A Psychophysically Validated Metric for Bidirectional Texture Data Reduction

Jiri Filip[1] Mike Chantler[1] Patrick Green[1] Michal Haindl[2]

[1]Heriot-Watt University

[2]Institute of Information Theory and Automation of the AS CR

SIGGRAPH ASIA 2008
NEW HORIZONS

Motivation

- Materials in real world

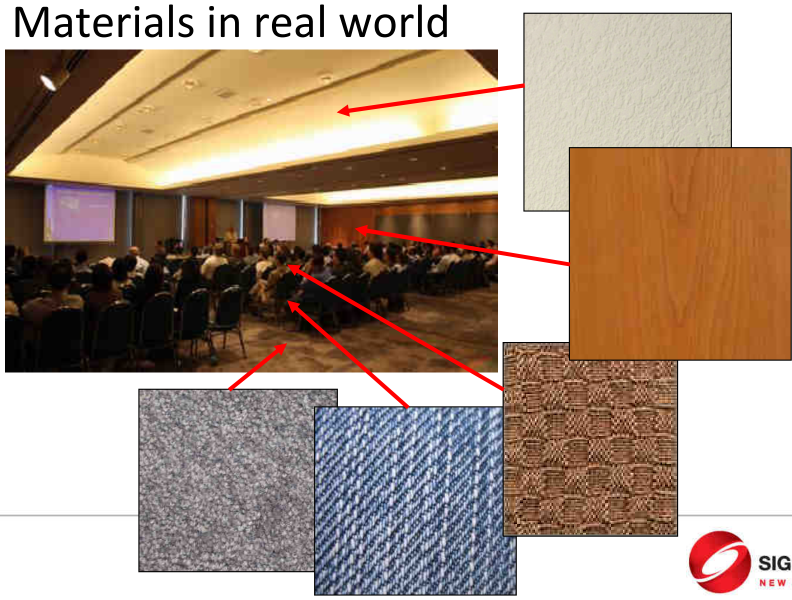

Any real-world environment contains a number of different materials. If we look more closely most of them have more-or-less apparent textures. Appearance of these textures often heavily depends on surrounding environment. Ultimate goal of computer graphics is to master real-time visualization of many natural effects and realistic appearance of textured materials is certainly one of them.

Although there is a number of ways to represent appearance of textured materials only some of them provide reasonable degree of realism. We can use for example displacement mapping methods or pixel-wise BRDF modelling, but without additional modelling or parameters setting we still struggle to capture important effects as translucency, inter-reflections or sub-surface scattering.

However, there are certain industrial applications requiring really accurate representation of all effects present in the real textured materials. One example of them is virtual designing or prototyping of interiors in architecture or in automotive industry, where the main goal is to find optimal materials or their surface finishes for specific environments and defined illumination.

Motivation

- Accurate simulation of material appearance
- **Safety Simulations**

Courtesy of RealReflect.org

Similar application are virtual visual safety simulations in car or airspace industry preventing unwanted lights reflections from materials covering certain parts of the environment.

If we want to capture all the mentioned effects accurately and generally for most of the applicable materials and without additional parameters tweaking, we do not have any other option than using bidirectional texture functions.

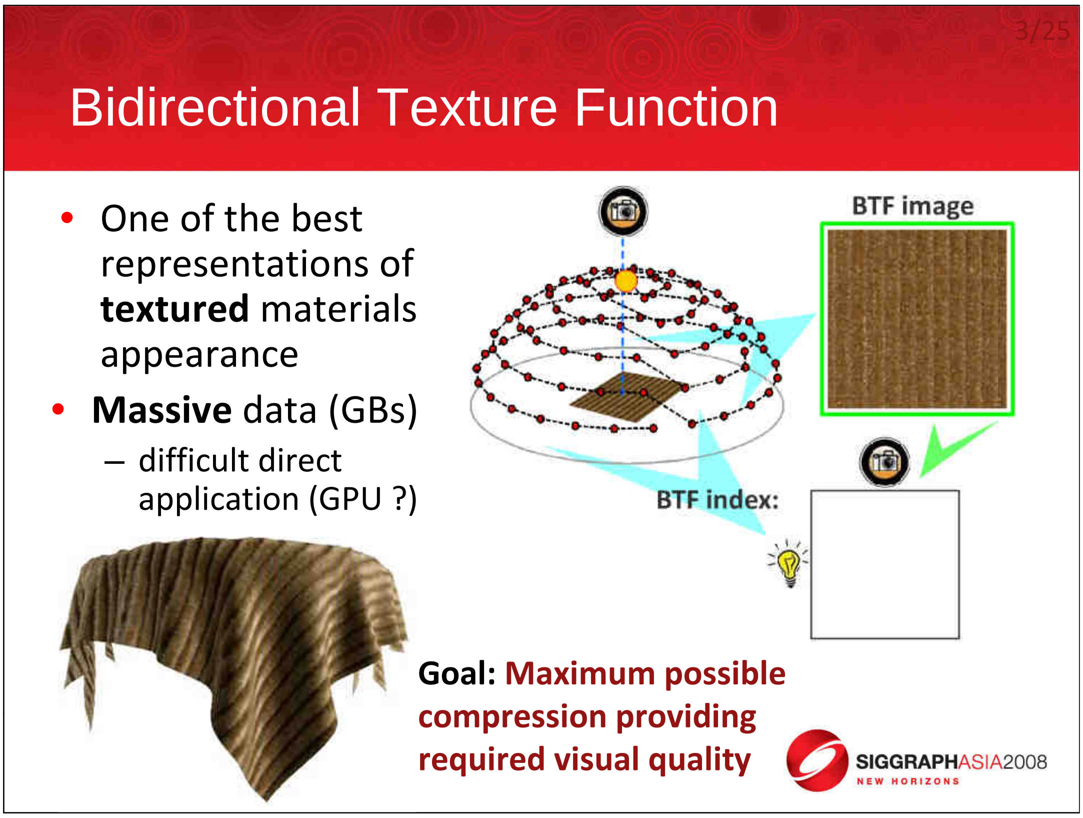

Bidirectional texture function, in short BTF, describe appearance of real material surface dependently on illumination and viewing directions. So BTF capture surface image for many combinations of light and camera positions. The word “many” means several thousands surface images.

This puts BTFs in position of one of the most accurate digital representations of textured material appearance, but on the other hand also one of the most memory consuming.

The storage costs of thousands of images per each material sample are in order of Gigabytes, which is still quite beyond the current graphics hardware capabilities.

So, our goal was to achieve maximal possible compression for each material represented by BTF sample, while still maintaining the required visual accuracy and fast data reconstruction.

Related Work

- **BTF Compression**
  - Reflectance models, PCA, wavelets….survey [Müller04]
  - Probabilistic models [Haindl07]
  - Clustering techniques [Leung01]
- **BTF** surface **dimensionality** estimation [Healey00]
- **Psychophysical aspects** of real measurements
  - BRDF [Pellacini00], [Matusik03], [tePas05], [Ramanarayanan07], [Vangorp07]
  - BTF [Meseth06], [Filip08]
- **Perceptual properties** of different BTF **samples** prior to their processing?

When we look back on research done on Bidirectional texture functions in the past, we can see substantial effort in compression of this data. Such methods were based on pixel-wise reflectance models, linear factorization methods, Markov random fields or pixel-wise clustering. Surprisingly, none of these methods adapted its parameters with respect to the type of the sample being compressed.

There was one attempt to analyse different BTF samples and evaluate their surface dimensionality. However, the authors didn't relate the obtained dimensionality to any visual quality controlling parameter.

As we suspected that setting of optimal BTF compression parameters might involve human evaluation we went through published works dealing with human perception of real view and illumination dependent measurements. We haven't found many references and most of them were related only to BRDF data – so without a texture information. There were just two BTF related references, where the first one was assessing the quality of different sample representations, and the second one was evaluating different down-sampling schemes.

We are not aware of any research analysing perceptual properties of different samples and using this knowledge for enhancement of BTF measurement and compression methods.

