



Perceptual Segmentation: Combining Image Segmentation with Object Tagging

Ruth Bergman, Hila Nachlieli, Gitit Ruckenstein, Mark Shaw and Ranjit Bhaskar
HP Laboratories
HPL-2008-185

Keyword(s):

object detection, image tagging, skin detection, sky detection, foliage detection

Abstract:

Most consumers do not want to edit their images, either because they do not have the time, or they do not have the know how. They do want to be able to press a button that will magically make the objects captured in their photo look better. At the heart of enabling such functionality, lay image analysis models. The more we know about the objects in the photo, the better we can enhance and modify it according to human preferences. We present a necessary piece of this puzzle, breaking the image into significant segments and finding important perceptual objects including skin, sky, snow and foliage.



Perceptual Segmentation: Combining Image Segmentation with Object Tagging

Ruth Bergman, Hila Nachlieli, Gitit Ruckenstein*, Mark Shaw and Ranjit Bhaskar†

October 6, 2008

Abstract

Most consumers do not want to edit their images, either because they do not have the time, or they do not have the know how. They do want to be able to press a button that will magically make the objects captured in their photo look better. At the heart of enabling such functionality, lay image analysis models. The more we know about the objects in the photo, the better we can enhance and modify it according to human preferences. We present a necessary piece of this puzzle, breaking the image into significant segments and finding important perceptual objects including skin, sky, snow and foliage.

Introduction

Today, even relatively simple image processing pipelines include mature algorithms for image denoising, image sharpening, contrast enhancement, color correction and more. In order to keep HP's leading position in the imaging market, HP imaging products are adding features that require some level of content understanding. Examples of such features that exist in current products are red eye removal, pet eye and automatic cropping. Each one of these algorithms requires the identification of relevant objects in the image. We propose a new algorithm in which a segmented image is tagged to identify specific object categories.

Our approach leverages recent advances in image segmentation [4, 6]. While there have been important advances in the quality of segmentation, improvement in speed has made this technology viable for some imaging applications. We also build on our prior work in the area of object detection [5]. Our current work unites these two image analysis techniques. Our algorithm to date tags blue sky, gray sky and snow, skin, and foliage. It produces demonstrable and compelling tagged segmentations, see, e.g., Figure 1. We discuss tests on natural outdoor images later in the paper.

Object detection uses as much domain knowledge about objects as is available. For example, Skin detection uses highly specialized features including face detection, skin color models and shape characteristics. Prior work developed some unified methods for identifying various objects in the image [5]. The approach evaluates regions in the image that match some expectation of the object, e.g., with regard to size, shape, location and color. For example, sky detection uses a color model that describes a range of blue and gray colors with high luminance, and evaluates additional characteristics we expect in sky regions, such as size, smoothness, gradients and overall luminance, to ascertain a final probability that the region depicts sky.

*Gitit Ruckenstein, our esteemed colleague and friend, passed away in August 2007

†ruth.bergman@hp.com, hila.nachlieli@hp.com, mark.q.shaw@hp.com, ranjit.bhaskar@hp.com

In the object detection approach, both skin and sky detection use object specific approaches to finding relevant regions. Skin regions arise from face detection, and sky regions are found using an open space method. For highly textured objects, for example, foliage, water and sand, there is no pixel-wise feature that we can use to define a region. The variety of appearance and placement of such objects in images is too unpredictable. In order to apply regional considerations for the detection of such objects, we need a global division of the image into regions, i.e., a segmentation.

Image segmentation is the process in which an image is divided into non-overlapping regions. The primary objective of any segmentation algorithm is to break an image up into a small number of contiguous regions of similar type. See, e.g., Figure 1(b). Once the image has been segmented, further analysis of those regions using specific heuristics will determine the features that one may want to tag within an image. Our algorithm currently tags blue sky, gray sky and snow, skin and foliage, as demonstrated in Figure 1(c).




Figure 1: Detection of skin, blue sky and foliage.

