Material Classification based on Training Data Synthesized Using a BTF Database (Supplementary Material)

Michael Weinmann, Juergen Gall, Reinhard Klein

Institute of Computer Science, University of Bonn

In the scope of the supplementary material, we summarize the main aspects of our material databases presented in the accompanying paper with some additional illustrations. First, we describe the acquisition process and the measured material database. Subsequently, we discuss our second database which is synthesized using the data from the measured database in combination with different environments and virtual viewpoints.

Table 1. Databases referenced in the paper: Our databases are highlighted in red (*: in principle, an arbitrary number of configurations could be considered in the synthesis)

	CURET [3]	KTH-TIPS [5]	KTH-TIPS2 [2]	MPI-VIPS [6]	spectral database [7]	measured database	FMD [9]	synthesized database
material samples	61	10	44	11	90	**84**	1,000	**84**
categories	61	10	11	11	8	**7**	10	**7**
samples per category	1	1	4	1	N.N.	**12**	100	**12**
illuminations	4 ... 55	3	4	4	150	**151**	100	**30***
illumination type	controlled	controlled	controlled & ambient	controlled & ambient	controlled	**controlled**	natural	**natural**
viewpoints	7	27	27	27	1	**151**	100	**42***
images per sample	205	81	108	108	150	**22,801**	1	**1,260***
total image number	12,505	810	4,752	1,188	13,500	**1,915,284**	1,000	**105,840***

1 BTF Database

In the context of our measured material database, we briefly explain the acquisition setup used for capturing the database and the material samples contained in the database. Subsequently, we discuss the main differences of this measured database to previous material databases.

1.1 Acquisition Setup

Modern acquisition devices (e.g. [8]) are capable of capturing BTFs at high spatial and angular resolution in a practical way. This way, thousands of images can be acquired automatically which allows us to consider the material characteristics under densely sampled view-light configurations.

For each of the material instances, we measured a full BTF with 22,801 high dynamic range images providing a bidirectional sampling of the utilized 151 view directions and 151 light directions. Here, the view directions and light directions are evenly sampled in the upper hemisphere above the sample as illustrated in Fig. 1. The spatial resolution of the BTFs is 512×512 texels. Furthermore, we also measured the surface geometry using a standard structured light approach.

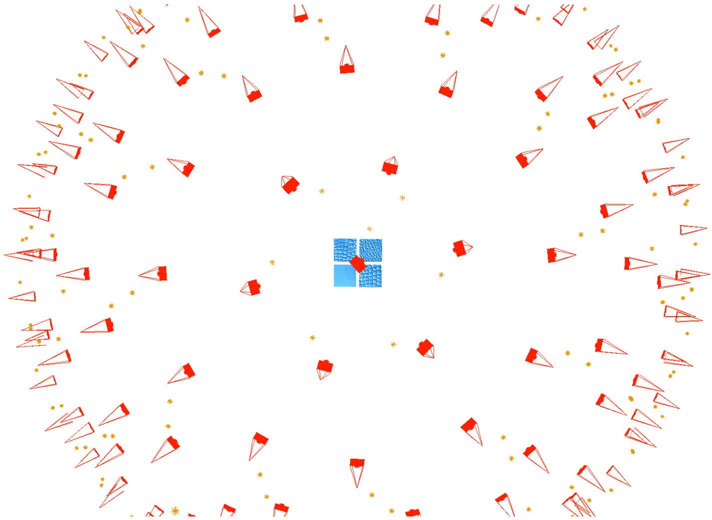

Fig. 1. Sketch of the acquisition setup (from top): The utilized 151 cameras (red) and 151 light sources (yellow) are arranged on a hemisphere above 4 samples

1.2 Measured Material Samples

For our database, we considered material categories which we expect to be relevant for scene analysis due to their presence in buildings, offices and streets. Therefore, the materials contained in our database can be categorized into the semantic classes *carpet*, *fabric*, *felt*, *leather*, *stone*, *wallpaper* and *wood*. For each of the individual material classes, we measured 12 different material instances to represent the intra-class variance. The respective material instances are illustrated in the Figures 2, 3, 4, 5, 6, 7, 8. As can be seen clearly, the variation in appearance differs for the different categories.

