

Segmentation using Meta-texture Saliency

Yaser Yacoob and Larry Davis
Computer Vision Laboratory-UMIACS
University of Maryland, College Park, MD 20742
yaser/lsd@umiacs.umd.edu

Abstract

We address segmentation of an image into patches that have an underlying salient surface-roughness. Three intrinsic images are derived: reflectance, shading and meta-texture images. A constructive approach is proposed for computing a meta-texture image by preserving, equalizing and enhancing the underlying surface-roughness across color, brightness and illumination variations. We evaluate the performance on sample images and illustrate quantitatively that different patches of the same material, in an image, are normalized in their statistics despite variations in color, brightness and illumination. Finally, segmentation by line-based boundary-detection is proposed and results are provided and compared to known algorithms.

1. Background

Image segmentation of scenes in which man-made objects are presented in their diversity of appearance is a challenging computer vision problem. For example, while clothing is, perhaps, the most diverse of such objects, its basic components can be simply reduced to: the material of the thread (i.e., the fiber such as cotton, wool, silk, etc.), the thickness of the thread and the particular weaving pattern. Figure 1 shows images we are interested to analyze and segment, (a) rug on a hardwood floor, (b) woman with colorful hair, (c) wool sweater laid on a wood surface and (d) textured painting hanging on a wall. Our objective is to delineate regions that have a salient surface texture despite significant variations in color, brightness and illumination attributes. Specifically, we seek the separation of the rug, hair, sweater and painting. We use the term *meta-texture* to convey and accentuate the *underlying* appearance of rough surfaces in the image plane to differentiate from texture that conveys the color patterns on a smooth surface and 3D-textons that are tuned to the *natural* appearance of monochromatic surface roughness under viewpoint and illumination variations. These images are taken at high res-



Figure 1. A rug on a wood floor, multi-color hair, wool sweater on a wood surface, and textured painting on a wall (respectively).

olution (1600x1200 for the face and 4368x2912 otherwise) to capture surface-detail.

Color and texture-based segmentation. There is a large body of research on texture-based segmentation of color images (e.g., [3, 5]). Color and texture features are typically extracted separately then clustering in a joint space is conducted. If applied to the images in Figure 1 they will result in the identification of multiple regions due to color or apparent texture variations and despite the similarity of the underlying surface texture.

Material perception and 3D Textons. Adelson et al. [1] studied the properties of materials under different illumination to determine classification by humans and study statistical texture measures. Leung and Malik [8] (also [5, 13]) proposed an approach for classifying materials based on 3D texture attributes, 3D textons, computed over small patches to capture local geometric and photometric properties of monochromatic images. Recognition of different materials under different lighting and viewing conditions were shown. Note that 3D texton-based approach can be adapted for image segmentation of Figure 1(c) since the color textured regions have clear boundaries and large sizes which in the monochromatic image may be amenable to detection, normalization and analysis. However, Figure 1(a,b,d) pose a challenge since the monochromatic image consists of small regions of constant texture in Figure 1(a) and gradual intensity changes in Figure 1(b,d). Consequently, we view the meta-texture image proposed in this paper as a potential input for a 3D-texton based recognition process and not as an equivalent approach for processing image texture.

Intrinsic Images. Intrinsic images (e.g., [7, 9, 12]) aim to reveal the underlying physical properties of a scene by estimating the shading (e.g., a function of lighting and surface

This work was sponsored by ONR grant N000140610102.

normals) and reflectance (e.g., surface color) images. Intrinsic analysis is suitable for coarse-level images (i.e., when surface roughness is not visible) but is less suitable for fine-level images where surface roughness is detrimental to the stability of the surface normal over small neighborhoods. Intrinsic images involve scenes which arise from albedo or color variations on smooth surfaces while Fig. 1 involves rough surfaces with complex dependencies on color, viewing and illumination directions that destabilize the albedo and color variations at fine scales. Measurement of rough surfaces have been proposed using the Bidirectional Texture Functions (BTF) [6] and analysis and recognition [5]. Analysis of rough surfaces in real-world images without BTF or learned models remains open for research.

This paper is focused on the diverse appearance of real-world rough-surfaces that defy smooth-surface assumptions and their surface texture functions (e.g., BTF, 3D Textons) are unknown. The paper’s contributions are characterizing the problem, proposing a discriminative approach to derive three intrinsic images from a color image (i.e., shading, reflectance, and meta-texture), proposing and developing the concept of salient meta-texture image (MTI) via transforming an image into a grey-level image in which the projected surface roughness is *preserved*, *equalized* and *enhanced* while other properties such as color, brightness variations, etc. are normalized. This MTI is evaluated by considering image segmentation by texture-boundary detection.

It is important to note that evaluating correctness of the derived three intrinsic images for real-world images cannot be done (a problem we share with existing work on intrinsic images). Moreover, the MTI is derived from intertwined imaging process and attributes of the scene. For example, in Figure 1(c) meta-textures of red regions appear less sharp than the yellow regions simply due to the imaging process while their surface-roughness are presumably equal. Nevertheless, we provide empirical results that quantify the performance with respect to similarity within an image and effectiveness of segmentation.

2. Approach

We propose that three derived images be computed, *shading*, *reflectance* and *surface-roughness*. The MTI is a logical extension of intrinsic images [9, 12] to account for surface-roughness. Deriving these images by extending known formulations of intrinsic images is not possible since surface-roughness violates the basic assumptions of derivative-based approaches (e.g., [9, 12]). Instead, we propose a patch-based approach to extract intrinsic information that is specifically-suited for image segmentation using surface-roughness appearance.