Tagged segmentation is a critical technology for a variety of image manipulation applications. For example, with the current algorithms, it is straightforward to develop tools for skin tone correction and sky correction. Figure 2 demonstrates some of the advantages of object aware image enhancement and manipulation. The original image is shown in Figure 2(a), with the corresponding tagged segmentation in Figure 2(b). Figure 2(a) was enhanced with an existing automatic contrast enhancement algorithm in Figure 2(c). Figure 2(a) shows the same image enhanced using object tagging. Contrast enhancement was applied in an object aware manner, specifically setting the skin to the expected luminance. The sky was reconstructed to have a more pleasing color and cloud texture.

Problem Statement

Given a natural image as input, our goal is to develop a set of algorithms that output a non-overlapping segmentation of the image, where each segment is tagged as a perceptual object. Both the segmentation and the perceptual tagging should coincide with human perception. For practical reasons, object tagging is partial. That is, some segments are labeled as known objects while others remain “unknown”.

Our Solution

The algorithm consists of two stages. First, the image is segmented based on location, color and texture characteristics. Next, we tag segments using object detection. To date the focus has been on detecting objects that most commonly appear in images. For all these objects, we can endow an algorithm with a large amount of prior knowledge, which may be based on our perception of the objects, their physical properties, or how they typically appear in images.

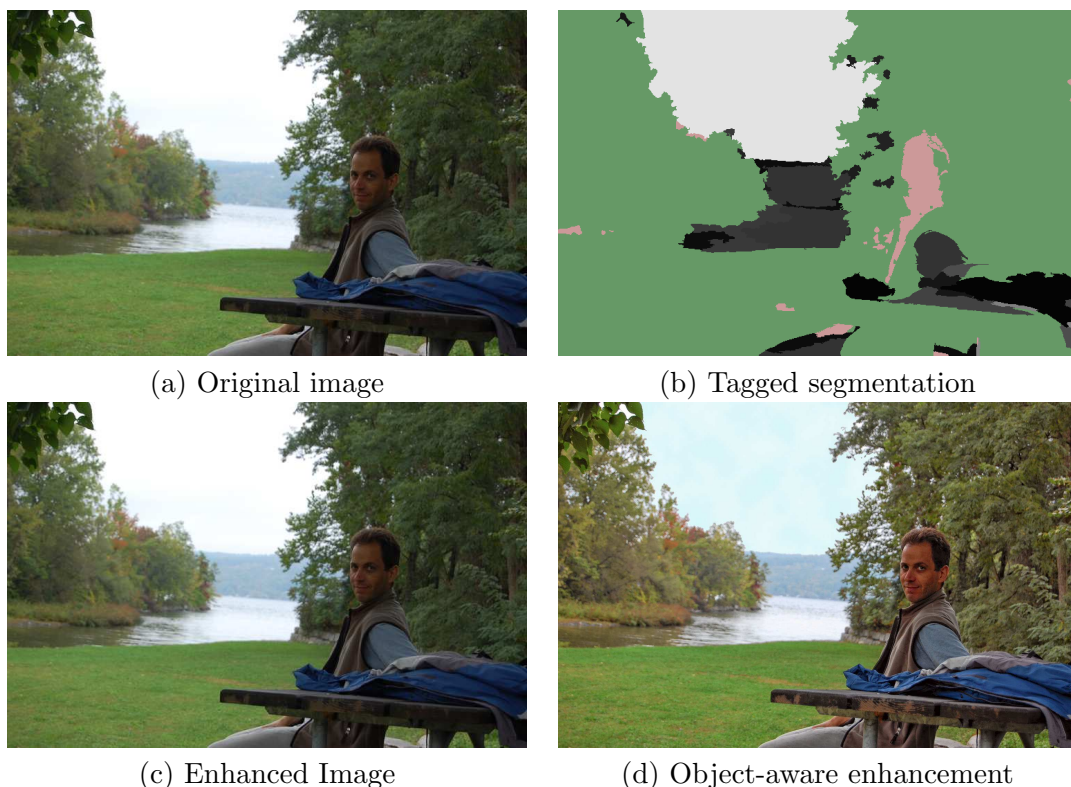


Figure 2: A demonstration of the benefit of tagged segmentation to image enhancement. The original image (a) is enhanced with an existing automatic contrast enhancement algorithm (c). The same image was enhanced in (d) using object tagging. Contrast enhancement was applied in an object aware manner. The sky was reconstructed with color and cloud texture. The tagged segmentation of (a) is shown in (b).

