

Intrinsic Image Decomposition with Non-Local Texture Cues

Li Shen
Microsoft Research Asia

Ping Tan
National University of Singapore

Stephen Lin
Microsoft Research Asia

Abstract

We present a method for decomposing an image into its intrinsic reflectance and shading components. Different from previous work, our method examines texture information to obtain constraints on reflectance among pixels that may be distant from one another in the image. We observe that distinct points with the same intensity-normalized texture configuration generally have the same reflectance value. The separation of shading and reflectance components should thus be performed in a manner that guarantees these non-local constraints. We formulate intrinsic image decomposition by adding these non-local texture constraints to the local derivative analysis employed in conventional techniques. Our results show a significant improvement in performance, with better recovery of global reflectance and shading structure than by previous methods.

1. Introduction

The color of scene points is determined by both its surface reflectance and shading. Reflectance describes the intrinsic color of a surface, which is invariant to illumination and imaging conditions. Shading arises from the amount of reflected light by the surface, and typically depends upon surface orientation and illumination condition. Intrinsic image decomposition addresses the problem of separating an image into its reflectance and shading components. Since each component represents a different physical element, many computer vision algorithms would benefit from such a decomposition, which isolates the information they wish to analyze. For example, shape from shading algorithms infer object geometry from shading variations, and will be misled by reflectance changes mistaken for shading. On the other hand, techniques based on segmentation, stereo and optical flow prefer a pure reflectance map as input to avoid appearance variations caused by shading.

Although intrinsic image decomposition offers distinct benefits, it remains a challenging, greatly underconstrained problem because for each pixel, there exists twice as many unknowns (reflectance, shading) as there are measurements.

Some previous works seek to make this problem more manageable by utilizing additional information from multiple registered images [21] [16] [15] [1], each having a different illumination condition. High quality results have been demonstrated by such methods, but their input requirements limit their application.

For general practicality, decomposition methods for single arbitrary images are favored. Previous work towards this end has mainly analyzed local derivatives for distinguishing image variations due to shading or reflectance. In general, these methods have assumed scene conditions similar to those addressed by the Retinex algorithm [12]: Mondrian-like images with piecewise constant reflectance and smoothly varying shading. For such scenes, large derivatives can be attributed to reflectance change, while smaller derivatives are due to shading. By further assuming reflectance and shading change not to coincide, intrinsic images can be recovered by integrating their respective derivatives across the input image. This approach has been adapted to color images [9] by associating reflectance derivatives to significant chromaticity changes, based on the property that shading variations do not alter chromaticity. A variational approach to the Retinex problem was proposed in [10] that subsumes many previous techniques. While these Retinex methods are intuitively simple, real-world scenes often do not adhere to their assumed conditions. Other methods based on local gradients classify edges or edge junctions according to heuristic rules [17] or trained classifiers [2] [19] [20]. It is, however, difficult to formulate all-inclusive rules or to comprehensively train classifiers over the possible range of shading and reflectance configurations. Furthermore, it is not always possible to disambiguate reflectance and shading changes based on only local appearance.

Besides local analysis, intrinsic image decomposition may also be guided by global considerations. In [17], decomposition is computed in a manner that aims for global consistency of local inferences with respect to a plausible 3D scene structure and illumination condition. In their case, scene structure is modeled by uniform-albedo polyhedra, and illumination is assumed to be diffuse ambient lighting plus a single distant light source. A related method was

presented based on patches from training images [8]. Such a high-level interpretation of the observed image can help to overcome ambiguities in local cues. However, generalization of this approach to real images containing complex surface geometries and lighting conditions, as well as object occlusions and cast shadows, is difficult. In [19], generalized belief propagation is used to broaden the influence of local cues to help resolve ambiguous local inferences. Our work also employs propagation for similar purposes.

In this paper, we introduce a new, non-local cue for intrinsic image decomposition. This cue is derived from the observation that a surface can often be depicted by a small number of local texture structures, or textons [13]. In other words, for each point on a surface, there generally exists a set of other points that share the same neighborhood texture configuration. Points within such a set can be considered likely to have the same underlying appearance, and such groupings of pixels have previously been used to reduce noise [3] and highlights [18] in images. Our work capitalizes on this property for intrinsic image decomposition by computing a shading-independent image and then identifying surface points with the same texture configuration as having the same reflectance value. Relating the reflectance values among pixels throughout an image provides important non-local decomposition constraints that complement information gained from local analysis.

