State of the Art in

Computational Fabrication and Display of Material Appearance

Matthias B. Hullin, Ivo Ihrke, Wolfgang Heidrich, Tim Weyrich, Gerwin Damberg and Martin Fuchs

---

Abstract

After decades of research on digital representations of material and object appearance, computer graphics has more recently turned to the problem of creating physical artifacts with controllable appearance characteristics. While this work has mostly progressed in two parallel streams – display technologies as well as novel fabrication processes – we believe there is a large overlap and the potential for synergies between these two approaches. In this report, we summarize research efforts from the worlds of fabrication display, and categorize the different approaches into a common taxonomy. We believe that this report can serve as a basis for systematic exploration of the design space in future research.

---

1. Introduction

In the decades since its inception, computer graphics research has developed rendering techniques so close to physical simulations that the resulting images of synthetic scenes are virtually indistinguishable from photos.

An important ingredient has always been the description of real-world material appearance. The term, as used within computer graphics, describes the space of all interactions of visible light with objects and materials, i.e. the space of all possible images that one could take of an object under arbitrary lighting and viewing conditions. In rendering it is now common practice to use either analytical models of appearance of sufficient quality to closely match gonioreflectometric measurements for a given material, or data-driven models that are directly based on such measurements. We refer readers who are interested in appearance modeling and acquisition to dedicated literature on these topics [Dorsey et al. 2008, Weyrich et al. 2009a].

In this report we focus instead on a relatively recent research theme in computer graphics; namely how to reproduce physical artifacts with a certain appearance derived from computational models. While there has obviously been a host of work on how to design specific products such as paints, textile fibers [Wada 1992], or coatings [Dobrowolski 1973] with specific optical and appearance characteristics, our focus in this report are fabrication and display processes capable of emulating and representing a large range of very diverse appearances. Within this focus of interest, we consider both dynamic technologies whose appearance can be reconfigured, as well as static technologies in which a physical artifact is fabricated with a certain appearance that then remains fixed for the lifetime of this artifact.

1.1. Reproducing Virtual Appearance in the Real World

While the virtual recreation of real objects is an important field of research, the opposite direction of transfer is at least of equal importance: the fabrication of objects that look a certain desired way. The usage scenarios are manifold and range from product design, pre-press proofing and the preservation of cultural heritage to a variety of medical applications (e.g., prostheses that look exactly like the body part they replace).

For many practical applications, a reliable virtual preview has been proven sufficient to guide the design process. In addition, however, traditional product design has always been relying heavily on material samples to preview the appearance of the resulting product. This requires that a sufficiently dense subset of materials from that space is sampled, which naturally becomes less tractable as the possible space of materials grows. Consequently, some of the efforts discussed in this report are rapid prototyping techniques that attempt to combine the visual properties of a given set of basis materials with the purpose of spanning a maximum appearance gamut.

So what does it take to recreate the appearance of a real or virtual object? Physically speaking, the goal is easily defined as the creation of a real object or device that closely recreates a given target gonioreflectometric distribution. In practice, perceptual similarity is often a good approximation to this rather strenuous requirement. Most of the works featured here

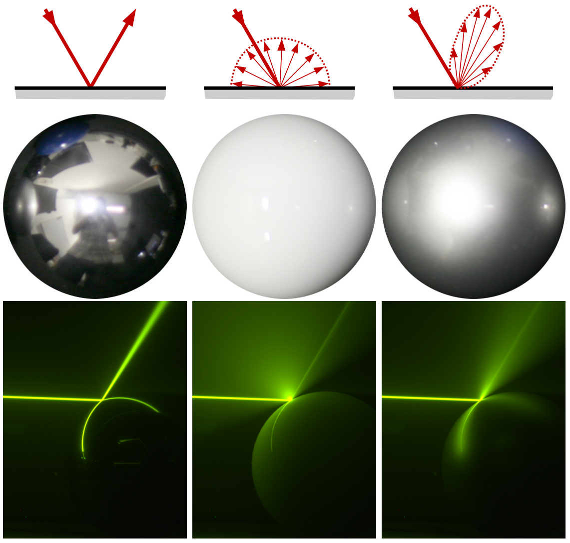

Figure 1: *The appearance of a surface is primarily defined by the angular distribution of light upon reflection, as shown using, from left to right, a mirroring (specular), a mostly diffuse and a glossy sphere. From top to bottom: idealized reflectance distributions, photo of sample spheres, light reflected off the surface in a scattering medium for visualization [Hullin et al. 2008]. Note that the glossy and diffuse spheres still contain weak specular contributions.*