Stage 1 - Segmentation

Pal and Pal[12] have provided a comprehensive review of traditional approaches to segmentation. More recently, the region growing approach based on color and texture features has been used with good results [2, 4, 6]. This region growing approach may combine segments that are not adjacent; see, for example, the mountains in Figure 1(b). This ability means the extracted features will be robust against occlusion; a common problem in natural images. The segmentation algorithm employs a multi-resolution scheme to attain significant speedup over the same approach without multi-resolution. See [6] for details about the multi-resolution implementation. This research takes advantage of the color gradient segmentation algorithm developed by HP in collaboration and our university partner RIT.

Stage 2 - Perceptual Tagging

The focus of the perceptual tagging work to date has been on the detection of skin, sky and snow, and foliage. The feature tagging algorithms that have been implemented to date are presented in the remainder of this section. They are presented in order of precedence. That is, if a segment is tagged as sky or snow, it will not be checked for skin tagging. Similarly, segments are tested for foliage only if they have not been tagged for either sky or skin.

Sky and Snow

Sky and Snow Detection uses several characteristics of the segment: size, location, color based probability of blue sky, color based probability of gray sky, color based probability of snow, texture (activity level), relative luminance, and correlation of color gradients (based on the Rayleigh scattering phenomenon). These features are combined to estimate the probability of blue sky/gray sky/snow for the segment. A segment will be tagged with the label that has the highest of these probability values, if this probability is above a threshold. Our algorithm, summarized in Figure 6, looks for such segments (step **1**). A more detailed discussion of the sky characteristics may be found in [5]. Step **2** of the algorithm analyzes all the probably sky and snow segments in the image to determine the sky color in the image.

Step **3** reassigns probabilities to each segment based on the observation that the sky is usually the most luminous object in the image. We use a histogram of luminance values in the image, to compute a regional measure of probability due to luminosity. However, this observation is invalid in images that contain snow. For example, the luminance of the sky in Figure 3(d) is lower than that of the snow region. This problem motivated our work on detection of snow. We have, in addition, added some logic to the treatment of luminance to deal with certain classes of images, in particular images that may contain snow. The algorithm examines a few common image scenarios, such as the blue sky and snow image of Figure 3(d).

As a final step (**4**), the algorithm uses the most probable sky region in the image to compute color and location models for the sky in the image. A sky probability is then re-computed for pixels in regions that were rejected because of their size or location in step 1 of the algorithm.

Figure 3 depicts several examples of sky and snow detection. Figure 3(a-c) demonstrate detection of blue sky. Figure 3(g-i) demonstrate detection of gray sky, and Figure 3(g-i) demonstrate detection of blue sky and snow.

We point out that snow is very difficult to distinguish from gray sky. We can separate snow from gray sky in images that contain very bright snow, such Figure 3(d). We have already stated that our motivation for adding snow detection came, in part, for improving the detection of blue sky. For this purpose, it suffices to detect very bright snow. In other images, that contain either very luminous gray sky or shaded snow, our tagging for gray sky and snow is interchangeable.

Face and Skin

Facial skin detection combines face detection, segmentation, and a global skin color model. Each of these modes of image analysis provides a different perspective and the combination enables the algorithm to accurately locate skin in the image. Our approach uses adaptive learning to estimate accurate skin models for the people in an image.