We define an MTI to be a grey-scale image derived from a color image so that scene surface-roughness similarity translates into similarity of the image texture regardless of

reflectance (e.g., color, absorbance, etc.), *lighting* (e.g., intensity, location) and *shading* (e.g., slant of the smooth surface as long as roughness detail remains visible) variations *within* the image (note image-centricity as opposed to inter-image variations in [5]).

2.1. Homogeneous Patch Selection

We first convert the RGB color image into (Hue, Saturation, Value) (HSV) space. The proposed transformations are applied to homogeneous patches that share underlying image attributes. We employ a conservative approach for patch delineation since some over-segmentation typically has little negative impact because the statistics of multiple patches sharing an underlying structure are similar. Each pixel, $P_{i,j}$, in HSV space will belong to a single patch R . Given a seed patch R with a single pixel $P_{i,j}$, a connected component expansion of R under constraints on the values $P_{i,j}^H, P_{i,j}^S, P_{i,j}^V$ can be performed. A pixel, $P_{k,l}$, 8-connected neighbor of a pixel, $P_{i,j}$, is added to R if

$$|P_{i,j}^H - P_{k,l}^H| < H_{thr}; |P_{i,j}^S - P_{k,l}^S| < S_{thr}; |P_{i,j}^V - P_{k,l}^V| < V_{thr} \quad (1)$$

where $H_{thr} = 10$, $S_{thr} = 40$, and $V_{thr} = 50$. These values were determined empirically and are applied to all images in this paper while taking into consideration the nonlinearity of the HSV space.

A two-pass algorithm for segmenting the image into patches is implemented. In the first pass each pixel is allowed to be part of as many patches as it conforms to, while in the second the optimal patch of a pixel is chosen based on the patch sizes and similarity of the pixel to patch statistics. Typically, thousands of patches are found in an image.

2.2. Deriving Intrinsic Images

We define our *variants* of intrinsic images that are related but not equal to existing formulations of intrinsic images. The visual information of a patch R is divided into shading, reflectance and meta-texture. Reflectance conveys the pure substrate color of the patch (i.e., independent of surface roughness and shading). Shading reflects illumination of the smooth-surface of the patch (i.e., excluding its surface roughness). Meta-texture reflects illumination variations due to surface roughness. Meta-texture and shading are related since both involve illumination interaction with scene surfaces, but they differ in capturing high and low 2D frequencies (if combined they equal shading in [12]).

Since reflectance is independent of surface-roughness and smoothness of R it can be computed by combining the hue at each point of the patch (or the average hue) with the average saturation and brightness over R (i.e., recombined (H, S_{avg}, V_{avg})). Since shading reflects illumination variations of the smooth-surface of R it can be computed by discarding the high-frequencies (by low-pass filtering) of the brightness V since these stem from surface roughness. It is critical to note that we associate shading with the *maximum*

amplitude of the light waveform as opposed to the *combined* RGB amplitudes of light waveforms (i.e., intensity; compare V and intensity images of the blanket in Fig. 2). We argue (in contrast to [7, 9, 12]) that since shading reflects the interaction of the illumination source with the surface normal it should reflect the maximum amplitude measured in light wavelengths discarding weaker wavelengths. The basis is that assuming scene material disperses all wavelengths in the same direction, and given homogeneity of the patch, all wavelengths reflect *equally* the surface shading with the maximum wavelength amplitude providing the strongest signal. Finally, meta-texture reflects the removal of the smooth-surface variations from the intensity image of the patch (e.g., by a high-pass filtering). Note that either of the intensity image (i.e., the combined wavelengths) or V can be used since surface-roughness is typically equally conveyed unless one of the RGB channels is saturated in which case the intensity image preserves the local texture while V eliminates it. Normalizing values across patches is necessary for the shading and MTI images (for now we use simple mean normalization).

Figure 2 shows the intrinsic images derived for three sample images: T-shirt and blanket with folds, and a wool sweater over an arm. The images convey color variations, surface-roughness and shading of smooth surface deformations. For each image, the input, intensity, V , reflectance, shading and MTI are shown, respectively. The reflectance images show a completely flat image with surface smoothness and roughness removed, the shading images show the global surface deformation revealed by the illumination and the MTIs show the surface-roughness excluding the shading information. Notice that patch boundary artifacts appear in the meta-texture and shading images. The bottom right side images show the reflectance and shading images as computed by [12]. Note that the shading images confound smooth and rough-surface information in comparison to our algorithm for separating them.

In the rest of the paper we focus on the MTI since it holds a great potential for image segmentation.

2.3. Deriving a Meta-texture Image

The derivation of MTI can be improved by adopting a construction-based approach. The objective is to transform a patch into a grey-scale patch in which surface-roughness is *preserved*, *equalized* and *enhanced* where needed. Preservation ensures that the meta-texture is not weakened. Equalization means that patches with the same underlying surface-roughness but perhaps different color, brightness or illumination are transformed into a similar meta-texture. Finally, where surface-roughness appears weak due to color, brightness or illumination it is enhanced (i.e., amplified) to reflect its prototypical appearance. Figure 3 shows these requirements computed for a colorful cotton sweater. Preservation reveals the meta-texture of the

