Workshop on Material Appearance Modeling (2015)
H. Rushmeier and R. Klein (Editors)

A Short Survey on Optical Material Recognition

M. Weinmann and R. Klein

Institute of Computer Science II, University of Bonn, Germany

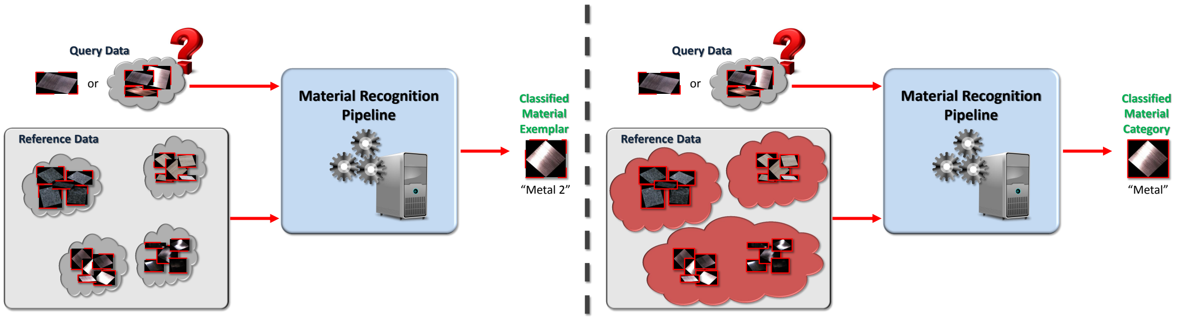

Figure 1: *Typical formulations of material recognition: Material instance recognition (left) is focused on finding the closest instance in the reference database, while material category recognition (right) aims at identifying corresponding semantic concepts. Depending on the application, the query data might consist of a single image or a set of images that show a certain, a-priori unknown material.*

Abstract

The complexity of visual material appearance as observed in the huge variation in material appearance under different viewing and illumination conditions makes material recognition a highly challenging task. In the scope of this paper, we discuss the facts that make material appearance that complex and provide a survey on technical achievements towards a reliable material recognition that have been presented in the literature so far. In addition, we discuss still open challenges that might be in the focus of future research.

Categories and Subject Descriptors (according to ACM CCS): I.2.10 [Artificial Intelligence]: Vision and Scene Understanding—Texture I.4.8 [Image Processing and Computer Vision]: Scene Analysis—Object recognition I.5.2 [Pattern Recognition]: Design Methodology—Pattern analysis

1. Introduction

Our interactions with the content of the surrounding environment in daily life are primarily guided by the rich information perceived via the human visual system. We not only perceive the presence of objects and their spatial arrangement in the scene, but are also able to infer their individual shapes and materials which e.g. tell us where and how careful to grasp a particular object. Even more, the perceived materials provide valuable information w.r.t. properties such as fragility, deformability and weight. These aspects are also important for industrial applications where objects should be handled automatically in an appropriate way as many tasks might have to be carried out depending on the material properties. Therefore, these applications rely on the availability of techniques that allow a reliable material recognition of individual material exemplars or semantic material classes in single query images or queries with several images (see Figure 1). In addition, automatic material retrieval techniques allow designers to find either a certain material or a similar material in the databases from suppliers.

DOI: 10.2312/mam.20151202

Unfortunately, material recognition is a rather complex problem due to the strong dependency of material appearance on surface geometry as well as viewing and illumination conditions. In the scope of this paper, we provide an analysis of the complexity of visual material appearance which is followed by a discussion of the key aspects that have to be considered by material recognition techniques and a survey on the advances in the related research domain. In particular, this includes surveys on:

- characteristic material attributes and how suitable descriptors can be derived,
- characteristic material "fingerprints", i.e. region-based material representations, as well as "material spaces", that can be derived from information w.r.t. material appearance under different view-light conditions, and
- recognition schemes used in the literature.

In addition, we review the improvements in establishing material databases. Finally, we discuss still remaining challenges in the context of material recognition.

2. The Complexity of Visual Material Appearance

A closer look at the materials of objects given in daily life scenarios immediately reveals the complexity of visual material appearance. These materials typically exhibit significant variations in properties such as color, texture, glossiness, specularities, translucency, transparency or surface profiles that determine their appearance. Furthermore, changes in the illumination conditions typically also lead to a changing material appearance. The observed colors and textures are a result of the complex interplay of surface material properties, surface geometry and illumination conditions which determines the visual complexity of surface appearance. Therefore, both the human visual system and acquisition devices are only capable of observing material appearance depending on the coupling of these three modalities.