The skin tagging algorithm is detailed in Figure 7. It first detects faces in the image using the HP Labs Multi-View Face Detector [7]. The second step in the algorithm tags face segments and computes a body skin map. To do so it examines each face in order. For each face it refines the skin color model repeatedly. The first refinement uses the most probable skin pixels in a central oval of the face box returned by face detection. The most probable skin pixels, and the least probable skin pixels are also used to learn a weight on the color features using information gain. This feature weight indicates how useful each feature is as a discriminator for this face. The algorithm computes a new skin probability by applying the estimated color model and feature weights to the face box. Next, face segments are found by matching the segmentation to this map. The caption in Figure 7 describe the statistical heuristics used to match segments with maps. With these face segments, the algorithm should have a very accurate location and outline of the face that arises from the segmentation. It now refines its color model further by estimating a model from the pixels




Figure 3: Examples of blue sky, gray sky and snow tagging.

in the face segments. This final color model is applied to the entire image to find the body skin associated with each face. The final step of the algorithm tags body skin segments. Again, segments that match the skin map are tagged as skin.

Examples of skin tagging results are shown in every figure in this paper. Figure 1 shows a segmentation map in which the face is divided into several segments, and tagging was able to associate all of the segments with the face. In this case, the combination of segmentation and feature tagging is much stronger than that of segmentation alone. Figure 4 shows two examples of tagged segmentation of images with several people and non-face skin. Note that, although some faces are missed by face detection in each image, the skin map gives a high probability to all the faces, as well as non-face skin regions. For Figure 4(a) segmentation unites all the skin areas to one segment, and it matches very well with the detected skin map of Figure 4(d). For Figure 4(g) there are several skin segments that match human expectation quite well. The skin map (Figure 4(j)) is not accurate on the boy's hands. Nonetheless, the combination of segmentation and skin map correctly tags the skin segments. These three examples demonstrate the advantages of using different

kinds of image analysis information.

Foliage

Unlike sky and skin colors which have a relatively narrow range of color possibilities, the possibilities for foliage are quite broad. It includes the bright green of grass, dark green of trees, browns which tend to appear within foliage regions and black which appears in shaded foliage regions. A simple application of selecting specific color regions within an image often includes too much of the image. In this case, segmentation enables us to evaluate each such region separately. For example, in Figure 5, the grass is evaluated separately from the trees. Similarly, in Figure 1, due to the correct segmentation of the chain, it does not get confused with the foliage regions around it.

The foliage tagging algorithm is outlined in Figure 8. The algorithm evaluates, for each segment, three color models: a general foliage model, a forest (or shaded foliage) color model, and a ground color model. Examples of the responses of these color models are shown in Figure 5(d-f). The ground color model is instrumental for reducing false detections of brownish segments, such as ground and walls. In addition, the algorithm uses expected texture characteristics of foliage and forest. In general, foliage has high texture, but there are some exceptions. Grass may appear quite smooth when viewed from far away, and shaded forest are often so dark that there is little texture. The algorithm, as can be seen in Figure 8, employs several cases in making a decision. For example, if the color of the segment is very foliage-like, it will be tagged as foliage. If the color is a less perfect match, but the segment matches foliage texture as well, then the segment will still be tagged as foliage.

An additional heuristic is related to the chaotic nature of natural objects. We have observed that a region containing foliage has elements in a variety of directions, whereas man-made objects tend to have a specific direction. We implement a heuristic using local gradient directions. Figure 5(g-i) depicts histograms of several segments. Note the relatively uniform distribution of directions for the forest segment in Figure 5(g), compared with the very peaked distribution of the man-made object in Figure 5(i). Figure 5(h) shows a the distribution of directions in a grass segment, it is less uniform than the forest segment, but less peaked than the man-made segment. The heuristic we implement, termed $p_{natural}$ attempts to capture this intuition in order to enabling us to distinguish between foliage and man-made green regions.

Training Images and Object Statistics

One of the key aspect of the current work is to use an extensible approach to the problem of object tagging. That is, an approach that we can leverage to tag additional objects. The primary ingredient for object tagging is statistics. We use statistics of object both at the pixel scale and at the region scale. We are accumulating a development dataset of images, which have been automatically segmented and manually tagged by the authors. Currently this dataset includes about 200 images. This dataset has supported our development for sky, skin, snow and foliage detection. (Although, many of the original color models originated from other image sets [3]. Typically, when we develop a tagging algorithm for a new type of object, we add images to the set in order to improve the statistical distribution of that object. We also use this dataset to evaluate algorithm modifications by comparing the accuracy before and after the change.