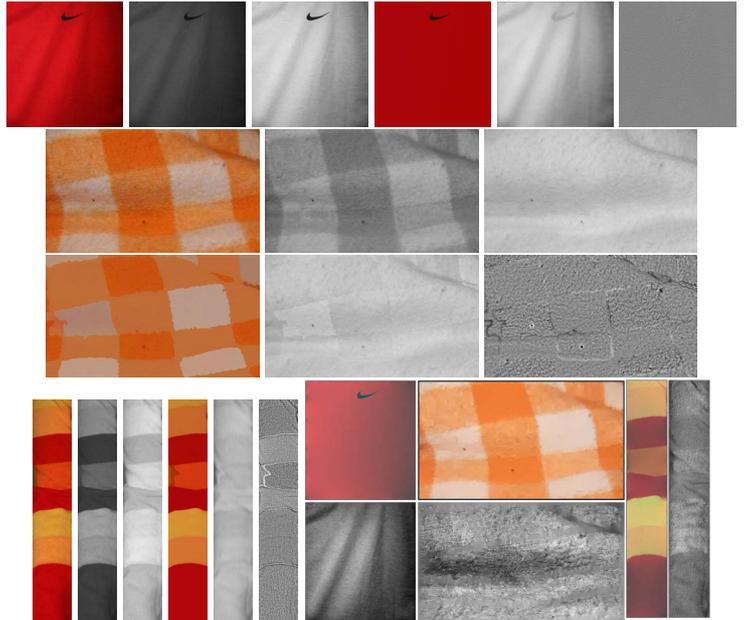


Figure 2. From top to bottom (raster-wise) input, intensity, V , reflectance, shading and MTI for each image followed by the respective reflectance and shading images computed by [12].

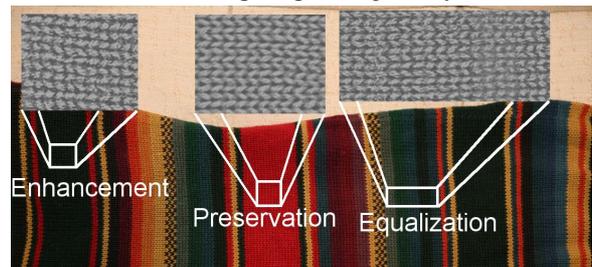


Figure 3. Achieving the requirements of MTI, preservation, equalization and enhancement on a cotton sweater.

red region patch. Equalization shows that the meta-texture is equalized across the five shades of green and blue. Enhancement is shown for the black region where the surface roughness is revealed, enhanced and equalized with respect to other meta-textures. The MTI of these patches are similar and thus could easily be classified into the same material.

We first consider the distribution of brightness values within homogeneous patches. Figure 4 (top two rows) shows sample patches and histograms of the brightness values of each patch. The patches were taken from the brightness images, V , of the rug, hair and wool and cotton sweater images, respectively (shown in figure 1 and 3). The histograms of the brightness coarsely fit a Gaussian distribution which is not surprising since the patches are assumed to be homogeneous. Each patch, however, has a distribution determined by its underlying texture.

The MTI of homogeneous patches should satisfy:

Preservation. The histogram of the brightness values of the MTI should qualitatively preserve the distribution of the brightness values.

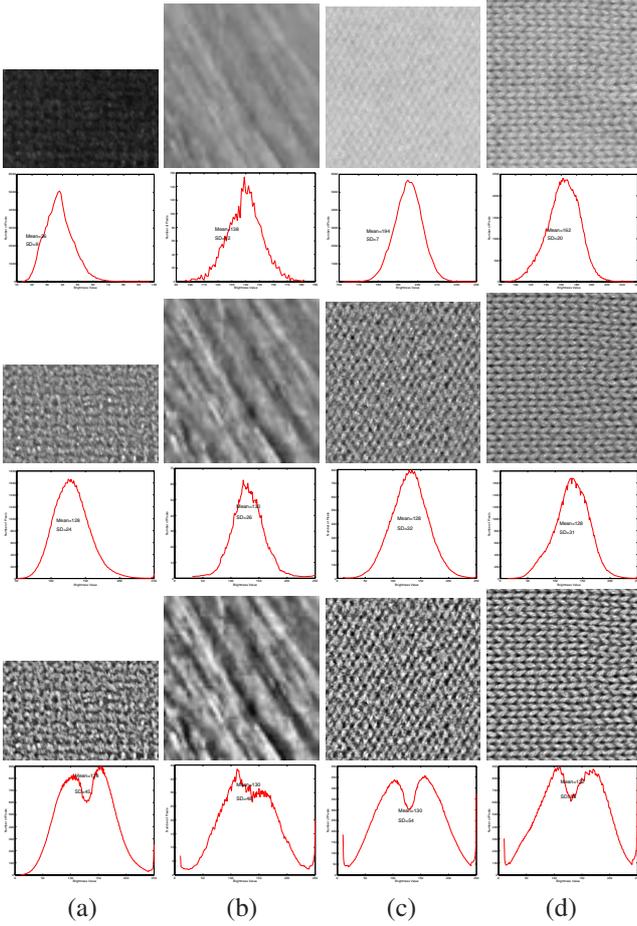


Figure 4. From top to bottom rows, (1) original brightness, (2) histogram of the brightness values, (3) MTI using $\gamma = 1$ (4) histogram of the brightness values for MTI $\gamma = 1$, (5) MTI using $\gamma = 0.8$ (6) histogram of the brightness values for MTI $\gamma = 0.8$. (a) rug (b) hair colored with different shades (c) a wool sweater and (d) cotton sweater.

Equalization. The mean of each meta-texture patch should be close to the means of other patches. Let MTI_{mean} denote the desired mean of MTI.

Enhancement. The standard deviation of any patch should be close (not necessarily equal) to other patches. Let MTI_{sd} be the desired standard deviation of MTI.