In this context, taking into account the scale-dependency of material appearance is inevitable. The structures on the *microscopic scale*, i.e. the scale of atoms and molecules, cannot directly be observed by the human visual system and yet they significantly contribute to material appearance. In particular, the appearance of materials such as metals, paper, plastics, etc. is determined on this scale. Furthermore, material appearance is also characterized by effects of light exchange happening on a *mesoscopic scale* at fine details in surface geometry such as scratches, engravings, weave-patterns of textiles or embossing of leathers. Such surface structures cause effects like interreflections or self-shadowing. While the effects on these aforementioned scales obviously represent the material characteristics and determine the material appearance, the 3D geometry of the object with the respective, considered material also influences the material appearance significantly. Considering this *macroscopic scale*, regular structures as e.g. given in woven cloth, brushed metal or surface textures of certain objects might appear distorted in the image because of the dependency on the object geometry.

Unfortunately, the consideration of these scales suffices only for a close distance between the surface material and the human observer. For an increasing distance, the effects of light exchange at fine surface details such as scratches, engravings, weave-patterns or embossing will become less visible and finally not be perceivable as mesostructures anymore. Hence, they might be treated as irregularities in a different kind of microscopic scale. In a similar way, some of the details in the surface geometry might not be perceived as macroscopic features anymore but rather as features on a novel mesoscopic scale.

To give a further example, shininess of specular objects or translucency might also depend on the distance between object and observer. When considering a highly specular surface with a rough surface profile from a close range, the resolution of the human visual system is sufficient to perceive the many surface patches with different surface normals, and the material will appear specular. With an increasing distance to the surface, the resolution of the visual system will become insufficient to perceive the appearance of all the individual surface patches with different orientations separately and, instead, perceive a superposition of the appearances of several of these patches. This will lead to a transition from specular to diffuse appearance perception. In contrast, for flat, highly specular surfaces, the surface will still appear highly specular with an increasing distance. In a similar way, the appearance of translucent objects with a rough surface profile is characterized by subsurface scattering effects when viewed from a close range. For an increasing distance, such objects might be perceived as opaque, if only the superposition of the appearances of the individual patches with the subsurface scattering effects is perceived by the visual system.

This clearly indicates that the definition of scale is of dynamic nature. Depending on the distance between the observer and the object of interest, the definition of microscopic scale, mesoscopic scale or macroscopic scale might have to be adapted. Therefore, material appearance involves a multitude of scales $\ldots \subset D_{i-1} \subset D_i \subset D_{i+1} \subset \ldots$ ranging from an atomic scale to the intergalactic scale [Kaj85, MMS*04].

3. Material Recognition Schemes

In order to allow an automatic image-based material recognition, the following challenging tasks have to be investigated. After segmenting images into regions for the different occurring materials, discriminative descriptors that reflect characteristic material traits have to be extracted. Based on such descriptors, efficient and appropriate models for the individual materials or material categories can be computed per region that represent characteristic material "fingerprints".

Typically, those compact representations are obtained based on assigning the individual descriptors to their closest match in a dictionary of textons that is derived by a clustering of the descriptors extracted from the training data. Then, e.g. either the weighted occurrences of the textons assigned to the individual descriptors per region or the offsets of the descriptors to their closest textons might be stored in the representations. Based on the availability of several fingerprints extracted under different view-light conditions, the more general concept of "material spaces" can be established in order to represent material exemplars or semantic material classes. Finally, material recognition can be performed based on these representations using adequate training data. Figure 3 provides an overview on such a typical recognition pipeline.