Evidence that the Solution Works

Figures throughout the paper demonstrate tagged segmentation results where the colored tagged segments indicate skin, sky, snow, and foliage, and gray-level segments indicate unknown objects. Figures 1, 3 and 5 demonstrates detection of blue sky, gray sky and snow. Skin detection for portraits and group images is demonstrated in Figures 1 and 4. Foliage detection examples are shown in Figures 1, 3 and 5.

The results demonstrated in the image capture, for the most part, human perception of the image. But, there are a few mistakes, which demonstrate the types of tagging errors made by the algorithm. For example, in Figure 3(c) the wood of the peer is tagged as foliage, which is not a surprising mistake. In Figure 3(f) a part of the skier is tagged as foliage, which is a mistake that stems from inaccurate segmentation.

It is expected that a classification algorithm, such as our tagging algorithm, will have classification errors. The right way to assess such an algorithm is to test tagging performance on classified data. It turns out that classifying data for this problem is not entirely straightforward. For example, one can classify the data at the level of an image, a region of the image or per pixel. There has been extensive evaluation of methods for validating segmentation [11]. Like the issues encountered by this research, the segmentation algorithm has similar issues, and its performance been evaluated using the Normalized Probabilistic Rand (NPR) index in [6].

For the tagging algorithm presented here, we might consider classification per segment or per pixel. However, a per segment evaluation has two advantages. The first advantage arises because tagging will not be likely to tag correctly where segmentation errors have occurred. Tagging at the segment level enables us to ignore segments that do not make perceptual sense. The second advantage is a practical one. Using the segmentation for each image, we developed a manual tagging tool that enables a user to tag an image within about a minute. With this approach we have assembled an independent test set of 196 outdoor images that was tagged by an unbiased observer, i.e., not one of the authors.

On this data set, our algorithm has the following performance. While performance is evaluated per segment, the statistics are reported by pixel count. These statistics diminish the effect of small segments. Blue sky is tagged correctly 90% of the time. About 2% of blue sky tags were false alarms. Gray sky and snow are detected correctly 92% of the time, and 15% of gray sky and snow tags are false alarms. These statistics do not count mislabels between blue sky/gray sky/snow, e.g., blue sky tagged as gray sky is a correct label. These results are quite similar to results for our development data set. The higher false alarm rate for blue sky is better for the test data set. For the development dataset we had a 10% false detection rate.

For skin tagging, on our development dataset, our algorithm tags about 60% of skin pixels correctly, with a 7% false detection rate. The statistics are not nearly as good on the test dataset with 25% correct pixel tagging and a 56% false detection rate. The foliage tagging algorithm tags foliage correctly 80% of the time, and has a 10% false detection rate on the test dataset. These statistics are better than on the development where we had a false detection rate of 43%. It is interesting to note that 8.5% of these incorrect detections are tagged as water segments.

These experimental results lead to several observations and conclusions. The sky tagging algorithm has satisfactory results. It also compares well to other approaches and works well in practice for image enhancement in [1]. The discrepancy in skin and foliage tagging bear further investigation. This difference may be due to the image content, or the tagging of the observer as compared with the author's tagging. The confusion between foliage and water suggests another item for future work.

Competitive Approaches

There exist several areas of prior work related to object detection. Many prior projects have investigated the detection of specific objects, such as faces, sky or foliage. Often the scope of the problem was reduced, for example, forward looking for faces or blue sky. Very good results for detecting blue sky were reported in [10]. They use a variety of physical insights and achieve a 95% success rate for detection of blue sky regions, with a low false detection rate. They do not, however, attempt to detect gray sky. It is not clear how they treat blue sky with clouds. A large scale experiment on skin detections was presented in [9]. They use a

skin color model and non-skin color model and some size and shape considerations and report about 83% detection rate of images with people, with about 30% false detection. Some work on foliage detection arises in the remote sensing community and uses hyperspectral imaging [14], to obtain highly accurate results. We are not familiar with results on foliage detection from RGB images. Our object tagging algorithms address the broader problem of tagging objects in natural images. Another area of related work is image classification, which aims to provide a general tagging of the image scene. Common image classification works include indoor/outdoor [8], natural/man-made [15], portrait [13], etc. In contrast, we try to capture the location of the important perceptual objects very accurately.