Let R_{mean}, R_{sd} denote the average and standard deviation of the histogram of the brightness values of pixels of R . Let P^v denote the brightness values of R . MTI will be constructed as a grey-scale image that takes on values between $(V_{min}, V_{max}) = (0, 255)$. Let $MTI_{mean} = (V_{max} + V_{min})/2$ which places the mean brightness at the center of the range of brightness values. The brightness value, $P_{i,j}^v$, of a pixel (i, j) , is transformed into a value, $P_{i,j}^{MTI}$, in the MTI,

$$P_{i,j}^{MTI} = ((P_{i,j}^v - P_{i,j}^D - R_{mean}) / (3R_{sd})) \cdot ((V_{max} - V_{min}) / 2) + MTI_{mean} \quad (2)$$

where $P_{i,j}^D$ is a floating point value of pixel (i, j) in the im-

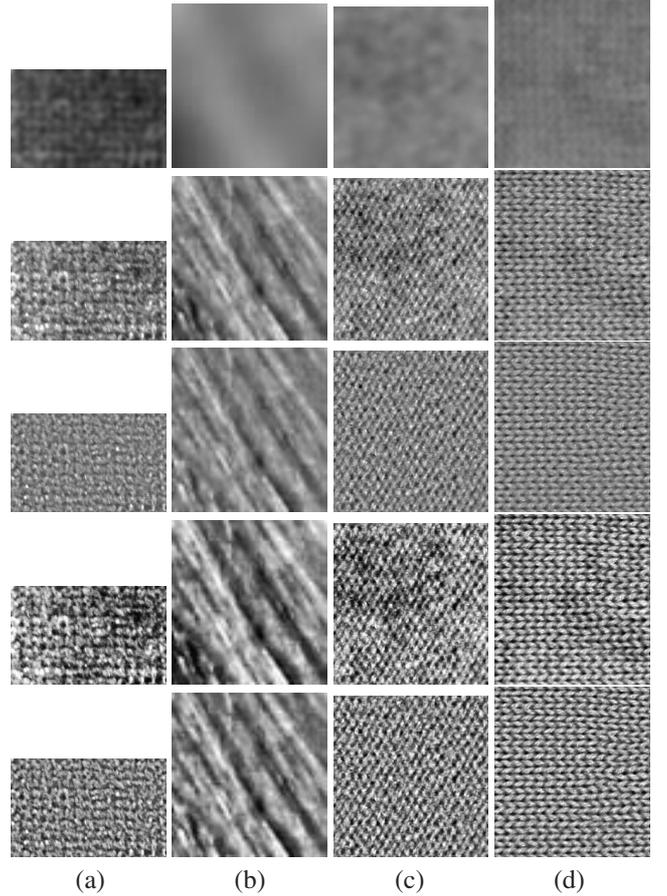


Figure 5. From top row to bottom, (1) the P^D image which is a smoothed version of the deviation from the average brightness of the patch (2) the MTI with $P^D = 0$ and $\gamma = 1$, (3) the MTI with P^D and $\gamma = 1$, (4) the MTI with $P^D = 0$ and $\gamma = 0.8$, (5) the MTI with P^D and $\gamma = 0.8$. (a) rug (b) hair colored with different shades (c) a wool sweater and (d) cotton sweater

age P^D . P^D is the Gaussian smoothed deviation of the brightness values of pixels in R from the average brightness in the patch R ,

$$P^D = \mathcal{S}(P^{v0}) \quad \text{and} \quad P_{i,j}^{v0} = P_{i,j}^v - R_{mean} \quad (3)$$

where \mathcal{S} is a Gaussian smoothing function. Since we approximate the brightness histogram as a normal distribution, the area within $\pm 3R_{sd}$ approximately covers 99.73% of the points. It should be noted that this normal-distribution assumption is an approximation and points beyond $\pm 3R_{sd}$ are treated as being at $\pm 3R_{sd}$.

In Figure 4, the third and fourth rows show the transformation of the four different patches into a MTI using Equation 2 and the histograms of the values of the MTIs. The appearance of the underlying textures is preserved, enhanced and equalized. Quantitatively, the means and standard deviations of the MTIs are close and thus readily reflect roughness differences between these patches.

Equation 2 can be modified to enhance the contrast

of the underlying brightness distribution by nonlinear amplification of the brightness instead of the linear scaling. A gamma-correction model accentuates brightness differences within the patch,

$$P_{i,j}^{MTI} = \text{sign}(P_{i,j}^G) \cdot |P_{i,j}^G|^\gamma \cdot ((V_{max} - V_{min})/2) + MTI_{mean} \quad (4)$$

where $P_{i,j}^G = (P_{i,j}^v - P_{i,j}^D - R_{mean})/(3R_{sd})$, the function sign is 1 if $P_{i,j}^G$ is positive and -1 otherwise, and $|P_{i,j}^G|$ is the absolute value of $P_{i,j}^G$. It is easy to see that if $\gamma = 1$ then Equation 4 is equal to Equation 2. If $\gamma < 1$ then $P_{i,j}^G$ values closest to R_{mean} are enlarged more than values far from R_{mean} and contrast is enhanced within the patch. In Figure 4, the last two rows show the images and histograms of brightness distributions for $\gamma = 0.8$. The enhancement in texture contrast is readily visible. The graphs in Figure 4 (bottom row) show that at the extreme low and high end of the brightness the number of pixels in the histogram have increased (the rising tails of both ends). This occurs since more points are shifted to the edges of the brightness range. Similarly, the mean area, which is the center of the transformation lost points that have moved farther to increase the brightness contrast. It is worth noting that $\gamma = 0.8$ does not preserve the normal distribution but it provides greater discrimination power. If a smaller γ is used these phenomena increase and overall it is not a desirable outcome since it can excessively distort the histogram. If $\gamma > 1$ a suppression of the texture contrast occurs (not relevant to our research objective). In the rest of the paper we use $\gamma = 0.8$ in all experiments.