3.1. Material Attribute Descriptors

Characteristic material traits such as shininess, roughness or homogeneity are manifested in characteristic local visual features with certain statistics of colors or textural patterns. Local color distributions can be described by using densely sampled color patches [VZ03, VZ09, LSAR10, SLRA13, WGK14, WK15]. Local texture characteristics are typically captured by considering local gradient information of the image intensities. This can be performed by densely sampled SIFT descriptors [LSAR10, SLRA13, WGK14, WK15], densely sampled Histogram of Oriented Gradients (HOG) descriptors [WK15], Local Binary Patterns (LBPs) [CHM05, LF12], kernel descriptors [HBR11] or filterbanks [LM99, LM01, VZ02, VZ04, CD04, CHM05, CHFE10, WK15] but other descriptor types such as basic image features [CG08, CG10] or sorted-random-projection descriptors [LFCK12] might also be used. In addition, several approaches use different combinations of these descriptor types to improve the accuracy of their classification frameworks [BG06, LSAR10, HBR11, LF12, SLRA13, WGK14, WK15]. Furthermore, learning features that are capable of capturing the characteristics of the individual categories has been investigated in [SN13, LF14, CMK*14, BUSB14] and attribute-based descriptions, that consider attributes such as e.g. bumpy, checkered, dotted, fibrous, knitted, porous, smeared, sprinkled, stained, striped, woven, or zigzagged, have been used in [CMK*14].

3.2. From Local Structures to Material "Fingerprints" and Material Spaces

After extracting such descriptors for the images contained in the training data and the query data, the next step typically consists in the computation of compact representations for the individual image regions covered by a certain material. For this purpose, the descriptors extracted from the training data are typically used to calculate a dictionary of representative descriptors denoted as textons. This allows assigning all of the extracted descriptors within a certain image region to the respective visual words in the dictionary to get a texton-based representation for image regions covered by a certain material as introduced in e.g. [LM99] and [LM01] and also followed in [VZ02, VZ04, VZ09, LSAR10, LF12, SLRA13, WGK14, WK15]. The computation of histogram-like fingerprints as well as the more sophisticated VLAD representation [JDSP10] are illustrated in Figure 2. Furthermore, instead of a hard quantization of the individual descriptors to their closest texton, it is also possible to use a soft quantization that introduces weights according to the distances of the descriptors to their closest texton [vGVSG10]. This results in soft-histograms or Fisher vectors [PD07]. Optionally, a dimensionality reduction technique such as PCA can be applied to compress the region-based representations. Based on such "fingerprints" of a certain material observed under multiple different view-light-conditions, a characteristic material space can be derived. Furthermore, the region representations of several different material samples belonging to the same semantic concept can be used to define representations for semantic categories.

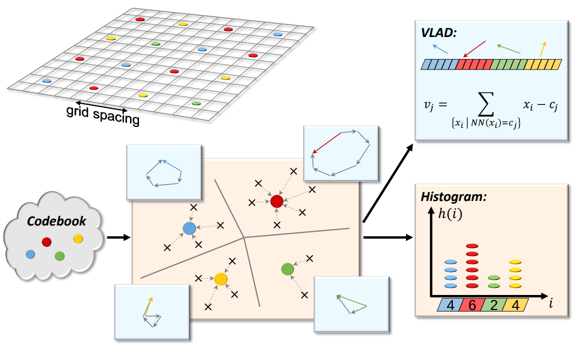

Figure 2: *Illustration of two widely used image region representations. Based on densely extracted descriptors a dictionary might be calculated. This allows to represent the descriptors extracted for a certain image region to be quantized in the form of histograms, where the occurrences are counted, or VLADs [JDSP10], where the offset vectors to the dictionary entries are stored.*

3.3. Recognition Schemes

The resulting texton-based image region representations can then be classified using nearest neighbor classifiers, Bayesian frameworks [VZ04, LSAR10], Markov random fields (MRFs) [VZ03], support vector machines (SVMs) [HCFE04, CHM05, LF12, LYG13, WGK14], random forests [Bre01], etc..

While most investigations focused on single-image-based material classification, some acquisition devices also offer the possibility to easily acquire several images under several view-light configurations, which might significantly facilitate material classification. In [LM99] and [LM01], histograms have been concatenated to form a single vector

