

Bidirectional Imaging and Modeling of Skin Texture

Oana G. Cula and Kristin J. Dana
Computer Science Department
Rutgers University
Piscataway, NJ, USA

Frank P. Murphy and Babar K. Rao
Dermatology Department
UMDNJ
New Brunswick, NJ, USA

Abstract—Human skin is a complex surface, with fine scale geometry and local optical properties that make its appearance difficult to model. Also, the conditions under which the skin surface is viewed and illuminated greatly affect its appearance. In this work, we capture the dependency of skin appearance on imaging parameters using bidirectional imaging. We construct a new skin texture database, containing bidirectional measurements of normal skin and of skin affected by various disorders. The complete database contains more than 3500 images, and is made publicly available for further research. Furthermore, we present two computational models for use in skin texture recognition. Both models are image-based representations of skin appearance that account for the varied appearance of the skin with changes in illumination and viewing direction. We employ these models in two contexts: discrimination between different skin disorders (e.g. psoriasis vs. acne), and classification of facial locations based on facial skin texture (e.g. forehead vs. chin). The classification experiments demonstrate the usefulness of the modeling and measurement methods.

I. INTRODUCTION

Modeling of human skin is an important task for both computer vision and graphics. For computer vision, accurate models of skin texture can greatly assist algorithms for human face recognition or facial feature tracking. In computer graphics, facial animation is an important problem which necessitates reliable skin texture models. In addition to computer vision and graphics, accurate skin models are useful in dermatology and several industrial fields. In dermatology, these skin models can be used to develop methods for computer-assisted diagnosis of skin disorders, while in the pharmaceutical industry, quantification is useful when applied to measuring healing progress.

However, human skin is a 3D textured surface, with local optical properties and fine scale geometry, that make its appearance difficult to model. The conditions under which the skin surface is imaged greatly affect its appearance, due to effects such as occlusions, foreshortening or shadowing. Recent work in texture recognition concerned with the dependency of surface appearance on the viewing and lighting conditions include [2][6][8][15][12][11][3][17][9][16][13][4][5]. The work in [7][8] has introduced the notion of bidirectional texture function (BTF) to describe image texture as a function of the viewing and illumination directions. While the BTFs of various real world surfaces have been of interest in the past, e.g. the CURET database [8], databases of skin texture BTFs



Fig. 1. Rows 1 and 2: examples of normal skin, corresponding to two subjects and to each of the four face locations (forehead, cheek, chin and nose) imaged for each subject. Notice that the skin features vary from one subject to another, and also from one region to another. Rows 3 and 4: examples from the clinical component of the database. Row 3 shows bidirectional images of psoriasis, located on the upper arm of a young female patient, while row 4 illustrates images of acute allergic contact dermatitis, located below the patient’s knee.

have not been available until now. In this work we present a novel skin texture database called Rutgers Skin Texture Database (URL http://www.rutgers.edu/rutgers_texture), which contains more than 3500 skin texture images. Each skin surface in the database is represented by a set of texture images, captured under different combinations of imaging parameters. The database has two components: (1) a normal skin component for recognition and rendering in computer vision and graphics; (2) a skin disorder component for quantitative imaging in dermatology.

In addition to the measurements and the bidirectional imaging methods employed to construct the database, we also present two texture modeling methods which we employ for texture recognition in two contexts: classification of skin disorders (e.g. psoriasis vs. acne), and classification of facial regions (e.g. forehead vs. chin). Both models are image based, and support recognition methods that have several desirable properties. Specifically, a single image can be used for fast non-iterative recognition, the illumination and viewing directions of the images need not be known, and no image




Fig. 2. Illustration of the face locations imaged during the construction of the normal skin component of the database. These facial texture images are employed during the experiments described in Section IV-B.

alignment is needed.

The rest of the paper is structured as follows: Section II presents details about the database, as well as the bidirectional imaging methods employed during construction of the database, Section III describes the two texture modeling methods we propose for use in skin texture classification, and Section IV describes the experimental results.

II. SKIN TEXTURE DATABASE

We present a skin texture database comprised of more than 3500 images, obtained under controlled imaging conditions. The database has two components.