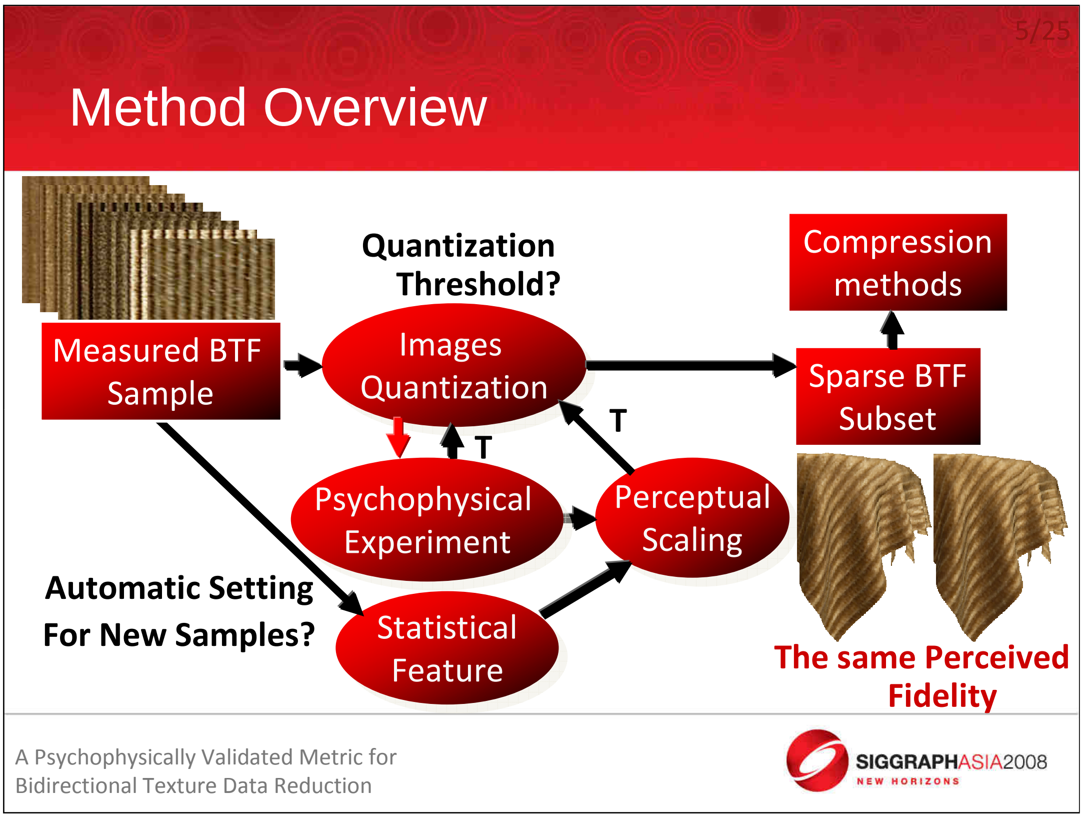

This fact motivated us try to develop a relatively simple technique for BTF data reduction outlined in this slide. We start with original data. We apply vector quantization of BTF images, but the question is how to set the optimal quantization threshold for each sample.

To get these thresholds we ran a psychophysical experiment on a test set of BTF samples. Well, this works, but when new sample arrives we have the same problem – how to set the threshold without doing the experiment again? So we were looking for computational feature correlating with the data from the experiment. We have found one so we then scaled its value with respect to data from the initial experiment and ended up with automatic BTF reduction metric preserving very sparse set of perceptually important BTF images. This small set produces visually indistinguishable results and can be used as an input to standard BTF compression methods.

Tested BTF Samples

- **Bonn BTF database**
  http://btf.cs.uni-bonn.de
- 8 samples
- 81 views x 81 illum.
- 6561 images/sample

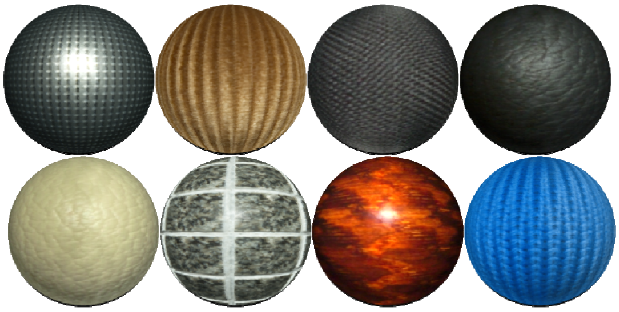

Point-light illumination

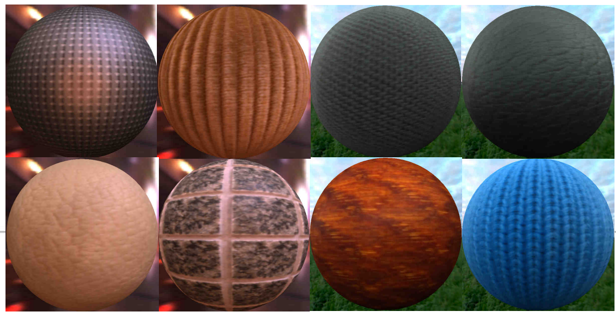

Two environment illuminations

SIGGRAPHASIA2008 NEW HORIZONS

We used eight samples from Bonn University BTF database, each having more than eighty illumination and view directions resulting into more then six thousands images per sample. In our work we used point-light and two environment illuminations.

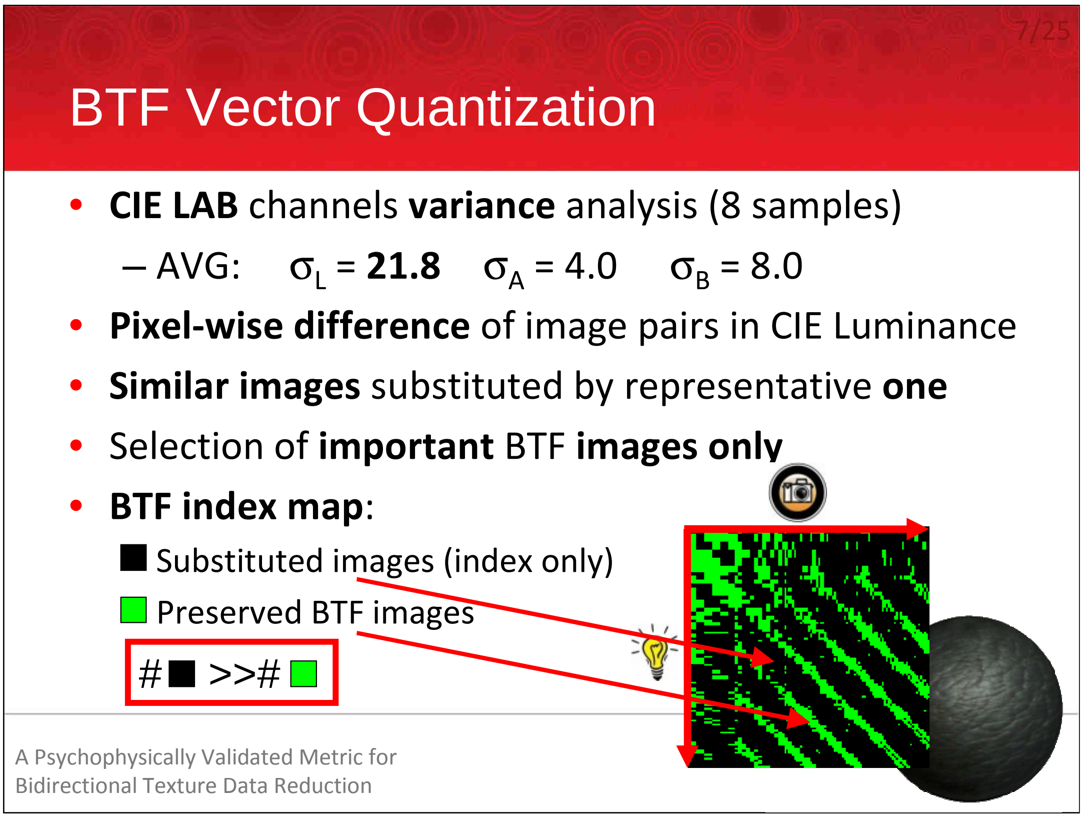

Prior to data quantization we performed data variance analysis in channels of LAB colour-space for all eight samples. We have found that the average variance in Luminance channel is far more higher then in the other two. So we decided to cut computational times and performed the quantization based on luminance information only. However, the method can be used in the same way using the full colour information.