1.3 Differences to Previous Databases

Whereas most standard databases such as CUReT [3], KTH-TIPS [5], KTH-TIPS2 [2], MPI-VIPS [6] and the spectral material database in [7] are restricted to only a sparse sampling of different view and light directions, our new material

Fig. 2. Material samples in the category *carpet*

database provides a significantly denser sampling. In particular, we consider all the 22,801 combinations of the utilized 151 view directions and 151 light directions for each acquired physical material sample as shown in Table 1. This is an essential prerequisite for preserving the mesoscopic effects in material appearance in an accurate way. In total, our database contains 1,915,284 million images. A second objective of our database is to better capture the intra-class variances of each individual material category by containing measurements from 12 different material instances. While these individual material instances share some common characteristics of the corresponding category, they also cover a large variability.

2 Synthesized Database

The description of our synthesized database is followed by a discussion of its main difference to other databases. Subsequently, we discuss the comparison of using our synthesized database to using previous synthesized databases for classifier training.

2.1 Generation of Renderings

The main objective of our synthesized database is to provide images where material appearance is depicted in high quality under complex real-world environment conditions. For this reason, this database is extremely valuable for applications

Fig. 3. Material samples in the category *fabric*

which focus on real-world scenarios beyond the controlled environment being present in a lab.

The central requirements for establishing a synthetic database are:

- a photo-realistic depiction of the digitized materials
- the selection of the conditions under which the materials should be rendered
- an efficient rendering technique to synthesize a sufficiently large amount of images

Using the acquired surface geometry and the densely sampled BTFs of the individual material instances in our database described in Section 1 allows synthesizing the materials under a huge range of almost arbitrary viewing and lighting conditions while still preserving characteristic material traits in an accurate way. The viewing and lighting conditions used for generating the synthetic database should be representative for the conditions under which the materials will appear during classification. To account for the illumination conditions, we utilize environment maps which are used for scene relighting (e.g. [4]). For our synthesized database, we considered five representative HDR environment maps as shown in Figure 9 in our synthesizing process to sample the illumination conditions one encounters in these settings. Please note that the choice of the utilized environment maps can be varied or extended to approach the expected illumination conditions.

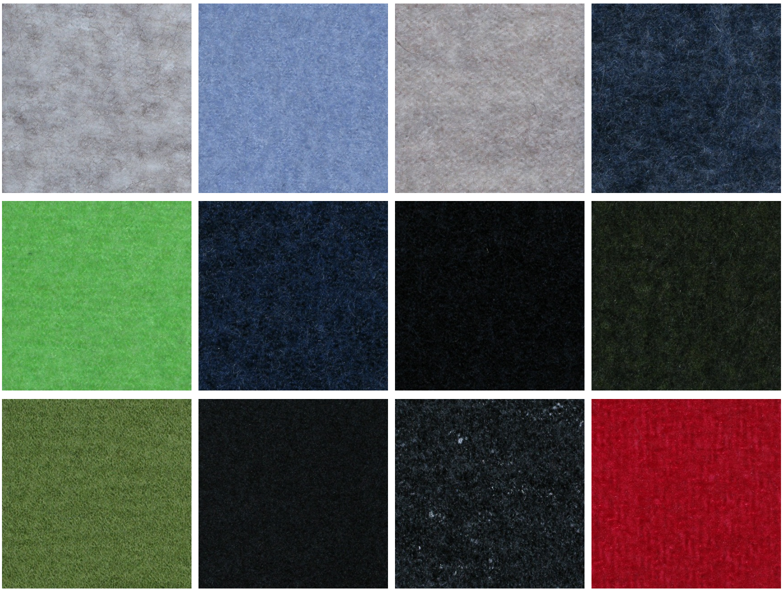

Fig. 4. Material samples in the category *felt*

Using 21 different rotations of the material sample ($\varphi = -67.5°, \ldots, 67.5° \times \theta = -45.0°, \ldots, 45.0°$) and two different distances to also consider the scale-induced variations in appearance of the materials, we synthesize images for each combination of material sample and environment map. To additionally increase the variance in the illumination conditions, we also rotated the utilized 5 environment maps (six different rotations) around the sample. In total, we obtain $21 \cdot 2 \cdot 5 \cdot 6 \cdot 12 = 15{,}120$ images for each of the semantic classes of our synthesized database. Example renderings are shown in Figure 10.

The rather high number of 105,840 synthesized images, for which still only a limited amount of viewpoints, scales and illumination conditions has been used, clearly indicates that an efficient rendering technique is mandatory for synthesizing all these images in a reasonable amount of time. To meet this requirement, we used an OpenGL based renderer and simulate the HDR environments in this renderer via approximating it in a similar way to the work in [1] with 128 directional light sources. These light sources are distributed representatively over the environment via a relaxation algorithm.

