SIGGRAPH 2008 Course on Projector-based Graphics

Projector-based Illumination for 3D Scene Modeling

Hendrik P.A. Lensch
MPI Informatik

Overview

- Scene appearance as higher dimensional reflectance fields
- Capturing (and removing) global versus local illumination effects
- Pattern projection for 3D geometry acquisition

max planck institut informatik

Hendrik Lensch

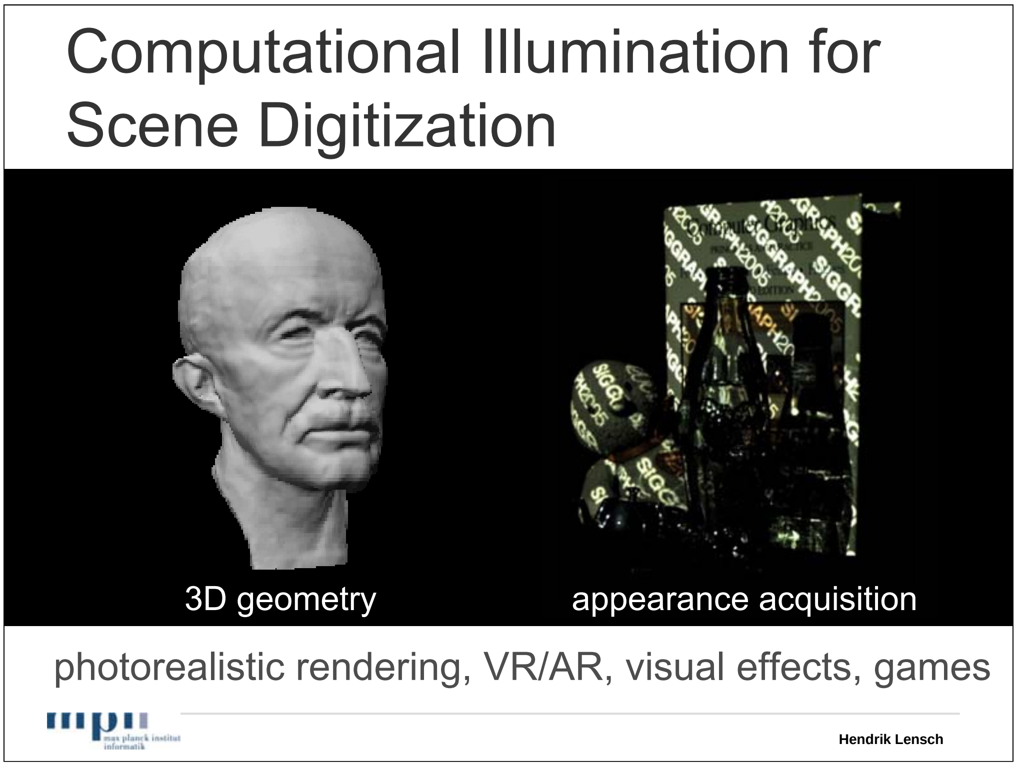

In this part of the talk we will discuss the design and the analysis of illumination patterns for scene digitization tasks, for acquiring 3D geometry to capturing the appearance of a scene. In the latter, we will discuss techniques for measuring the light transport on a ray-to-ray basis which further allows to distinguish local illumination effects, where the incident light is directly reflected at the scene surface, from global effects where light might be scattered or interreflected multiple times before arriving at the observer. We will first address the problem of appearance acquisition and then discuss the implications of a separation into direct and global components for 3D range scanning.

There are of course lots and lots of other fields where computational illumination plays a role, e.g. confocal microscopy,

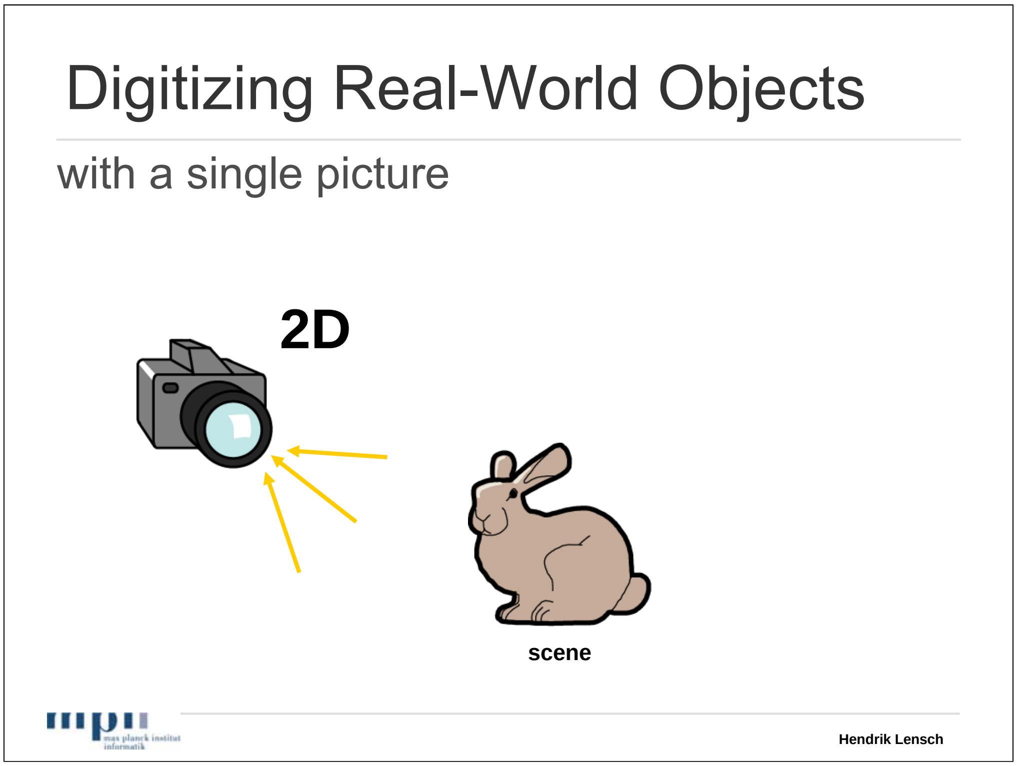

The simplest way to digitize a scene is to acquire a 2D photograph. Unfortunately, besides zooming in and out, the operations that can be performed on a single image are rather limited.

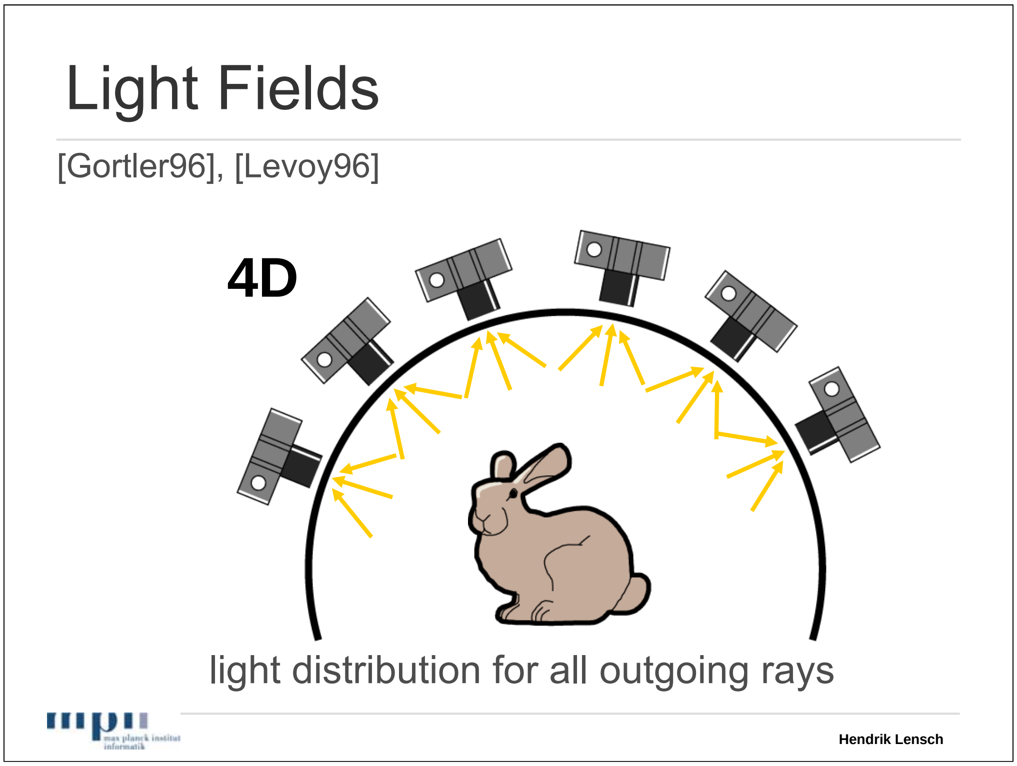

Lumigraphs or Light Fields represent a collection of photographs from a set of different view points. In essence, a fully sampled light field captures the outgoing radiance for any ray leaving the surface.

Light Fields

[Gortler96], [Levoy96]

novel view synthesis but no relighting

Hendrik Lensch

With light fields, applying view interpolation (maybe geometry assisted) arbitrary views of the object can be generated correctly as long as the virtual camera stays outside the visual hull. Direct and indirect reflections are reproduced in the rendered images, i.e. one can observe moving highlights, but the object will always be shown in the lighting that was present during acquisition.

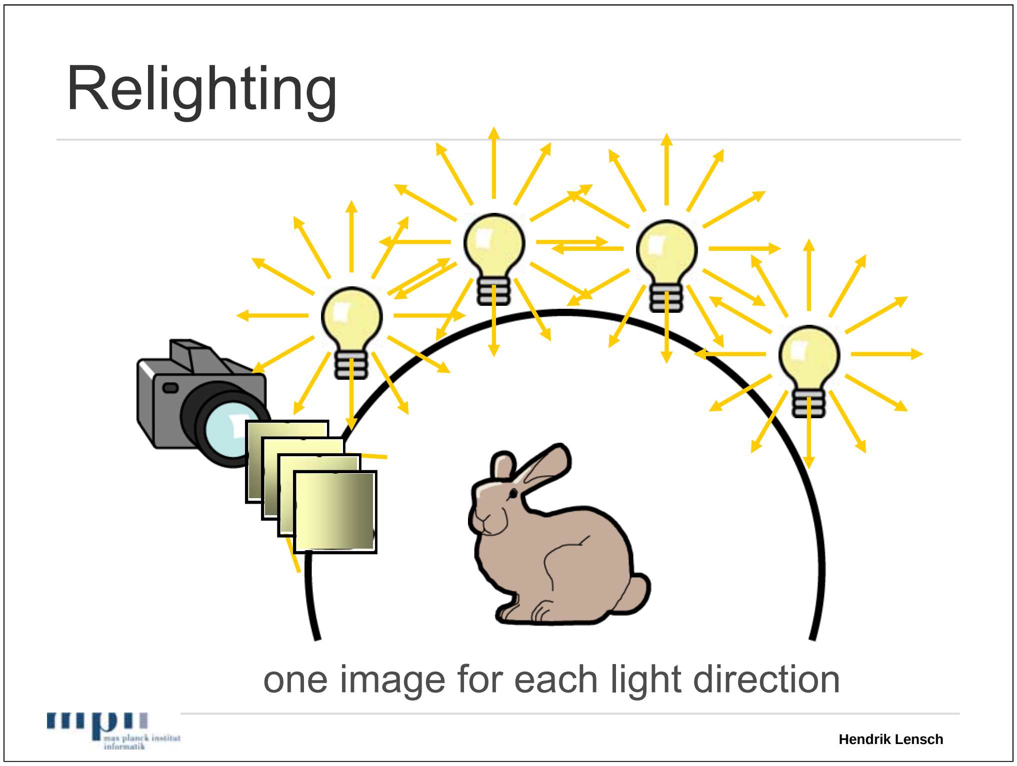

In order to obtain renderings in a novel lighting condition, one needs to capture a data set of images in a variety of illumination conditions. The simplest is to illuminate the scene subsequently from a number of directions, e.g. from all positions on a hemisphere.

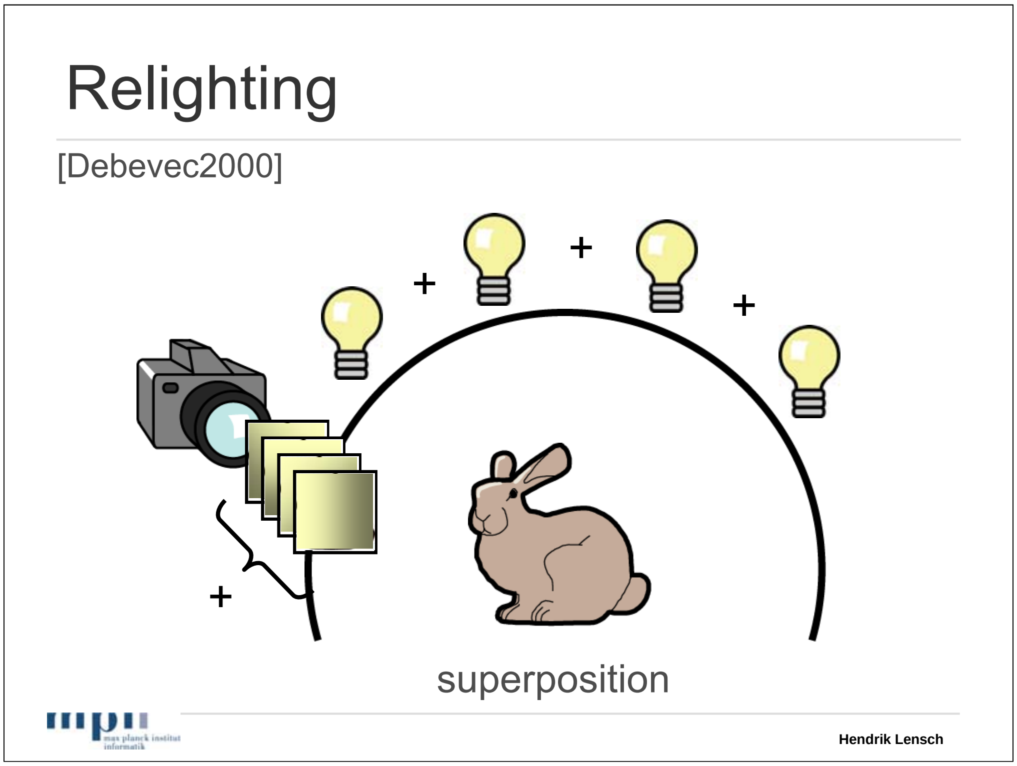

Due to the superposition principle of light these captured images can simply be added up to produce an image of the scene as if illuminated from all light sources at the same time. The data structure is called a reflectance field.