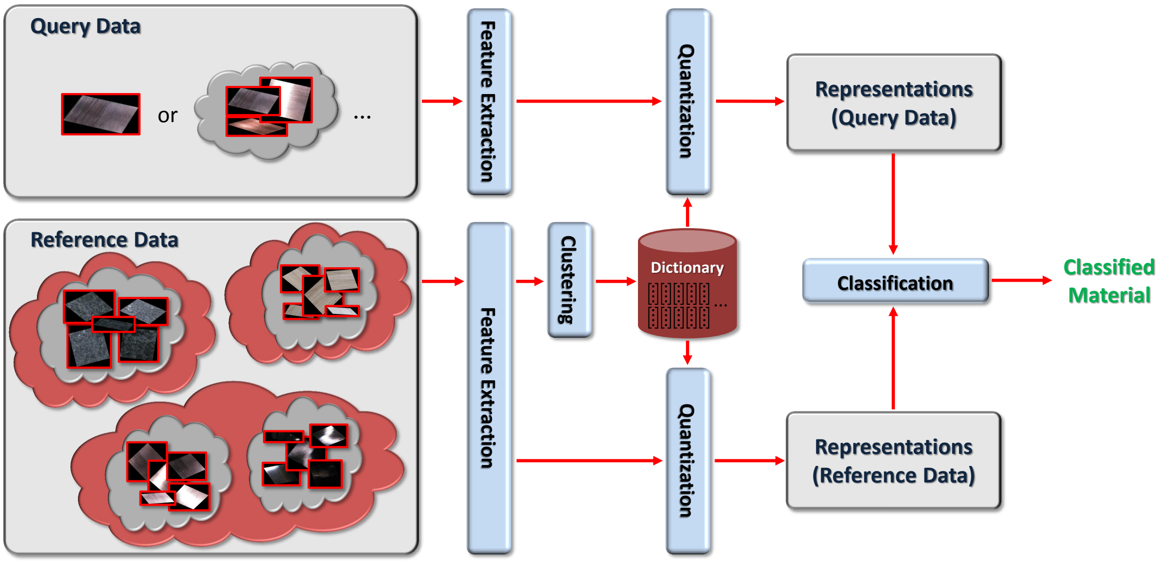

Figure 3: *Typical material recognition scheme: Descriptors are extracted for both the training and the query data. The descriptors extracted from the reference data can be used to compute a dictionary that allows the quantization of the descriptors into a compressed representation. Material recognition is typically carried out based on these compressed representations. The reference data consists of several image collections that capture the appearance variations of either individual material exemplars (gray) or semantic material categories (red).*

for each particular material, which imposes that materials are represented by a fixed ordering of the configurations within the combined vector where all the individual image representations have to be carefully registered. Comparing materials based on these vector-based representations hence requires that exactly the same view-light configurations are considered in each vector with the same fixed ordering. Furthermore, material classification based on BRDF slices has been proposed in [WGSD09]. Bidirectional feature histogram manifolds as introduced in [CD04] overcome the need for considering exactly the same view-light configurations for all materials but still rely on a densely sampled set of view-light configurations. The method described in [WK15] aims at classifying material instances using only a few images by representing materials based on convex hull models or affine hull models similar to [CT10]. Recognition can then be performed based on the distances between the convex hulls or affine hulls of the individual materials. Their method yields significantly better recognition rates than previous methods while using smaller numbers of view-light configurations.

Alternative approaches include learning optimal illumination for material classification [JSJ10], material classification based on learning coded illumination to directly measure discriminative features such as projections of spectral BRDFs [GL12, LG14] or learning discriminative illumination patterns and texture filters to directly measure optimal projections of BTFs [LYG13].

4. Databases for Material Recognition

Training an appropriate classifier requires having adequate training datasets which representatively cover the appearance variations of a huge multitude of different material exemplars under a large variety of different viewing conditions, illumination conditions and surface geometries expected to be encountered in the query data. If material instances are to be recognized, a class is defined by images depicting the appearance of the respective sample under the aforementioned varying conditions. In contrast, when the objective consists in recognizing semantic categories instead of single exemplars, data of several exemplars that adequately define the intra-class variations of the respective semantic class has to be taken into account. Consequently, defining a single category adequately might easily require several thousands of images.

The CUReT database [DvGNK96] is probably the first database with a large set of 61 material samples that have been systematically acquired under 205 different view-light configurations. In the scope of the KTH-TIPS database [HCFE04], this database has been extended by also adding scale information for material appearance by varying the distance of the acquired samples to the camera. Furthermore, the ALOT database [BG09] offers significantly more and also a wider range of different material types, which have additionally been observed under illumination with different colors. However, only a few images have been taken per material sample. In addition, in all of these databases the individual material categories are defined based on appearance variations per material exemplar. Aiming for a generalization to classifying object categories, the KTH-TIPS database has been further extended by adding measurements of different samples of the same material category and also considering ambient lighting in the KTH-TIPS2 database [CHM05]. However, taking only four samples per category still limits the representation of the intra-class variance of materials observed in real-world scenarios. More recently, a spectral material database has been presented in [LYG13] for multi-spectral material recognition. However, the samples are imaged from only one single viewpoint. A common limitation of all these databases is the rather limited number of measurements, which are furthermore acquired in a lab environment. Hence, the influence of the complexity of real-world environment conditions is not taken into account, and, therefore, material recognition under natural illumination cannot be performed based on such training data.