2.2 Differences to Previous Databases

The CUReT database [3], the KTH-TIPS database [5], the KTH-TIPS2 database [2], the MPI-VIPS database [6] and the spectral material database in [7] only consider controlled illuminations and a single ambient illumination (KTH-TIPS2 [2], MPI-VIPS [6]). However, considering only this severely limited fraction of possible illumination conditions within a lab environment is not sufficient for

Fig. 5. Material samples in the category *leather*

material classification under natural illumination. In contrast, our synthesized database incorporates different natural illuminations and, in the accompanying paper, we show the benefit of this data for material classification in complex real-world scenarios. While natural illumination has also been considered in the Flickr Material Database (FMD) [9], our approach of synthesizing data does not rely on manually acquiring and segmenting the images which severely limits the number of images included in the FMD. Following our synthesis technique, it is easily possible to synthesize thousands of images fully automatically.

2.3 Using our Synthesized Database vs Using Previous Synthesized Training Data for Classifier Training

Due to lack of space, we report more details on this experiment here in the supplementary material. Unfortunately, a comparison to other approaches using synthesized data, such as [6], is not directly possible. While the material shaders and the selected illumination conditions utilized for the generation of synthetic data in [6] are chosen to correspond up to some degree to the conditions during the acquisition of the KTH-TIPS2 database, our synthesized data considers different material categories which we expect to be more relevant for scene analysis due to their presence in offices, buildings and streets. Our data does not focus on the controlled illumination conditions in a lab environment and, instead, approaches the more complex real-world conditions in arbitrary environments. The only class we directly share with the MPI-VIPS and the KTH-TIPS2 databases

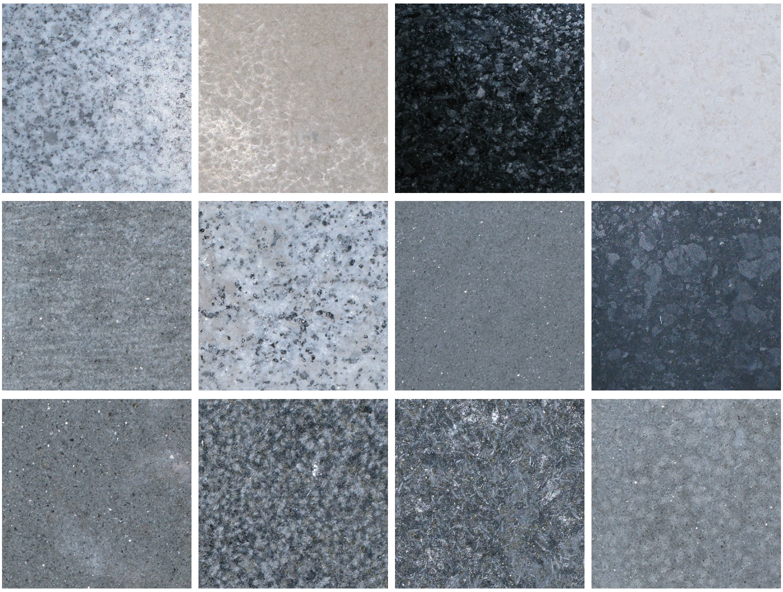

Fig. 6. Material samples in the category *stone*

is the class wood. For analyzing the difference of using shaders with the reproduction of the illumination conditions present in the test dataset and our data synthesized from several real-world wood samples under more complex illumination conditions, we first train a wood-vs-rest classifier on the synthesized MPI-VIPS database and perform a classification on the KTH-TIPS2 database. Here, 61.57% of the images in the wood category of the KTH-TIPS2 database are classified correctly. In contrast, we also perform an experiment where we replaced the wood images of the MPI-VIPS database in the training set with images taken from our OpenGL-synthesized data with environmental illumination for the class wood. Even though the illumination conditions in our data are rather different in comparison to the ones present during the acquisition of the KTH-TIPS2 database, we obtain a correct classification of 76.16% for the images with the label wood of the KTH-TIPS2 database which represents a significant improvement. This clearly demonstrates the benefit of using accuratly digitized materials for material classification from synthesized data and taking the intra-class variances into account.