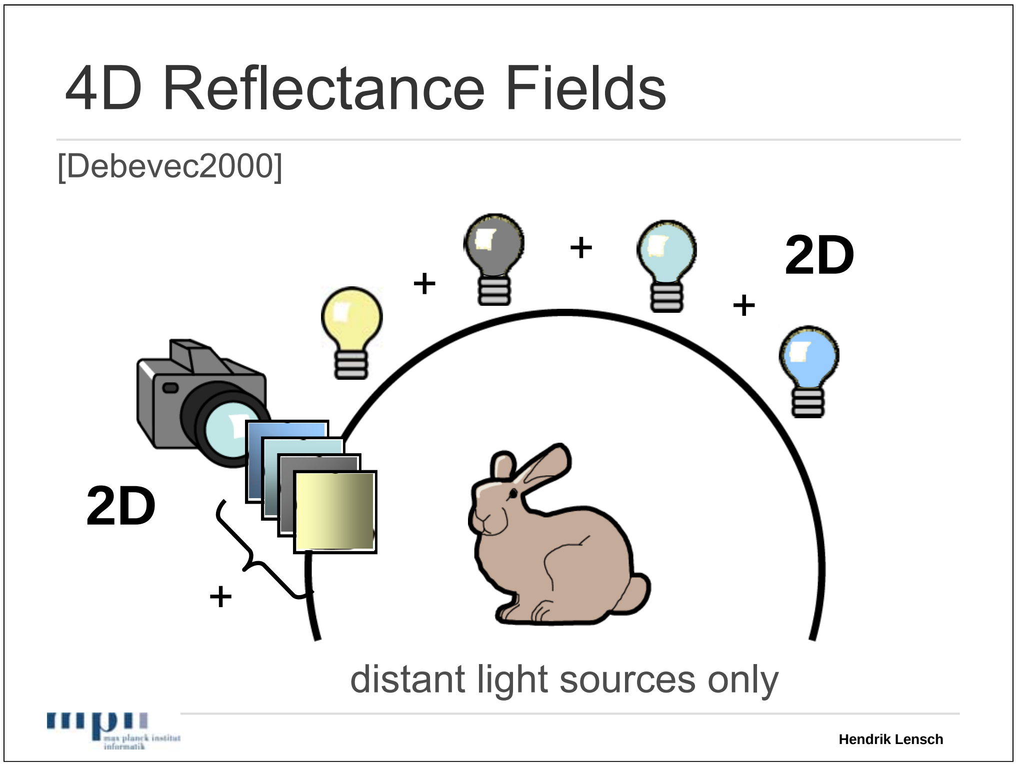

It is even possible to assign arbitrary weights to the virtual light sources by simply multiplying the individual images with a different color before adding up. This way, the appearance of the scene can be reproduced in arbitrary environments.

One restriction however is that the incident illumination is constrained to originate from the positions of the capturing light sources, for example from the sampling positions on a sphere. It is not possible to change the distance of the light source.

As all light sources are assumed to be directional light sources, reflectance fields captured this way cannot reproduce the appearance according to a spatially varying illumination pattern formed, for example, by a spot light or by projected shadows.

The direct reflection might be correctly reproduced, but the indirect reflections or subsurface scattering will be wrong.

Far and Near-Field Illumination

Hendrik Lensch

Far and Near-Field Illumination

Hendrik Lensch

This is demonstrated in these two slides. The projected light pattern alters the incident illumination for every point on the lamp shade. In the back on even sees how the incident light pattern is slightly blurred due to the scattering within the cloth. These effects cannot be reproduced with a far-field reflectance field.

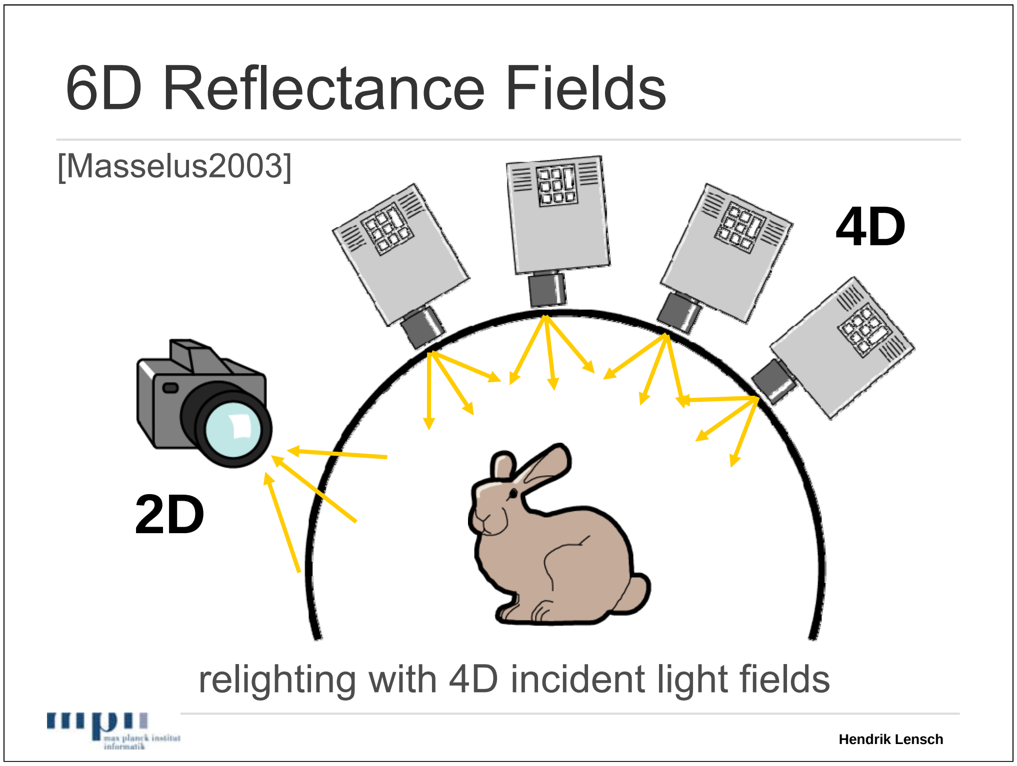

In order to cover global illumination effects correctly one needs to acquire a reflectance field for individual incident rays. Now, one image is captured for every possible incident light ray, instead of for each light position. The collection of incident rays form a 4D incident light field. Capturing one image per ray is of course expensive. Masselus et al. therefore only employed projectors with sixteen pixels. In the remainder of the course we will show how different illumination patterns can speed up this process.

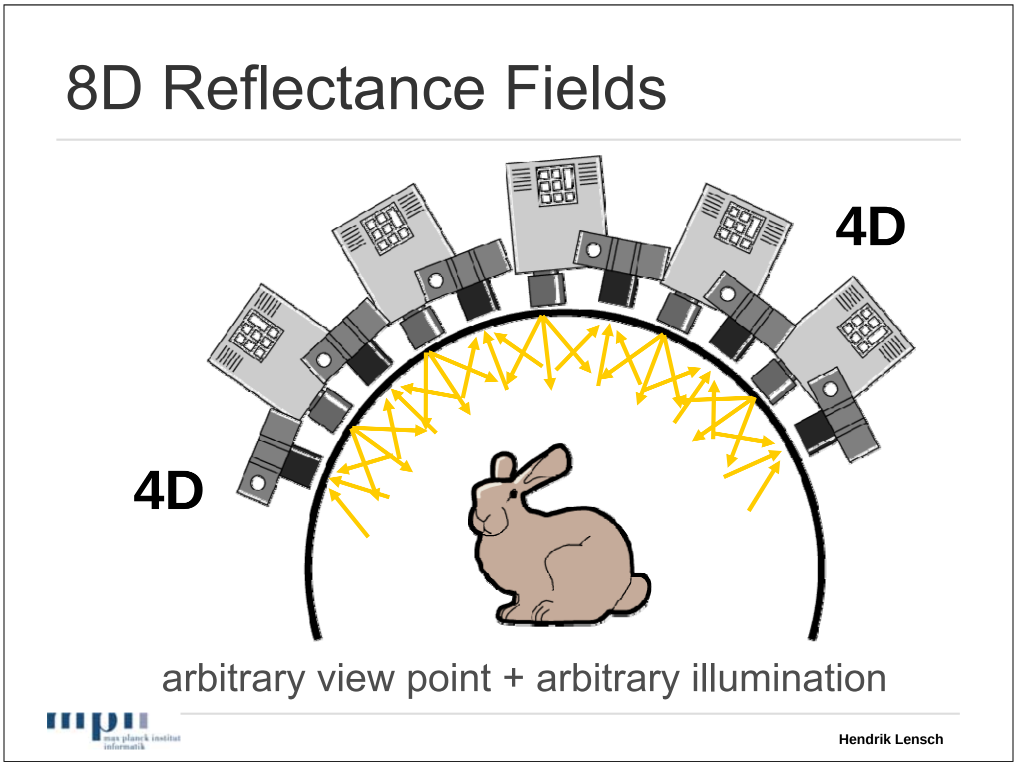

The Holy Grail of appearance acquisition is to capture a full 8D reflectance field where the look of the object is captured from all possible direction from every ray in the incident light field. One can think of a reflectance field as an operator transforming the incident light field into the reflected light field.

Main Problem

- sampling an ***8D function***
  - spending 100 samples/dimension
    → 10^{16} samples
  - hi-res 3D geometry: 10^8 vertices

- coherence in reflectance fields
  → reduced data complexity
- no complete solution yet

max planck institut informatik

Hendrik Lensch

Sampling, storing or rendering an 8D function is a tough problem. Approaches therefore often reduce the dimensionality of the problem, e.g., assuming a single view point, distant illumination, etc. In addition, using structured illumination, it is possible to capture multiple samples in the same shot.

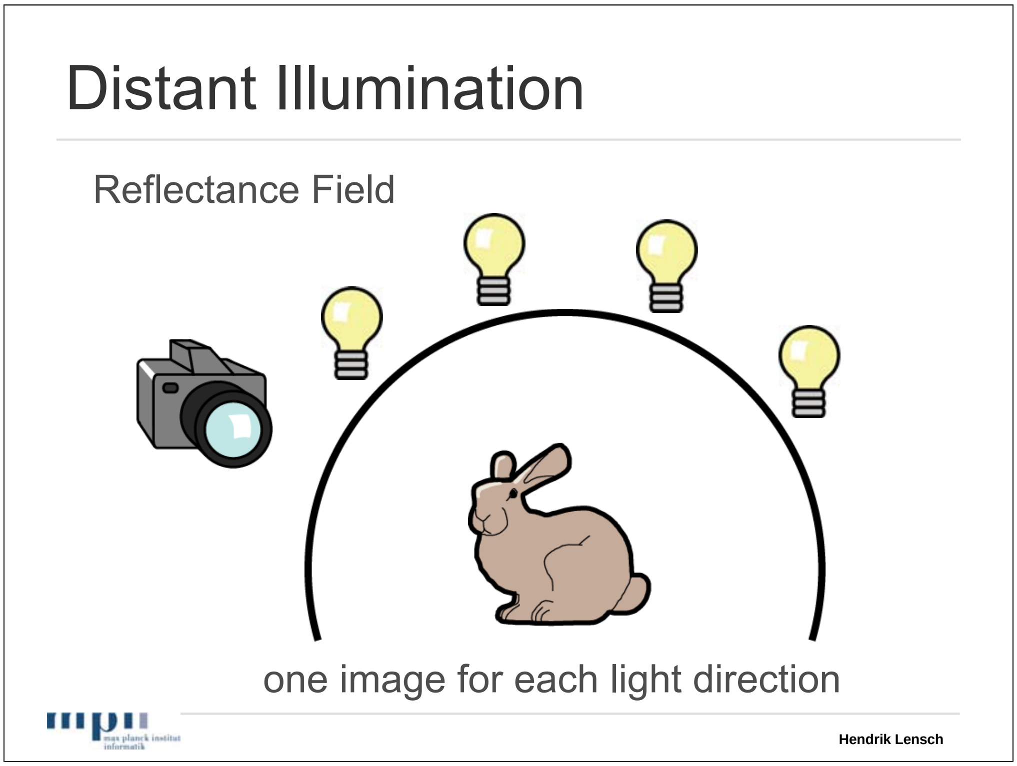

Let's start with a reflectance field for distant illumination.

4D Reflectance Field

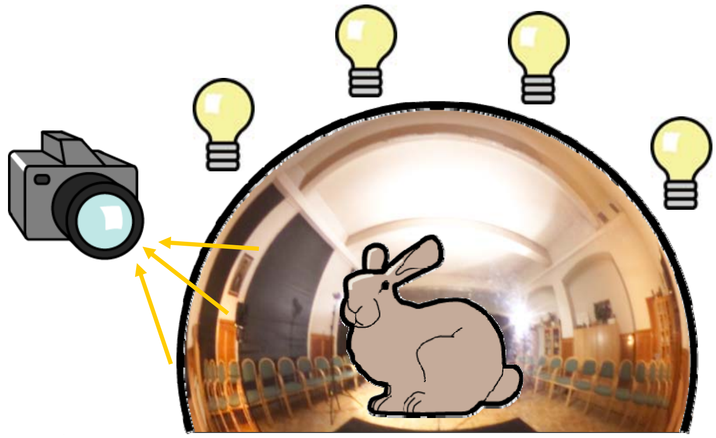

Hendrik Lensch

As already said, one can reproduce the appearance of an object in the illumination of an environment map, assuming infinitely far away light sources.

Light Stage

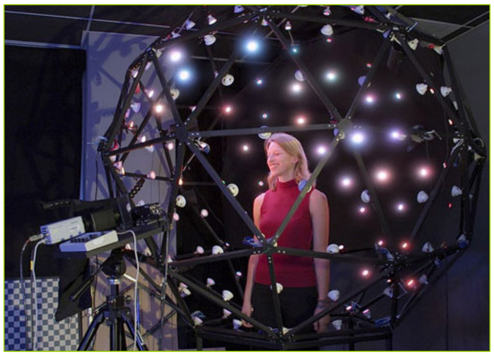

- [Debevec2000]
- single view point
- assumes distant light sources
- video

The light stage is a device that is able to capture this. A dome of switchable light sources light sources illuminates the actor in sequence. The video demonstrates how the collected images can then be combined to reproduce the appearance in arbitrary environments.

I would like to mention two properties of this particular setup:

-The light source images are captured one at a time, i.e. one image per light source. This corresponds to scanning through the space of light sources and is relatively slow. In this setup. it has been accelerated using a high-speed camera.
A positive point of this scanning illumination pattern is that the images can be directly used for rendering. In principle, no further analysis is necessary.

-The second issue is that the sphere of directions is relatively sparsely sampled from a set of fixed positions. As with any sampling process, one might observe sampling artifacts due to undersampling

4D Reflectance Field

typical: a sparse sampling of light directions

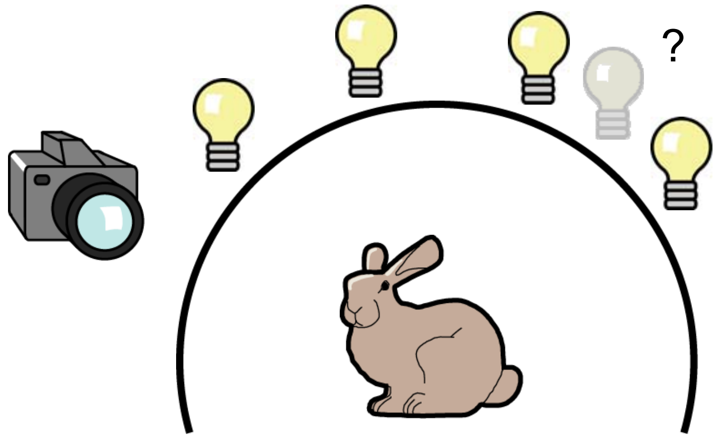

Hendrik Lensch

Generating the appearance for a non-captured light source position therefore requires a modified acquisition system or additional processing in the form of light source interpolation.