To demonstrate the benefits of non-local texture constraints, we formulate an algorithm that utilizes them in conjunction with color-based Retinex cues (significant chromaticity derivatives indicating reflectance changes, and other derivatives resulting from shading). We show that decomposition with the addition of these texture constraints can considerably outperform decomposition using Retinex cues alone. Our results also show that global shading and reflectance structures can be better preserved by employing this non-local information.

2. Non-Local Texture Analysis

In the intrinsic image literature, an image I is modeled as a product of a shading image S and a reflectance image R :

$$I_{x,y} = S_{x,y}R_{x,y}$$

where x, y are image coordinates. Intrinsic image decomposition is clearly seen to be ill-posed, since at each pixel there are two unknowns (S, R) with one known variable (I). For an image with M pixels, this amounts to $2M$ unknowns with only M observations. To solve this problem, different priors on the solution have been employed [21] [10] [19]. Here, we seek to reduce the number of unknowns before applying any priors to solve it.

This reduction of unknowns is achieved through an analysis of texture, the spatial configuration of colors in a local

neighborhood. In our work, texture may be represented by any model, such as filter responses [14] or textons [13], but for simplicity we directly represent the texture at a pixel as a vector of concatenated pixel values from its surrounding neighborhood. For a pixel p , its texture vector for an $N \times N$ neighborhood can be expressed as

$$\begin{aligned} \mathbb{I}(p) &= (I_{1,1}, I_{1,2}, \dots, I_{N,N}) \\ &= (S_{1,1}R_{1,1}, S_{1,2}R_{1,2}, \dots, S_{N,N}R_{N,N}). \end{aligned}$$

In a color image, each element in a texture vector is composed of an rgb triple, i.e., $I_{x,y} = (I_{x,y}(r), I_{x,y}(g), I_{x,y}(b))$, and $R_{x,y} = (R_{x,y}(r), R_{x,y}(g), R_{x,y}(b))$. $S_{x,y}$ remains a scalar that indicates the amount of shading.

To remove the influence of shading on texture, texture vectors are constructed using intensity-normalized color values (chromaticity) of pixels, which are used in [9] to distinguish shading and reflectance changes. These shading-independent texture vectors are then expressed as

$$\hat{\mathbb{I}}(p) = (\hat{R}_{1,1}, \hat{R}_{1,2}, \dots, \hat{R}_{N,N})$$

where $\hat{R}_{x,y}$ corresponds to an intensity-normalized reflectance, i.e., $\hat{R}_{x,y} = R_{x,y} / \|R_{x,y}\| = I_{x,y} / \|I_{x,y}\| = \hat{I}_{x,y}$. We note that although the normalized image \hat{I} is free of shading, it is not equivalent to the reflectance image R because of the missing intensity components.

According to the theory of Markov Random Fields [22], if the surrounding neighborhoods of two pixels are the same, the two pixels are expected to have the same value. This property has been successfully used for texture synthesis [4]. In this work, we employ a similar constraint by considering two pixels to have identical reflectance if their intensity-normalized texture vectors are the same. We group pixels in the input image according to this criterion to identify pixels that have the same reflectance. This non-local analysis thus yields intrinsic image constraints among pixels throughout the image.

Given these non-local texture constraints where $R_{p^1} = R_{p^2} = \dots = R_{p^n}$ for pixels p^1, p^2, \dots, p^n within each group, the number of unknowns for reflectance values $R_{x,y}$ drops from M to K , where K is the number of distinct groups in the input image. As a result, the number of unknowns in the intrinsic image decomposition is reduced from $2M$ to $M + K$. It is often the case that $K \ll M$ because of the abundant textures found in natural images. Although the decomposition problem remains ill-posed, the significant reduction of unknowns generally leads to a more accurate solution. In particular, the non-locality of these constraints can help to overcome ambiguities in local inference, as done in global algorithms [17] [8].

3. Decomposition Algorithm

Our method for intrinsic image decomposition is performed in two steps. We first search the intensity-normalized (shading-independent) image for pixels embedded within the same texture configuration. These pixels are grouped together and treated as having the same reflectance. To account for improper matches that may exist within a group, we associate with each pixel a soft grouping weight that reflects the confidence that this pixel indeed has the same reflectance as others in the group. These non-local constraints are joined with local Retinex constraints within an energy formulation to solve for the intrinsic images.