3 Future Work

Currently, we are extending our database with the categories *artificial leather*, *granules*, *laminate*, *metal*, *sponge* and *tile* depicted in Fig. 11, Fig. 12, Fig. 13, Fig. 14, Fig. 15, Fig. 16. Again each of the categories is composed of 12 material samples to account for the intra-class variance within a material category. This

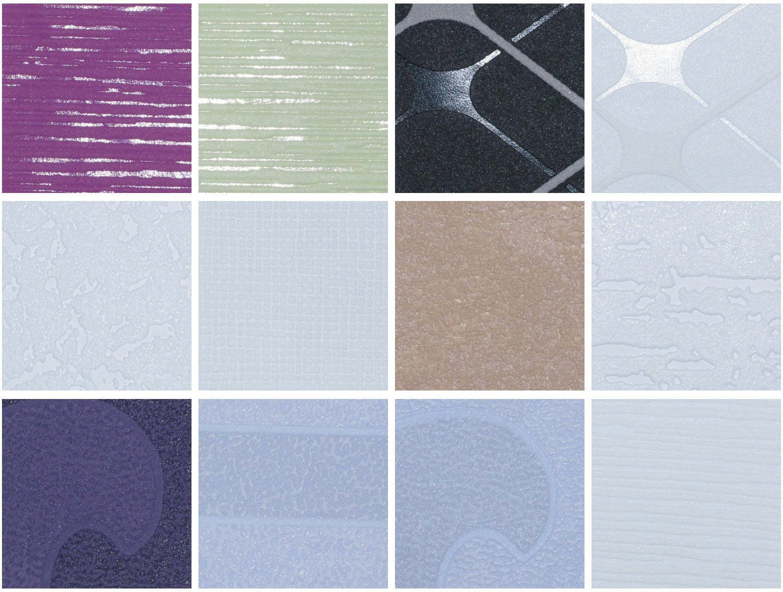

Fig. 7. Material samples in the category *wallpaper*

data will be published together with the BTF database introduced in the paper. As a result, the database will contain dense BTF measurements of 156 material samples containing $156 \cdot 22{,}801 = 3{,}556{,}956$ images in total.

References

1. Ben-Artzi, A., Ramamoorthi, R., Agrawala, M.: Efficient shadows for sampled environment maps. J. Graphics Tools 11(1), 13–36 (2006)
2. Caputo, B., Hayman, E., Mallikarjuna, P.: Class-specific material categorisation. In: ICCV. vol. 2, pp. 1597–1604 (2005)
3. Dana, K.J., van Ginneken, B., Nayar, S.K., Koenderink, J.J.: Reflectance and texture of real world surfaces. Tech. rep. (1996)
4. Debevec, P.: Rendering synthetic objects into real scenes: bridging traditional and image-based graphics with global illumination and high dynamic range photography. In: SIGGRAPH. pp. 189–198 (1998)
5. Hayman, E., Caputo, B., Fritz, M., Eklundh, J.O.: On the significance of real-world conditions for material classification. In: ECCV. pp. 253–266 (2004)
6. Li, W., Fritz, M.: Recognizing materials from virtual examples. In: ECCV. pp. 345–358 (2012)
7. Liu, C., Yang, G., Gu, J.: Learning discriminative illumination and filters for raw material classification with optimal projections of bidirectional texture functions. In: CVPR. pp. 1430–1437 (2013)
8. Schwartz, C., Weinmann, M., Ruiters, R., Klein, R.: Integrated high-quality acquisition of geometry and appearance for cultural heritage. In: The 12th International

Fig. 8. Material samples in the category *wood*

Symposium on Virtual Reality, Archeology and Cultural Heritage VAST 2011. pp. 25–32 (2011)
9. Sharan, L., Rosenholtz, R., Adelson, E.H.: Material perception: What can you see in a brief glance? Journal of Vision 8 (2009)

Fig. 9. Three of the environments used during the synthesizing process. These three environment maps are publicly available as mentioned in the paper