The first component of the database is the normal skin component, with more than 2400 texture images of normal facial skin corresponding to 20 human faces, 4 locations on the face (forehead, cheek, chin and nose, as illustrated in Figure 2) and 32 combinations of imaging angles for each imaged surface (four camera poses, and eight light directions for each camera pose). The images in the database are acquired from both female and male subjects (7 females and 13 males), while the subjects age ranges from 24 to 53. The size of the imaged skin patch is about $3 \text{ cm} \times 3 \text{ cm}$. Rows 1 and 2 of Figure 1 show examples of texture images corresponding to each of the four imaged regions, for two different subjects. Notice the visual differences between texture images across various regions and from one subject to another. Skin features like pores, freckles, wrinkles or moles can differ in both size and shape, making skin modeling a difficult task. Also, the appearance of skin is influenced by factors like age and gender of the subject. This component of the database can be used as a testbed for a variety of applications of interest for computer vision and graphics, i.e. classification and synthesis methods.

The second component contains texture measurements of skin affected by various disorders. We present 75 clinical cases, which include conditions like acne, rosacea, psoriasis, sebhorreic keratosis, dermatomyositis, actinic keratosis, congenital melanocytic nevus, keratosis pilaris, and dermatofibroma. Each case may correspond to a different patient and a different body location. Depending on the location of the disorder, there are images from the face, arms, legs, back and abdomen. Each case has been measured with multiple texture images, and each image is characterized by a certain

combination of imaging conditions. More specifically, 28 cases are captured using imaging method 1, as described in Section II-A, and each case is represented by 10 images. The remaining 47 cases are imaged with method 2, also described in Section II-A, and each case is imaged under approximately 24 imaging conditions. The varying number of images per case is due to the difficult positioning of some lesions, affecting the visibility for some combinations of imaging conditions. The size of the imaged skin patch is dependent on the size of the lesion, however is not larger than $5 \text{ cm} \times 5 \text{ cm}$. As an example, consider row 3 of Figure 1, which illustrates a case of psoriasis. The images show that the bidirectional imaging captures the visual characteristics of this disorder, i.e. the inflammation, redness, and thickness of skin areas with silvery scale. Row 4 of Figure 1 shows bidirectional images of a case of acute allergic contact dermatitis. Again, the image set as opposed to a single image captures features of the disease, e.g. fine scale bumpiness, redness, and shininess.

The complete database is made publicly available for further research. While the normal skin component is meant to further stimulate skin texture research and algorithm development for computer vision and graphics, the clinical component will constitute an effective teaching tool for the medical community and an informative reference for patients.

A. Skin Imaging Methods

To obtain bidirectional measurements, there are four imaging parameters: the polar and the azimuthal angles of the viewing direction, and the polar and the azimuthal angles of the illumination direction. Bidirectional measurements can be quite difficult to obtain due to mechanical requirements of the setup. When the sample is non-planar, non-rigid and not-fixed, as is the case for human skin, measurements are even more cumbersome. Consequently, our methods for bidirectional imaging are developed with convenience of imaging human subjects in mind. The sampling of the full space of imaging angles is not dense, but the range of angles is sufficiently large that the observed images show complementary detail. For example, consider row 3 of Figure 1, which presents images of the same skin surface affected by psoriasis, images obtained with different imaging conditions. Notice how well the color variation appears when the light falls frontally to the surface (the first and third images from the left), and how the surface geometric detail becomes apparent when the light direction is more oblique (the second and fourth images from the left).

We present two techniques to achieve bidirectional imaging. Both methods use a Sony DFW-V500 IEEE-1394 digital camera equipped with InfiniMini video lens with variable focus from 750mm to 150mm, and video magnification 2.9-15X. Illumination intensity and color is monitored by imaging a white diffuse reflectance target providing an extremely flat spectral response.

Imaging method 1 uses a light arc and a camera mounted on a tripod. The various viewing directions are obtained by




Fig. 3. The illumination setup for imaging method 1. Various light poses are obtained by mounting the light source on a rotating arc. The light source is allowed to slide along the arc such that the polar angle of illumination could be varied, while the arc can be rotated around the vertical axis, providing variation of azimuthal angle of illumination. The setup allows scanning of horizontal surfaces.