Then we performed BTF images quantization. This process is based in computation of pixel-wise difference between each pair of images. When the difference between pair of images is lower than the predefined quantization threshold one of them is substituted by an index pointing to the other one. Given the threshold a certain number of images is removed.

As a result of the quantization we obtain an index map, representing each combination of camera and light directions by one dot. The black dots correspond to the images substituted by index pointing to one of the preserved images in green.

We expect that number of substituted images will be considerably higher than number of the preserved images, but that visually salient features of the material sample will be still preserved.

8/25

BTF Vector Quantization – demo

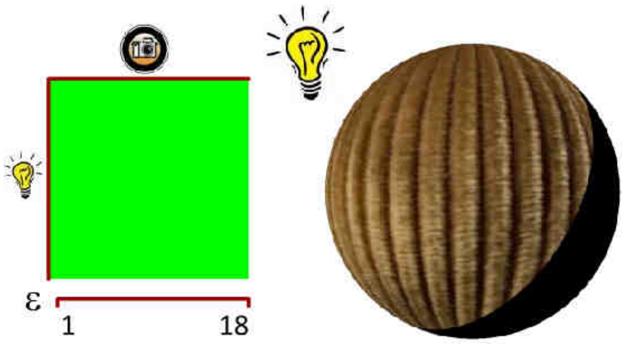

A Psychophysically Validated Metric for
Bidirectional Texture Data Reduction

Following demo for corduroy sample shows distribution of removed images (in black), and corresponding visual degradation of the sample as the threshold is continually increased.

BTF Vector Quantization

- No. of **preserved images** vs. **quantization threshold**

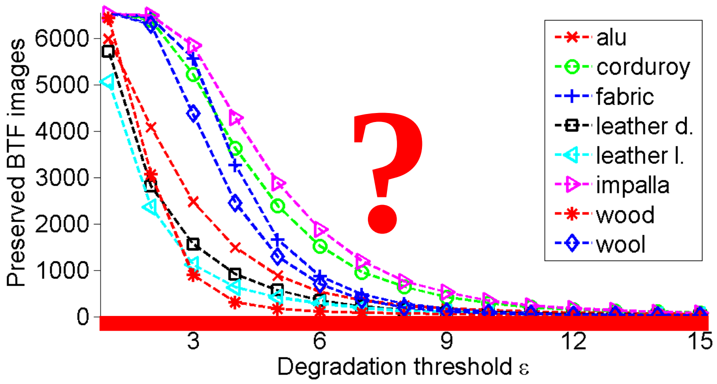

How to set thresholds for different samples?

This graph, shows dependency of number of preserved images on the quantization threshold. The shape of all graphs is similar for different samples, but they are all shifted apart.

This confirms our assumption that each sample requires different threshold setting. Our task is to set the thresholds appropriately for all the samples.

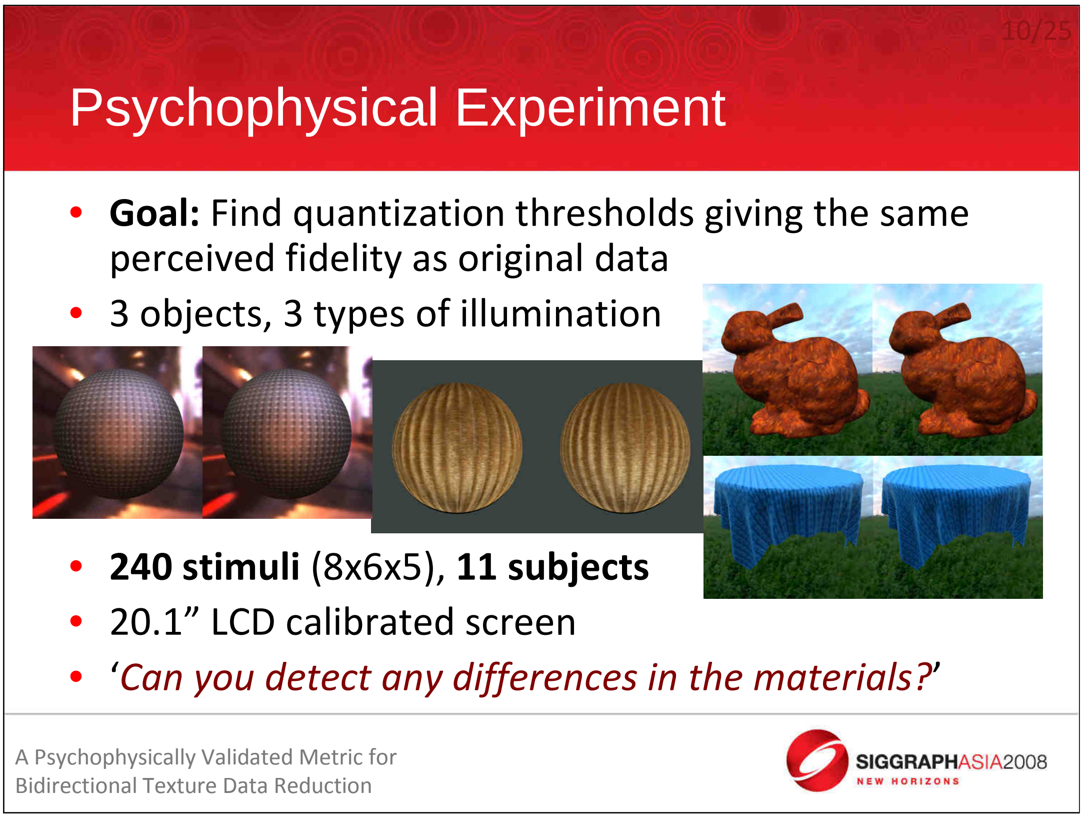

To achieve this, we prepared a simple psychophysical experiment. The goal of the experiment was to find the quantization thresholds providing the same perceived fidelity as rendering from original data. The experiment comprised six quantization thresholds applied on rendering of eight tested BTF samples using three different shapes and three different illuminations. This configuration resulted into 240 stimuli showing in random order side-by-side renderings of original and quantized data.

Eleven participants observed the stimuli on calibrated LCD screen and responded to yes-no question: “Can you detect any difference in the materials on the objects?”.

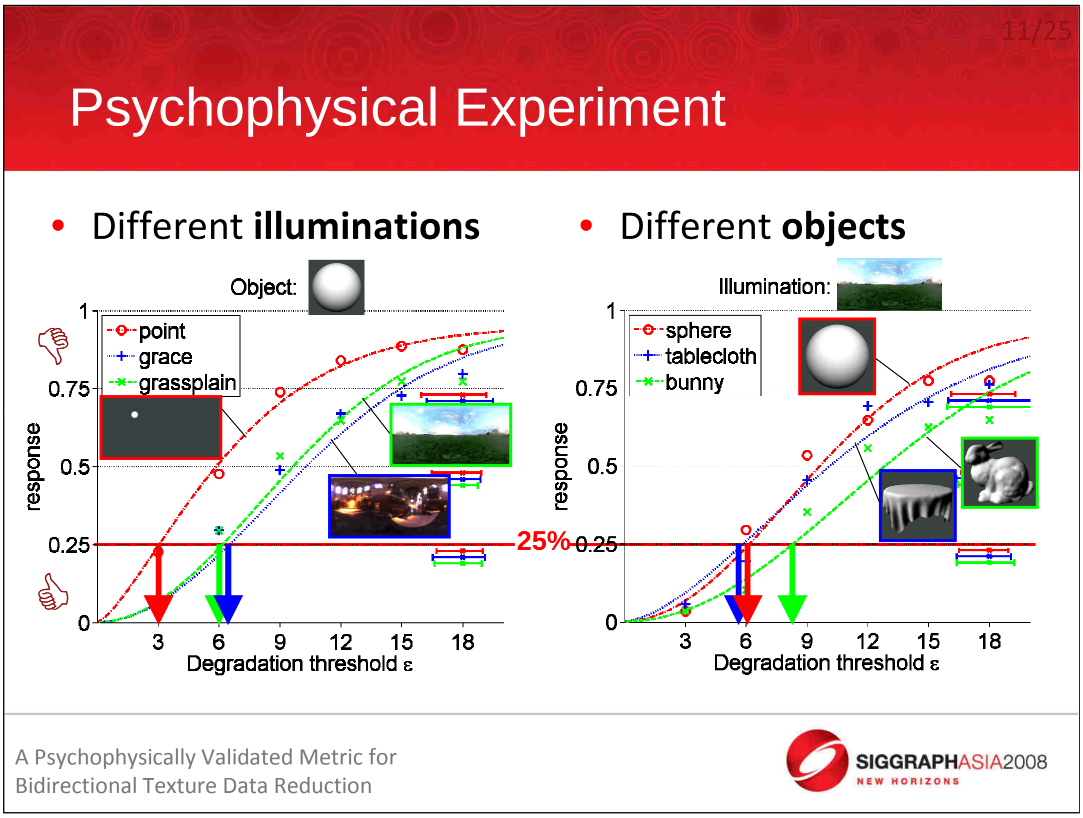

Afterwards responses of all participants were averaged and psychometric functions were fitted to the averaged data. These functions describe smoothly dependency of perceived difference on quantization thresholds.