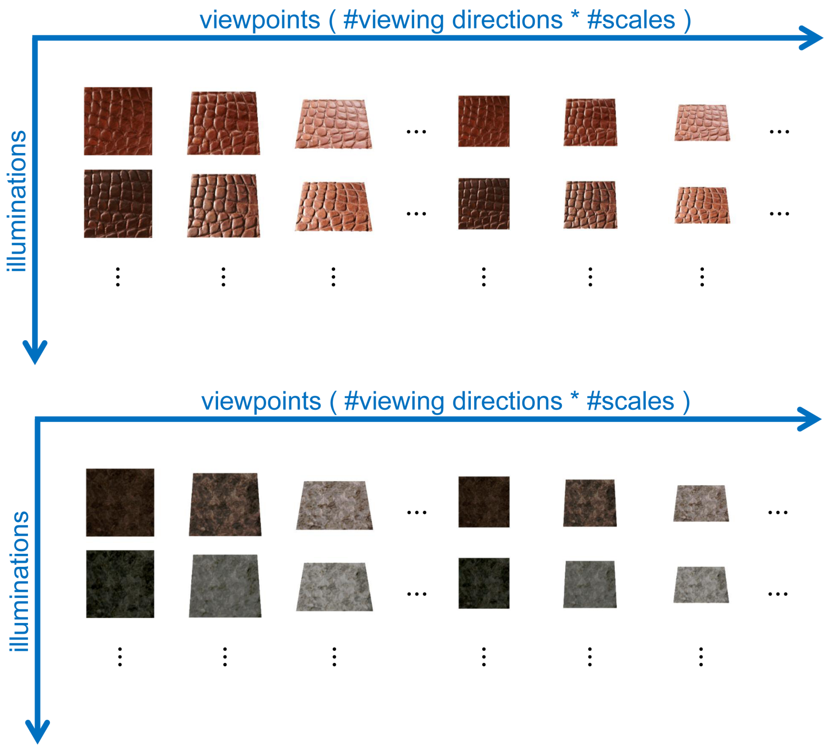

Fig. 10. Exemplary illustration of the images in the synthesized database for 2 of the 84 different material samples: The number of viewpoints is determined by the number of different viewing directions and relative distances between the virtual camera and the sample. In our current database, we incorporated two different scales. However, incorporating more different scales can be achieved easily with our rendering pipeline. Per row, the environment remains the same and only the sample is rotated. For the illustrated examples, the main part of the light comes from the top, i.e. when the samples are rotated in a way so that more light hits the surface the surface appears brighter. Please note that only the rotation around one axis is illustrated in the figure. We also combined this rotation with a rotation of the sample by up to 67.5° to the left and right respectively. To further increase the variance of the illumination conditions in our database, we additionally rotated the utilized 5 environment maps.