3.1. Soft Grouping

In the grouping process, we examine small (3x3) windows centered on each pixel in the image. Groups are formed by iteratively selecting an unmatched pixel and finding all matches in the image with a sum of squared differences (SSD) less than a specified threshold (0.01 in our implementation). To promote matching, rotated windows may be considered in comparing pixels. Our implementation specifically accounts for rotations of 90°, 180°, and 270°, though additional rotation angles may be included.

Inaccuracies may arise in this threshold-based grouping. To alleviate the problem caused by incorrect grouping, we compute a match weight at each pixel, which indicates the confidence of that pixel sharing the same reflectance as others in the group. We consider matches with greater local support to be more indicative of a proper reflectance match. For each pixel (x, y) , we compare windows of larger sizes (5x5, 7x7) to its group's median window of the same size, where the entries of the median window are computed by taking the median values among corresponding window entries in the group. The largest matching window size, $N \times N$, is recorded for the pixel as $m_{x,y} = N$. The closeness of a match, given by the normalized SSD of the original 3x3 match ($\frac{1}{9}SSD$), is also used as evidence of proper reflectance grouping. These two factors are jointly used to determine the grouping weight of each pixel:

$$w_{x,y} = m_{x,y} \cdot \left(1 - \frac{1}{9}SSD_{x,y}\right).$$

While this grouping scheme allows us to determine which pixels in the image have the same reflectance, the reflectance color of each group is yet unknown. At this point, we know only their intensity-normalized colors, and need to infer the reflectance intensity r_k of each group k in order to compute the intrinsic image decomposition:

$$R_{x,y} = \hat{R}_{x,y} \cdot r_{g(x,y)}, \quad S_{x,y} = I_{x,y}/R_{x,y} \quad (1)$$

where $g(x, y)$ indexes the group to which pixel (x, y) belongs, and the computation of $S_{x,y}$ involves element-wise division.

3.2. Energy Formulation and Minimization

In principle, this texture-based grouping of pixels could be incorporated as additional constraints within any previous intrinsic image decomposition algorithm. In this paper, we introduce these constraints into a Retinex-based method that is similar to that in [10], for simplicity of optimization. The effectiveness of our non-local texture cues is tested by comparing the results of the Retinex algorithm with and without this texture-based grouping.

We compute intrinsic images by solving for the reflectance intensities r_k of the K different groups in the image. In determining these values, we utilize the local derivative constraints of color-based Retinex methods, namely that significant chromaticity differences result from reflectance changes while other variations are due to shading. These constraints are combined within an energy function that we wish to minimize:

$$\begin{aligned} \arg \min_{r_1, r_2, \dots, r_K} \sum_{x,y} & \left[\left(\frac{I_{x+1,y}}{\hat{R}_{x+1,y} \cdot r_{g(x+1,y)}} - \frac{I_{x,y}}{\hat{R}_{x,y} \cdot r_{g(x,y)}} \right)^2 \right. \\ & + \left(\frac{I_{x,y+1}}{\hat{R}_{x,y+1} \cdot r_{g(x,y+1)}} - \frac{I_{x,y}}{\hat{R}_{x,y} \cdot r_{g(x,y)}} \right)^2 \\ & + \alpha_{x,y}^X \left(\hat{R}_{x+1,y} \cdot r_{g(x+1,y)} - \hat{R}_{x,y} \cdot r_{g(x,y)} \right)^2 \\ & \left. + \alpha_{x,y}^Y \left(\hat{R}_{x,y+1} \cdot r_{g(x,y+1)} - \hat{R}_{x,y} \cdot r_{g(x,y)} \right)^2 \right] \end{aligned}$$

where

$$\alpha_{x,y}^X = \begin{cases} 0.1 & \text{if } |\hat{I}_{x+1,y} - \hat{I}_{x,y}| > 0.01 \\ 10 & \text{otherwise,} \end{cases}$$

$$\alpha_{x,y}^Y = \begin{cases} 0.1 & \text{if } |\hat{I}_{x,y+1} - \hat{I}_{x,y}| > 0.01 \\ 10 & \text{otherwise.} \end{cases}$$

The first two terms in this energy function penalize large shading derivatives, while the last two terms are for reflectance derivatives. When there exists a significant change in intensity-normalized color, the coefficients α^X, α^Y are set low to allow for large reflectance changes. Otherwise, they are set high and discourage large reflectance derivatives. A pixel with lower confidence of membership in its group should have less influence on the solution of the group's reflectance intensity. This is achieved from its lower weight $w_{x,y}$ in the energy function.

