cut through the shadow edges. We can see that $con_{shd,lit}(p)$ quantifies how heavily a patch p is connected to the shadow/bright boundaries in a local area ($\sigma_{clr} = 5$ in the experiments.)

Hence, we define the local shadow/bright measure as:

$$\gamma_{shd,lit}(p) = \exp\left(-\frac{con_{shd,lit}(p)}{2\sigma_{con}^2}\right), \quad (2)$$

and local shadow/bright probability can be computed as $Pr_{loc}^{shd,lit}(p) = 1 - \gamma_{shd,lit}(p)$ which is close to 1 when shadow/bright connectivity is large, and 0 when it is small. Note that $Pr_{loc}(p)$ is only affected by the local shadow edges in the region p belong to. If a patch is hardly connected to any of the shadow edges in the image, $Pr_{loc}(p)$ is low, which indicates that we cannot get a correct prediction with the local information.

We define the global shadow/bright measure as:

$$\Gamma_{shd,lit}(p) = \frac{\sum_{i=1}^N w_{app}(p, p_i) w_{spa}(p, p_i) \gamma_{shd,lit}(p_i)}{\sum_{i=1}^N w_{app}(p, p_i) w_{spa}(p, p_i)} \quad (3)$$

where $w_{app} = \exp(-\frac{d_{app}^2}{2\sigma_{app}^2})$, and $w_{spa} = \exp(-\frac{d_{spa}^2}{2\sigma_{spa}^2})$. d_{app} and d_{spa} are the Euclidean distance between average colors and locations, respectively. It is based on the observation that if two patches in an image are of the same colour and near to each other, they usually are both in shadow or bright regions. Global shadow/bright probability at p can be computed as $Pr_{glb}^{shd,lit}(p) = 1 - \Gamma_{shd,lit}(p)$.

3.2. Global shadow optimization

The input image is first abstracted as a set of nearly regular superpixels using the Quick Shift segmentation method [23]. The shadow and bright boundary patches can be extracted from the local shadow edge detection results. We only select the most reliable patches as the shd and lit

boundary regions. Specifically, if a superpixel at shadow edge is darker(brighter) than all the adjacent superpixels, we set it as a shadow(bright) boundary patch. If a superpixel is brighter than some of its neighbours but darker than some others, we consider the patch to be ambiguous and it would be discarded.

After obtaining **shd** and **lit**, the local shadow/bright measures $\gamma_{shd,lit}$ can be computed at each superpixel respectively, as described in Eq. 2. The global shadow/bright measures $\Gamma_{shd,lit}(p)$ can be computed as in Eq. 3, consequently.

We formulate the problem of shadow detection as optimizing the shadow values of all image superpixels. The objective cost function is designed to assign shadow regions to value 1 and bright regions to value 0, respectively. The optimal shadow map is then obtained by minimizing the cost function. Let the shadow values of N superpixels be $\{s_i\}_{i=1}^N$. The cost function is thus defined as

$$E = \sum_{i=1}^N w_i^{shd} s_i^2 + \sum_{i=1}^N w_i^{brt} (1 - s_i)^2 + \sum_{i,j} w_{ij} (s_i - s_j)^2 + \lambda \sum_{i \in \{shd, lit\}} (s_i - \tilde{s}_i)^2 \quad (4)$$

where $w_i^{shd,lit} = \Gamma_{shd,lit}(p_i)$ defined in Eq. 3. $\tilde{s}_i$ is the initial values for $p \in \{\mathbf{lit}, \mathbf{shd}\}$, where $\tilde{s}_i = 1$ for $p_i \in \mathbf{shd}$ and 0 for $p_i \in \mathbf{lit}$. We set $\lambda = 0.001$ in the experiments.

The four terms are all squared errors and the optimal shadow map is computed by least-square. Fig. 5 shows the optimized results.

4. Experiments

4.1. Datasets

UCF Shadow Dataset: This dataset contains 355 images with manually labeled region-based ground truth. Only 245/355 images were used in [25, 6]. The split of the train/test data is according to the software package provided by [6] as the original authors did not disclose the split.

CMU Shadow Dataset: This dataset contains 135 images with manually labeled edge-based ground truth for shadow on the ground. As our algorithm is not restricted to ground shadows, we create the ground plane masks and augment the edge-based ground truth to region-based ground truth. The authors did not report train/test data split therefore we follow the procedure in [12] where even order images are used for training and odd images for testing.

UIUC Shadow Dataset: This dataset contains 108 images (32 train images and 76 test images) with region-based ground truth.