This left graph shows observers' performance on shape sphere illuminated by different illuminations. We can see that the point light (in red) has higher responses than environment illuminations (blue, green) where the imperfections are hidden in convolution of pixel with set of lights representing the environment. The right graph shows observers' performance on different shapes illuminated by environment illumination, and suggest that the more complex the shape is the more difficult it is to spot the difference.

Of course, response level for which we estimate the thresholds depends on requirements of the intended application. In context of our work we have used 25% level to obtain our thresholds, which means that 75% of observers were not able to spot any difference between the original and quantized data.

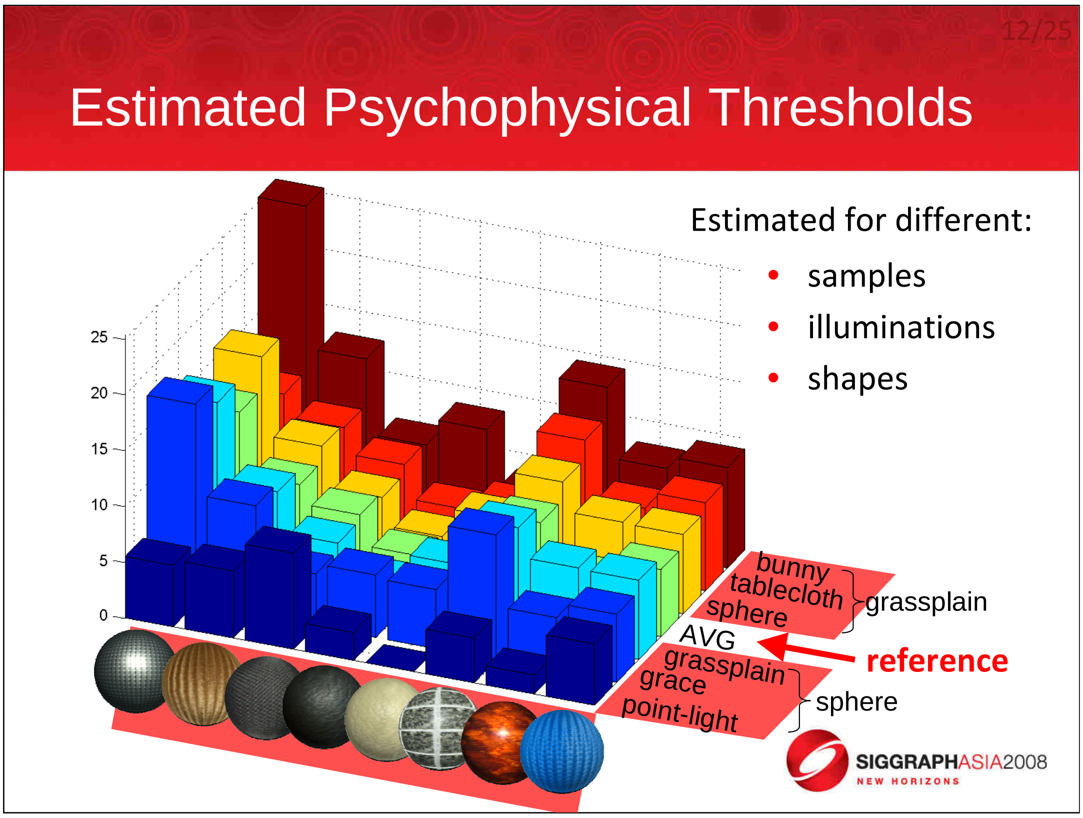

In this way we obtained thresholds for different samples, illuminations, and shapes. As mentioned before the thresholds for point light are considerably lower than for the environment illuminations. Also the thresholds across different environments and shapes are more-or-less similar. So we have chosen thresholds averaged across all illuminations for shape sphere as a reference for our further investigation.

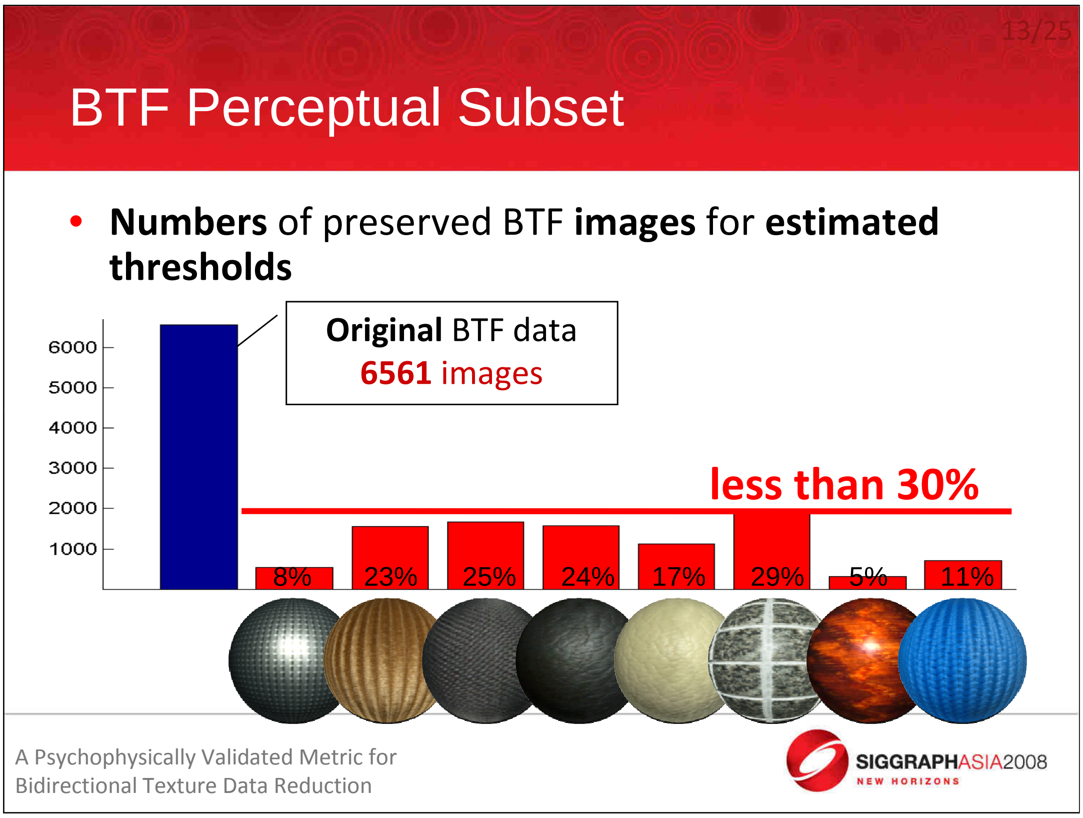

For this reference thresholds we got the following numbers of images preserved for individual BTF samples. This numbers confirms our initial assumption that majority of images can be removed without perceptible impact on visual fidelity. It turns out, that we need in all cases less than 30% of images, what is actually quite encouraging.