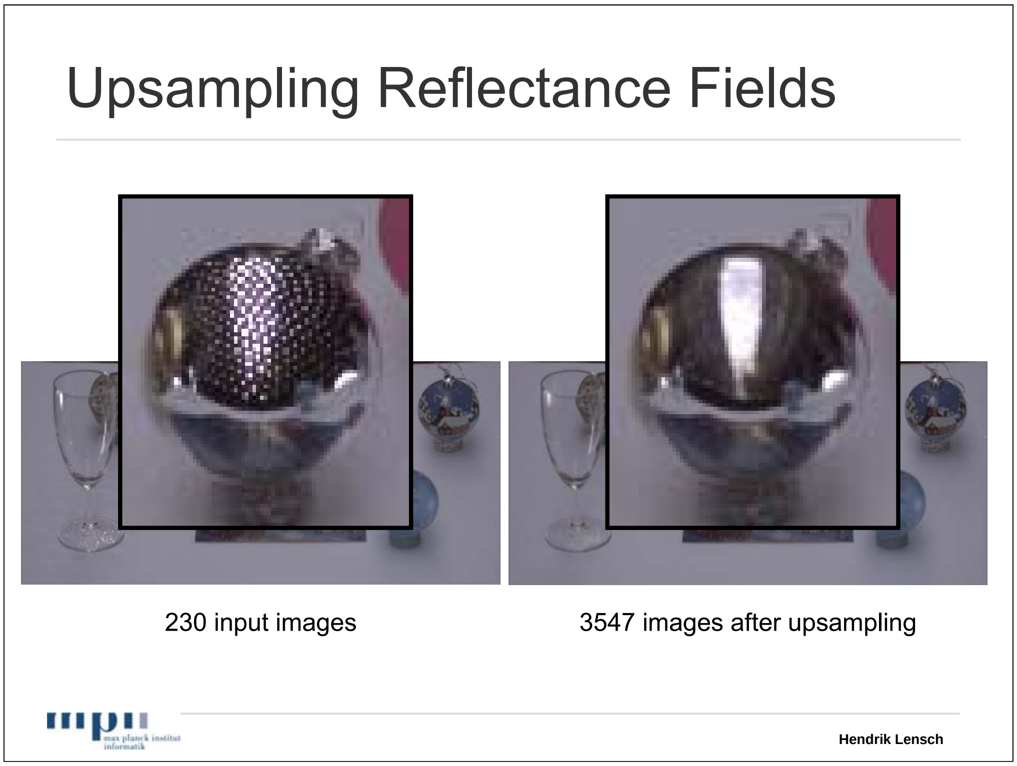

Too few directional samples lead to artifacts most noticeable at high frequency effects such as specular highlights or shadows. Employing a non-linear upsampling scheme, it is possible to obtain a much clearer reflection of the environment in the sphere.

4D Reflectance Fields

no problem at glossy surfaces

max planck institut informatik

Hendrik Lensch

The problem gets clearer when taking a close up view on the scene: For a point on a glossy surface the reflectance function, the dependence on the incident light ray direction (theta,phi) is rather smooth. A coarse sampling does not lead to any problem.

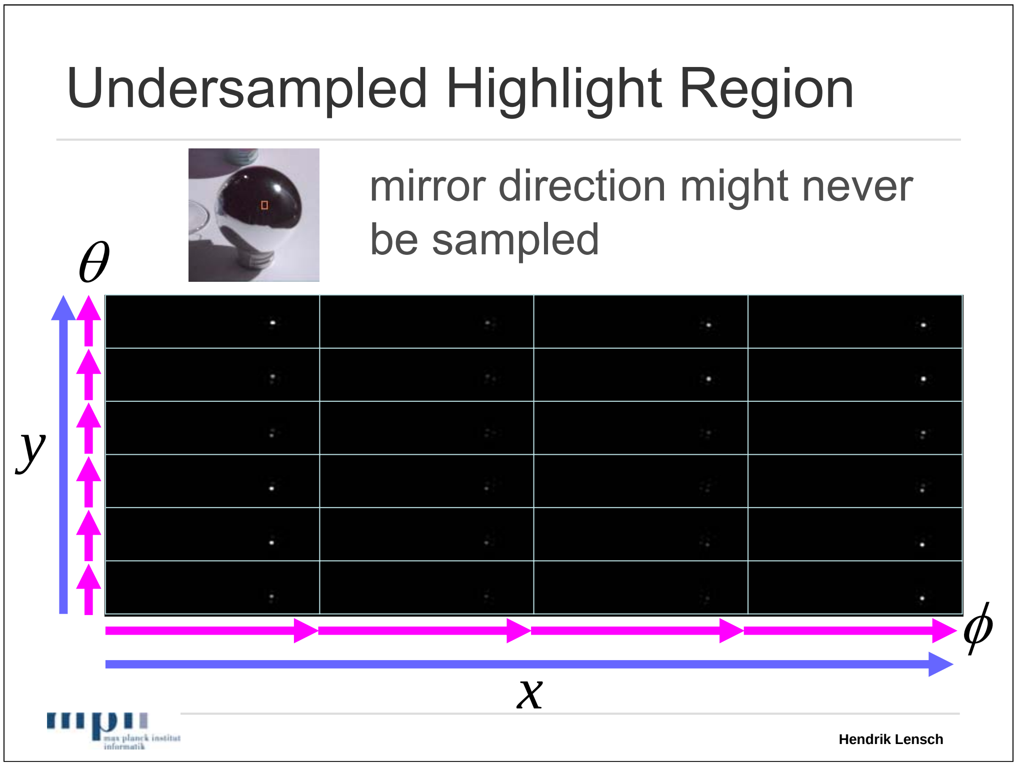

On a mirroring sphere however, the reflection function changes drastically with the incident or the viewing angle. In a neighborhood of pixels on a sphere, it might be that only for a very view pixels, the coarse sampling locations actually contain a mirror direction. A highlight generated by a moving light source would generate an incorrect intensity variation.

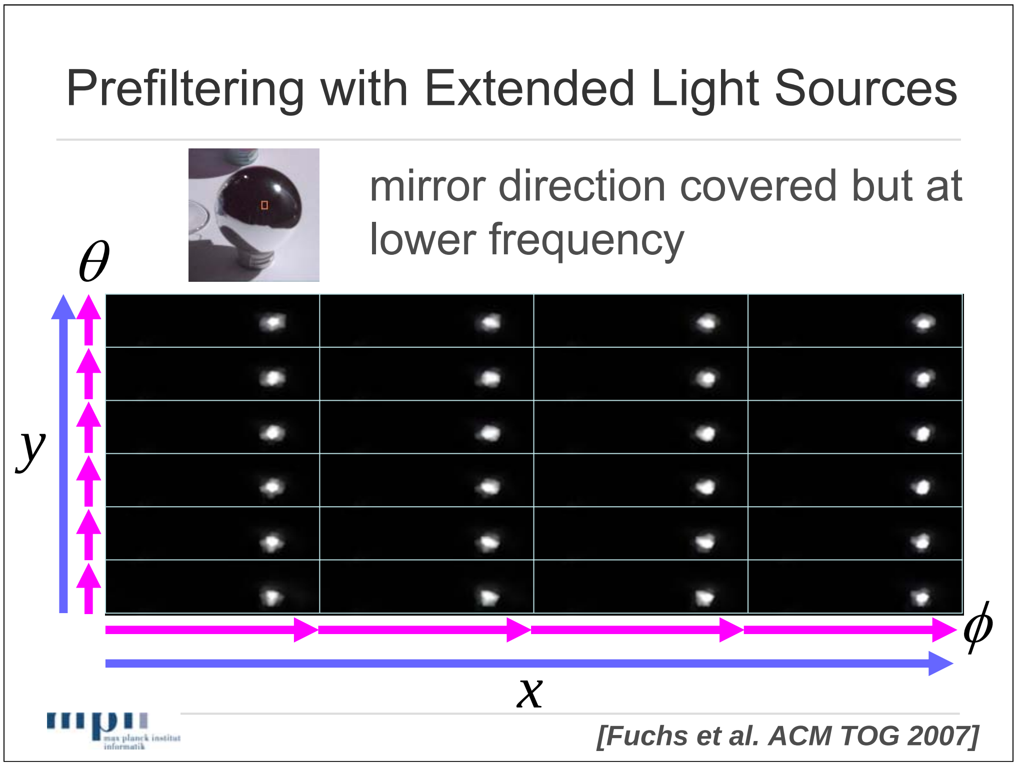

One can mitigate this problem by prefiltering the incident illumination, limiting the maximum frequency in the environment map. This has the effect the acquisition is done with extended light sources rather than point light sources. For each surface point a highlight will be captured.

Prefiltering with Extended Light Sources

[Fuchs et al. ACM TOG 2007]

This can for example be achieved using an indirect light stage where disco spot lights project an extended spot onto the wall of a tent indirectly illuminating the scene in the middle.

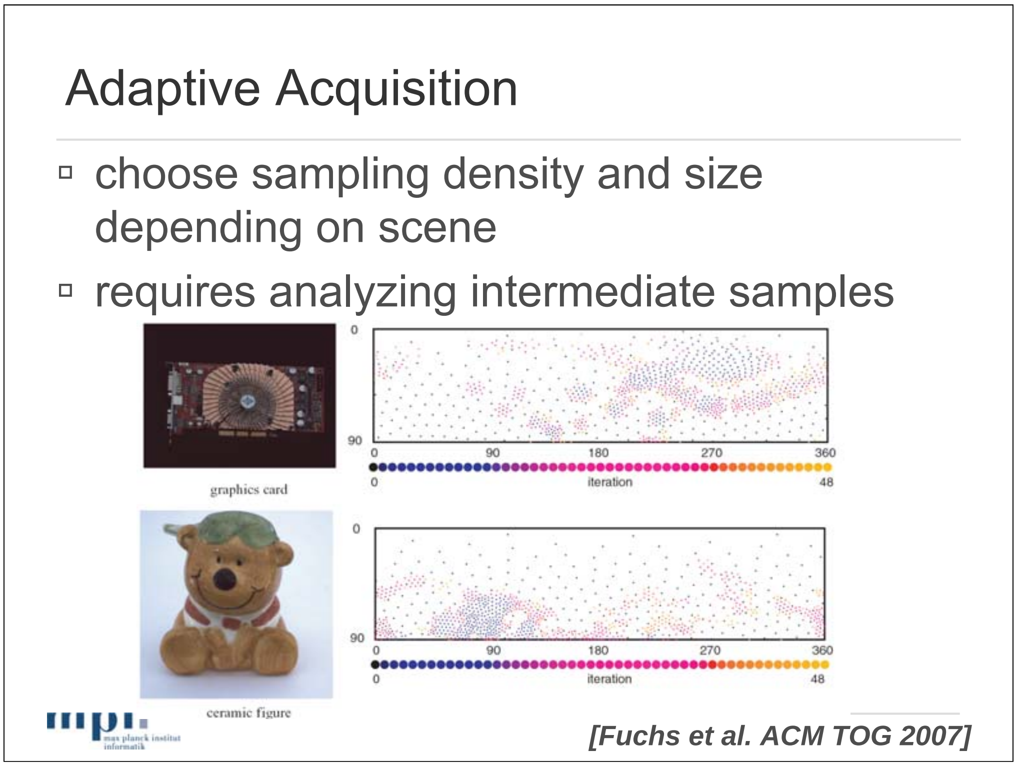

This can for example be achieved using an indirect light stage where disco spot lights project an extended spot onto the wall of a tent indirectly illuminating the scene in the middle. With this setup, one can even run an adaptive sampling pattern where the size of the light sources and the sampling density is adapted to meet different sampling requirements for different section of the illuminating sphere.

This adaptive scheme provides a better sampling only at places where required. Compared to a sampling at full resolution a lot of sampling position can be well approximated from a sparser sampling. However, this adaptivity comes at the cost of a more complex acquisition setup, where the size and the location of the light sources need to be controllable. In addition, after each acquisition step the captured images need to be analyzed in order to predict, which samples to acquire next.

Subdivision Hierarchy

- take few input pictures
- analyze the reflectance field
- perform pair-wise upsampling, subdividing the mesh of light directions

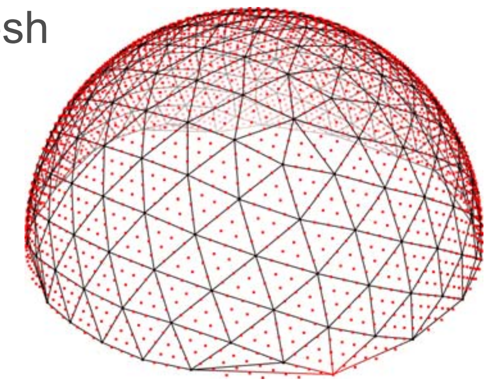

Hendrik Lensch

Another way is to upsample the acquired data by performing non-linear interpolation between the originally captured samples.

Processing Pipeline

- separate in different effects:
  - highlights
  - shadows
  - low frequency effects
- separate upsampling
- refine result with texture priors

We apply different upsampling schemes for highlights, shadows and low frequency effects.

Result Videos

- continuous highlights
- smooth movement of reflections

linear

superresolution

[Fuchs et al. EG 2007]

The benefit of this approach is that from the same set of input images much more faithful renderings can be produced.

The environment is nicely reflected in the silver sphere. Shadows, highlights and even caustics move quite smoothly

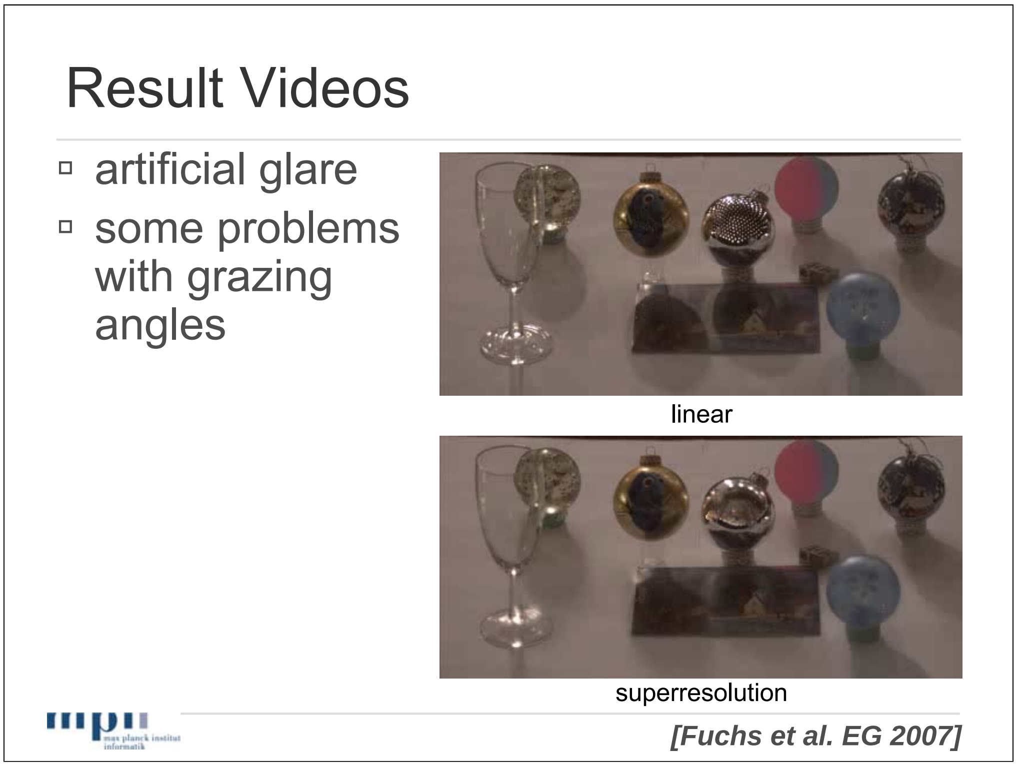

Here the performance on shadows is demonstrated.

Note the correct effect of refraction in the glass. Some artifacts remain though.