To effectively optimize this energy function, we first express it in a graph structure, where there is a node for each group, and links between nodes that represent the shading and reflectance derivative constraints between pixels. Associated with each node is the reflectance intensity r_k of the



Figure 2. Two pixel groups for the box shown in Figure 1. The pixels in each group are taken to have the same reflectance. In our algorithm, matches with less neighborhood support and less consistency with the group are down-weighted in the energy formulation.

group, which we optimize using tree-reweighted message passing [11] with the values of r_k discretized to 30 levels. Once the values of r_k have been optimized, they are used in Eq. (1) to obtain the decomposition solution.

4. Results

To demonstrate the benefit of using non-local texture cues in intrinsic image decomposition, we compare the results of our algorithm to results obtained without the use of texture analysis. The method without texture analysis resembles color-based Retinex algorithms, and was implemented by removing the texture processing from our technique, essentially making each pixel its own group. All of the presented examples except for Figure 1 were downloaded from Flickr (<http://www.flickr.com>). Our algorithm performs well on these images with general illumination conditions, unknown camera settings and relatively low image quality.

In Figure 1, we present an illustrative example of an object with a simple shape. The image contains a rectangular box with stronger illumination from the right side. This example represents a common case where Retinex scene assumptions do not hold, since it contains pixels that have both shading and reflectance changes, as seen on the foremost edge of the box. With conventional Retinex constraints, pixels that contain significant reflectance derivatives should be smooth in shading. The result of this is exhibited in (b) and (c), where the shading is seen to be approximately even across the box edge, even though it is actually not. With non-local texture analysis, our method identifies pixels with the same reflectance on both sides of the box, as shown in Figure 2. That these pixels must have the same reflectance provides an important constraint in computing correct intrinsic images, and allows our algorithm to find the shading discontinuity as shown in (d) and

(e). With this non-local information, our method can recover the global shading and reflectance structure that cannot easily be inferred using local cues alone.

The second example is a ball of colored string, shown in Figure 3. This ball exhibits rapid changes in both reflectance and shading. At the same time, the global shading over the ball gradually changes from darker to brighter, from left to right. Without the texture cue, the recovered shading image misses much of the global and local shading structure, as demonstrated in Figure 3 (b), because the algorithm misinterprets many image gradients as purely reflectance changes due to the large color differences. In contrast, Figure 3 (d) with the non-local texture cues correctly captures the gradual shading change across the image, as well as the local shading variations.

Figure 4 presents a challenging scene. The sky is divided by the window frame into three regions. In the decomposition without texture cues, shading and reflectance in each of these regions are thus computed separately, which results in estimated reflectances that are inconsistent, as seen in Figure 4 (c). With non-local texture cues, divided portions of the sky are grouped together to share the same reflectance value, which leads to a more consistent decomposition in Figure 4 (e). The image also contains a blanket that exhibits fine-scale variations in shading and shadows that arise from the 3D weave pattern. As demonstrated in Figure 4 (b), the Retinex-based algorithm produces very noisy shading decomposition due to the overlap of many shading and reflectance edges. In contrast, our method can better deal with this problem and more accurately recovers the shape of the blanket surface in Figure 4 (d). This can be seen more clearly for a closeup of a small blanket region in Figure 4 (f). In contrast to the purely Retinex-based method, the inclusion of non-local texture cues can correctly recover the subtle shading caused by 3D surface textures.

Further results are included in Figures 5 and 6. Figure 5 exemplifies problems with decomposition in very dark regions, e.g., the deep shadow at the lower-right of the image. In dark regions, the signal-to-noise ratio becomes smaller in the image, which makes texture-based matching less reliable. Even with correct matching, the effects of quantization errors in such areas can be magnified and lead to artifacts in the reflectance image.