Current Status

We presently have the capability of detecting a small set of perceptual objects: skin, sky, snow and foliage. The combination of image segmentation and feature extraction usually correlates with human perception of the content of the image much better than either segmentation or specific object detection alone. Some of the object detection algorithms are implemented in the HP Indigo Photo Image Enhancement (HIPIE)[1].

Next steps

We plan to develop detection algorithms for additional objects including hair, water and sand, for which we have some insight already. We believe the segmentation and object detection algorithms should be more tightly coupled, and intend to pursue this approach over the long term. The expected benefit of this technology to the consumer is demonstrated in the object-aware image enhancement and manipulation as shown in Figure 2.

1 References

- [1] Hp indigo photo image enhancement (hipie), 2008. was developed by HP-Labs Israel, under the coordination of Renato Keshet. Shlomo Harus and his group integrate the tool into the Indigo press.
- [2] T. Asano, D.Z. Chen, N. Katoh, and T. Tokuyama. Polynomial-time solutions to image segmentation. In *Proceedings of the 7th Annual SIAM-ACM Conference on Discrete Algorithms*, pages 104–113, January 1996.
- [3] O. Martinez Bailac. *Semantic retrieval of memory color content*. PhD thesis, Universitat Autònoma de Barcelona, 2004.
- [4] G. Balasubramanian, E. Saber, V. Misić, E. Peskin, M. Shaw, and R. Bhaskar. Unsupervised color image segmentation using a dynamic color gradient thresholding algorithm. Master’s thesis, Dept. Electrical Engineering, RIT, 2006.
- [5] R. Bergman, H. Nachlieli, and G. Ruckenstein. Perceptual segmentation: per-face skin map, eye location, body skin map and sky map. Technical Report HPL-2007-135, HP Laboratories, 2007.
- [6] L. Garcia, E. Saber, V. Amuso, M. Shaw, and R. Bhaskar. Automatic image segmentation by dynamic region growth and multiresolution merging. Master’s thesis.
- [7] D. Greig. Hp labs multi-view face detection. In *Proceedings of hPICS’06*, 2006.
- [8] G.H. Hu, J.J. Bu, and C. Chen. A novel bayesian framework for indoor-outdoor image classification. In *Proceedings of the Second International Conference on Machine Learning and Cybernetics*, November 2003.
- [9] M.J. Jones and J.M. Rehg. Statistical color models with application to skin detection. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition*, volume 1, pages 274–280, 1999.
- [10] J. Luo and S. P. Etz. A physical model-based approach to detecting sky in photographic images. *IEEE Transactions on Image Processing*, pages 201–212, 2002.
- [11] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proc. 8th Int’l Conf. Computer Vision*, volume 2, pages 416–423, July 2001.

- [12] N.R. Pal and S.K. Pal. A review on image segmentation techniques. *Pattern Recognition*, 26(9):1277–1294, September 1993.
- [13] J. Quyang. An image classification algorithm for automatic enhancement of both global contrast and local contrast in hp’s real life technologies, 2007.
- [14] S. W. Running, D. L. Peterson, M. A. Spanner, and K. B. Teuber. Remote sensing of coniferous forest leaf area. *Ecology*, 67(1):273–276, February 1986.
- [15] A. Torralba and A. Oliva. Statistics of natural image categories. *Network: Computation in Neural Systems*, 14:391–412, 2003.