Other databases are designed to capture the large intra-class variation in the appearance of materials in complex real-world scenarios. The Flickr Material Database (FMD) [SRA09] contains images that have been downloaded from Flickr.com and show different associated material samples under uncontrolled viewing and illumination conditions and compositions. Even larger collections are given by the OpenSurfaces dataset [BUSB13] or the Materials in Context Database (MINC) [BUSB14]. However, annotations and segmentations of the images of these collections require plenty of work in a time-consuming process and are typically obtained by costly crowdsourcing services such as Amazon Mechanical Turk (AMT) [BUSB13, BUSB14, CMK*14]. In addition, while manual segmentations are available, these masks are not always accurate, leading to the inclusion of background appearance and problematic artifacts for material classification. Obviously, the significantly more complex variations of material appearance encountered under natural illumination make material classification much more challenging and only recognition rates far below the ones obtained for databases acquired under controlled lab conditions have been reached so far [LSAR10, SLRA13]. The main reason for this is that it is more complex to include the possibly encountered variations on material appearance in the training data than for material classification under controlled illumination data, where a smaller subset of training data might already be sufficient. Furthermore, a different approach has been presented with the Describable Textures Dataset (DTD) [CMK*14]. While the aforementioned databases establish classes for different material instances or more general semantic material categories, this database considers semantic material attributes as classes. This allows to represent materials in terms of how well they match the individual attributes.

The required manual processes for capturing exemplars as well as for segmenting and annotating materials in images severely limit the number of images per material category in all of the above-mentioned databases. As an alternative, the potential of computer graphics has been investigated to introduce a new promising trend of using synthesized training data. In seminal work, the virtual MPI-VIPS database has been introduced to approach material recognition based on synthetic data [LF12]. This database contains images of virtual materials that are synthesized under view-light configurations similar to the ones given in the KTH-TIPS2 database. The renderings are created based on rather simple BRDF shaders and the local mesostructure of the material surface is simulated via bump maps to improve the shading effects. Unfortunately, the use of such approximate material models results in a less realistic depiction of several materials as the complexity of the reflectance characteristics of many materials has not been adequately considered. In particular, mesoscopic effects that contribute to the appearance materials such as textiles, bread or cork are not modeled. As demonstrated in [WGK14], the approach for synthesizing virtual samples matters. The UBO2014 database (see Figure 4 and Figure 5) presented in the scope of this investigation is based on BTFs to also model mesoscopic effects. In addition, the intra-class variance of semantic material categories is covered in a better way and significantly more viewing and lighting configurations are included than in any of the other systematically acquired databases. These dense measurements are required for the realistic depiction of many materials with their characteristic traits in a virtual scene via BTFs to preserve the mesoscopic effects in the synthesized data. Furthermore, using synthesized training data automatically provides annotations and segmentations per image and, hence, overcomes the need for time-consuming and costly annotations and segmentations as performed for the FMD [SRA09], the OpenSurfaces dataset [BUSB13] or the MINC database [BUSB14].

5. Open Challenges

Despite the remarkable progress that has already been achieved in the domain of material recognition several aspects still require further investigations.

Further effort might be spent on the investigation of better material-specific descriptors. In this context, it might also be worth to analyze the subjective perception criteria such as warm/cold, rough/soft, etc. in addition to semantic, visual features such as glossiness or roughness as these attributes guide the material selection process of designers as well as the material editing. Therefore, even more effort has to be spent on the development of attribute-based datasets such as the Describable Texture Dataset [CMK*14] which consider