4.2. Results

We extensively evaluated our proposed algorithm on three publicly available single image shadow datasets. The evaluation results on UCF in [12] were based on the full dataset. To be comparable to their results, we reported both the UCF results using the full dataset and the subset 245 images. As shown in Table 2, our shadow detection method (SCNN-LinearOpt) achieves the best performances for all three datasets. In particular, we achieve almost 2% and 5% gain over state-of-the-art results for UCF and CMU datasets. Table 1 shows the comparisons of class-specific detection accuracies. We take the shadow accuracy to be the number of pixels correctly detected as shadow divided by the total number of pixels marked as shadow in the ground truth. Likewise, non-shadow accuracy is obtained in a similar manner. Our approach achieves significantly higher shadow accuracies. This is consistent with the finding from Fig. 7 where our approach delivers highest AUC.

Fig. 6 shows some of the qualitative results obtained with our method. The results suggest that our shadow detectors performed robustly under various cases ranging from indoor images to outdoor and aerial images that exhibit soft shadow, low light condition, colour cast, and complex self-shading regions. In Fig. 8, we compare our approach with Zhu’s work [25]. Our method can correctly recover the shadow regions in the complex scene. In Fig. 9, we show that our approach outperforms Guo’s work [6] in the ambiguous situation that the object material has the similar colour of the shadows in the image. In Fig. 10, we compare our shadow edges results with the Lalonde’s results [15]. Our method can accurately detect the shadow edges of the image which Lalonde’s method fails with. We also compare our method with Khan et al.’s very recent work [12] in

Datasets\Methods	Shadows	Non-Shadow
UCF Dataset		
BDT-BCRF [25]	63.9%	93.4%
Unary-Pairwise [6]	73.3%	93.7%
CNN-CRF [12]	78.0%*	92.6%*
SCNN-LinearOpt	91.1%(91.6%*)	93.5%(93.4%*)
UIUC Dataset		
Unary-Pairwise [6]	71.6%	95.2%
CNN-CRF [12]	84.7%	95.5%
SCNN-LinearOpt	91.3%	95.03%
CMU Dataset		
BDT-CRF-Scene [15]	73.1%	96.4%
CNN-CRF [12]	83.3%	90.9%
SCNN-LinearOpt	91.6%	97.7%

Table 1: Pixel-wise Shadow/Non-Shadow Detection Accuracy. *: The result is produced using the full dataset.



Figure 8: Comparison with Zhu’s work [25]. Left: input image. Middle: Zhu’s results. Right: our result.



Figure 9: Comparison with Guo’s work [6]. Left: input image. Middle: Guo’s results. Right: our result. Our method can correctly recovery the shadow regions.

Fig. 11.

5. Conclusions

In this paper, we propose an efficient structured labelling framework for shadow detection from a single image. We show that the structured CNN Networks framework can capture the local structure information of shadow edge. Moreover, we present a novel global shadow/bright measures to model the complex global interactions based on spatial layout of image regions. The non-local constraints on shadow/bright regions help to overcome ambiguities in local inference. Using these non-local region constraints, we formulate the shadow detection as a least-square optimization problem which can be solved efficiently. Our method can be easily extended to other low-level problems

Methods	UCF dataset	UIUC dataset	CMU dataset
BDT-BCRF [25]	88.7%	-	-
BDT-CRF-Scene [15]	-	-	84.8%
Unary-Pairwise [6]	90.2%	89.1%	-
CNN-CRF [12]	90.7%*	93.2%	88.8%
SCNN-LinearOpt	93.1 % (92.3%*)	93.4%	94.0%

Table 2: Performance Comparisons of Shadow Detection Methods




Figure 6: Shadow optimization results from detected shadow edges. Top: input shadow images. Bottom: recovered shadow regions.




Figure 10: Shadow edge detection results comparing with Lalonde's work [15]. Left: input image. Middle: Lanlonde's results. Right: our result. Our method can accurately detect the shadow edges.

such as object edge detection, smoke region detection etc..