Results – psychophysical thresholds

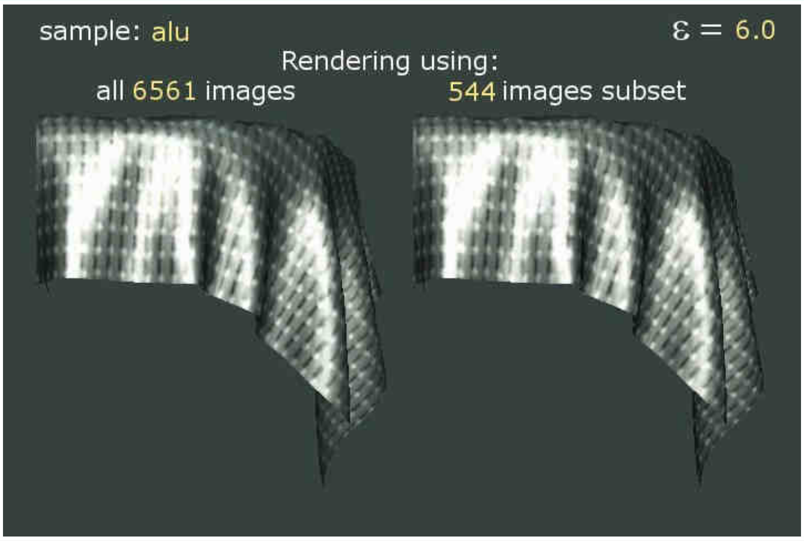

This movie compares rendering using all measured BTF images (on the left) with renderings from psychophysically estimated subsets (on the right).

As you can see even when sizes of the subsets are quite small the perceived quality is still very good, preserving all important features of the tested samples.

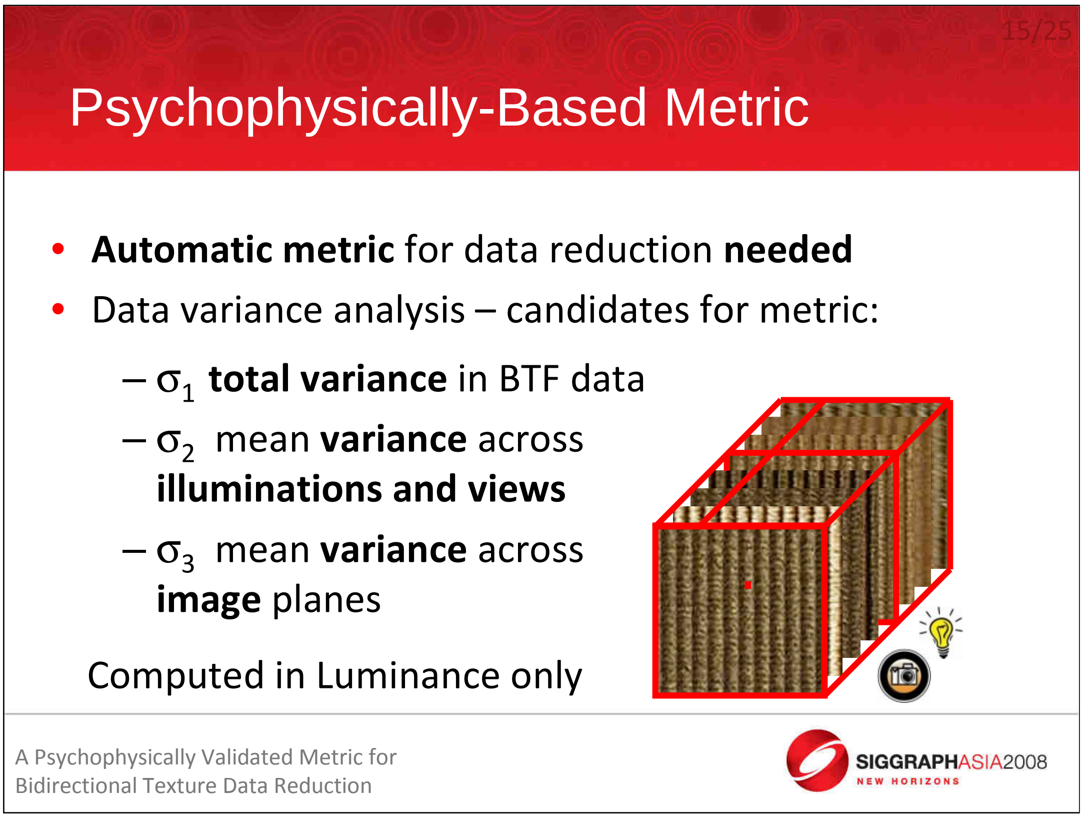

This works quite well, but performing the psychophysical experiment for any new sample is not very practical, so we need an automatic data-driven metric for settings of the thresholds.

As we believe that the key difference between individual material samples lies in their variance, we tested several features of BTF data variance:

The first one computes total variance in the data, the second one computes mean variance over all view and illumination dependent pixels and the third computes mean variance over all image planes. Again, all these features were computed in luminance channel only.

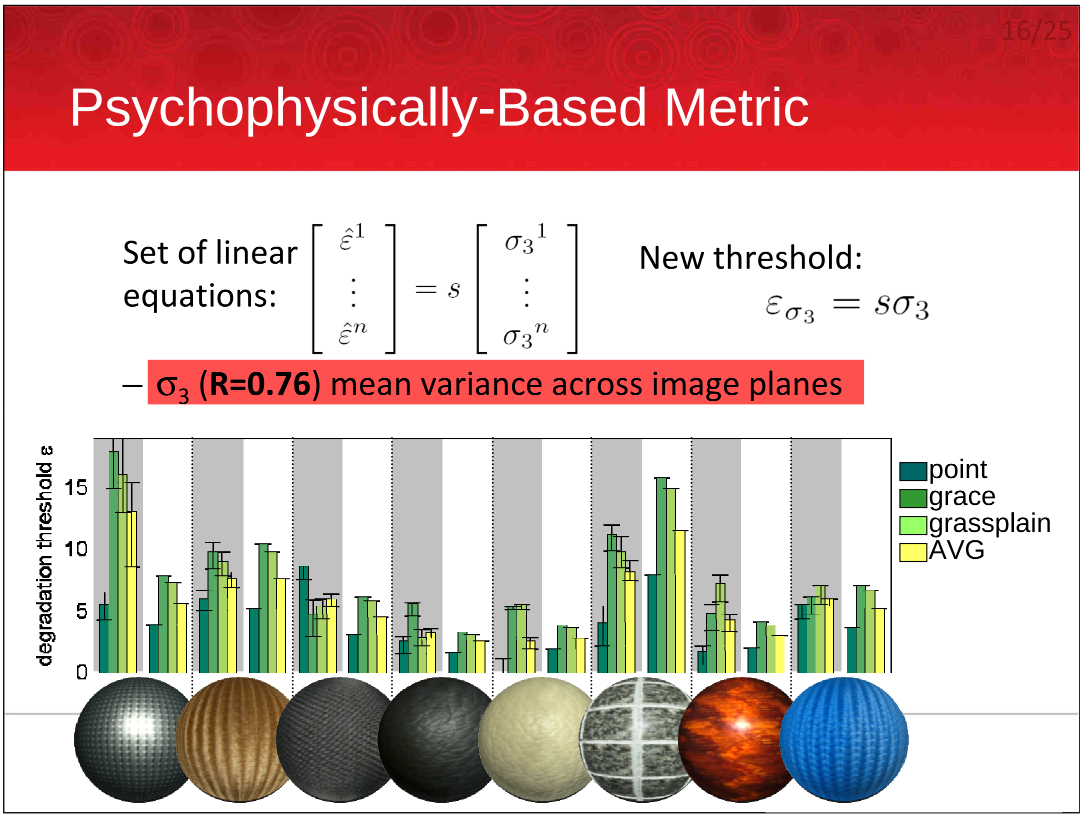

To find which of these three features corresponds the best to the thresholds obtained from the experiment (that you can see in the graph below), we computed correlation coefficient between the estimated thresholds and each group of data features from the previous slide. It turned out that the highest correlation corresponds to mean variance across all image planes. Now we have a feature that is relative to the estimated threshold. But to get absolute value we have to apply some kind of scaling. To do so we solved a set of linear equations using estimated thresholds and the feature values for all the tested samples and obtained a scaling factor “s” . This factor was estimated for each illumination type separately. Our new threshold is then obtained by multiplication of the computed variance by the scaling factor “s”. Then we obtained the following thresholds, that were set based on data variance. Most of these thresholds match quite reliably to those obtained from the psychophysical experiment (so those with grey background). One exception is the first aluminium sample, where the high difference for environment illuminations was caused by the improper fitting of psychometric functions to quite inconsistent observes’ data.