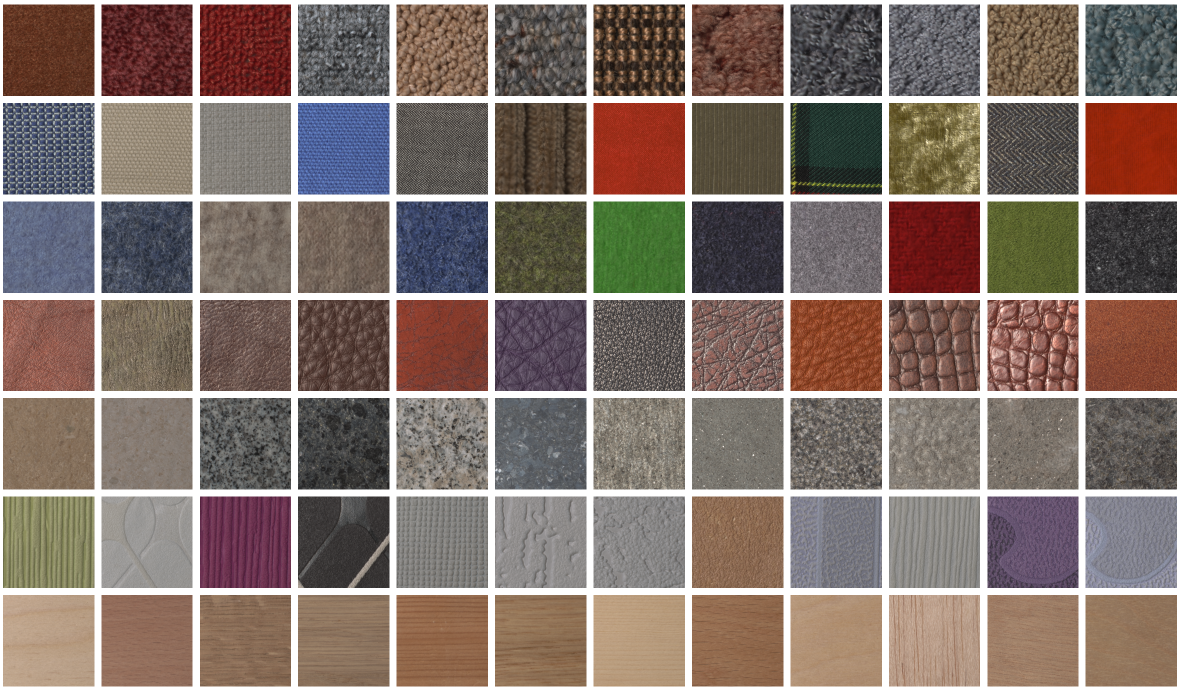

Figure 4: *Material samples in the UBO2014 database [WGK14].*

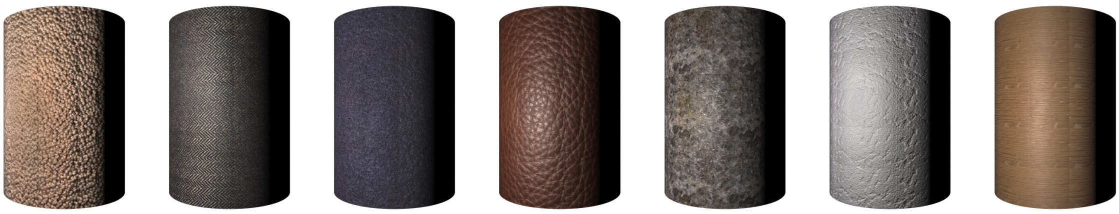

Figure 5: *Renderings of cylinders with some of the virtual materials in the UBO2014 database [WGK14].*

the variations in the appearance of the attributes in a better way. Considering the material reflectance in different spectra such as the near infrared [SFS09] or the ultra-violet spectral range as well as illuminating material samples with different wavelengths [LYG13] might additionally contribute to a more robust material recognition, as materials might be distinguished more easily than in the RGB channels.

Furthermore, even larger material databases are required to recognize the typical materials we encounter in daily life. This might involve either a costly manual acquisition, segmentation and annotation of huge masses of data via crowd-sourcing or creating larger datasets via synthesis following approaches such as in [WGK14]. The latter approach might become more practical with the advances towards more efficient automatic acquisition of material samples. Having huge masses of data, i.e. many images, there is an additional need for efficient large-scale learning techniques that can train per-class models based on a high number of images in a reasonable time. The aforementioned approach of synthesizing data might also be an important practical step towards learning which of the view-light-configurations are most informative regarding material recognition.

A further challenge is the automatic segmentation of materials within images which reduces the involved manual work significantly. In this context, more investigations towards color segmentation strategies with robustness w.r.t. shadows, highlights, and textures such as the one in [VB-vdWV11] will have to be carried out.

Another important objective for future research can be identified in the development of suitable material metrics to efficiently assess similarity or dissimilarity of materials based on distinctive material characteristics. This might allow a more practical material retrieval and material recognition. While similar efforts have long reached maturity in color science for comparing colors, the massive increase in physical degrees of freedom imposes significant challenges for the generalization of color metrics to general material appearance.