The choice of sampling pattern for the incident illumination depends on the scene properties.

Prefiltering, adaptive sampling or an advanced interpolation scheme can significantly improve over a sparse sampling.

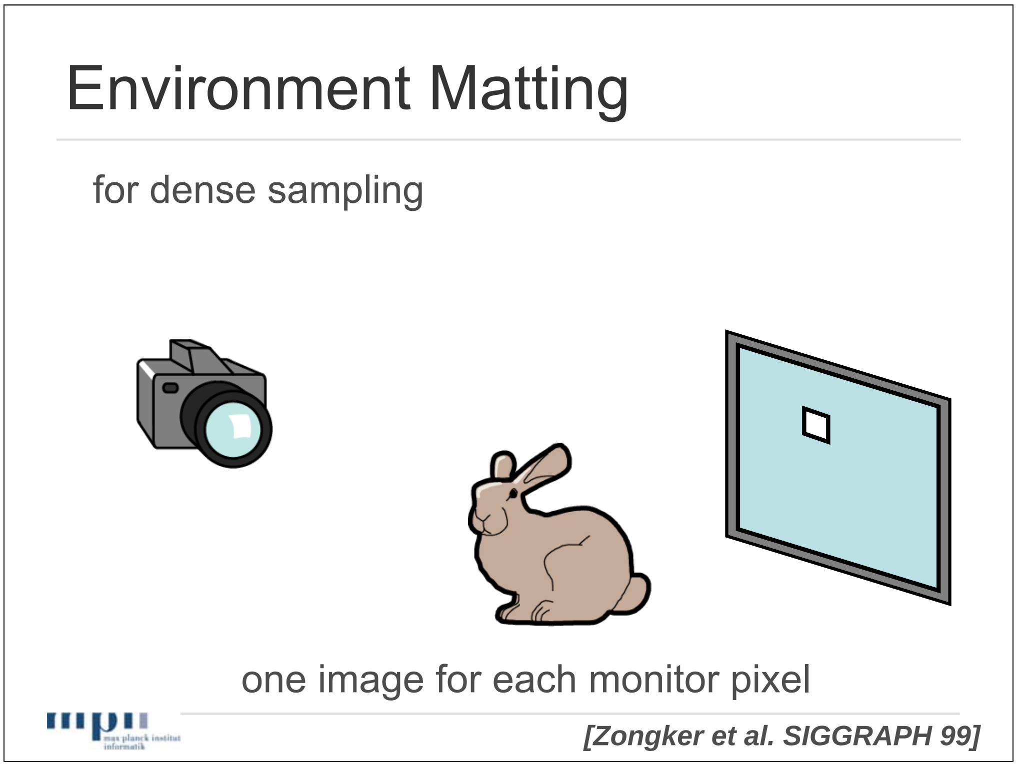

Environment mattes are designed to solve this task by densely sampling the light source positions. Especially geared to capture specular reflections or refractions.

In principle, one could scan through all monitor pixels, but this would amount to 1 million images.

Environment Matting

- Extension of Alpha Matting capable of capturing transparent and specular objects for one view.
- Allows for reproduction with arbitrary backdrops.
- A high-resolution 4D reflectance field.

[Zongker et al. SIGGRAPH 99]

Traditional Alpha Matte

$$C = F + (1 - \alpha)B$$

- Composite color C
- Foreground color F
- Background color B
- Pixel coverage α

- Acquired by blue/green screening
- missing: dependence on light direction

max planck institut informatik

Hendrik Lensch

Environment Matting - Definition

- Add reflected and refracted rays (sum over backdrops)

$$C = F + (1-\alpha)B + \sum_{i=1}^{m} R_i M(T_i, A_i)$$

- Reflectance R
- Texture T
- Axis-aligned area A
- Averaging operator $M(T, A)$

max planck institut informatik

Hendrik Lensch

Environment matting extends traditional blue screen matting to incorporate a directionally dependent part. In this model, it is expressed as the dependence on a rectangular patch of the environment M(T,A).

In order to acquire an environment matte one needs to determine the size and the location of the region of the backdrop/sides that influences each camera pixel.

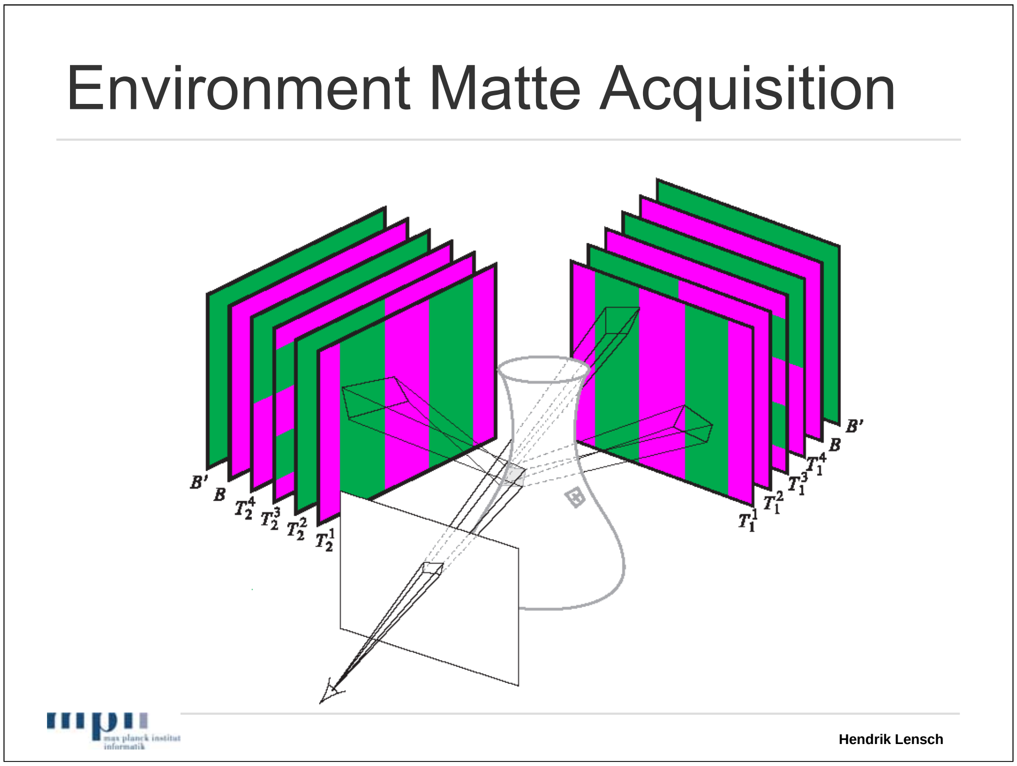

The process is visualized here. On each monitor side a set of hierarchical stripe patterns is displayed. From the recorded images for each individual displayed pattern the algorithm determines the direction of the sub-cone that contributes to the pixel's color.

In addition, the algorithms determines the spread of the cone.

Compared to scanning with individual pixels this hierarchical patterns are less precise as only an approximation of the actual beam is estimated. However, the acquisition time is significantly reduced.

Environment Matting Acquisition

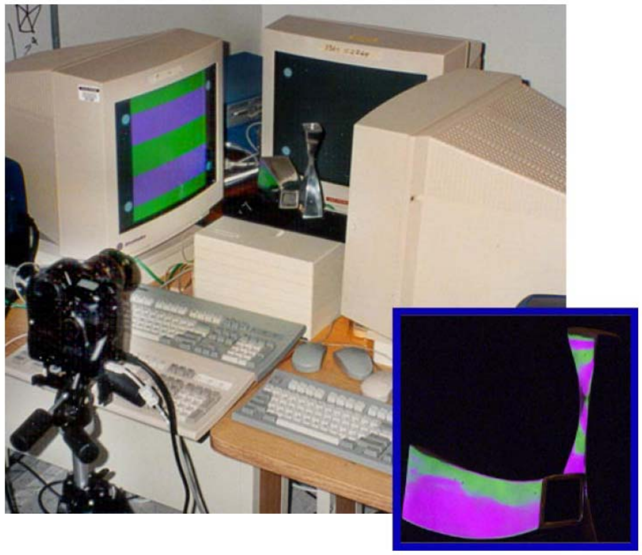

mpi max planck institut informatik

Hendrik Lensch

Environment Matte Extensions

- Gaussian Filter kernel [Chuang et al. 2000]
- Real-time acquisition [Chuang et al. 2000]
- Wavelets in acquisition [Peers et al. 2003]
- Multiple View Points (Opacity hulls) [Matusik 2002]

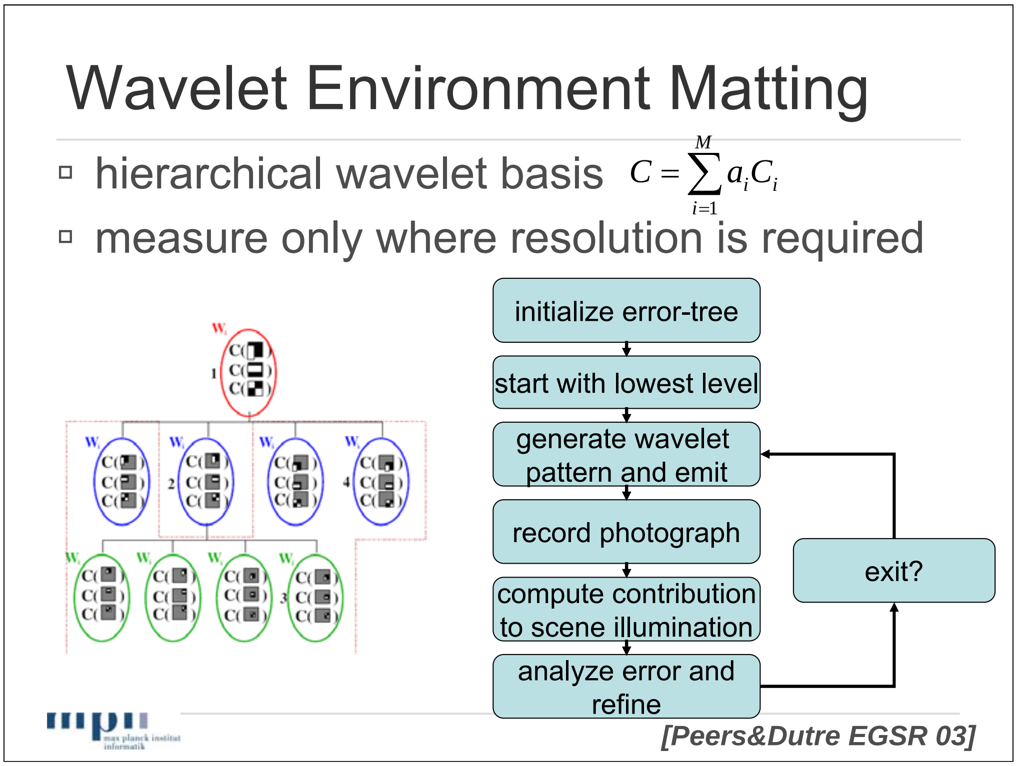

Peers and Dutre suggested an adaptive approach based on wavelets. The illumination domain is only refined if there are pixels whose reflection show high frequency responses in this region. For rather diffuse reflection, sub-division is not required and a lot of images can be saved. On the other hand, very sharp reflections can be correctly reproduced as well.

In principle, the output quality is as good as with a scanning approach, the number of images required is however drastically reduced. As can be seen in the flow diagram. the adaptive process requires to interleave pattern generation, acquisition and analysis.

Wavelet Noise

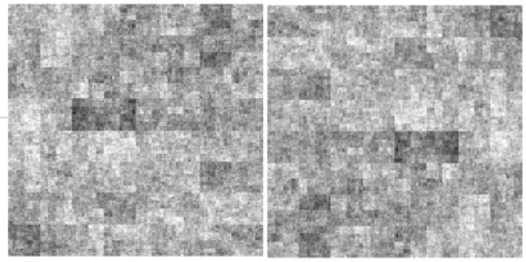

- randomly mix all wavelet bases for illumination (generate about 600 patterns)
- infer per-pixel reflection $T(x, y)$
- measurements $C_i(x, y) = T(x, y) W_i$
- minimize $E = \sum_i \left| C_i(x, y) - T(x, y) W_i \right|_2$ subject to a sparse $T(x, y)$
- (compare to compressed sampling)

[Peers&Dutre EGSR 05]

In their wavelet noise approach Peers and Dutre no longer illuminate with individual wavelet basis but rather combine a random collection (with random weights) together in order to form a fixed set of illumination patterns.

From this set of patterns the reflection functions now can only be inferred. They are no longer measured directly but need to be determined based on an optimization process. In this particular case the reflectance function of each pixel is hierarchically estimated minimizing the error between the current prediction and the measured samples while keeping the structure of the wavelet tree as simple as possible.

This process bears some resemblance with compressed sampling approaches.

The benefit of this technique is, that one obtains a hierarchical basis representation from a fixed set of patterns. No explicit sampling or adaptation is necessary. The acquisition is significantly simplified but the analysis at the end is rather complicated.

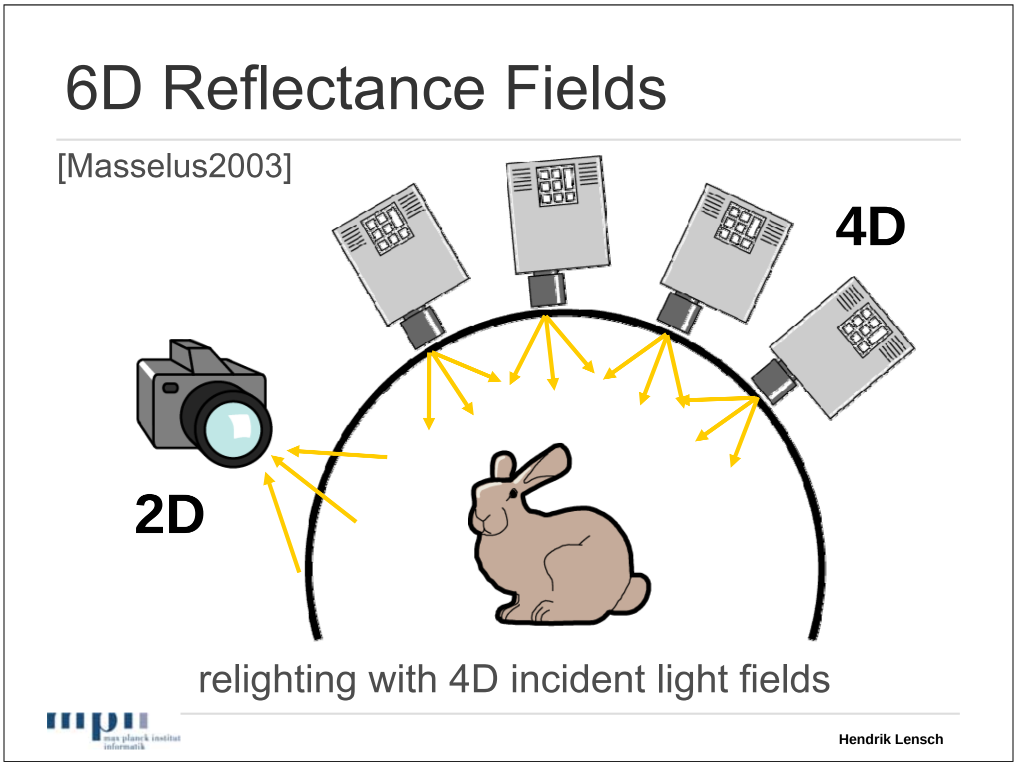

Now let’s look at illumination pattern for a different type of reflectance fields where the illumination is recorded not for individual light source positions but rather for individual pairs of rays.