As defined in Equation 3, P^D is a Gaussian-smoothed image of the normalized brightness values (i.e., average intensity subtracted from the brightness value). In our experiments we use a Gaussian smoothing filter of 41 units. P^D removes low-frequency brightness differences (e.g., shading) within a patch whether these were a result of material brightness or illumination differences. In Figure 5, top row, the value of P^D is shown as a grey-scale image although the actual values are positive and negative numbers. The second and fourth rows show the MTI images for $P^D = 0$, while the third and fifth rows show the MTI images using the actual P^D ($\gamma = 1$ and $\gamma = 0.8$, respectively).

2.4. Evaluation of MTI Transform

We compare MTI to histogram equalization and low-frequency elimination and normalization using Fourier Transform. The latter approach eliminates the lowest-frequencies and normalizes values. Figure 6 shows the original V patches in (a) columns, MTIs computed using Equation 4 in (b) columns, histogram equalized patches (c) columns, and low-frequency eliminated and normalized images in (d) columns. The MTIs achieve far more uniformity of texture appearance within images and between images. Brightness variations due to illumination differences, material creases, or other anomalies are removed or reduced



Figure 7. Left, a colorful region in the sweater and its MTI. Right, the distance between the textures on the sides of A is larger than the distances between the textures on the sides of B.

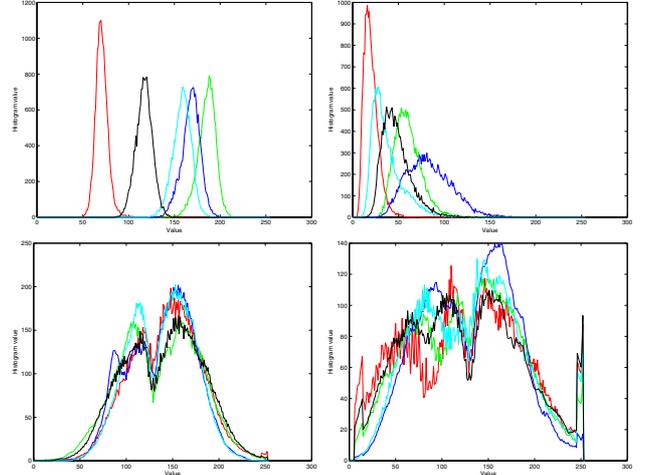


Figure 8. The top row shows the histogram of brightness values for 5 wool (left) and cotton (right) patches, the bottom graphs show the histograms for the MTI values.

significantly so that the underlying texture appearance is uniform. Note that the low-frequency removal weakens the sharpness of the texture. Robustness of MTI to linear transformations can be shown mathematically but is omitted.

Figure 7(left) shows MTI of an area taken from Figure 3. The MTI reflects appropriately the underlying salient weave of the sweater. Very small regions such as the black/yellow interweaved thread conform to the general texture pattern. The meshing of patch boundaries remains an open challenge; Boundary pixels are occasionally a blend of two colors, and depending on whether they are associated with one patch or another (or form a third patch) their meta-texture value will be different.

While it is not possible to evaluate MTI with respect to the ideal MTI, we can compare the similarity of MTIs of different patches in which surface roughness is presumed to be equal. Two sets of five patches of different colors were selected from both the wool and cotton sweaters (Fig. 1 (c) and Fig. 3). Fig. 8 shows the histograms of the original brightness and the MTI values (top and bottom row, respectively) for the wool and cotton sweater patches. An ideal transformation should bring each set of 5 patches to perfect alignment (if the original patches were indeed of identical surface-roughness). The MTI comes close to this ideal given likely intrinsic differences between the patches.