Fig. 4. The illumination setup for imaging method 2. Various light poses are obtained by mounting the light source on a rotating arm. The rotation of the arm around the horizontal or vertical axis allows variation of the azimuthal angle of lighting. Also, the light source is allowed to slide onto the arm, providing variation of the polar angle of lighting. Moreover, the axis around which the arm is rotated, can be folded by 90° such that the light can span both vertical (left) and horizontal (right) surfaces.

mounting a digital camera on a manually articulated arm supported by a tripod. Positioning the light source on sample points of the hemisphere of all possible directions is obtained by mounting the lamp on a rotating arc, as illustrated in Figure 3. The light source is allowed to slide along the arc such that the polar angle of illumination could be varied, while the arc can be rotated around the vertical axis, providing variation of azimuthal angle of illumination. The illumination setup ensures constant distance between the surface of interest and the light source. The lamp is positioned distant enough from the surface of interest such that a quasi-collimation of light is ensured.

Imaging method 2 uses a modified light arc and a camera manipulator, utilized for the face texture database. These two methods are essentially the same because in both the camera and illumination are carefully controlled and a series of images is obtained. The main difference between the imaging methods is that our equipment for method 2 can also image a vertical surface and is more convenient for facial imaging, so that the subject does not have to lie horizontally under the imaging equipment. Figure 4 shows the light source of method 2 mounted on a rotating arm which approximates the light arc. The rotation of the arm around the horizontal or vertical




Fig. 5. (a) The fiducial marker used for both positioning the camera and for controlling the illumination direction. (b) The projections in the camera of the fiducial marker, corresponding to the four camera poses we seek for facial imaging.

axis allows variation of the azimuthal angle of lighting. Also, the light source is allowed to slide onto the arm, providing variation of the polar angle of lighting. Moreover, the axis around which the arm is rotated, can be folded by 90° such that the light can span both vertical and horizontal surfaces. For more convenient positioning of the viewing direction, the camera is mounted on an articulated arm boom stand allowing six degrees of freedom. Both the light setup and the camera manipulator are augmented with mobile carts, allowing the light and camera to be easily moved.

During both imaging methods, for each surface area, we take a sampling of camera and lighting positions at approximate directions and employ fiducial markers so that the camera and lighting position can be recovered more precisely if needed. Specifically, the fiducial marker we use is a circular pattern presenting a height marker, as illustrated in Figure 5 (a). The height marker has length equal with the disk radius. This pattern is attached to the surface of interest such that it is comprised in the field of view of the camera. Positioning the camera in one of the desired poses is achieved by inspecting the projection of the pattern in the image. As an example, consider the projections corresponding to the four camera poses we seek for facial imaging, illustrated in Figure 5 (b). The camera calibration is completed by imaging a known 3-dimensional object during each imaging session, allowing the computation of the intrinsic parameters of the camera.

III. SKIN TEXTURE MODELING

Many standard approaches in texture modeling characterize an image texture with a feature distribution. Bidirectional texture necessitates a model which should account for the appearance variation with the imaging conditions. Our modeling methods use a feature distribution which is a function of the imaging conditions, i.e. the model is a bidirectional feature histogram. We develop two models which both employ the bidirectional feature histogram, but differ in defining the image feature. While the feature in the first model is an image texton, obtained by clustering the output of oriented multiscale filters, the feature in the second model is a symbolic texture primitive, defined as a spatial configuration of indices of filters with maximal response at a pixel. In this section we give the details of these two methods.



Fig. 6. Illustration of the set of five dermatological disorders we classify during the experiments described in Section IV-A: acne, congenital melanocytic nevus (medium sized), keratosis pilaris, psoriasis, and acute allergic contact dermatitis.

A. Image Textons as Features

Due to various evidences of similar processing in human vision, a classical computational approach for encoding the local texture feature employs multichannel filtering [1][10][14][11]. In our work, we obtain a computational description of the local texture attribute by using a multiresolution filter bank consisting of oriented derivatives of Gaussians and center surround derivatives, on three scales. Each pixel of a texture image is represented by three multidimensional feature vectors obtained by grouping the corresponding filter responses over scale. To simplify our discussion, we will refer to a single scale, however the processing is done in parallel for all three scales. Our method differs from the method in [11] because the resulting feature vector does not encode the variation of local appearance as the viewing and illumination directions change. Instead, we account for the change in appearance with imaging conditions globally, by populating the feature space with feature vectors from sampled BTFs. We cluster the feature space to determine the representatives among the population, called the image textons. Specifically, we invoke k-means algorithm, which is based on the first order statistics of data, and finds a predefined number of centers in the data space. Empirical results suggest that the resulting set of representatives in the space spanned by the local structural feature vectors, namely the image texton library, is generic enough to represent a large set of texture samples. The histogram of image textons is used to encode the global distribution of the local structural attribute over the texture image. We represent the surface using a collection of image texton histograms, acquired as a function of viewing and illumination direction.