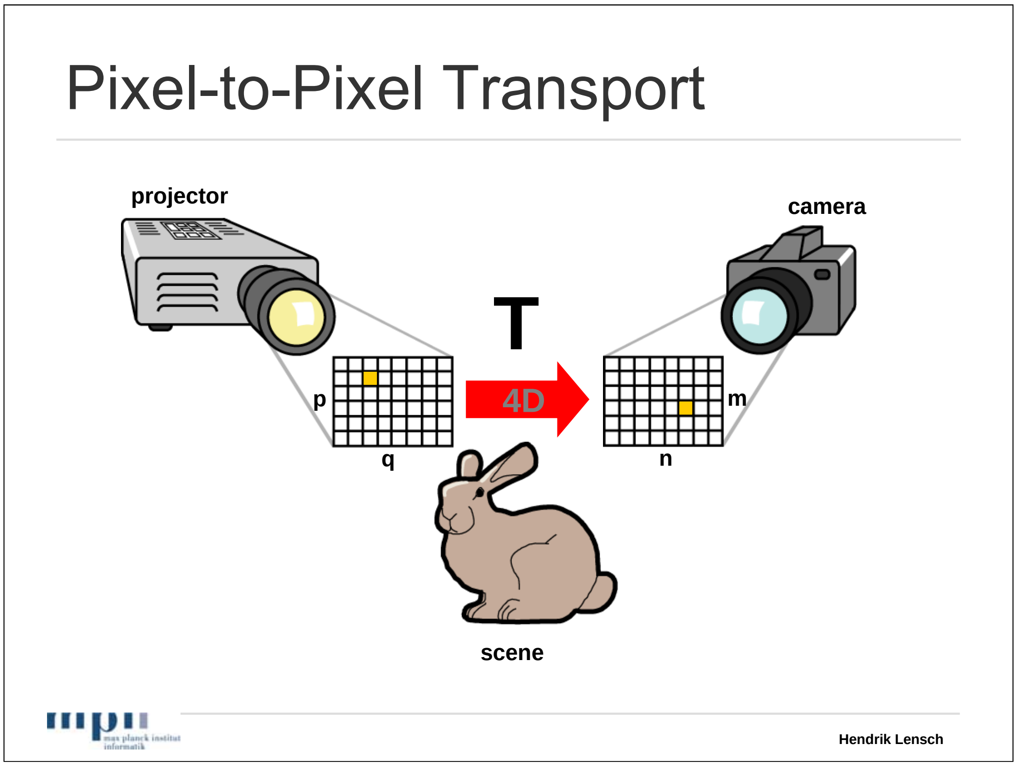

Focusing on a single camera projector pair we need to determine for each projector pixel (p,q) what is the resulting camera image. The fourth-order tensor T stores the reflection coefficient for every pair of projector/camera pixel.

It fully describes the scenes reflection properties.

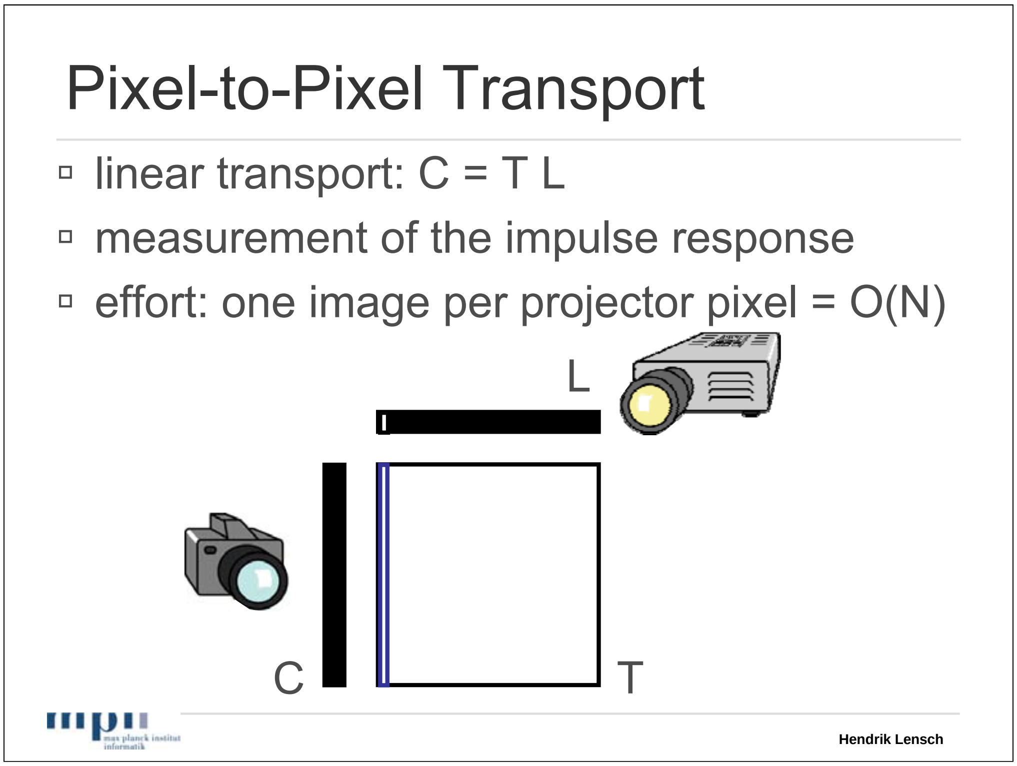

This matrix can again be acquired by scanning, i.e. turning on each projector pixel individually, recording one image per projector pixel.

This is of course extremely expensive, but doable.

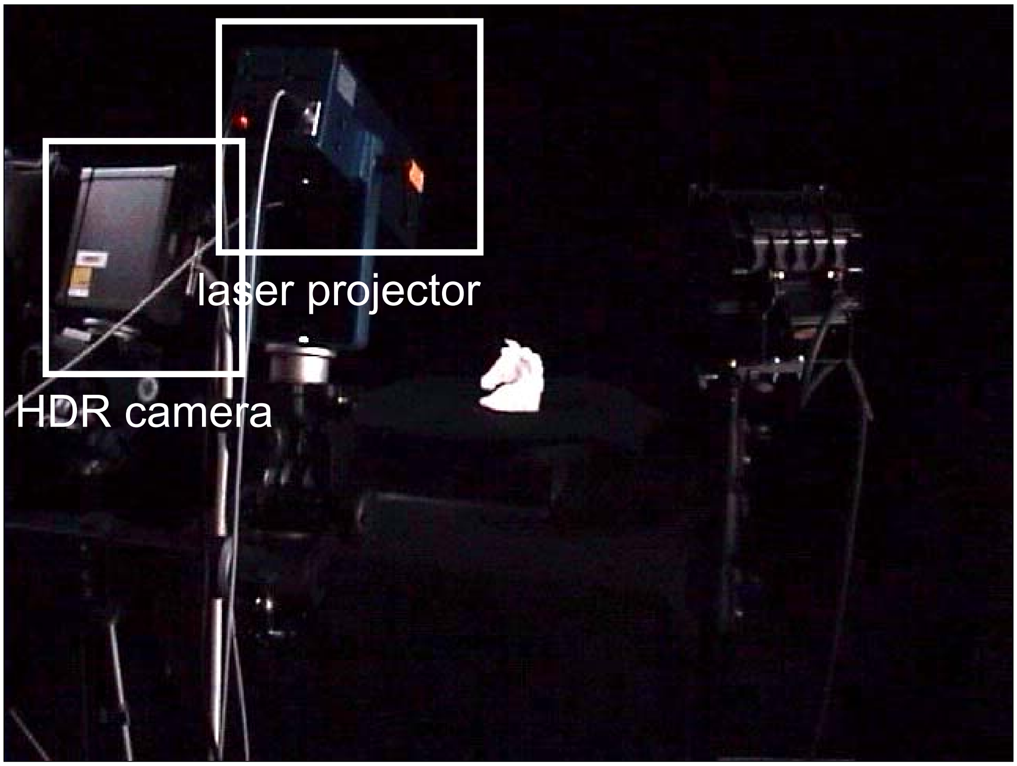

Here, one sees a corresponding measurement setup. Since a single projector pixel is not very bright compared to the projectors black level, we apply a laser projector with galvanic mirrors. The laser projector provides a significantly better contrast compared to video projectors.

The laser beam sweeps over the surface and a HDR video camera records every sample.

When the laser beam hits a translucent surface, the area surrounding the laser point lights up due to subsurface scattering. The footprint varies drastically with the incident location.

Video

1.000.000 images, 22 hours → model - 800MB

[Goesele et al. SIGGRAPH 2004]

After one day and after capturing about one million images the entire matrix has been acquired. One can apply some compression for interactive rendering. Hole-filling is applied.

The data set now allows for relighting with arbitrary light patterns from arbitrary directions. The very specific properties of the light transport in translucent objects is correctly captured and reproduced.

One sees for example the drastic difference between front and back illumination as well as the light bleeding into shadowed regions.

Adaptive Parallel Acquisition

- assumption: sparse matrix,hierarchical basis
- radiometrically independent blocks can be sensed in parallel

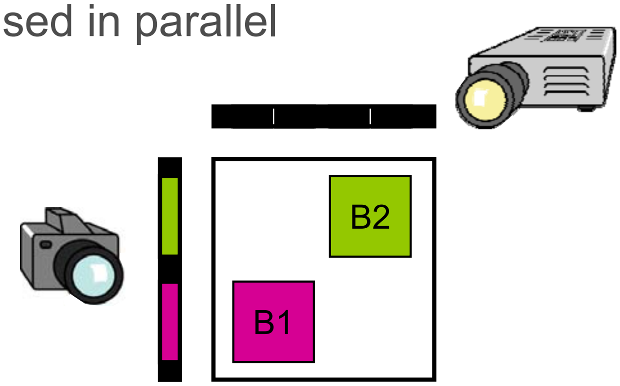

Adaptive Parallel Acquisition

- assumption: sparse matrix,hierarchical basis
- radiometrically independent blocks can be sensed in parallel

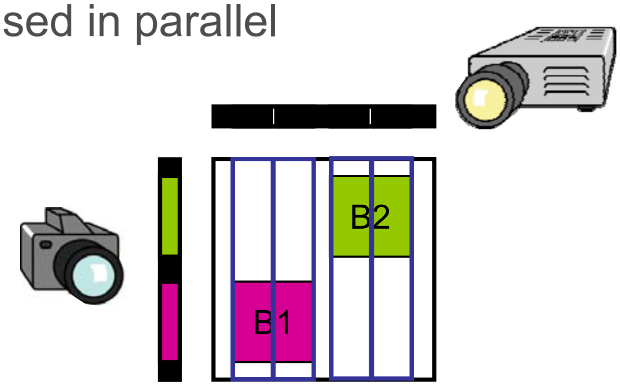

Hendrik Lensch

We now apply an adaptive scheme to this problem, again using a hierarchical wavelet bases, or equivalent hierarchically nested blocks.

In the case that the transport matrix is only sparsely populated, we can even parallelize the acquisition of multiple blocks. As long as they influence well separated camera regions only, they can be illuminated at the same time, and we can afterwards determine, based on the pixel location, which block in the reflectance tensor has been measured. The parallelization drastically reduces the acquisition time compared to a sequential adaptive scheme.

Adaptive Parallel Acquisition

parallelized acquisition of regions which do not overlap in the camera image

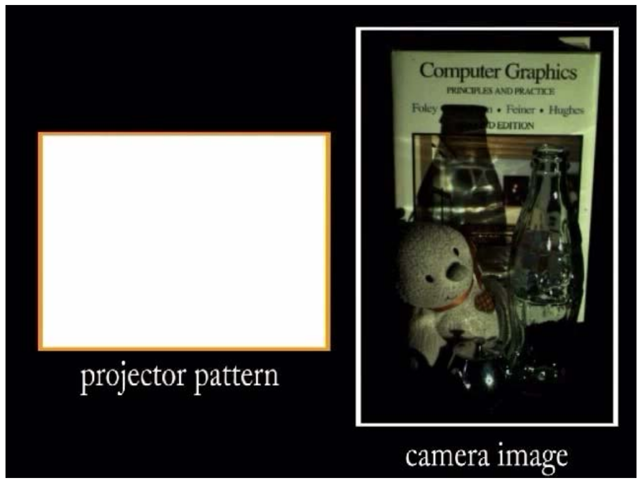

Here you can see the acquisition sequence for a quite complex scene. Initially the subdivided regions are sequentially acquired. At some point the algorithm detects that the corresponding footprints (areas of influence) in the camera image are well separated. All subsequent measurements of radiometrically independent blocks can further be captured in parallel. This process starts at some specific level of subdivision.

Adaptive Parallel Acquisition

parallelized acquisition of regions which do not overlap in the camera image O(log N)

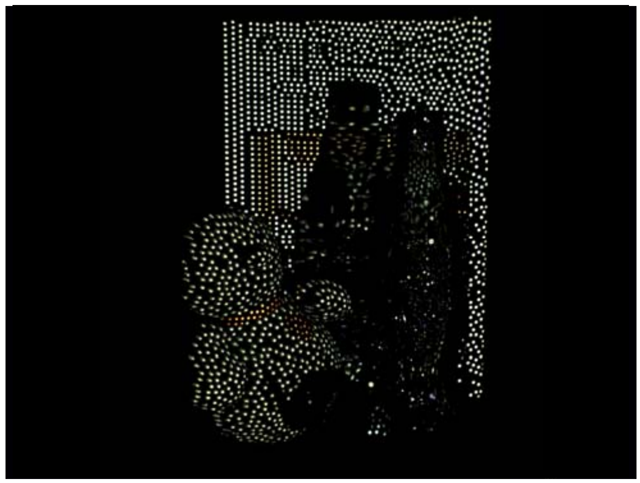

At the pixel-level, lots of samples can be acquired in parallel, especially in areas where direct reflections dominate the light transport. If the light transport is more complicated, as for example within the bottle, where refraction spreads the extent of the incident light rays, less samples can be captured at the same time.

Still the acceleration is dramatic. For most scenes, the acquisition time is in the order of O(log N) for N projector pixels

Relighting with Arbitrary Patterns

1.200 images, 2 hours → model - 220MB

The recorded reflectance field can now be relit with arbitrary illumination patterns. The tensor captured all light transport paths: one can see direct diffuse and specular reflections, interreflections. The glass bottle furthermore features refraction and caustics.

Global Light Transport

Hendrik Lensch

In the Cornell Box we demonstrate even diffuse color bleeding and mirror reflections.