Validation Experiment

- **6 additional BTF samples**, **18 subjects**, one shape
- scaling value s from the previous experiment.
  - Automatic thresholds setting metric
  - No need to perform psychophysical experiment again

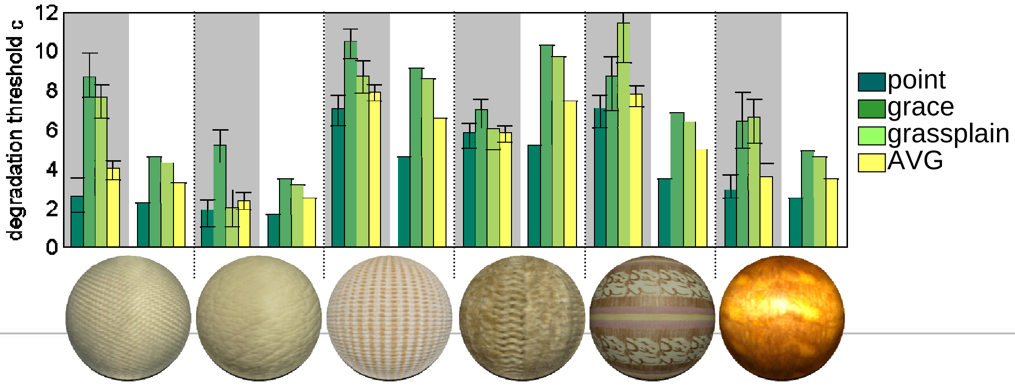

34% 24% 24% 32% 15% 5% of images

This was quite encouraging, so we decided to verify our results by additional validation experiment. We used another set of six BTF samples and ask 18 observers for the same task as in our main experiment. We have used the same set of illuminations, but just the shape of sphere. As a result of the experiment we obtained the psychophysical thresholds for each sample and illumination.

Then we computed our automatic variance-based thresholds, but using scaling factors obtained for the previous set of materials. The resulted thresholds shows again good proportional match with the estimated ones. They are mostly conservative, except highly complex and structured sample of knitted wool, where the values for environment illuminations were slightly higher.

The numbers of preserved images for all samples are is still reasonable low.

Results – derived thresholds

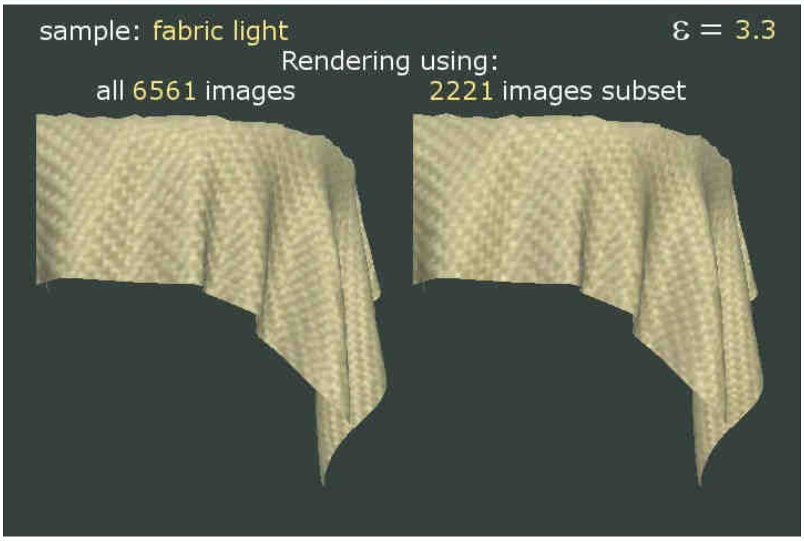

This movie compares again rendering of original BTF measurements (on the left) with rendering of subset obtained for our variance-based thresholds (on the right).

Still there are no apparent perceptual differences and all major samples characteristics are preserved.

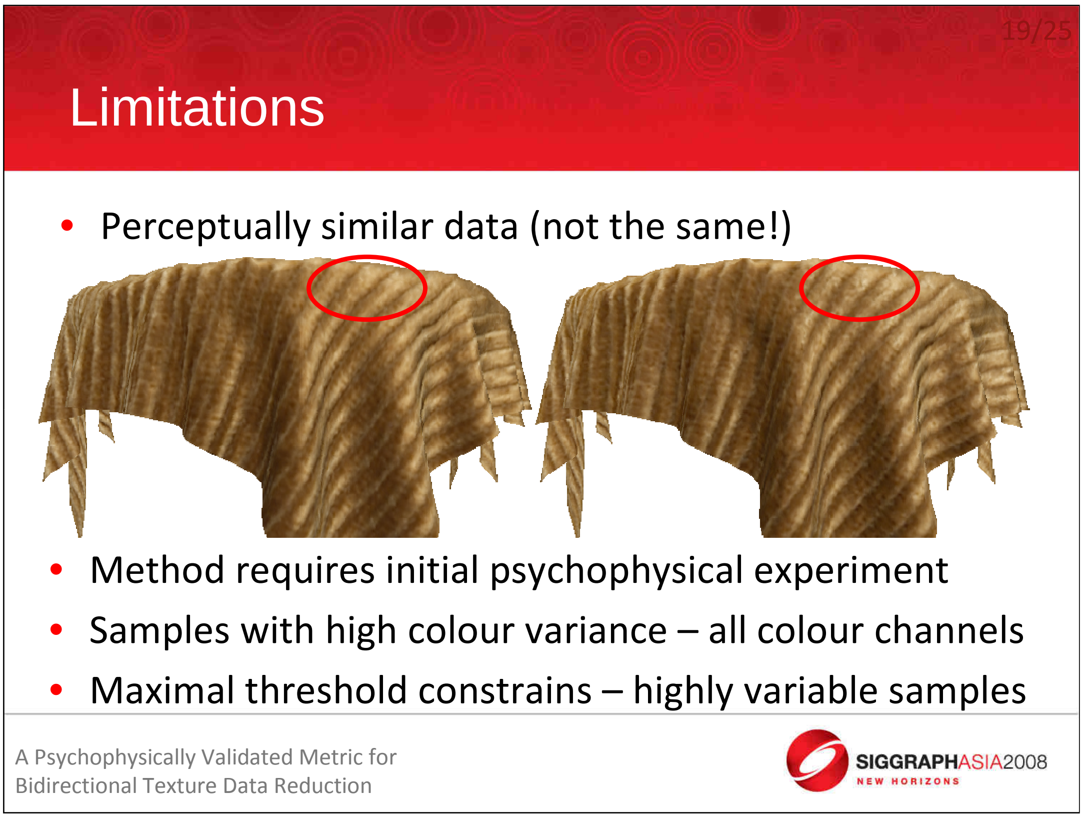

Now I'd like to briefly mention limitations of our approach:

First: our methods is not lossless so the rendered data are not exactly the same, but only perceptually very similar. You can see one difference nicely here.

Second: Our method requires to run initial psychophysical experiment with representative samples and illuminations to obtain appropriate scaling factors, that can be then used for automatic settings of thresholds for any new sample.

Third: Samples having wide range of colours should be handled more carefully. So, the quantization and data variance should be computed using all colour channels not just luminance one.

And the last one: Highly structured samples exhibit extensive masking and occlusions effects and so very high variance. The maximal applicable thresholds for such samples have to be limited to certain maximum level to obtain appropriate visual quality. The proper understanding of this behaviour is currently a subject of our research.