B. Symbolic Texture Primitives as Features

Though the image texton method works well when inter-class variability is large, there are several drawbacks to this approach. Clustering in high dimensional space is difficult and the results are highly dependent on the prechosen number of clusters. Furthermore, pixels which have very different filter responses are often part of the same perceptual texture primitive.

We seek a simplified representation that preserves the commonality of edges with the same orientation regardless of the strength of the filter response. The important entity is the label of the filter that has a large response relative to the other filters. If we use a filter bank F , consisting of N filters $\{F_1, \dots, F_N\}$, the index of the filter with the maximal response is retained

as the feature for each pixel. In this sense, the feature is symbolic. While a single symbolic feature is not particularly descriptive, we observe that the local configuration of these symbolic features is a simple and useful texture primitive. The dimensionality of the texture primitive depends on the number of pixels in the local configuration and can be kept quite low. No clustering is necessary as the texture primitive is directly defined by the spatial arrangement of symbolic features.

C. The Recognition Method

The recognition method consists of three main tasks: (1) creation of the library of representative features (either image textons or symbolic primitives), (2) training, and (3) classification. A subset of texture images is processed to determine the collection of representative features. Each training image provides a feature histogram and several training images obtained with different imaging parameters are collected for each texture class. The histograms from all training images for all texture classes are used to create an eigenspace, then the primitive histograms from a certain class are projected to points in this eigenspace, which represent a sampling of the manifold of points for the appearance of this texture class. In theory, the entire manifold would be obtained by histograms from the continuum of all possible viewing and illumination directions. For recognition, the feature histograms from novel texture images are projected into the eigenspace and compared with each point in the training set. The class of the nearest neighbor is the classification result.

IV. EXPERIMENTAL RESULTS

A. Classification of Skin Disorders

We employ the texture modeling approach based on image textons as texture features in the context of a simple classification experiment, where we attempt to discriminate between five instances of skin disorders: acne, congenital melanocytic nevus (medium sized), keratosis pilaris, psoriasis, and acute allergic contact dermatitis. Figure 6 shows images corresponding to each disorder considered for classification. Each class is characterized by a set of 24 bidirectional images captured from a certain patient. To compute the model, i.e. the bidirectional feature histogram, each image is converted to grayscale, and is manually segmented such that only the largest quasi-planar surface is considered during the experiments. We employ a filter bank consisting of 45 oriented derivatives of Gaussian and Laplacian of Gaussian filters. These filters are chosen to efficiently identify the local oriented features evident in skin structures.

Each image is processed to compute the corresponding image texton histogram, which encodes the statistical distribution over the image of conspicuous spatial features, such as oriented edges or spot-like features. As the image texton library we employ a collection of image textons obtained from images from the CURET database [7][8], which is a rich collection of 61 real world surfaces, each characterized




Fig. 7. (a) The set of five filters ($F_i, i=1\dots5$) used during the face skin modeling: four oriented Gaussian derivative filters, and one Laplacian of Gaussian derivative filter. (b) The set of five local configurations ($P_i, i=1\dots5$) used for constructing the symbolic primitives.

by a BTF sampled in more than 200 points. The reason for using this image textron library and not one derived from skin images, is to verify the assumption that our image textron library is generic enough to represent skin images, though it has been derived from images corresponding to different types of surfaces.

During training, the size of the train set of each class is varied from 12 to 20, therefore the size of the testing set varies from 12 to 4. The recognition rate as a function of the size of the training set is illustrated in Figure 8. Three of the classes (congenital melanocytic nevus - medium sized, keratosis pilaris, psoriasis) are 100% recognized, independent of the size of the training set, while the other two cases (acne and acute allergic contact dermatitis) are somewhat more difficult, attaining a recognition rate of 75%. By visually inspecting the images we observe that indeed the common characteristics of cases 1 and 5 are more subtle, and this is reflected by the reduced recognition rate. Globally, we obtain an excellent recognition rate of 90%, and this result shows that our modeling method successfully captures skin features. This result is even stronger when one considers that the textron library is computed based on non-skin images.