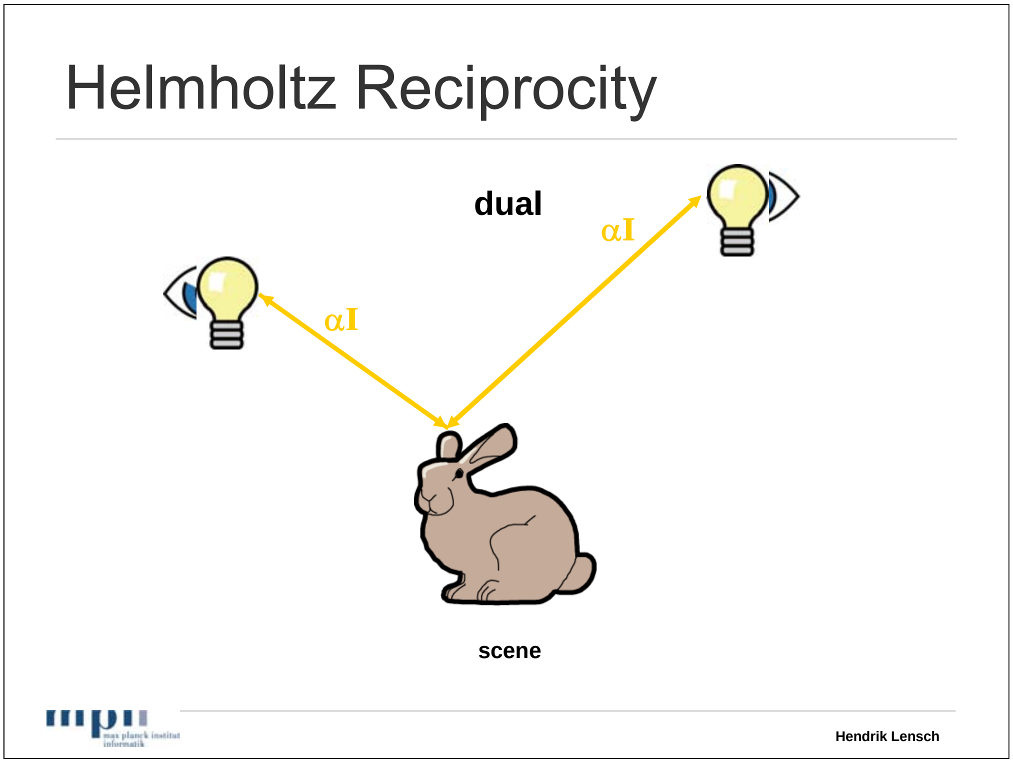

Given the reflectance measured on a ray-to-ray basis, i.e. along one path, we can apply Helmholtz reciprocity.

It states that the reflectance along a path is the same no matter in which direction the path is followed. We can follow the light from the source to the receiver or we can swap their location and follow exactly the same path in the other direction, the observed reflectance coefficient will be exactly the same.

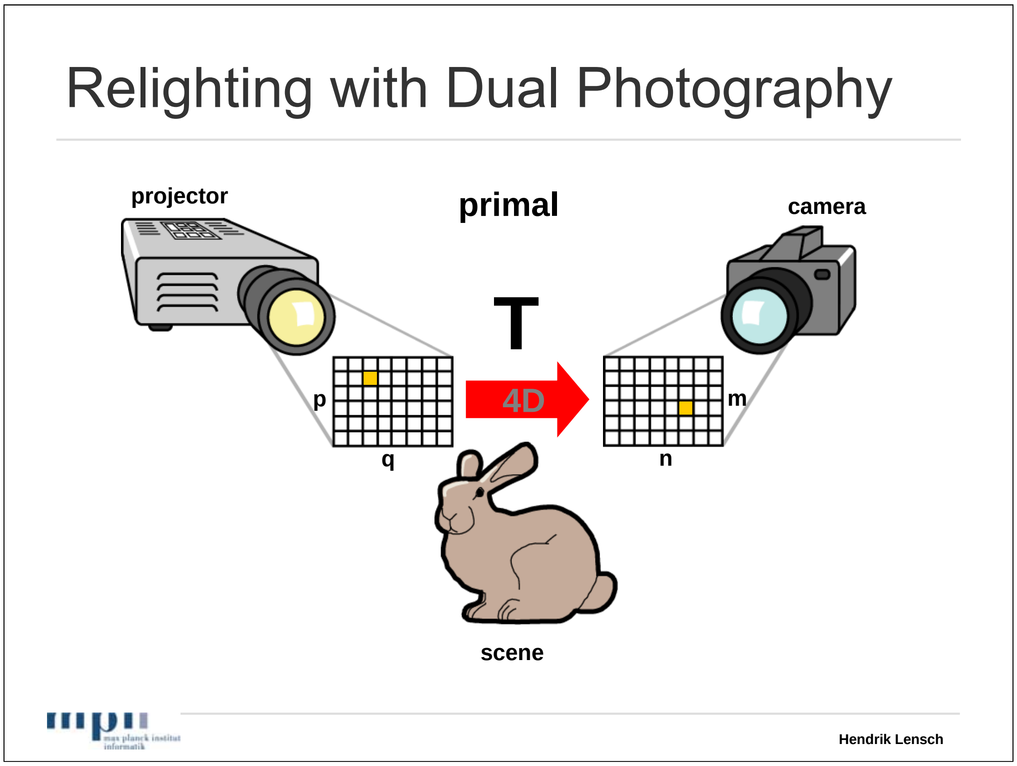

We can apply Helmholtz reciprocity as the tensor T captures the reflectance for any path between the projector and the camera. Helmholtz reciprocity holds for any collection of paths as long as the light source and the camera are within the same optical medium.

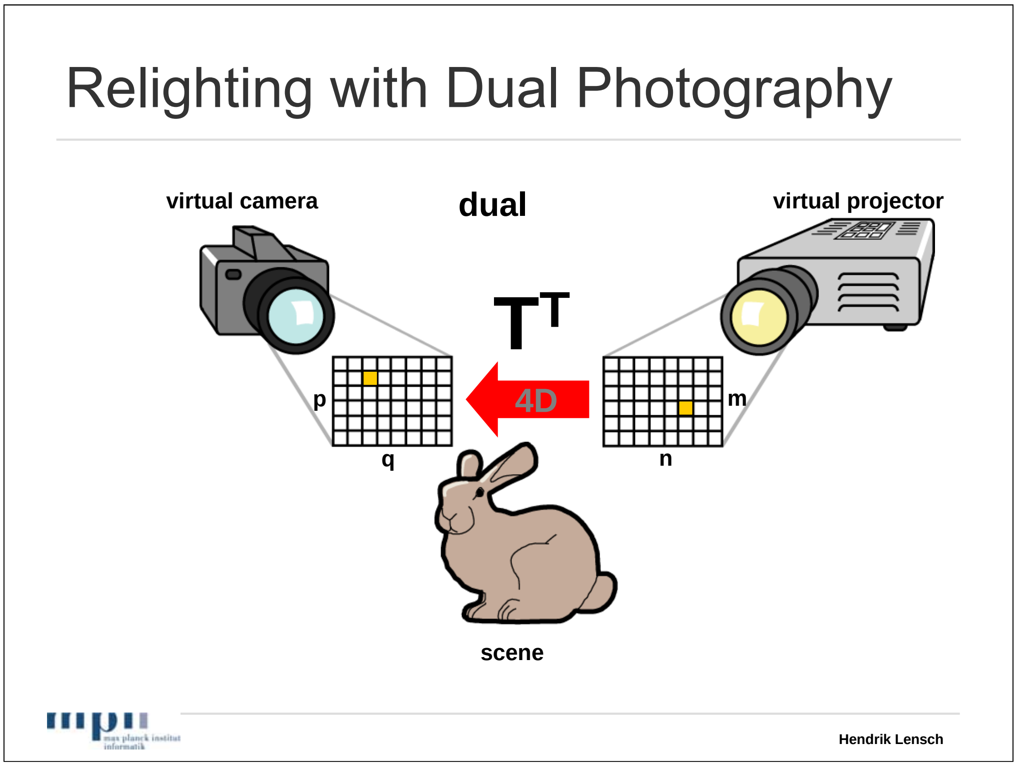

We can now swap the role of the camera and the projector, turning the original projector into a virtual camera, the original camera into a virtual projector. Given T one can easily compute an image from the point of view of the original projector as if illuminated from the camera. Here, one can even apply arbitrary patterns from the virtual camera. The process is surprisingly simple as the reflectance coefficients stay exactly the same. They just need to be reordered in order to account for the novel ray configuration. This can be done by simply transposing the captured tensor T.

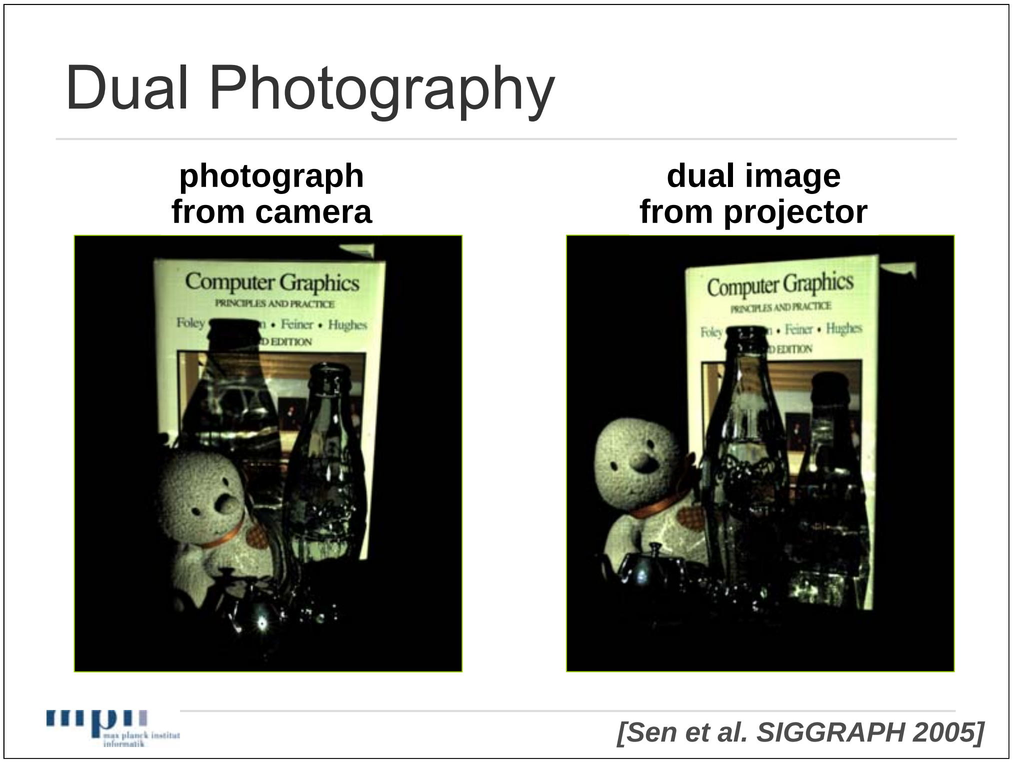

We can now generate different views from the same transport tensor, one from the original camera, illuminated from the original projector, and one from the projector illuminated from the camera.

All light transport paths are correctly incorporated, leading to a faithful reproduction of the bottles refractions and the caustics of the teapot.

Due to the duality, objects and shadows typically change place in the dual image.

Famous Examples

primal

dual

Hendrik Lensch

Another set of rather famous objects.

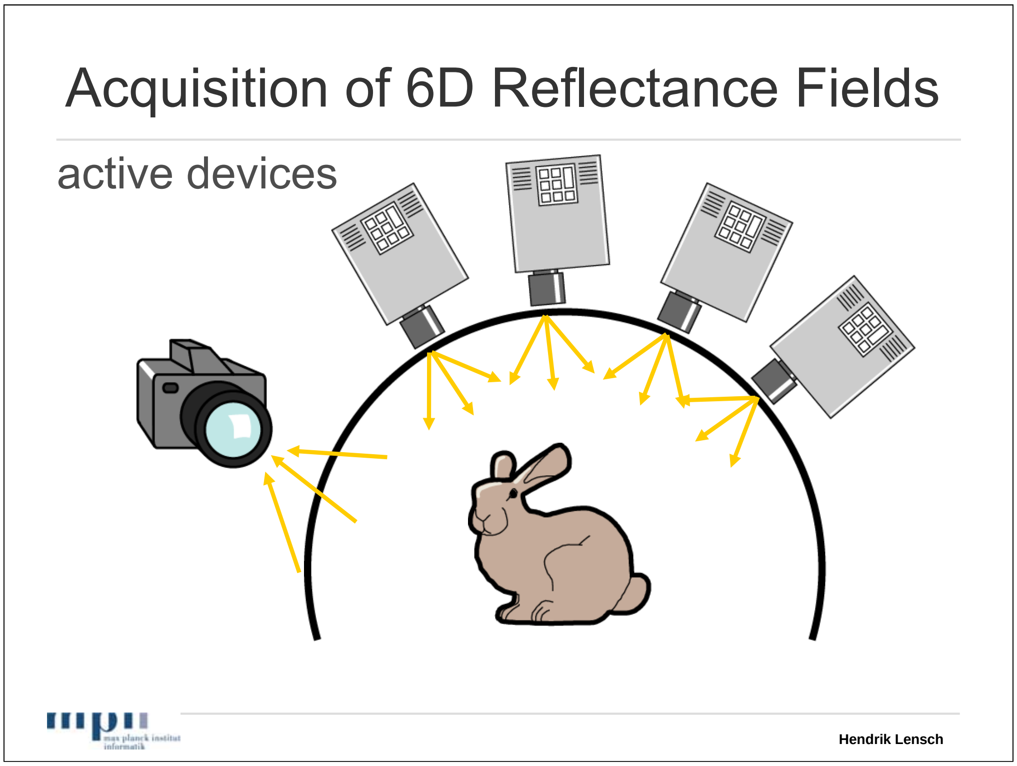

Besides being a nice ‘trick’, Dual photography can be used to accelerate the acquisition of higher order reflectance fields.

In this configuration, a set of projectors produce an incident 4D light field. In order to measure the corresponding reflectance field, the reflectance field of each projector needs to be captured in sequence. Otherwise the illumination patterns of multiple projectors would overlap. It would be rather difficult to disentangle the contributions of the individual projectors.

Dual Acquisition Process

parallel acquisition by passive devices

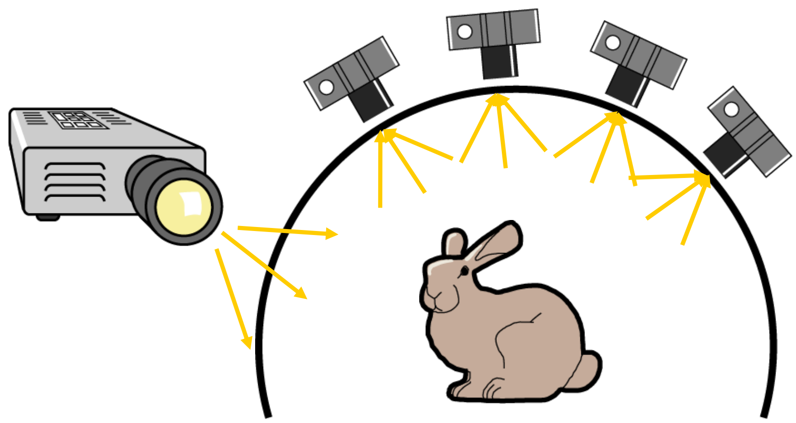

Hendrik Lensch

With dual photography we can swap the role of cameras and projectors both during acquisition and during rendering. Since cameras are passive devices they can easily operate in parallel without interfering with each other.

This way, the acquisition process of a 6D reflectance field for relighting with 4D incident light fields can be accelerated to the time cost of acquiring the reflectance field between a single camera/projector pair.

Smooth Interpolation

100.000 images, 12 hours → model - 4.5GB

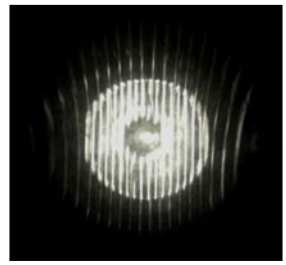

[Chen&Lensch VMV2005]

Thus it is possible, to acquire a rather dense reflectance fields on a ray-to-ray basis in significantly reduced time. It allows for relighting with arbitrary patterns from arbitrary directions.