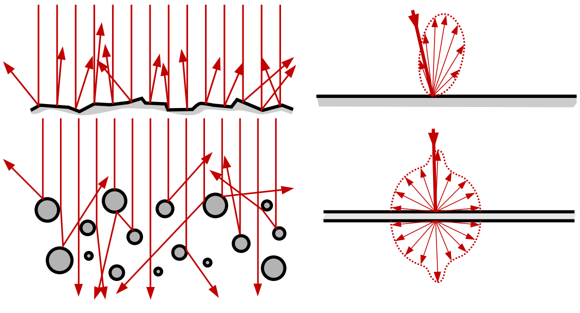

Figure 2: *The presence of microscopic geometric detail on a surface or in a volume (left) gives rise to macroscopic distributions of scattered light in reflection and transmission(right).*

are situated between both worlds in that their approaches are motivated by the physical definition, but they evaluate their results mainly for visual plausibility.

1.2. Scope of this Report

In this report, we look at attempts to synthesize real-world material appearance based on purely virtual model parameters, effectively embodying the appearance representations from Figure 3 into actual objects and devices. In recent years, the output of reflectance models in extension of flat pictures and videos has emerged as a new research direction. Since then, we have seen a variety of publications dealing with the fabrication of physical objects, made from simple materials so as to emulate the look and/or feel of a visually complex computer graphics model.

At the same time, researchers have also attempted to achieve similar goals in a dynamic, re-programmable fashion. We argue that these two strands of research are in fact closely related. This report aims at introducing a framework that connects fabrication and display by investigating the underlying principles and design approaches.

The remainder of this report is structured as follows: in Section 2, we lay the theoretical foundations by summarizing the most common theoretical representations of material appearance. Section 3 introduces efforts within the graphics community to translate these representations into real-world objects and devices. A discussion of the individual approaches and working principles is then given in Section 4, also pointing out some directions that might be particularly interesting to pursue in the future. In Section 5, we argue why in our context, “display” and “fabrication” have more in common than what separates them.

While we acknowledge the vast amount of work that is somewhat connected to the reproduction of material appearance, this review on reproduction of material appearance is much more focused. The central criterion for work to be featured here is that it must close the loop from real-world illumination via digital processing or computational optics back to real-world output. Additional literature on appearance models in physical and perceptual space, image-based lighting and rendering, plenoptic imaging, light/reflectance fields and light transport theory, general display technology and autostereoscopic, but illumination-insensitive displays [Holroyd et al. 2011, Wetzstein et al. 2011, Ito et al. 2010, Jones et al. 2007], mechanical fabrication [Bickel et al. 2010], human-computer interaction, as well as most of augmented and mixed reality literature, falls outside the scope of this report, but we refer to it whenever we deem it useful for understanding of a concept.

2. The Dimensionality Problem

The main purpose of appearance fabrication and display is the creation of a facsimile, something which looks like something else, and the appearance of which may be determined by a computer graphics model. Appearance being the optical response to illumination as observable for different viewing conditions, we may classify methods according to the *dimensions of the light interaction space* they provide. The key considerations are whether a technology can independently model spatial and/or directional variation in both the incident illumination and the reflected/observed light. Since each spatial and directional component accounts for 2 dimensions, variability of all these aspects already accounts for up to 8