B. Classification of Facial Regions Based on Skin Texture

We construct the skin texture representation using symbolic texture primitives as described in Section III-B. We define the symbolic texture feature by employing a filter bank consisting of five filters: 4 oriented Gaussian derivative filters, and one Laplacian of Gaussian filter. These filters, illustrated in Figure 7 (a), are chosen to efficiently identify the local oriented patterns evident in facial skin structures, such as wrinkles or pores. Each filter has size 15×15 . We define several types of symbolic texture primitives by grouping symbolic features corresponding to nine neighboring pixels. Specifically, we define five types of local configurations, denoted by $P_i, i=1\dots5$, and illustrated in Figure 7 (b). As described in Section III-B, a symbolic feature corresponding to a pixel in the image is the index of the maximal response filter at that pixel. The pertinent question is how a featureless region in the image would be characterized in this context. To solve this issue we

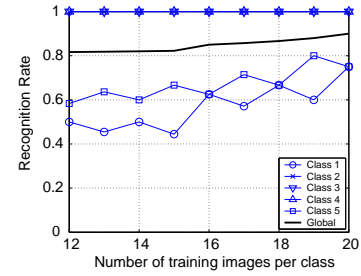


Fig. 8. Classification of skin disorders. Recognition rate as a function of the number of training images for each class.

define a synthetic filter, denoted by F_0 , which corresponds to pixels in the image where features are weak, that is, where the maximum filter response is not larger than a threshold. Therefore the symbolic texture primitive can be viewed as a string of nine features, where each feature can have values in the set $\{0, \dots, 5\}$. A comprehensive set of primitives can be quite extensive, therefore the dimensionality of the primitive histogram can be very large. To avoid this problem we prune the initial set of primitives to a subset consisting of primitives with high probability of occurrence in the image. Also, this reasoning is consistent with the repetitiveness of the local structure that is characteristic property of texture images.

We construct the set of representative symbolic primitives by using 384 images from 3 randomly selected subjects and for all four locations per subject. We first construct the set of all symbolic primitives from all 384 skin images, then we eliminate the ones with low probability. The resulting set of representative symbolic primitives is further employed for labeling the images, and to construct the primitive histogram for each image.

By examining the skin texture images in the database, we observe that there are local features of skin which vary with the location on the face. For example, the pores on the chin seem smaller than the ones on the forehead. Also, there are wrinkles on the forehead which are not present on the chin, or they appear different on the nose. However, we notice that there is a great similarity between the appearance of forehead and cheek, especially if they are imaged from the same subject. Consequently, we design a recognition experiment which discriminates between three locations on the face: forehead, chin and nose. We use skin texture images from all 20 subjects, therefore the total number of skin images employed for this experiment is 1920 (20 subjects, 3 locations per subject, 32 imaging conditions per location). Each face region to be classified is represented by a set of $20 \times 32 = 640$ texture images. We achieve classification of face regions by using individually each type of local configuration $P_i, i=1\dots5$, as well as using all $P_i, i=1\dots5$ combined. The training set is constructed with a subset of images from all subjects. The size of the training set is varied from 320 to 520 texture images for each face location. Classification is achieved with the remaining images, which are characterized by completely different imaging conditions than the ones for the training

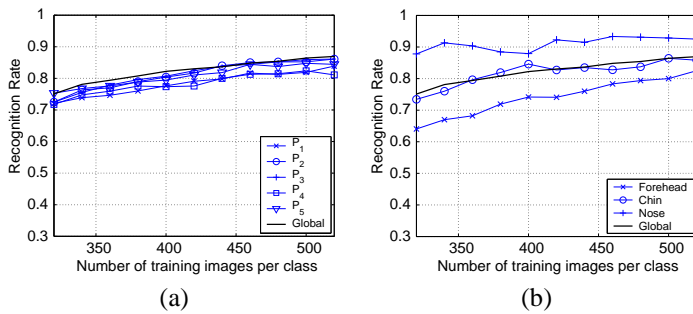


Fig. 9. Classification of face regions. (a) Recognition rates vs. the size of the training set plotted for six scenarios: using all five types of primitives for recognition (global), and using each configuration separately (P_i , $i = 1, \dots, 5$). (b) The global recognition rate as well as the recognition rates for individual classes (forehead, chin, nose) vs. the size of the training set.

images. The recognition rate as a function of the size of the training set is plotted in Figure 9. We achieve a global recognition rate of 87%. Notice from Figure 9 (a) that there is a slight difference in the performance of various types local configurations. However, the combination of P_i , $i=1\dots5$ has more descriptive power than a single configuration. The global recognition rate as well as the individual recognition rates for each region to be classified are shown in Figure 9 (b). Notice that while the classification performance for nose skin is quite good, the forehead skin is less recognizable. By visually inspecting the images we observe that indeed the common characteristics of forehead skin images are more subtle than those of nose skin images.