6. Acknowledgements

The research leading to these results was partially funded by the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 323567 (Harvest4D); 2013-2016.

References

[BG06] BURGHOUTS G. J., GEUSEBROEK J.-M.: Color textons for texture recognition. In *Proceedings of the British Machine Vision Conference (BMVC)* (2006), pp. 1099–1108. 3

[BG09] BURGHOUTS G. J., GEUSEBROEK J.-M.: Material-specific adaptation of color invariant features. *Pattern Recognition Letters 30*, 3 (2009), 306–313. 5

[Bre01] BREIMAN L.: Random forests. *Machine Learning 45*, 1 (2001), 5–32. 3

[BUSB13] BELL S., UPCHURCH P., SNAVELY N., BALA K.: Opensurfaces: A richly annotated catalog of surface appearance. *ACM Transactions on Graphics 32*, 4 (2013), 111:1–111:17. 5

[BUSB14] BELL S., UPCHURCH P., SNAVELY N., BALA K.: Material recognition in the wild with the materials in context database. *CoRR abs/1412.0623* (2014). URL: http://arxiv.org/abs/1412.0623. 3, 5

[CD04] CULA O. G., DANA K. J.: 3d texture recognition using bidirectional feature histograms. *International Journal of Computer Vision (IJCV) 59*, 1 (2004), 33–60. 3, 4

[CG08] CROSIER M., GRIFFIN L. D.: Texture classification with a dictionary of basic image features. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2008), pp. 1–7. 3

[CG10] CROSIER M., GRIFFIN L. D.: Using basic image features for texture classification. *International Journal of Computer Vision (IJCV) 88*, 3 (2010), 447–460. 3

[CHFE10] CAPUTO B., HAYMAN E., FRITZ M., EKLUNDH J.-O.: Classifying materials in the real world. *Image and Vision Computing 28*, 1 (2010), 150–163. 3

[CHM05] CAPUTO B., HAYMAN E., MALLIKARJUNA P.: Class-specific material categorisation. In *Proceedings of the International Conference on Computer Vision (ICCV)* (2005), vol. 2, pp. 1597–1604. 3, 5

[CMK*14] CIMPOI M., MAJI S., KOKKINOS I., MOHAMED S., VEDALDI A.: Describing textures in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2014), pp. 3606–3613. 3, 5

[CT10] CEVIKALP H., TRIGGS B.: Face recognition based on image sets. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2010), pp. 2567–2573. 4

[DvGNK96] DANA K. J., VAN GINNEKEN B., NAYAR S. K., KOENDERINK J. J.: *Reflectance and Texture of Real World Surfaces*. Tech. rep., Department of Computer Science, Columbia University, Columbia, USA, 1996. Columbia University Technical Report CUCS-048-96. 4

[GL12] GU J., LIU C.: Discriminative illumination: Per-pixel classification of raw materials based on optimal projections of spectral brdf. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2012), pp. 797–804. 4

[HBR11] HU D., BO L., REN X.: Toward robust material recognition for everyday objects. In *Proceedings of the British Machine Vision Conference (BMVC)* (2011), pp. 1–11. 3

[HCFE04] HAYMAN E., CAPUTO B., FRITZ M., EKLUNDH J.-O.: On the significance of real-world conditions for material classification. In *Proceedings of the European Conference on Computer Vision (ECCV)* (2004), pp. 253–266. 3, 4

[JDSP10] JEGOU H., DOUZE M., SCHMID C., PÉREZ P.: Aggregating local descriptors into a compact image representation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2010), pp. 3304–3311. 3

[JSJ10] JEHLE M., SOMMER C., JÄHNE B.: Learning of optimal illumination for material classification. In *Pattern Recognition* (2010), vol. 6376 of *Lecture Notes in Computer Science*, pp. 563–572. 4

[Kaj85] KAJIYA J. T.: Anisotropic reflection models. In *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)* (1985), pp. 15–21. 2

[LF12] LI W., FRITZ M.: Recognizing materials from virtual examples. In *Proceedings of the European Conference on Computer Vision (ECCV)* (2012), vol. 4, pp. 345–358. 3, 5

[LF14] LI W., FRITZ M.: Learning multi-scale representations for material classification. *[arXiv]* (2014). URL: http://arxiv.org/abs/1408.2938http://arxiv.org/pdf/1408.2938v1.pdf. 3