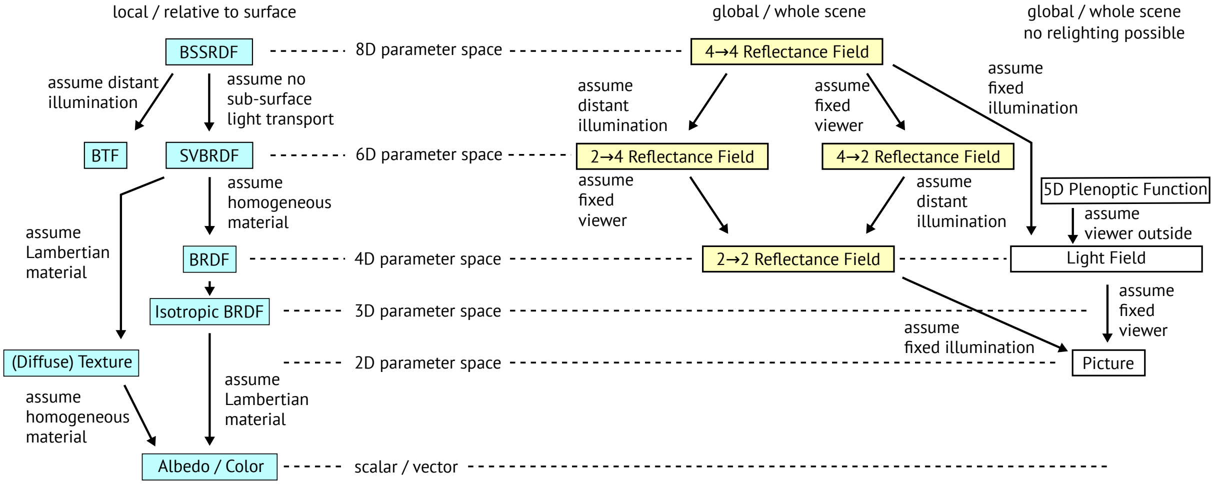

Figure 3: *A taxonomy of visual appearance representations, extended from [Fuchs 2008], [Lensch 2003], and [Rusinkiewicz and Marschner 2000]: methods for fabrication of material appearance and its interactive display are tightly related to methods for describing material appearance in computer graphics contexts, and follow the same patterns. Simplifying the material types to reduce the parameter space (blue) maintains full interactivity with viewer and light, while restricting the interactions between viewer, material, and illumination maintains a full material gamut (yellow). For reference, illumination-invariant representations are in white.*

dimensions. Color adds an additional dimension if treated as a separable effect in individual color channels, as is the case with current fabrication techniques. Adding control over fluorescence and/or phosphorescence would require introducing additional dimensions. to the interaction space. Generally speaking, a higher dimensional interaction space can capture more physical properties of an object, and therefore produces a more realistic experience.

On the other hand, we may look at how many parameters a technique can (or even must) control, and consider the *controllable parameter dimensionality*, the number of dimensions of the variable space (such as surface color, optical density, ...) which must be independently controlled by the technique. As every additional parameter space dimension enlarges the problem size (data structures, manufacturing / computation time, etc.), a low-dimensional parameter space dimensionality is paramount for the feasibility of a practical implementation.

This creates conflicting goals for which a compromise must be found, and, as Figure 3 illustrates, the same problem is addressed by the classical computer graphics methods for describing appearance. Most generally (on the very top of the figure), the Bi-Directional Scattering-Surface Reflectance Distribution Function [Nicodemus et al. 1977] and the reflectance field [Debevec et al. 2000] both provide for general viewer, general illumination interaction with arbitrary materials. (As we will look at specific simplifications of the general reflectance field, we label the general reflectance field “$4 \rightarrow 4$”, as it describes the transformation of a 4D incident light field [Levoy and Hanrahan 1996] to a 4D outgoing light field). BSSRDF and reflectance field differ only in that the BSSRDF is defined relative to a scene surface, while the reflectance field is defined relative to some more arbitrary geometry; both require an 8D parameter space when stored in tabulated form, and the re-creation of the appearance of arbitrary materials in arbitrary ways of interaction hence requires the control of eight dimensions of variables.

2.1. Reducing Dimensions to Enable Feasibility

Reducing the unwieldy number of parameter dimensions involves two basic strategies: reducing the expressivity of the representation to simpler materials (blue) maintains full interactivity, and lends itself to a fabrication process, which may shape the synthesized surface into desired geometry. For instance, one may consider only materials which do not transport light below the surface; then, only light that enters and leaves the surface at the same point in space must be modeled – and a 6D spatially varying bi-directional reflectance distribution function (SVBRDF) is sufficient. If the surface may be assumed to appear the same in every point, two more dimensions are lost and a BRDF is sufficient, and so forth.

Reducing the possible interaction (yellow) either with the viewer (varying view points) or the illumination motivates a construction of a display-type material appearance reproduction technique; for instance, if all illumination may be assumed to be distant from the scene, it may be modeled as a 2D environment map, a $2 \rightarrow 4$ reflectance field models the material appearance exhaustively, and the incident illumination can be observed with a single, wide-field-of-view camera [Nayar et al. 2004].