Figure 4: Examples of skin detection and tagging for two images (a) and (g). Face detection results are shown in (b) and (h), respectively. (c) and (i) show the corresponding global skin maps. Face detection and the global skin map are used by the per-face learning algorithm to compute a per-face skin model. This model is applied to the rest of the image to find the improved skin maps in (d) and (j). Segmentation for each image is shown in (e) and (k). Finally, (f) and (l) depict the skin tagging for each segment.

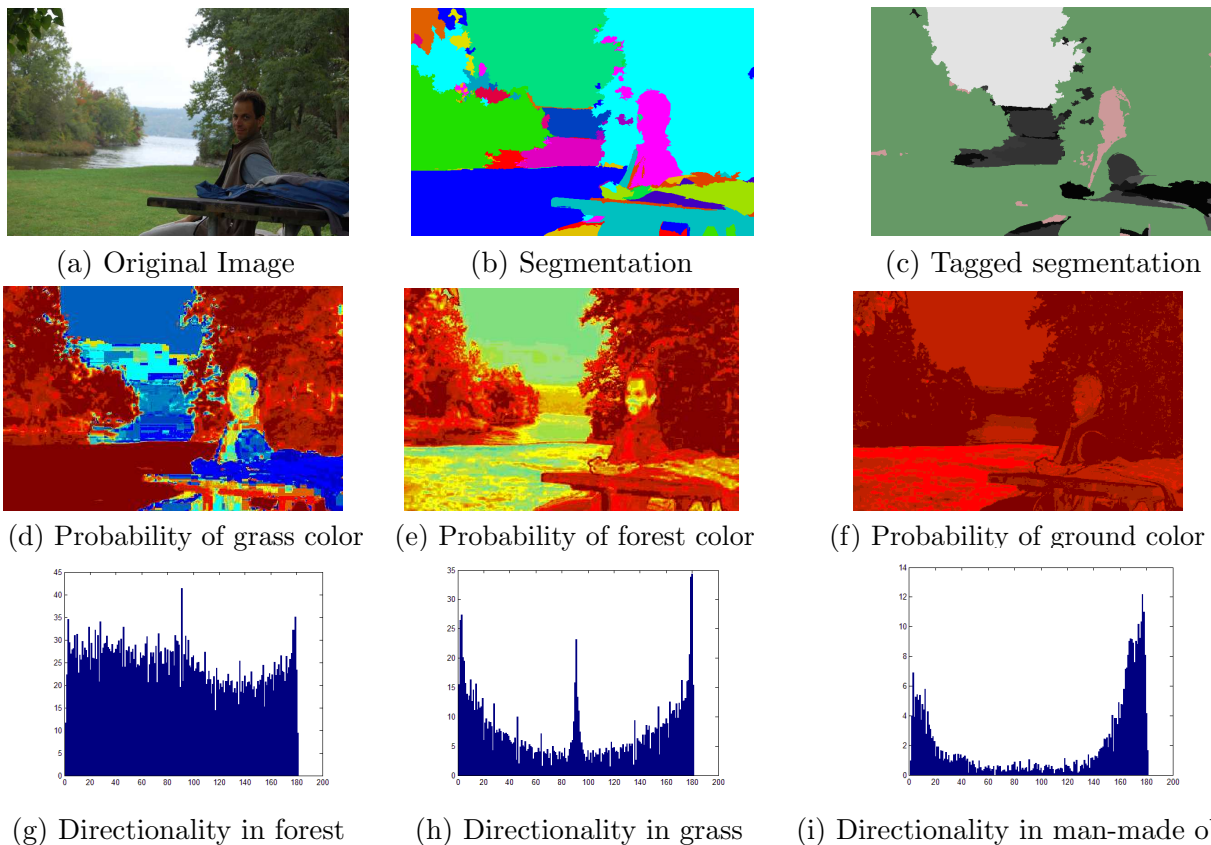


Figure 5: Example of foliage tagging. The original image is shown in (a), and the corresponding segmentation in (b). The tagging output by the algorithm is shown in (c). (d), (e) and (f) show the response to the three color models the algorithm uses, i.e., foliage, forest and ground, respectively. (g)–(i) illustrate the intuition captured by the natural heuristic used by the algorithm. (g) shows the histogram of local directions in a forest segment. Note the relatively uniform distribution. (i) is the histogram on the picnic table, which has a well-defined direction. (h) shows the directions in the grass segment.