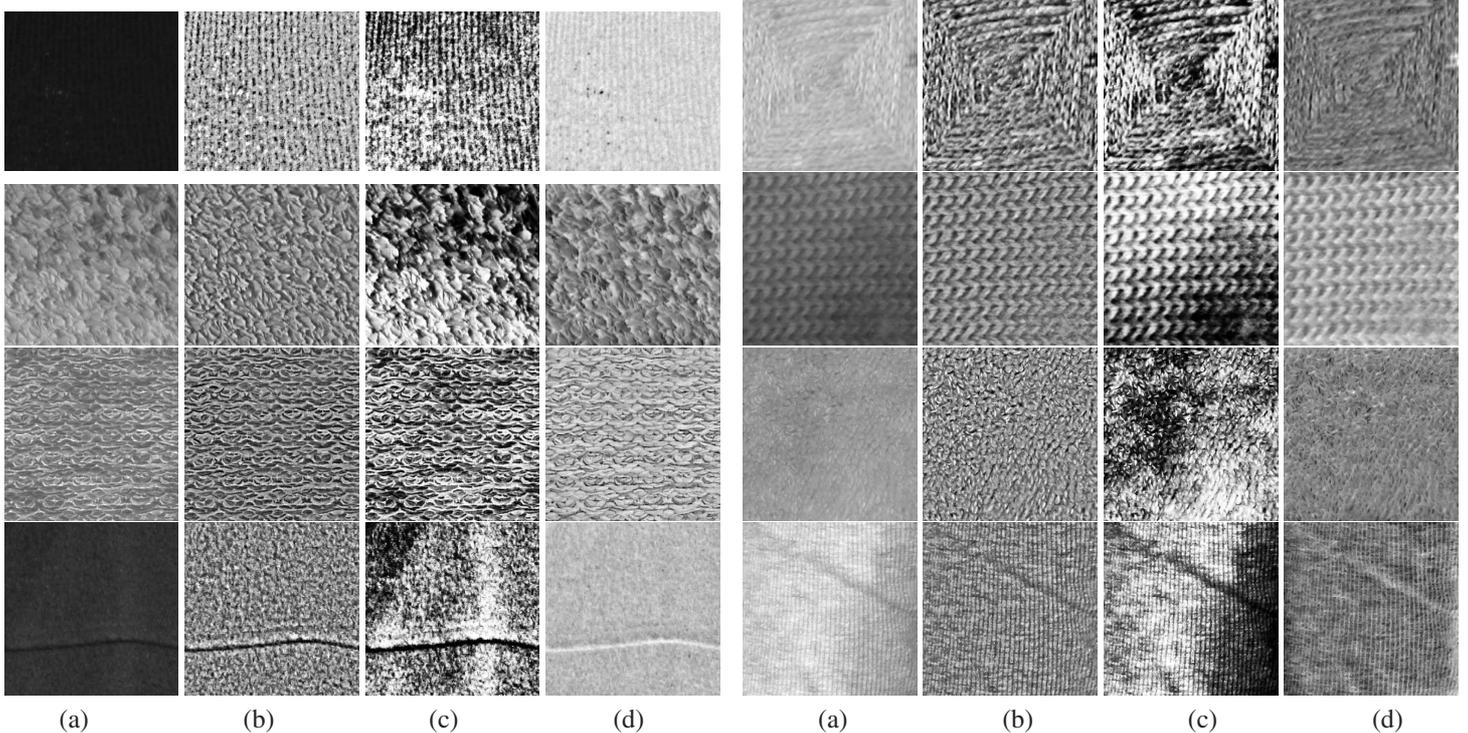


Figure 6. Columns show (a) the original V image, (b) the MTI of the patches using Equation 4, (c) the transformed patches using histogram equalization, (d) the low-frequency removal and normalized filtering.

3. Segmentation of the Meta-texture Image

There is a variety of approaches for computing texture features [1, 3, 5, 8, 10, 13]. We adopt the Gabor wavelets method (details can be found in [10]). We re-pose the image segmentation problem as a line-based texture-boundary detection. We make a critical assumption that the true boundary between two different textures *coincides* with the maximum distance between textures on the sides of the boundary. In Figure 7(right) it means that the distance between the textures on the sides of line *A* is larger than the distance between the textures on the sides of any line such as *B*. The basis for this assumption is that the difference between the textures is maximized when the textures are as different from each other as possible. However, there are situations in which the assumption is violated since textures are not normally as distinct as illustrated. Note that since line-boundaries are computed, inaccuracies and jagged boundaries appear at curved boundaries. The texture-boundary detection consists of the following:

Coarse texture representation. Given the MTI, we compute all $N \times N$ texture descriptors for a sliding square window (in raster scan). The vertical regions do not overlap, while the horizontal regions overlap by $N/2$. We compute the Gabor wavelet features for each square rotated between ± 75 degrees by 15 degrees steps.

Coarse texture-boundary detection. We evaluate the evidence of a texture boundary between any two neighbor-

ing squares (both horizontal and vertical neighbors). If the minimum distance between their texture features (across orientations) exceeds T then a boundary is assumed to exist.

Fine texture-boundary detection. Since the boundary between textures is unlikely to fall at the edge between adjacent squares it is necessary to estimate the optimal *line location* and *orientation* of the boundary. The optimal boundary is the one that maximizes the distance between the textures on its sides even when all image orientations are considered. In practice, a scan of all possible boundary locations and orientations (± 45) is done, where for each candidate boundary the minimum distance between the textures on the two sides of the candidate boundary is computed, and the maximum of these minima is selected as the optimal boundary. To reduce computations, this process is done raster-wise skipping N rows for vertical line boundaries and N columns for horizontal line boundaries.

The texture-boundary detection is followed by delineation of regions of uniform texture and merger of regions that have similar textures. Delineation of regions is achieved by connecting proximate edges texture boundaries. Two regions are merged if the distance between their textures is less or equal than $T + 3$.

4. Results

Two parameters need to be set in the image segmentation approach: the window size for texture feature extraction is $N = 128$ and the distance threshold between texture fea-