References

- [1] A. Abrams, I. Schillebeeckx, and R. Pless. Structure from shadow motion. In *ICCP*, pages 1–8, May 2014. **1**
- [2] Y. Caspi and M. Werman. Vertical parallax from moving shadows. In *CVPR*, volume 2, pages 2309–2315, 2006. **1**
- [3] P. Dollar and C. Zitnick. Structured forests for fast edge detection. In *ICCV*, pages 1841–1848, Dec 2013. **2**
- [4] G. Finlayson, M. Drew, and C. Lu. Intrinsic images by entropy minimization. In *ECCV*, volume 3023, pages 582–595. 2004. **2**
- [5] G. Finlayson, S. Hordley, and M. Drew. Removing shadows from images. In *ECCV*, volume 2353 of *Lecture Notes in Computer Science*, pages 823–836. 2002. **2**
- [6] R. Guo, Q. Dai, and D. Hoiem. Single-image shadow detection and removal using paired regions. In *CVPR*, pages 2033–2040, June 2011. **1, 2, 5, 6, 7**
- [7] J.-B. Huang and C.-S. Chen. Moving cast shadow detection using physics-based features. In *CVPR*, pages 2310–2317, June 2009. **2**
- [8] X. Huang, G. Hua, J. Tumblin, and L. Williams. What characterizes a shadow boundary under the sun and sky? In *ICCV*, pages 898–905, Nov 2011. **1**
- [9] C. Jiang and M. Ward. Shadow identification. In *CVPR*, pages 606–612, Jun 1992. **2**




Figure 7: ROC curves on (a) UCF dataset, (b) UIUC dataset, and (c) CMU dataset



Figure 11: Comparison with Khan's work [12]. Left: input image. Middle: Khan's results. Right: our result.

- [10] A. Joshi and N. Papanikolopoulos. Learning to detect moving shadows in dynamic environments. *IEEE TPAMI*, 30(11):2055–2063, Nov 2008. 1, 2
- [11] I. Junejo and H. Foroosh. Estimating geo-temporal location of stationary cameras using shadow trajectories. In *ECCV*, volume 5302 of *Lecture Notes in Computer Science*, pages 318–331. 2008. 1
- [12] S. Khan, M. Bennamoun, F. Sohel, and R. Togneri. Automatic feature learning for robust shadow detection. In *CVPR*, pages 1939–1946, June 2014. 1, 2, 3, 6, 7, 8
- [13] P. Kotschieder, S. Rota Bulò, H. Bischof, and M. Pelillo. Structured class-labels in random forests for semantic image labelling. In *ICCV*, pages 2190–2197, Nov 2011. 2
- [14] J.-F. Lalonde, A. Efros, and S. Narasimhan. Estimating natural illumination from a single outdoor image. In *ICCV*, pages 183–190, Sept 2009. 1
- [15] J.-F. Lalonde, A. Efros, and S. Narasimhan. Detecting ground shadows in outdoor consumer photographs. In *ECCV*, volume 6312 of *Lecture Notes in Computer Science*, pages 322–335. 2010. 1, 2, 6, 7
- [16] J. Lim, C. Zitnick, and P. Dollar. Sketch tokens: A learned mid-level representation for contour and object detection. In *CVPR*, pages 3158–3165, June 2013. 2
- [17] Z. Liu, K. Huang, T. Tan, and L. Wang. Cast shadow removal combining local and global features. In *CVPR*, pages 1–8, June 2007. 2
- [18] T. Okabe, I. Sato, and Y. Sato. Attached shadow coding: Estimating surface normals from shadows under unknown reflectance and lighting conditions. In *ICCV*, pages 1693–1700, Sept 2009. 1
- [19] A. Panagopoulos, D. Samaras, and N. Paragios. Robust shadow and illumination estimation using a mixture model. In *CVPR*, pages 651–658, June 2009. 1
- [20] F. Porikli and J. Thornton. Shadow flow: a recursive method to learn moving cast shadows. In *ICCV*, volume 1, pages 891–898 Vol. 1, Oct 2005. 2
- [21] A. Prati, R. Cucchiara, I. Mikic, and M. Trivedi. Analysis and detection of shadows in video streams: a comparative evaluation. In *CVPR*, volume 2, pages II–571–II–576 vol.2, 2001. 2
- [22] I. Sato, Y. Sato, and K. Ikeuchi. Illumination from shadows. *IEEE TPAMI*, 25(3):290–300, Mar. 2003. 1
- [23] A. Vedaldi and S. Soatto. Quick shift and kernel methods for mode seeking. In *ECCV*, volume 5305, pages 705–718. 2008. 5
- [24] T. Yago, C.-P. Yu, and D. Samaras. Single image shadow detection using multiple cues in a supermodular mrf. In *Proceedings of the British Machine Vision Conference*. BMVA Press, 2013. 2
- [25] J. Zhu, K. Samuel, S. Masood, and M. Tappen. Learning to recognize shadows in monochromatic natural images. In *CVPR*, pages 223–230, June 2010. 1, 2, 5, 6, 7
- [26] W. Zhu, S. Liang, Y. Wei, and J. Sun. Saliency optimization from robust background detection. In *CVPR*, pages 2814–2821, June 2014. 2, 5