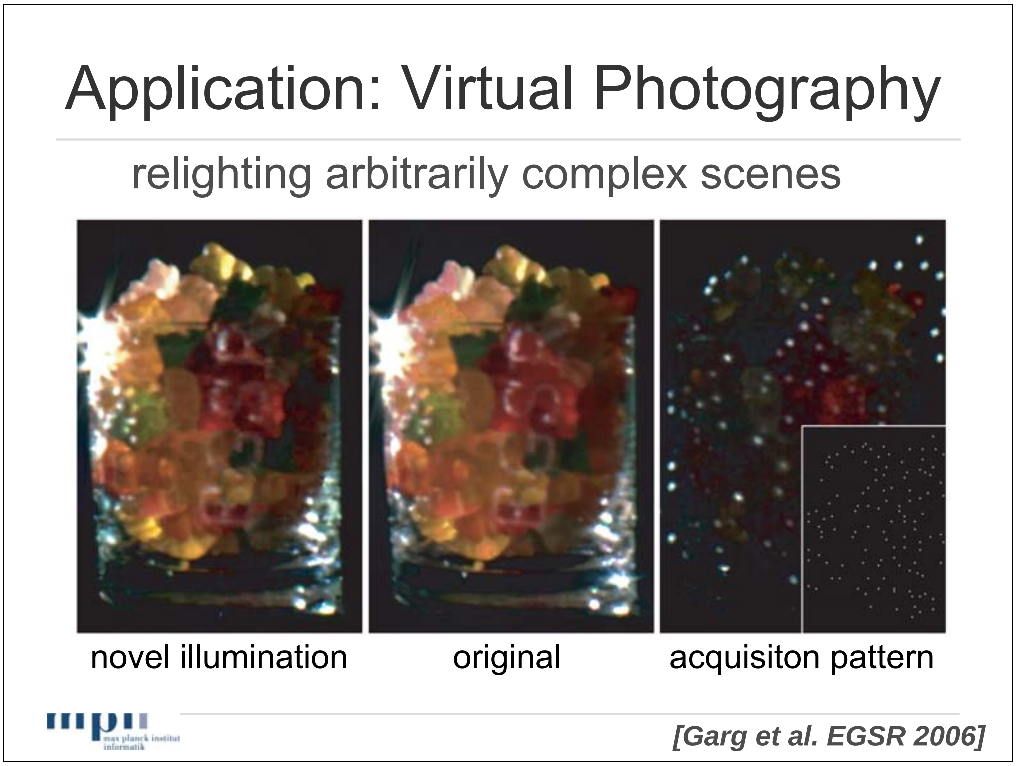

Slightly extending the previous adaptive parallel acquisition approach the reflectance fields of arbitrarily complex scenes can be acquired at moderate time cost. Here you see the complex light pattern formed in a glass of gummy bears illuminated from two different directions with a high frequency pattern.

One can acquire reflectance field that reproduce correctly all local and global illumination effects.

Local vs. Global Reflections

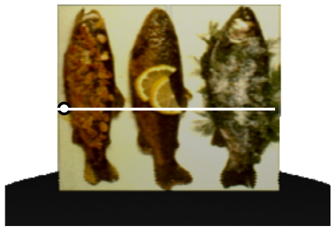

Hendrik Lensch

Looking at the structure of the resulting light transport tensors the difference between global and local illumination effects are clearly revealed.

In this case, we just concentrate on the slice which is formed by illuminating every point along the white line and looking at how much light is reflected from all points along the line.

In the left case, the reflections are due to the direct reflection off a planar, textured sheet of paper.

As one can see, light will only be reflected at the place where the sheet is illuminated, leading to the (band) diagonal structure of the matrix. The tensor does not contain any contribution in off-diagonal elements.

In the right case, the light transport between the two pages of an open book are captured.

The transport is still dominated by the direct reflection, i.e. the entries along the diagonal. However, there is additional indirect light transport from the left to the right side of the book. as well as tertiary reflections. Those populate the off-diagonal elements in the matrix.

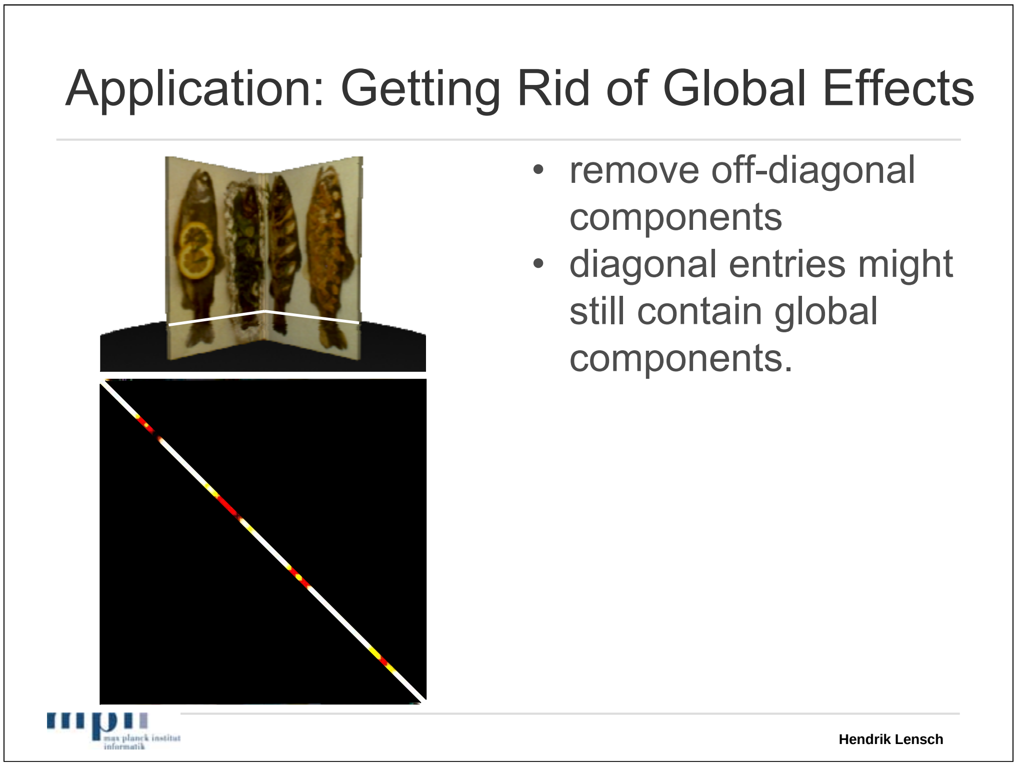

With this knowledge at hand, one can now try to separate direct and global illumination effects.

Simply removing off-diagonal elements from the transport tensor removes most of the indirect reflections. However, even the diagonal elements will to some extent be influenced by global illumination effects, e.g. 3rd order reflections back to the point of incidence.

Fast Separation of Direct and Global Effects

- main idea: global illumination effects dampen high frequencies
  - illuminate with shifted high frequency patterns
  - only the local illumination will change
  - global illumination will be invariant to phase shifts

image direct global component

mpii max planck institut informatik

[Nayar et al. SIGGRAPH 2006]

A principled method for a fast separation of direct from global illumination effects in images captured from a single view point has been proposed by Nayar et al.

In this example, the direct component contains all specular highlights. The global component is strongest for the translucent surfaces and for the interreflections between the walls.

For this problem again, structured illumination is applied.

The idea is that only the direct reflection component will respond to high frequency variation in the incident illumination while the global component will remain fix.

Fast Separation of Direct and Global Effects

- shifted periodic patterns (e.g. checker board)
- record per-pixel minimum and maximum
- approximation: $L_g = L_{\min}$ $L_d = L_{\max} - L_{\min}$

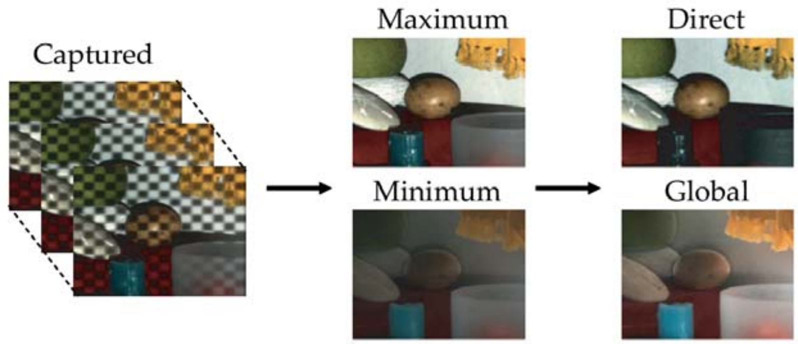

[Nayar et al. SIGGRAPH 2006]

This effect can be observed when illuminating the scene with a set of shifted periodic high frequency patterns.

In the recorded image sequence one just need to determine the minimum and maximum per pixel in order to quickly compute the direct and global component.

In the paper a more precise approximation based on the projector contrast and black-level can be found.

This tool of quickly removing any global illumination effect from measurements performed with structured light can help to make computer vision tasks more robust.

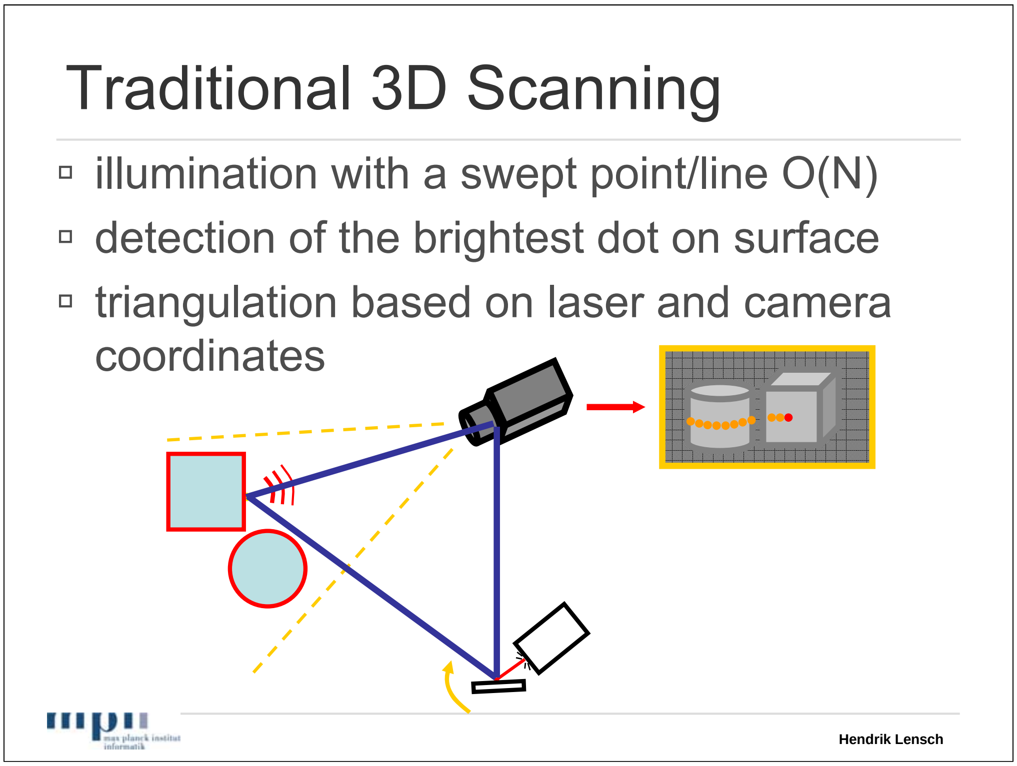

For example, 3D range scanning, where structured illumination has been initially been applied.

Here, we see the principle of a traditional laser range scanner. Sweeping a line or a dot over the scene the camera records the brightest point along the scan line.

This detection can be performed with sub-pixel precision.

Given the location of the brightest point in the camera image and the rotation angle of the laser, one can now easily determine the distance between the camera and the surface point by triangulation.

Structured Illumination

- reduced number of acquired images O(log N)
- analyze the captured images to determine projector coordinates

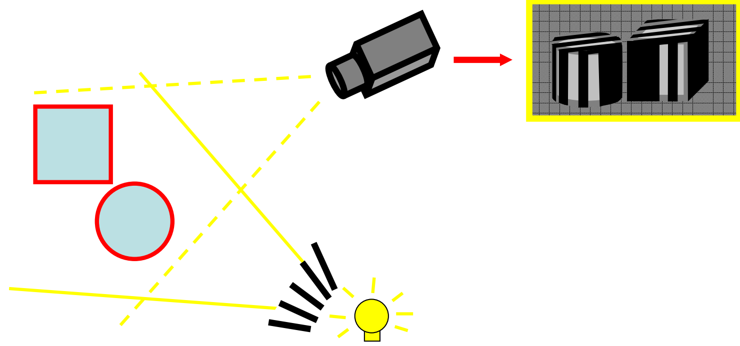

Hendrik Lensch

Instead of sweeping a line, structured hierarchical codes can reduce the acquisition effort to O(log N). The recorded camera images now need to be analyzed in order to determine the projector row that illuminated the visible surface point in each camera pixel.

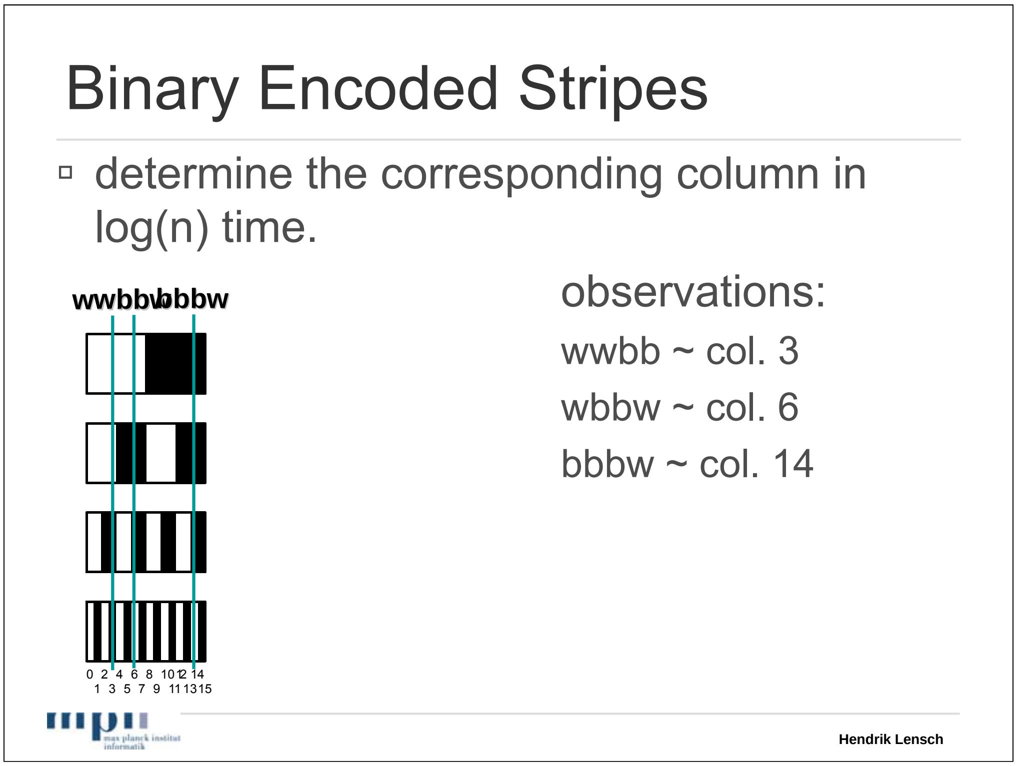

The easiest code simply encodes the column number in a set of bit planes, producing one illumination pattern for the most significant bit and so forth.