V. CONCLUSION

In this work we introduce a novel skin BTF database, called Rutgers Skin Texture Database¹, which contains more than 3500 bidirectional texture images of normal facial skin, as well as of skin affected by various disorders, like psoriasis, acne or keratosis pilaris. The images in the database are captured under controlled view/lighting conditions. Also, we describe the bidirectional imaging methods employed during the construction of the database. In addition, we present two image-based surface modeling methods, which we employ in two contexts: classification of skin disorders, and recognition of facial regions based on skin texture. The recognition rates are good, demonstrating the usefulness of our modeling techniques in capturing textural features of the skin surface.

The ability to recognize or classify skin textures has a wide range of useful multidisciplinary applications. For example, skin texture as a biometric can be used to assist face recognition in security applications. Biomedical evaluation of skin appearance with this method can provide quantitative measures to test topical skin treatments. Computer-assisted diagnosis in dermatology is another potential application.

Future work includes expanding the current texture model to support efficient synthesis. The difficulty in using image-based representations for synthesis is that even large measurements

are still a sparse sampling of the continuous six-dimensional bidirectional texture function. A combined approach of geometry or geometric transformations and image-based rendering is needed, although exactly how this combination is achieved is an open issue.

ACKNOWLEDGMENT

This material is based upon work supported by the National Science Foundation under Grant No. 0092491 and Grant No. 0076585.

REFERENCES

- [1] A. C. Bovik, M. Clark, and W. S. Geisler. Multichannel texture analysis using localized spatial filters. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12, January 1990.
- [2] M.J. Chantler. Why illuminant direction is fundamental to texture analysis. *IEE Proceedings Vision, Image and Signal Processing*, 142(4):199–206, 1995.
- [3] O. G. Cula and K. J. Dana. Compact representation of bidirectional texture functions. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1:1041–1047, December 2001.
- [4] O. G. Cula and K. J. Dana. Image-based skin analysis. *Proceedings of Texture 2002 - The 2nd international workshop on texture analysis and synthesis*, pages 35–40, June 2002.
- [5] O. G. Cula and K. J. Dana. 3D texture recognition using bidirectional feature histograms. *to appear in International Journal of Computer Vision*, 2003.
- [6] K. J. Dana and S. K. Nayar. Histogram model for 3D textures. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 618–624, June 1998.
- [7] K. J. Dana, B. van Ginneken, S. K. Nayar, and J. J. Koenderink. Reflectance and texture of real world surfaces. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 151–157, June 1997.
- [8] 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):1–34, January 1999.
- [9] J. Dong and M. Chantler. Capture and synthesis of 3D surface texture. *Proceedings of Texture 2002 - The 2nd international workshop on texture analysis and synthesis*, pages 41–46, June 2002.
- [10] A.K. Jain, S. Prabhakar, and L. Hong. A multichannel approach to fingerprint classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(4):348–369, 1999.
- [11] T. Leung and J. Malik. Representing and recognizing the visual appearance of materials using three-dimensional textons. *International Journal of Computer Vision*, 43(1):29–44, 2001.
- [12] G. McGunnigle and M. J. Chantler. Rough surface classification using first order statistics from photometric stereo. *Pattern Recognition Letters*, 21:593–604, 2000.
- [13] Sylvia C. Pont and Jan J. Koenderink. Bidirectional texture contrast function. *Proceedings of the European Conference on Computer Vision*, 2353(IV):pp. 808–822, 2002.
- [14] T. Randen and J. H. Husoy. Filtering for texture classification: a comparative study. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(4):291–310, 1999.
- [15] B. van Ginneken, J. J. Koenderink, and K. J. Dana. Texture histograms as a function of irradiation and viewing direction. *International Journal of Computer Vision*, 31(2-3):169–184, 1999.
- [16] M. Varma and A. Zisserman. Classifying images of materials. *Proceedings of the European Conference on Computer Vision*, pages 255–271, 2002.
- [17] A. Zalesny and L. van Gool. Multiview texture models. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1:pp. 615–622, December 2001.

¹http://www.caip.rutgers.edu/rutgers_texture