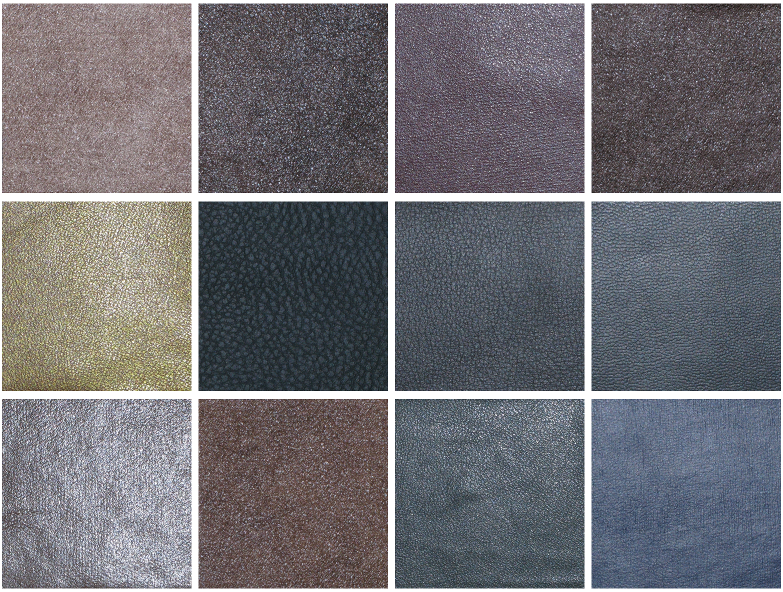

Fig. 11. Material samples in the category *artificial leather*

Fig. 12. Material samples in the category *granules*

Fig. 13. Material samples in the category *laminate*

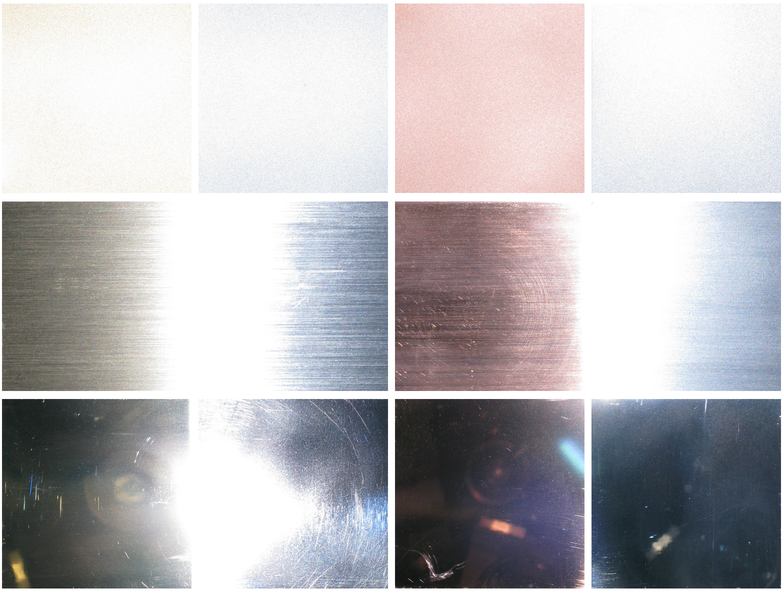

Fig. 14. Material samples in the category *metal*

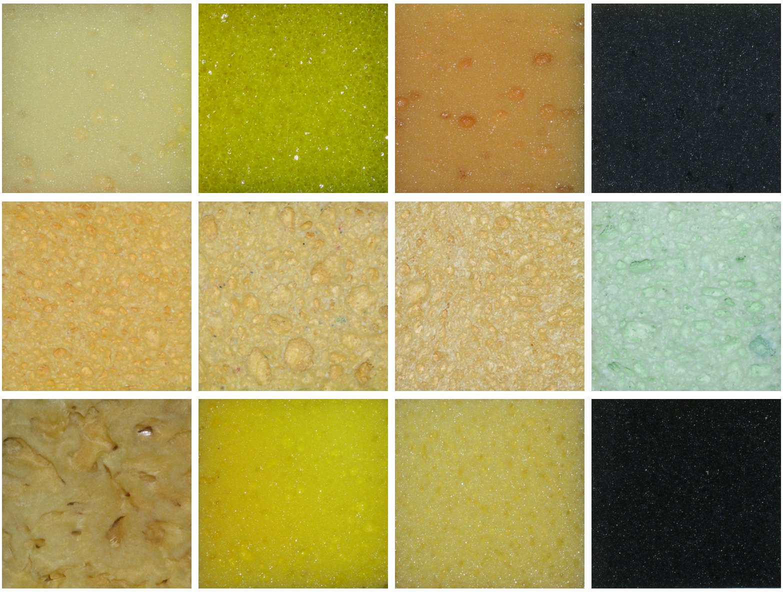

Fig. 15. Material samples in the category *sponge*

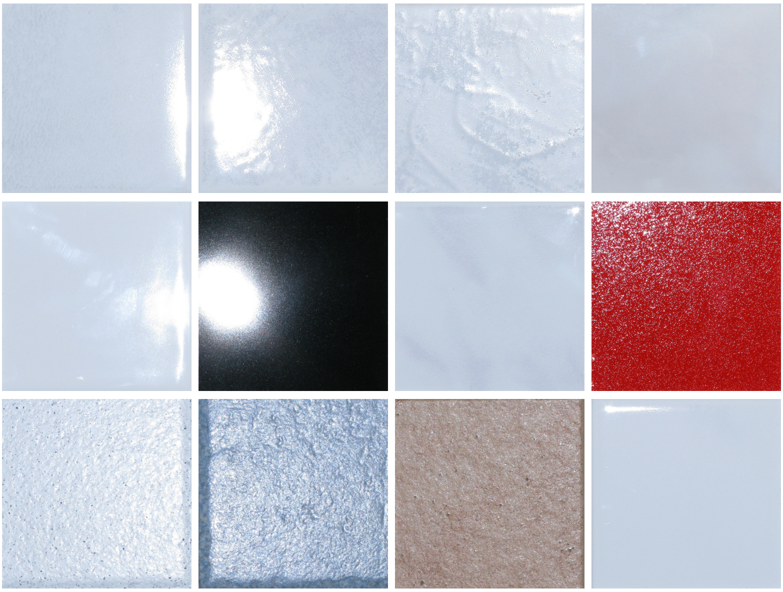

Fig. 16. Material samples in the category *tile*