[LFCK12] LIU L., FIEGUTH P., CLAUSI D., KUANG G.: Sorted random projections for robust rotation-invariant texture classification. *Pattern Recognition 45*, 6 (2012), 2405–2418. 3

[LG14] LIU C., GU J.: Discriminative illumination: Per-pixel classification of raw materials based on optimal projections of spectral BRDF. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 36*, 1 (2014), 86–98. 4

[LM99] LEUNG T., MALIK J.: Recognizing surfaces using three-dimensional textons. In *Proceedings of the International Conference on Computer Vision (ICCV)* (1999), vol. 2, pp. 1010–1017. 3

[LM01] LEUNG T., MALIK J.: Representing and recognizing the visual appearance of materials using three-dimensional textons. *International Journal of Computer Vision (IJCV) 43*, 1 (2001), 29–44. 3

[LSAR10] LIU C., SHARAN L., ADELSON E. H., ROSENHOLTZ R.: Exploring features in a bayesian framework for material recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2010), pp. 239–246. 3, 5

[LYG13] LIU C., YANG G., GU J.: Learning discriminative illumination and filters for raw material classification with optimal projections of bidirectional texture functions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2013), pp. 1430–1437. 3, 4, 5, 6

[MMS*04] MÜLLER G., MESETH J., SATTLER M., SARLETTE R., KLEIN R.: Acquisition, synthesis and rendering of bidirectional texture functions. In *Eurographics 2004 State of the Art Reports* (2004), pp. 69–94. 2

[PD07] PERRONNIN F., DANCE C. R.: Fisher kernels on visual vocabularies for image categorization. In *Proceedings of the*

IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2007), pp. 1–8. 3

[SFS09] SALAMATI N., FREDEMBACH C., SÜSSTRUNK S.: Material classification using color and NIR images. In *Proceedings of the Color and Imaging Conference (CIC)* (2009), pp. 216–222. 6

[SLRA13] SHARAN L., LIU C., ROSENHOLTZ R., ADELSON E. H.: Recognizing materials using perceptually inspired features. *International Journal of Computer Vision (IJCV) 103*, 3 (2013), 348–371. 3, 5

[SN13] SCHWARTZ G., NISHINO K.: Visual material traits: Recognizing per-pixel material context. In *Proceedings of the Color and Photometry in Computer Vision Workshop (Workshop held in conjunction with ICCV 2013)* (2013), pp. 883–890. 3

[SRA09] SHARAN L., ROSENHOLTZ R., ADELSON E. H.: Material perception: What can you see in a brief glance? *Journal of Vision 8* (2009). 5

[VBvdWV11] VAZQUEZ E., BALDRICH R., VAN DE WEIJER J., VANRELL M.: Describing reflectances for color segmentation robust to shadows, highlights, and textures. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 33*, 5 (2011), 917–930. 6

[vGVSG10] VAN GEMERT J., VEENMAN C., SMEULDERS A., GEUSEBROEK J.-M.: Visual word ambiguity. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 32*, 7 (July 2010), 1271–1283. 3

[VZ02] VARMA M., ZISSERMAN A.: Classifying images of materials: Achieving viewpoint and illumination independence. In *Proceedings of the European Conference on Computer Vision (ECCV)* (2002), pp. 255–271. 3

[VZ03] VARMA M., ZISSERMAN A.: Texture classification: Are filter banks necessary? In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2003), vol. 2, pp. 691–698. 3

[VZ04] VARMA M., ZISSERMAN A.: Unifying statistical texture classification frameworks. *Image and Vision Computing 22*, 14 (2004), 1175–1183. 3

[VZ09] VARMA M., ZISSERMAN A.: A statistical approach to material classification using image patch exemplars. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 31*, 11 (2009), 2032–2047. 3

[WGK14] WEINMANN M., GALL J., KLEIN R.: Material classification based on training data synthesized using a BTF database. In *Proceedings of the European Conference on Computer Vision (ECCV)* (2014), vol. 3, pp. 156–171. 3, 5, 6

[WGSD09] WANG O., GUNAWARDANE P., SCHER S., DAVIS J.: Material classification using BRDF slices. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2009), pp. 2805–2811. 4

[WK15] WEINMANN M., KLEIN R.: Material recognition for efficient acquisition of geometry and reflectance. In *Computer Vision - ECCV 2014 Workshops* (2015), vol. 8927 of *Lecture Notes in Computer Science*, pp. 321–333. 3, 4