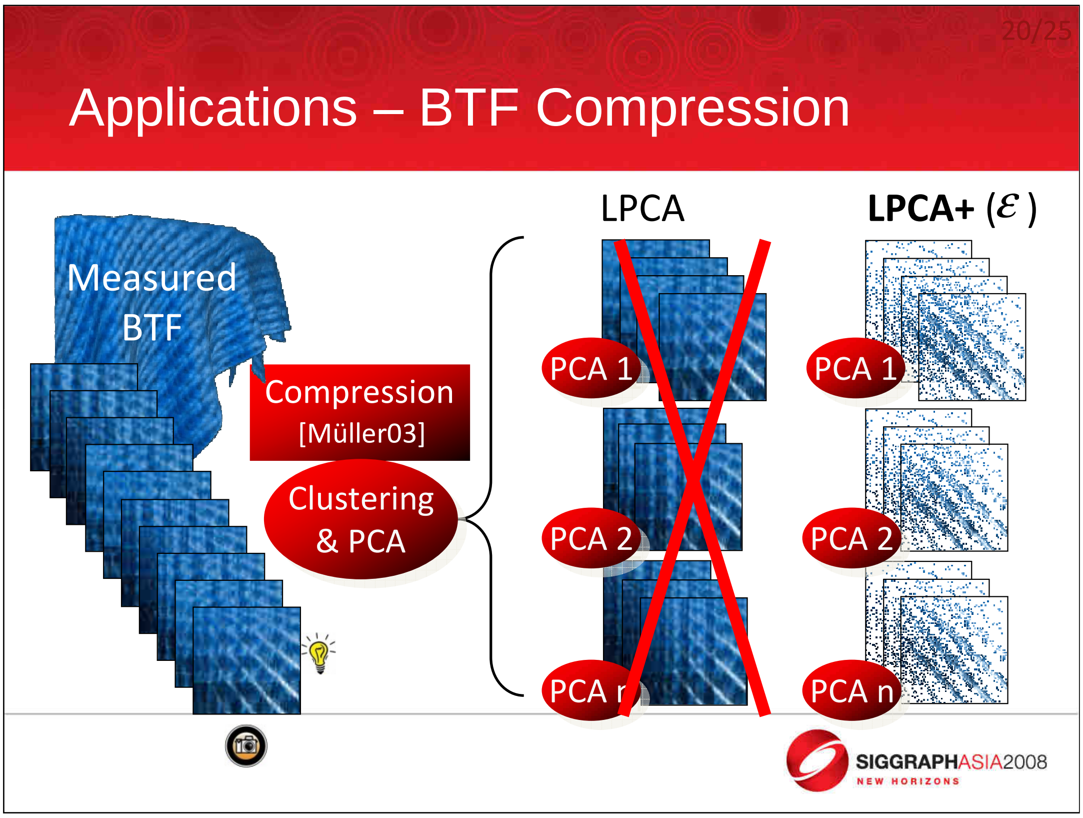

In the rest of my time I'd like to mention two major applications of the proposed data reduction method. First of them is BTF data compression. Our technique allows significantly increase efficiency of many BTF compression methods, just by compressing not all, but only the relevant data. We did an experiment with one of the most efficient methods published by Muller five years ago. This method is based on BTF clustering and simultaneous compression of each cluster locally by PCA. We represented BTF sample by set of illumination and view dependent pixels as recommended by the authors. However, instead of compressing all data as the original method did, we compressed only the selected subset of BTF images.

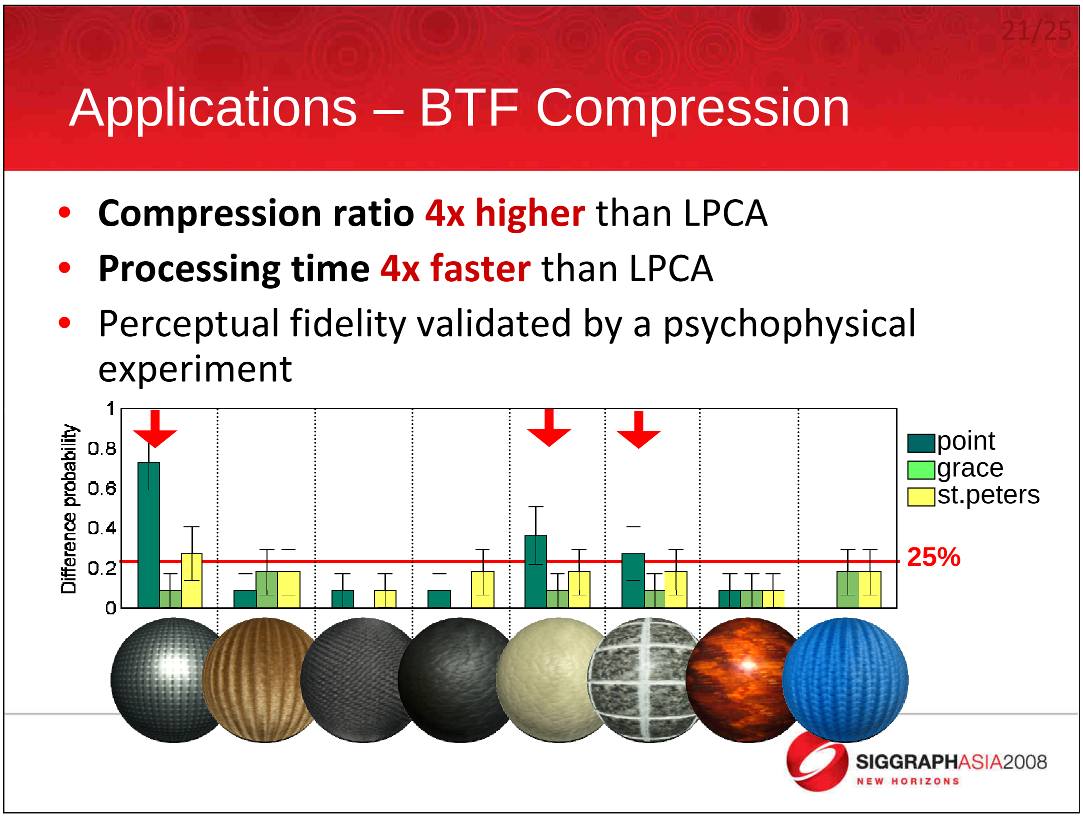

This approach proved to be much more efficient as we in average obtained 4 times higher compression rates and compression time was reduced 4 times in comparison with the original method. To validate our results we again run a psychophysical study with 11 observers comparing the performance of the original and the modified method and got the following results. This graph shows perceived differences between both methods for different samples and illuminations. Most of the tested samples were actually under the 25% level, which was set previously during estimation of data quantization thresholds. Only three samples exceeded this level and those are shown in the following slide.

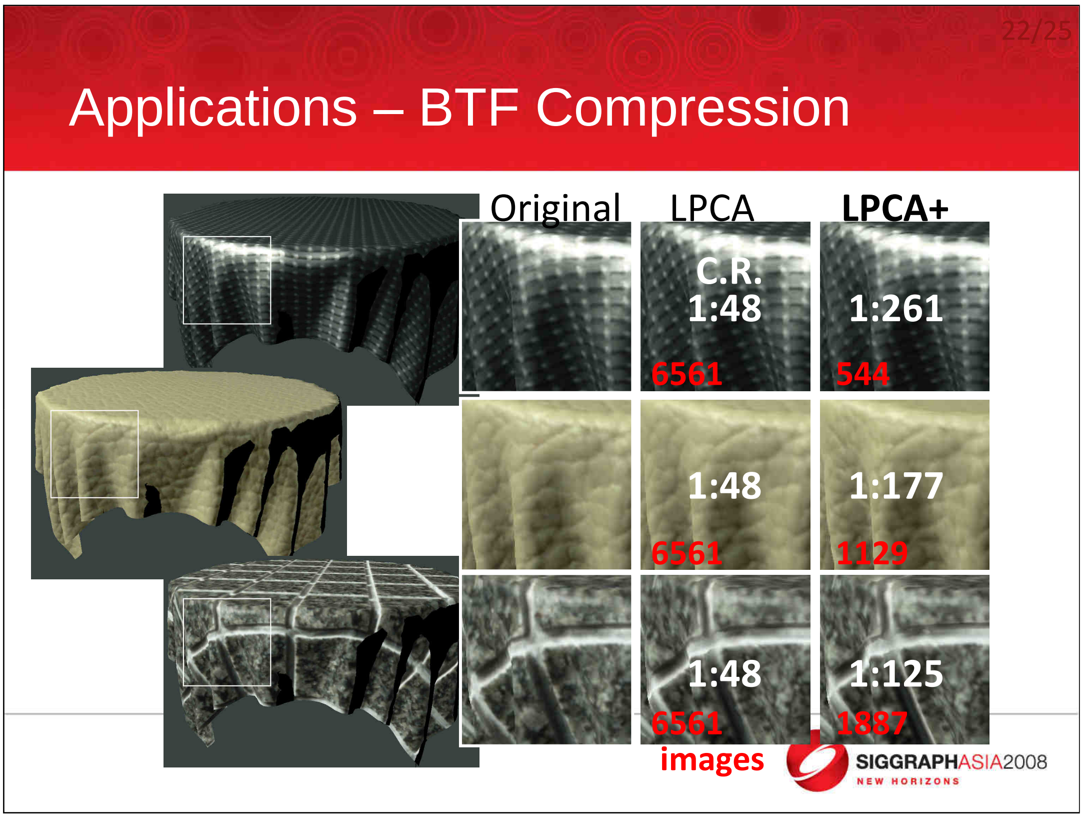

The samples are mapped on tablecloth shape and illuminated by a point-light from the left. On the right side you can compare performance of the original data mapping with results of both compression methods. The results are still very similar, however, the number of BTF images required is quite different. So while the original method has fixed compression ratio regardless the input data our modification produces much higher compression rates dependent on actual sample variance. We believe that many other BTF compression methods can be enhanced in the same way.