By interpreting the recorded on/off patterns as a binary number one directly can read of the projector column again for each pixel.

One limitation of this approach is however, that the encoding is constant within each bar of the highest resolution.

Stripe Patterns - Extensions

- Use Gray-code for more reliable detection
- Different pattern for interactive scanning [Rusinkiewicz ’02]
- Color patterns for faster, texture invariant acquisition

Hendrik Lensch

Various extensions have been proposed to make the encoding more robust or more efficient.

Phase Shifting

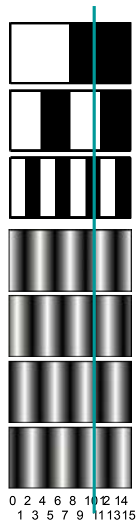

- precision of binary encoding limited by highest resolution
- use shifted sine pattern to obtain subpixel precision
- requires precise gray levels
- period needs to be determined using coarser levels (phase unwrapping)

One way to overcome the problem of the discrete locations it to apply sinusoidal functions shifted in phase for the last few iterations. Given some photometric calibration, the exact phase of a surface point in these shifted pattern can be reconstructed with subpixel precision. The output is the phase within one period of the highest frequency. The binary codes of coarser resolution are necessary to locate the period. Here, the constant location is sufficient.

Robust 3D Range Scanning

- use patterns that inherently filter out global illumination effects
  - high frequency sinusoids do this
- use additional polarization – multiply scattered light is typically unpolarized

Hendrik Lensch

With the shifted sinusoidal patterns we can do two things at the same time:

-3D range scanning using phase shifting

-separation into direct and global reflections

This produces a 3D range scanning approach that is robust to global illumination effects.

Robust Sine Patterns

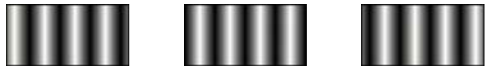

- illuminate with phase shifted sinusoids: e.g.

$$I = \sin(x + \delta_i), \qquad \delta_i \in \{-2\pi/3, 0, 2\pi/3\}$$

- estimate phase -> 3D position:

$$\Phi = \tan^{-1}\left(\frac{-\sum L_i \sin\delta_i}{\sum L_i \cos\delta_i}\right)$$

- or perform separation

$$L_g = \tfrac{2}{3}(L_0 + L_1 + L_2) - L_d\,, L_d = \tfrac{2}{3}\sqrt{3(L_0 - L_2)^2 + (2L_1 - L_0 - L_2)^2}$$

max planck institut informatik

Hendrik Lensch

The equations for obtaining the phase and for performing the separation are rather simple.

It is worth noting that estimating the direct component using sine pattern using the above equations is much more robust than detecting the minimum and maximum from a sequence of shifted binary patterns.

Using sinusoids all measurements are incorporated.

Another way to do descattering is using polarization.

There exists a vast literature on polarization for separation of reflection components.

Rather than the algorithmic approaches presented early that use high frequency patterns

polarization filtering makes use of the physical properties of scattering events. Each scattering event or reflection changes the polarization of the reflected light. Multiple scattering and subsurface scattering decorrelate the resulting polarization state from the incoming orientation while direct reflections typically keep them.

These images show the comparison in 3D reconstruction.

The underlying structures, here the straws, destroy the reconstruction if pure phase-shifting is used. Its descattering capabilities are not strong enough in this case.

Using polarization difference imaging on top of the shifted sine patterns we can obtain a clear 3D scan of the first surface.

Application: 3D Scanning

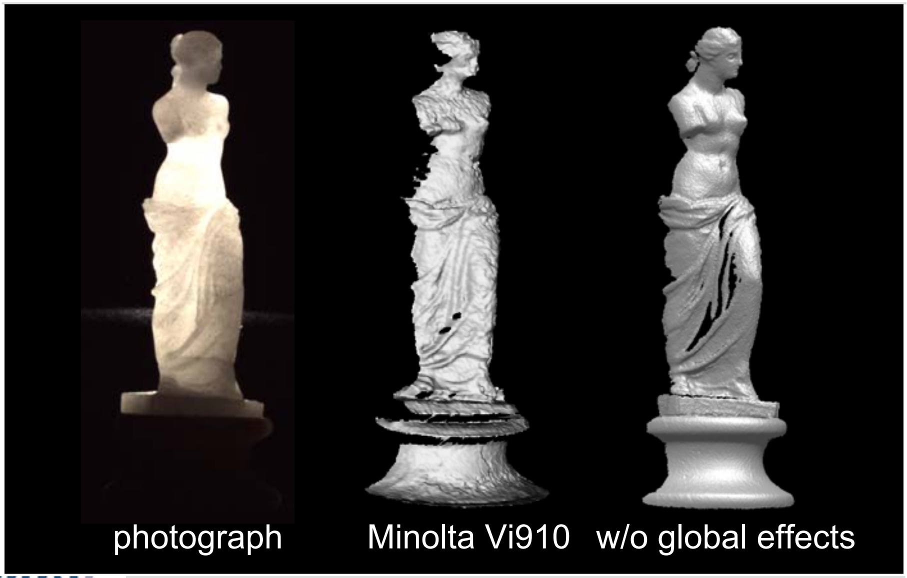

[Chen et al. CVPR 2007]

A commercial laser range scanner has sever difficulties in scanning translucent objects. The subsurface scattering offsets the location of the peak, resulting in a corrupted 3D scan.

Removing the global effects a significantly better 3D scan can be obtained.

Modulated Patterns in 2D

- use sinusoids in both dimensions
  - results in descattering in both dimensions

$$I = \sin(x + \delta_i) \cdot \sin(y + \delta_j)$$

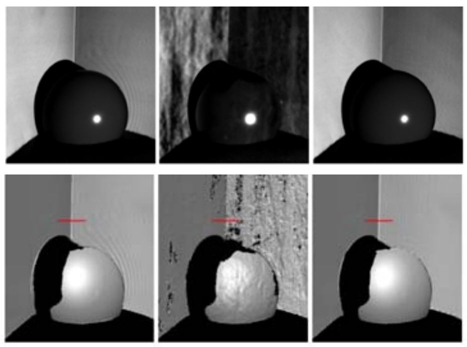

1D polarization 2D

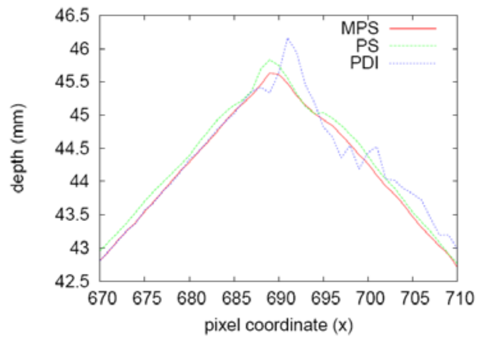

[Chen et al. CVPR 2008]

Modulating a horizontal sine pattern with a vertical one and capturing M*N phase shifts for both directions, the separation capabilities can be drastically increased. Modulation and shifting in two dimensions outperforms even polarization difference imaging.

Illumination Patterns and Codes

- scan
- periodic
- hierarchical
- multiplexed / Noise

- modified by the properties of light
  - wavelength
  - polarization
  - time of flight / phase (not discussed here)

Hendrik Lensch

In the different fields we have seen the same codes and illumination patterns over and over again.

Scanning

- moving a pixel or a line
- simple
- slow
- low light efficiency, consider black frame

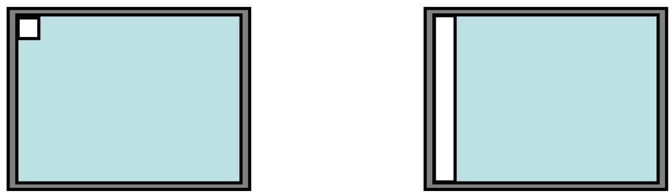

O(N^2) O(N)

Hendrik Lensch

Scanning is the slowest but delivers quite robust measurements. The benefit is that the acquired images can be used directly without any further processing.

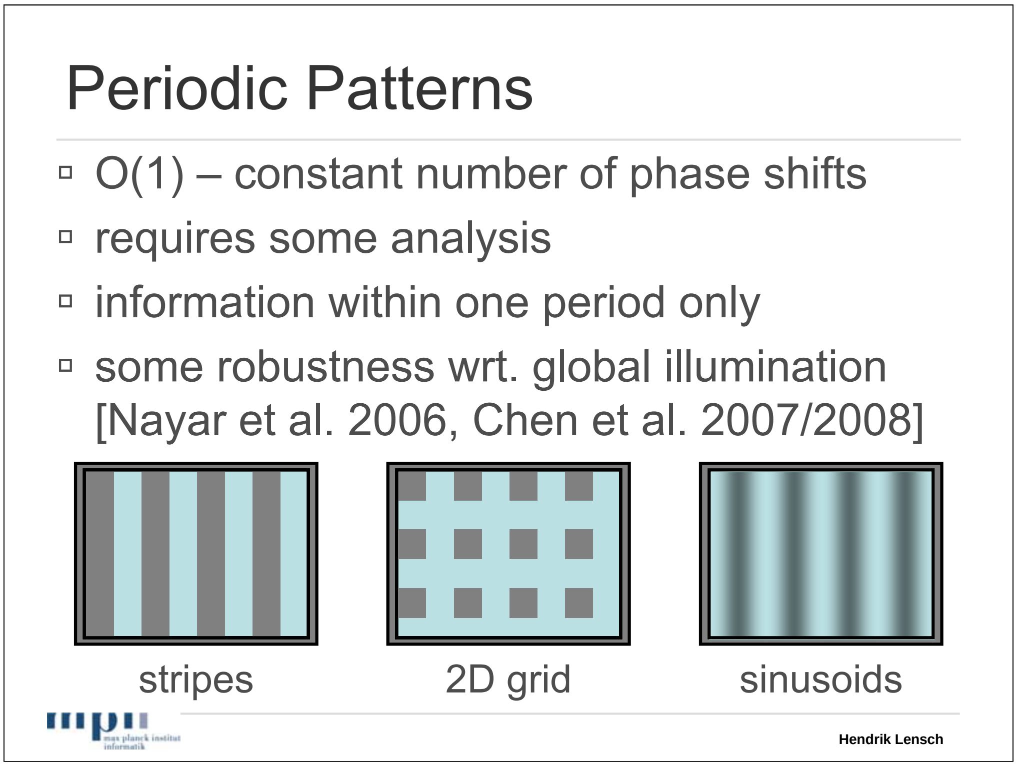

Periodic patterns can be acquired much faster. As the number of phase shifts is fixed the acquisition effort is O(1). Of course, periodic patterns are note always applicable, e.g. not for measuring reflection properties.

They require some analysis which is typically rather simple.

The second benefit of periodic patterns is that the are inherently robust against global illumination effects if analyzed appropriately.

Hierarchical Bases

- O(log N)
- deliver global localization
- stripes or sinusoids (phase unwrapping), spherical harmonics

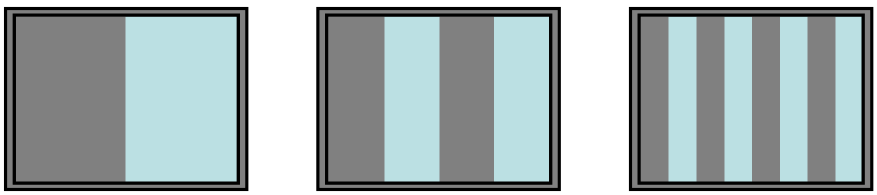

Hendrik Lensch

Hierarchical bases are used to compactly localize features globally. Furthermore, when acquiring reflectance fields or environment mattes the can be used to directly sample into a hierarchical representation of the otherwise large data structures.

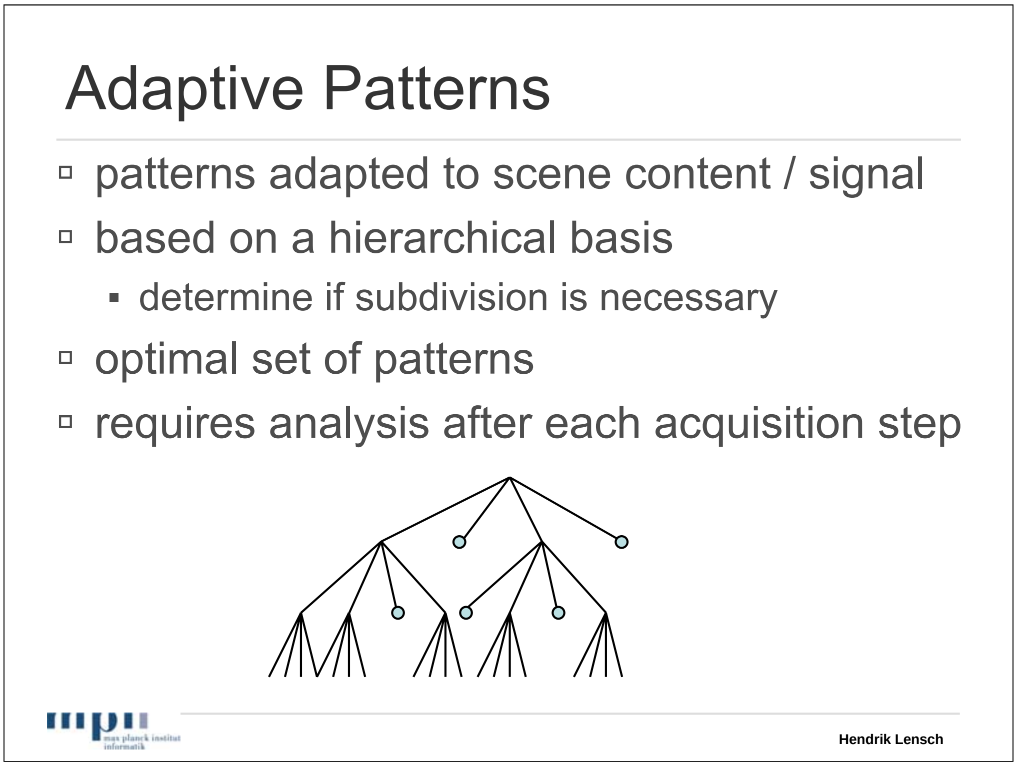

Combined with an adaptive acquisition scheme, hierarchical bases can be used to accelerate the acquisition. Sub-trees that do not require further subdivision will be culled already during acquisition, unnecessary acquisition steps are saved.

Multiplexed Illumination

- random distribution of basis elements
- same (or less) effort than scanning
- each sample is obtained by combining multiple measurements – increased SNR
- requires analysis
  - multiplication with inverse measurement matrix
  - optimization (compressed sampling)

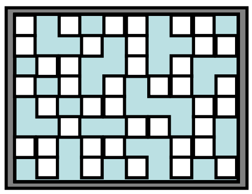

Hadamard

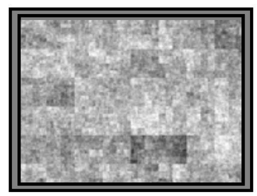

wavelet noise

Multiplexing multiple basis functions, it is possible to either improve the SNR slightly or to accelerate the acquisition sacrificing some quality.

Multiplexed illumination requires some additional analysis to obtain the measurements for individual bases.

Summary

- computational illumination used for measuring
  - 3D shape
  - object appearance
- choice of ***codes, acquisition*** and ***analysis*** are always interlinked

Hendrik Lensch

Challenges

- robust acquisition under strong ambient illumination
- dynamic scenes
- large scenes
- uncontrolled environments