5. Discussion

Although intrinsic image decomposition has long been recognized as an important problem in computer vision, this area has seen limited progress due to the severely ill-posed nature of this problem. In this work, we have introduced a new non-local texture constraint that can significantly reduce the number of unknowns to be solved. Our approach identifies pixels that have the same reflectance, as inferred from their neighboring image contexts. We test



Figure 1. Box. (a) Original image. (b)-(c) Shading and reflectance image without the use of texture cues. (d)-(e) Shading and reflectance image with our non-local texture algorithm.



Figure 3. Ball of string. (a) Original image. (b)-(c) Shading and reflectance image without the use of texture cues. (d)-(e) Shading and reflectance image with our non-local texture algorithm.

these non-local texture constraints within the conventional Retinex framework. By combining the Retinex algorithm and our non-local texture cue, we obtain a method that relies less on weak priors such as shading smoothness. We have demonstrated that the inclusion of these non-local texture constraints can bring significant improvement in intrinsic image decomposition, especially in the recovery of global shading and reflectance structure. In instances where no texture matches are found within an image, this method becomes equivalent to standard Retinex techniques.

Although our experiments have focused on Retinex methods, non-local texture constraints also have the potential to elevate the performance of other derivative-based methods, such as those which utilize trained classifiers [2] [19] [20]. Texture provides decomposition cues different from those of classifiers, in that it does not attempt to distinguish between reflectance and shading edges, but rather seeks pixels that have the same reflectance. In principle, intrinsic image information from both texture and classifiers could be jointly used within a common framework.

In our texture descriptor, the effects of shading are removed by the use of intensity-normalized color values, or chromaticity. This facilitates texture matching on surfaces with complex shading patterns, such as those with small-

scale 3D structure. On the other hand, removal of the intensity component reduces the distinctiveness of texture patterns, which can possibly lead to some incorrect matches. This tradeoff of distinctiveness for shading invariance is a fundamental issue that affects most chromaticity-based algorithms.

In our method, we have included soft matching based on window size and match similarity in order to reduce the impact of incorrect matches on the decomposition solution. But even with this safeguard, there may nevertheless exist grouped pixels with large match weights that do not actually have identical reflectance values. This ambiguity could potentially be reduced by considering the uniqueness or complexity of their texture patterns. A match between pixels with a more unique or complex texture could be expected to have a higher likelihood of having the same reflectance, whereas matches with simple texture (e.g., a single straight edge) are more susceptible to error. Incorporating such an indicator of reflectance match confidence could potentially improve the soft matching weights. Investigation of a suitable texture complexity measure is a potential direction for future work.

While intensity normalization of texture can remove the effects of shading prior to our matching process, it may not




Figure 4. Example with 3D surface texture. (a) Original image. (b)-(c) Shading and reflectance image without the use of texture cues. (d)-(e) Shading and reflectance image with our non-local texture algorithm. (f) A cropped region from the blanket. (g)-(h) Shading and reflectance on the cropped region without the use of texture cues. (i)-(j) Shading and reflectance on the cropped region with our non-local texture algorithm.




Figure 5. Yarn. (a) Original image. (b)-(c) Shading and reflectance image without the use of texture cues. (d)-(e) Shading and reflectance image with our non-local texture algorithm.

necessarily eliminate the effects of shadows. In contrast to shading, shadows may produce a chromaticity shift (instead of just an intensity shift) in observed colors, due to the absence of illumination from a particular colored light source. This chromatic shift in shadows has been modeled in [6] [5] [7], which have moreover used this information in removing shadows from images. With the described algorithm, the chromatic shift can cause problems for decomposition as seen for the lower-left shadowed area of the blanket in Figure 4. The red and pink stripes within the shadow are not matched with those in other areas because of the chromatic shift, and consequently they are not identified as having the same reflectance. Additionally, the chromatic shift within the shadow causes the red and pink stripes to appear similar in chromaticity, resulting in some reflectance changes being mistaken for shading changes. In future work, we plan to examine how chromaticity models for shadows could be incorporated into our matching procedure to improve handling of shadows from colored illumination.

6. Acknowledgement

The work is supported by Singapore FRC Grant R-263-000-477-112.