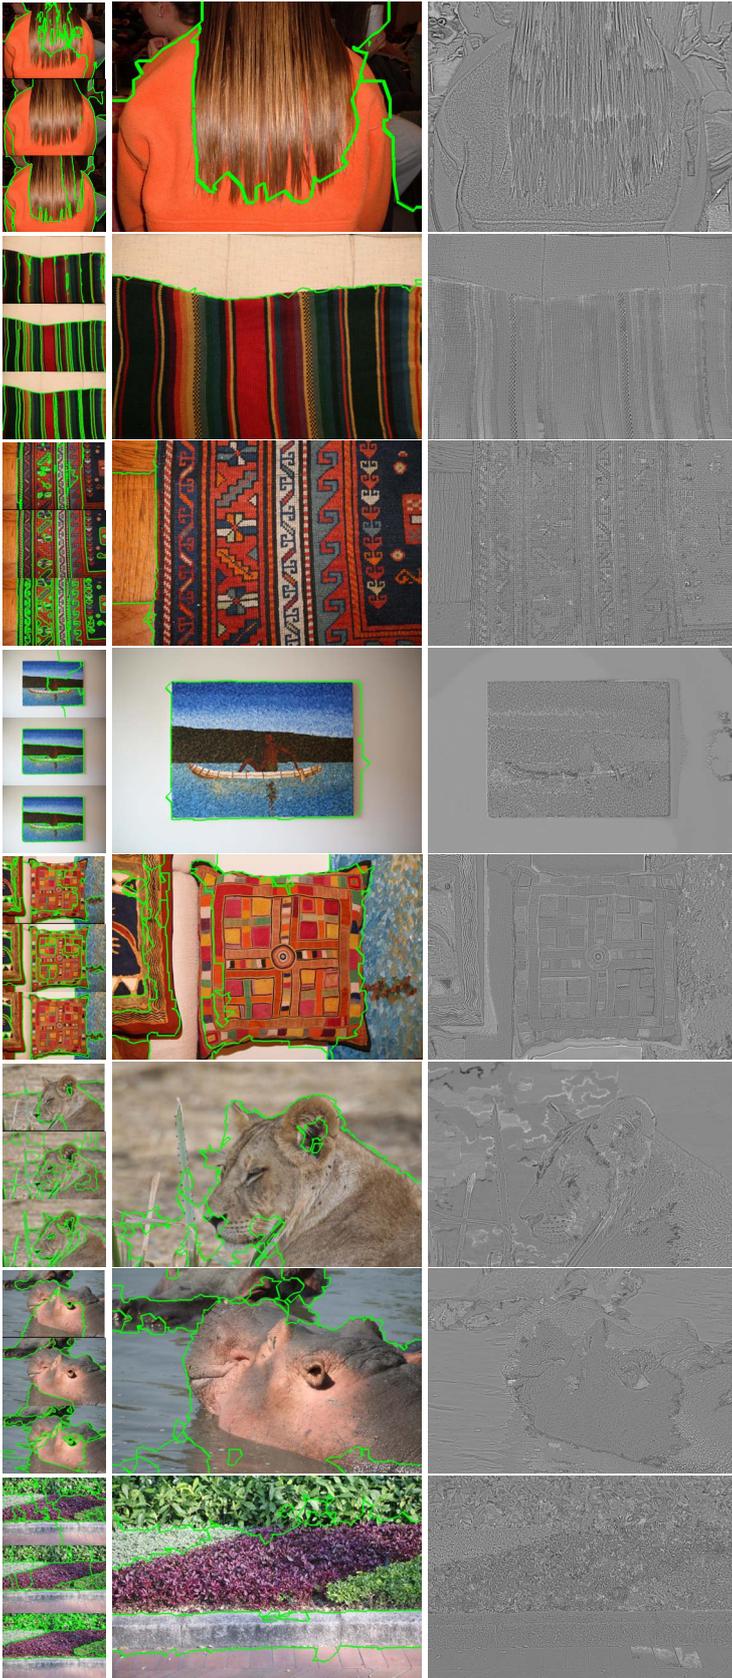


Figure 9. MTI (right), segmentation of MTI (middle) and segmentation using [4] [11] [2], top to bottom respectively.

tures determines whether two texture regions should have a line-boundary detected, $T = 7$ ($N = 64$ and $T = 6$ for the lower resolution images).

Figures 9,10 show MTIs (right column), segmentation using MTI (center column), and the output of three segmentation algorithms, Cour et al [4], Martin et al [11] and Arbelaez [2] (top to bottom, respectively, left column). The top image shows hair with highlights and shaded regions. The MTI clearly reveals the hair as having a uniform texture and accentuates the texture of the fleece. The cotton sweater, rug and painting show near optimal segmentation (including texture differences within the wood). The multi-colored pillow is in front of the textured painting and to its left the cream color fabric has visible texture while the other colorful pillow region (on the left edge) has no meta-texture visible at this resolution. The segmentation result is reasonable as the pillow, painting, wall and cream-color fabric are properly segmented. Finally, three outdoor images show appropriate segmentation of a lion, hippopotamus, and vegetation where surface detail is visible. The reflection of hippopotamus appears close to its surface roughness and therefore it was merged with the hippopotamus, also the vegetation is segmented by the size of leaves not color.

In Figure 10, the roughness of the sweaters is sufficient for segmentation. Parts of the background scenes mostly lack surface detail so the segmentation typically follows color edges (except the carpet in the third image where surface roughness is visible). Parts of the sweater are not merged properly since the 3D creases of the texture are not accounted for by our texture analysis algorithm. The second image from the top is notable since the background and foreground share the same colors, and the top of the pants at the waist is separated despite having the same white color.

The experiments suggest that MTI-segmentation provides better results than the three approaches since surface-roughness saliency transcends edges and color attributes which lead to over-segmentation in other approaches. Our approach suffers from inaccurate boundary detection brought about by using a single cue, line-boundary constraint and sparse boundary search. These are not core issues and can be remedied in future research. The absence of roughness detail leads to reflecting differences in color and as a result the segmentation suffers.

5. Summary

We addressed the problem of normalizing surface-roughness of image patches. A transformation of an image into MTI was proposed by the constraints: preservation, equalization and enhancement of meta-texture properties of patches across color, brightness or illumination differences. The MTI was used as input to a texture boundary segmentation process and experiments on realistic scenes were reported with respect to three state-of-the-art algorithms.