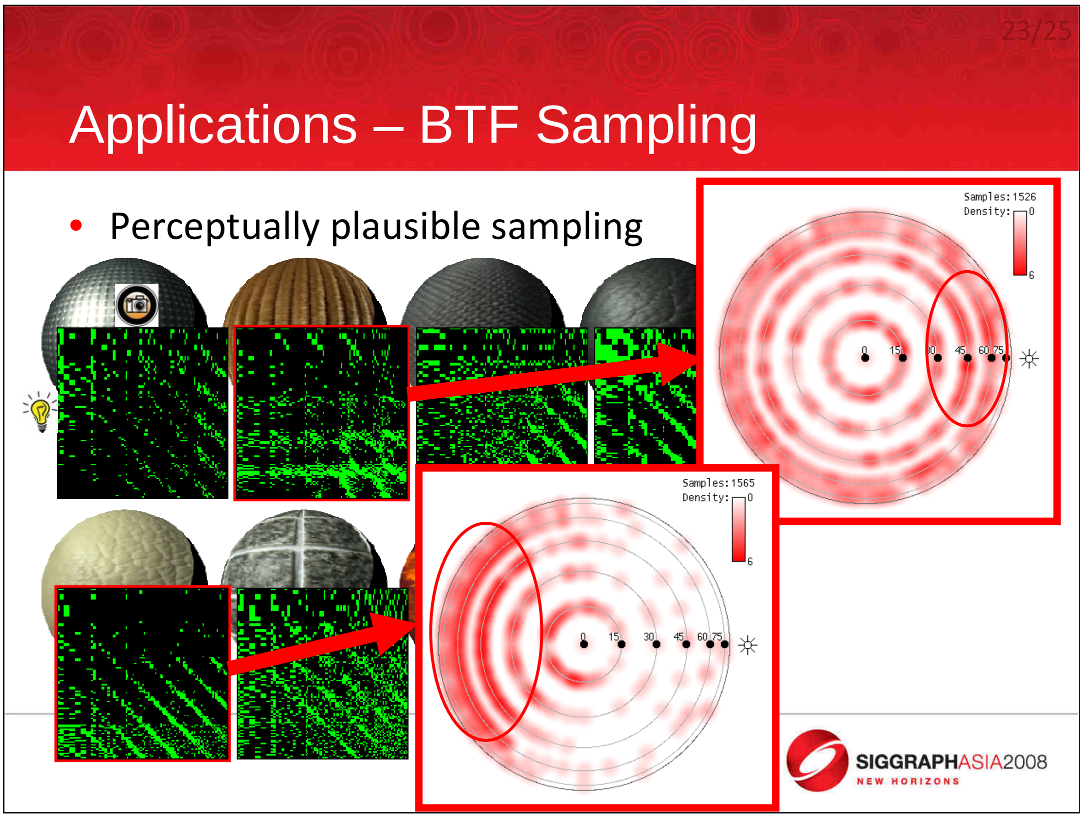

The second application of our technique is related to optimal BTF sampling. The BTF measurement is very demanding process, often requires hours of acquisition and post-processing of huge amount of data. Any method suggesting which camera and light directions should be preferred during sampling would be quite handy.

Our technique produces index maps showing that preserved BTF images (in green) form patterns which seems to be more-or-less typical for groups of similar materials. When we visualize the information from the index maps more practically, we can for example obtain the appropriate distribution of view directions with respect to fixed illumination on the right. For the “leather” sample we can see that the sampling of material around specular peak is the most important, while for the “corduroy” sample it seems that we should rather sample the material more uniformly with particular care about retro-reflections back to direction of illumination.

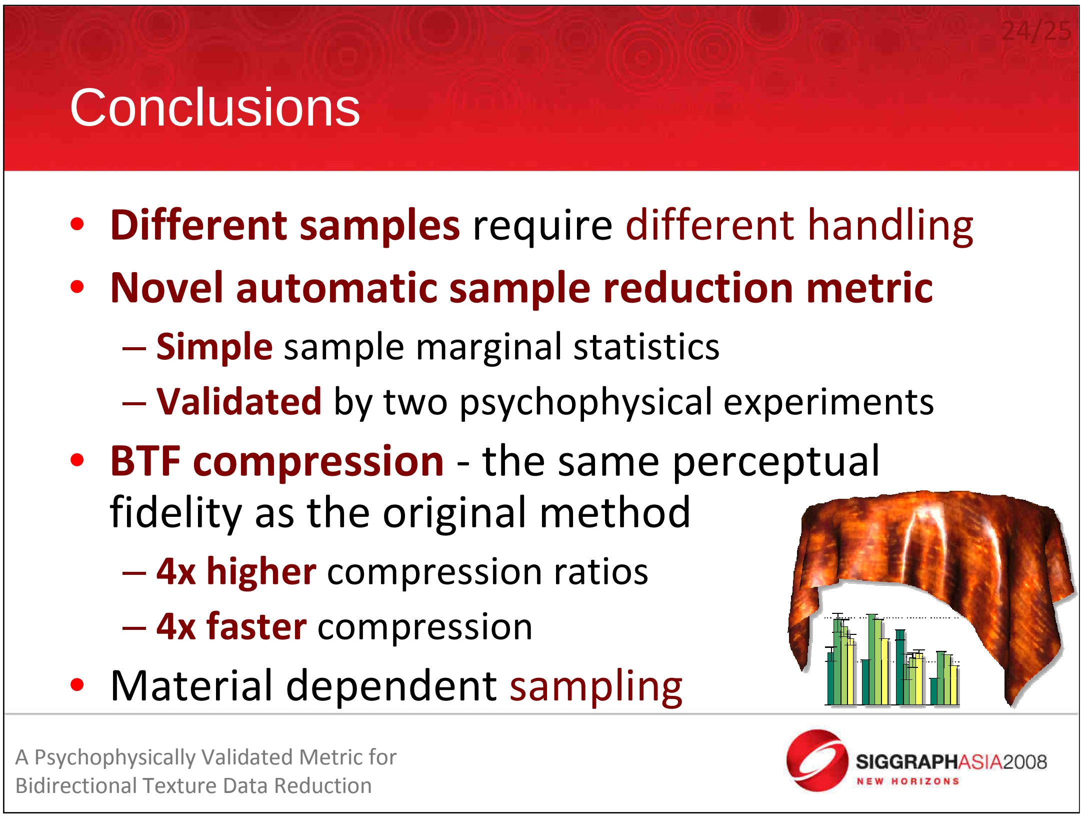

To sum up this talk I would like to point out four main results.

First, although it seems to be quite straightforward, this work is the first one explicitly showing that different BTF samples require different treatment during their processing, to achieve predefined visual fidelity.

Second, and probably the most important result is a new simple BTF reduction metric, which we believe can be extended to any other view or illumination dependent data not just to Bidirectional Texture Functions.

Third, the proposed data reduction metric allows us to determine perceptually important subset of original data, which produces the same perceived fidelity, but allows significantly increase efficiency of BTF compression methods.

And the last one, the method provides information about appropriate material dependent sampling.

Acknowledgements

- Bonn University
- Vlastimil Havran
- Subjects in the experiments
- Anonymous reviewers

Funding:

- EC Marie Curie Intra-European Fellowship No.41358
- EPSRI SMI grants GR/S12395, EP/F02553X/1
- GACR grant 102/08/0593
- Grant 1ET400750407

Finally, let us say big thanks to Bonn University for providing the BTF samples, Vlastimil Havran for help with environment illumination, participants in the experiments for their time, and anonymous reviewers for their great contribution to the paper.