References

- [1] A. Agrawal, R. Raskar, and R. Chellappa. Edge suppression by gradient field transformation using cross-projection tensors. In *IEEE Computer Vision and Pattern Recognition (CVPR)*, volume 2, pages 2301–2308, 2006. 1
- [2] M. Bell and W. T. Freeman. Learning local evidence for shading and reflectance. In *IEEE Int’l Conf. on Computer Vision (ICCV)*, volume 1, pages 670–677, 2001. 1, 5
- [3] A. Buades, B. Coll, and J. M. Morel. A non-local algorithm for image denoising. In *IEEE Computer Vision and Pattern Recognition (CVPR)*, volume 2, pages 60–65, 2005. 2




Figure 6. Cushion. (a) Original image. (b)-(c) Shading and reflectance image without the use of texture cues. (d)-(e) Shading and reflectance image with our non-local texture algorithm.

- [4] A. A. Efros and T. K. Leung. Texture synthesis by non-parametric sampling. In *IEEE Int'l Conf. on Computer Vision (ICCV)*, volume 2, pages 1033–1038, 1999. 2
- [5] G. D. Finlayson, M. S. Drew, and C. Lu. Intrinsic images by entropy minimization. In *European Conf. on Computer Vision (ECCV)*, volume 3, pages 582–595, 2004. 6
- [6] G. D. Finlayson, S. D. Hordley, and M. S. Drew. Removing shadows from images. In *European Conf. on Computer Vision (ECCV)*, volume 4, pages 823–836, 2002. 6
- [7] G. D. Finlayson, S. D. Hordley, C. Lu, and M. Drew. On the removal of shadows from images. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 28:59–68, 2006. 6
- [8] W. Freeman, E. Pasztor, and O. Carmichael. Learning low-level vision. *International Journal of Computer Vision*, 40:24–57, 2000. 2
- [9] B. V. Funt, M. S. Drew, and M. Brockington. Recovering shading from color images. In *European Conf. on Computer Vision (ECCV)*, pages 124–132, 1992. 1, 2
- [10] R. Kimmel, M. Elad, D. Shaked, R. Keshet, and I. Sobel. A variational framework for retinex. *International Journal of Computer Vision*, 52:7–23, 2003. 1, 2, 3
- [11] V. Kolmogorov. Convergent tree-reweighted message passing for energy minimization. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 28(10):1568–1583, 2006. 4
- [12] E. Land and J. McCann. Lightness and retinex theory. *Journal of the Optical Society of America A*, 3:1684–1692, 1971. 1
- [13] T. Leung and J. Malik. Representing and recognizing the visual appearance of materials using three-dimensional textons. *International Journal of Computer Vision*, 43:29–44, 2001. 2
- [14] J. Malik and P. Perona. Preattentive texture discrimination with early vision mechanisms. *Journal of the Optical Society of America A*, 7:923–932, 1990. 2
- [15] Y. Matsushita, S. Lin, S. B. Kang, and H.-Y. Shum. Estimating intrinsic images from image sequences with biased illumination. In *European Conf. on Computer Vision (ECCV)*, volume 2, pages 274–286, 2004. 1
- [16] Y. Matsushita, K. Nishino, K. Ikeuchi, and M. Sakauchi. Illumination normalization with time-dependent intrinsic images for video surveillance. In *IEEE Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages 3–10, 2003. 1
- [17] P. Sinha and E. Adelson. Recovering reflectance and illumination in a world of painted polyhedra. In *IEEE Int'l Conf. on Computer Vision (ICCV)*, pages 156–163, 1993. 1, 2
- [18] P. Tan, S. Lin, and L. Quan. Separation of highlight reflections from textured surfaces. In *IEEE Computer Vision and Pattern Recognition (CVPR)*, volume 2, pages 1855–1860, 2006. 2
- [19] M. F. Tappen, W. T. Freeman, and E. H. Adelson. Recovering intrinsic images from a single image. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 27:1459–1472, 2005. 1, 2, 5
- [20] M. F. Tappen, W. T. Freeman, and E. H. Adelson. Estimating intrinsic component images using non-linear regression. In *IEEE Computer Vision and Pattern Recognition (CVPR)*, volume 2, pages 1992–1999, 2006. 1, 5
- [21] Y. Weiss. Deriving intrinsic images from image sequences. In *IEEE Int'l Conf. on Computer Vision (ICCV)*, volume 2, pages 68–75, 2001. 1, 2
- [22] S. Zhu, Y. Wu, and D. Mumford. Filters, random fields and maximum entropy (frame): Towards a unified theory for texture modeling. *International Journal of Computer Vision*, 27:107–126, 1998. 2