SkyTag

1. Initial sky tagging
 - For each segment
 - Compute sky probability based region characteristics
 - Color: blue sky, gray sky, or snow
 - Size, Location, Texture, Relative luminance, and
 - Correlation of color gradients (Rayleigh scattering)
 2. Find probable sky segments
 - Tag each as *blue sky*, *gray sky* or *snow* according to the sky color characteristic
 3. Determine **sky color**
 - If there is a blue probable sky region set the **sky color** to blue
 - Otherwise if there is a gray probable sky region set the **sky color** to gray
 4. Reassign probabilities based on luminance
 - Consider several scenarios of images
 - Based on these scenarios a probability based on luminance is assigned to each segment
 4. Tagging from per-image sky map
 - Compute a color model from the most probable sky region
 - Apply this color model to the rest of the image
 - For each segment tagged as *unknown*
 - If the mean probability of sky according to this color model $> .6$
 - Tag it as the **sky color**
-

Figure 6: A sky detection algorithm

SkinTag

1. Detect faces
 2. Tag face skin segments and compute body map
 - Initialize the body skin map to be 0 everywhere
 - For each face
 - Compute Skin probability in face rectangle using the global skin color model
 - Label *skin* pixels - the 1/4 most probable skin pixels in a central area of the face rectangle
 - Label *non-skin* pixels - the 1/4 least probable skin pixels of the face rectangle
 - Compute features weights for LCH features based on information gain
 - Estimate a color model from the *skin* pixels and weights
 - Apply the color model to the face rectangle to get a face skin map
 - For each segment that matches the face skin map, tag it as **face**
 - Estimate a color model from the *face* pixels and weights
 - Apply the color model to the image
 - Set the body skin map at each pixel to the maximum of the current body skin map and the map just computed
 3. Tag body skin segments
 - For each segment that matches the face skin map, tag it as **skin**
-

Figure 7: A skin tagging algorithm. In the above algorithm, a segment matches a map if it passes the following test. Consider all the pixels in the segments with probability $\geq .1$. If the number of such pixels is more than 1/4 of the region size and the median skin probability of these pixels is $\geq .7$, then the segment matches the map.

FoliageTag

For each segment

 Compute segment statistics

$Mean_{Lum}, Median_{foliage}, Median_{forest}, Median_{ground}$

$Perc75_{foliage}, Perc75_{forest}, Perc75_{foliageTexture}, Median_{forestTexture}$

 Compute $p_{foliage} = Perc75_{foliage}Perc75_{foliageTexture}$

 Compute $p_{forest} = Perc75_{forest}Perc75_{forestTexture}$

 Compute $p_{natural}$ from the distribution of local directions in the segment

 If $Median_{foliage} < Median_{ground} - T_{ground} \wedge Median_{forest} < Median_{ground} - T_{ground}$

 tag segment as *unknown*

 Elseif $Median_{foliage} > T_{foliageDef}$

 tag segment as *foliage*

 Elseif $((p_{foliage} > T_{foliageHi} \wedge Lum < T_{foliageLum}) \vee (p_{Forest} > T_{foliageHi} \wedge Mean_{Lum} < T_{forestLum})) \wedge p_{natural} > T_{natural}$

 Set the segment probability of foliage to $\max(p_{foliage}, Median_{foliage}, p_{forest})$

 tag segment as *foliage*

 Else tag segment as *unknown*

Figure 8: A foliage detection algorithm. In the above, $Mean_{Lum}$ is short notation for the mean luminance of a segment. The notation $Median_{foliage}$ is short for the median response of a segment to the foliage color model. Likewise, $Perc75_{foliage}$ is the 75th percentile of the response to the foliage color model. The algorithm uses a number of thresholds, which are noted as $T_{foliageDef}$, for example. These threshold have been tuned by hand.