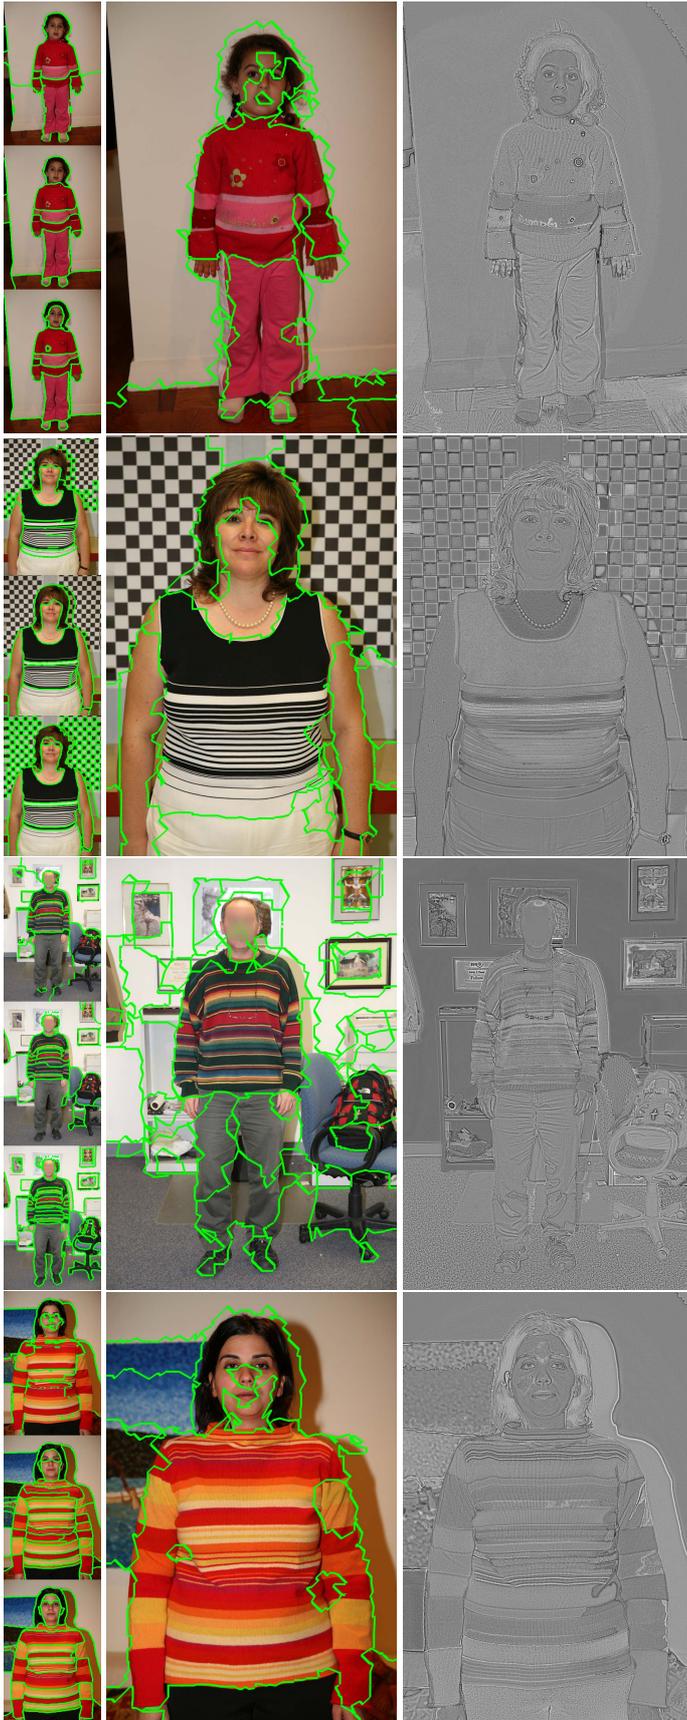


Figure 10. MTI (right), segmentation of MTI (middle) and segmentation using [4] [11] [2], top to bottom respectively.

While our results are better than these algorithms, they suffer from jagged and inaccurate line-boundaries since we use only surface-roughness (as opposed to edges and color and texture) and we approximate the boundary by lines and perform sparse computation to reduce complexity.

Acknowledgment. We thank the authors of [2, 4, 11, 12] for facilitating comparing our results to their work.

References

- [1] E.H. Adelson, Y. Li and L. Sharan. Image statistics for material perception. *Journal of Vision*, Volume 4, Number 8, Abstract 123, 2004. 1, 6
- [2] P. Arbelaez. Boundary Extraction in Natural Images Using Ultrametric Contour Maps. Workshop on Perceptual Organization in Comp. Vision, 2006. 7, 8
- [3] J. Chen, T. N. Pappas, A. Mojsilovic, B. E. Rogowitz. Adaptive perceptual color-texture image segmentation, *Tr. Im. Processing*, 14, 2005, 1524-1536. 1, 6
- [4] T. Cour, F. Benezit, J. Shi. Spectral Segmentation with Multiscale Graph Decomposition. CVPR, 2005, 1124-1131. 7, 8
- [5] O. Cula and K.J. Dana, 3D Texture recognition using bidirectional feature histograms, *IJVC*, (59)1, 2004, 33-60. 1, 2, 6
- [6] K.J. Dana, B. van Ginneken, S.K. Nayar, and J.J. Koenderink. Reflectance and texture of real world surfaces. *ACM Transactions on Graphics*, 18(1):134, 1999. 2
- [7] G.D. Finlayson, M.S. Drew, and C. Lu. Intrinsic Images by Entropy Minimization. *ECCV 2004*, 582-595. 1, 3
- [8] T. Leung J. Malik. Representing and recognizing the visual appearance of materials using three-dimensional textons. *IJCV*, 43(1),2001. 1, 6
- [9] E.H. Land and J.J. McCann, Lightness and Retinex Theory. *J. of Optical Soc. Am.*, 61, 1971, 1-11. 1, 2, 3
- [10] B.S. Manjunath and W.Y. Ma. Texture features for browsing and retrieval of image data. *PAMI*, (18)8, 1996, 837-842. 6
- [11] D. Martin, C. Fowlkes, J. Malik. Learning to Detect Natural Image Boundaries Using Local Brightness, Color and Texture Cues, *TPAMI* 26 (5) 530-549. 7, 8
- [12] M.F. Tappen, W.T. Freeman, and E.H. Adelson. Recovering Intrinsic Images from a Single Image. *TPAMI*, (27)9, 2005, 1459-1472. 1, 2, 3, 8
- [13] M. Varma and A. Zisserman. Classifying images of materials: Achieving viewpoint and illumination independence. *ECCV*, 2002, 3, 255-271. 1, 6