In a sense, a bi-directional texture function (BTF, [Dana et al. 1999]) bridges the two approaches of dimensionality reduction. Just like an SVBRDF, it is defined relative to an object surface, and, in tabulated form, requires 6D storage. Practically, it is however often used to describe composite materials of complex micro-geometry, which may have self-shadowing and local-scale global light transport, such as woven fabric. The dimensionality reduction against a full BSSRDF is achieved by assuming that the illumination comes from a far greater distance than the local light transport, and that its spatial variation, consequently, is of low frequency. This can be interpreted as either restricting the interaction space (for sufficiently distant illumination, any material can be modeled) or the material type (for a certain local variation of illumination, the material must restrict local light transport to a smaller scale).

For completely static, pre-determined illumination, we end up with classical representations such as the plenoptic function [Adelson and Bergen 1991], light field [Levoy and Hanrahan 1996] or a simple 2D picture of a scene in this light.

Figure 4 shows fabrication and display techniques grouped by their interaction dimensionality, plotting the interaction with the viewer (up axis) against the interaction with the light (right axis).

At the end of the day, it is most likely the limited supply of hardware resources, be it memory, compute power, display resolution, or our drawer of optical parts, that forces us to identify a "sweet spot" within the space of appearance models that provides an optimal viewing experience, but is still feasible for implementation. An important trade-off in this respect is the one between the dimensionality of our interaction space and how coarsely we discretize it. For instance, Fuchs et al. [2008] implemented a passive display prototype which permits $2 \to 4 = 6$D interactions, but offers programmability only in $7 \times 7 \times 6 \times 5$ locations for four of the dimensions.

Another compromise can be found in limiting the interaction scenarios. If we can fix the observer position, we also may neglect two dimensions entirely, as done in the $2 \to 2$ design of Fuchs et al. [2008] and Malzbender et al. [2012].

Finally, the depicted scenes or materials can be constrained so as to collapse some dimensions (see, for instance, Matusik et al. [2009]). Ultimately, for a Lambertian surface, the appearance would be uniform for all observation directions – hence, a Lambertian material can be plausibly reproduced by a printing process which spatially varies albedo along two dimensions, while maintaining the full 8D interaction space.

3. Embodiments of Appearance Models

The availability of *relightable* representations has had an enormous impact for realistic rendering of real-world materials and scenes under arbitrary illumination, and spawned a wide variety of applications. The most prominent field of use and of significant economic impact is arguably the movie industry with its need to composite real actors into virtual environments [Debevec et al. 2000]. As detailed above, these relightable representations can either be global (such as the reflectance field, originally introduced under a different name by Nimeroff et al. [1994]) or defined with respect to a surface (the BRDF and its extensions [Nicodemus et al. 1977]). They all are inherently high-dimensional.

It is obvious that traditional display technologies can only give an insufficient viewing experience for such relightable representations, because they lack the capability of interacting with surrounding light. Therefore, researchers from the graphics community have recently started to implement devices that reproduce computerized appearance models in the real world. In many cases, it is hard to draw a line between efforts from rendering and related fields, namely virtual and augmented reality.

In the following, we outline work that seeks to translate the taxonomy in Figure 3 into objects and devices that can be viewed and relit in the real world. We observe how researchers strive to increase the dimensionality and expressivity of light transport scenarios that these devices can reproduce. A more detailed discussion of practical challenges encountered along the way is given in Section 4.

3.1. Global (Scene) Representations

Global representations treat the scene as a "black box" that receives and emits light. As such, reflectance fields lend themselves to implementations using cameras and displays, although some passive devices have been demonstrated as well.

Among the first efforts to introduce relightable objects as a display modality is a device that measures 2D illumination conditions using a wide-angle camera embedded in the display frame, and renders and displays the corresponding 2D image in real time [Nayar et al. 2004]. Since the device uses standard 2D imaging devices and computes the image formation in software, this approach in principle allows for the display of 2→2 reflectance fields at a full resolution that is the product of the individual resolutions of the camera and display. The rendering approach used by the authors is image-based and made feasible by exploiting the local coherence in the reflectance field in order to compress the large amounts of data.

Malzbender et al. demonstrate flat, optically passive objects that react to surrounding light in a visually plausible way [Malzbender et al. 2012]. By combining an array of curved concave mirrors with a printed transmissive layer stacked on top, the method employs an idea from the world of integral imaging (multiplexing while trading spatial for angular resolution [Lippmann 1908]) to achieve the illusion

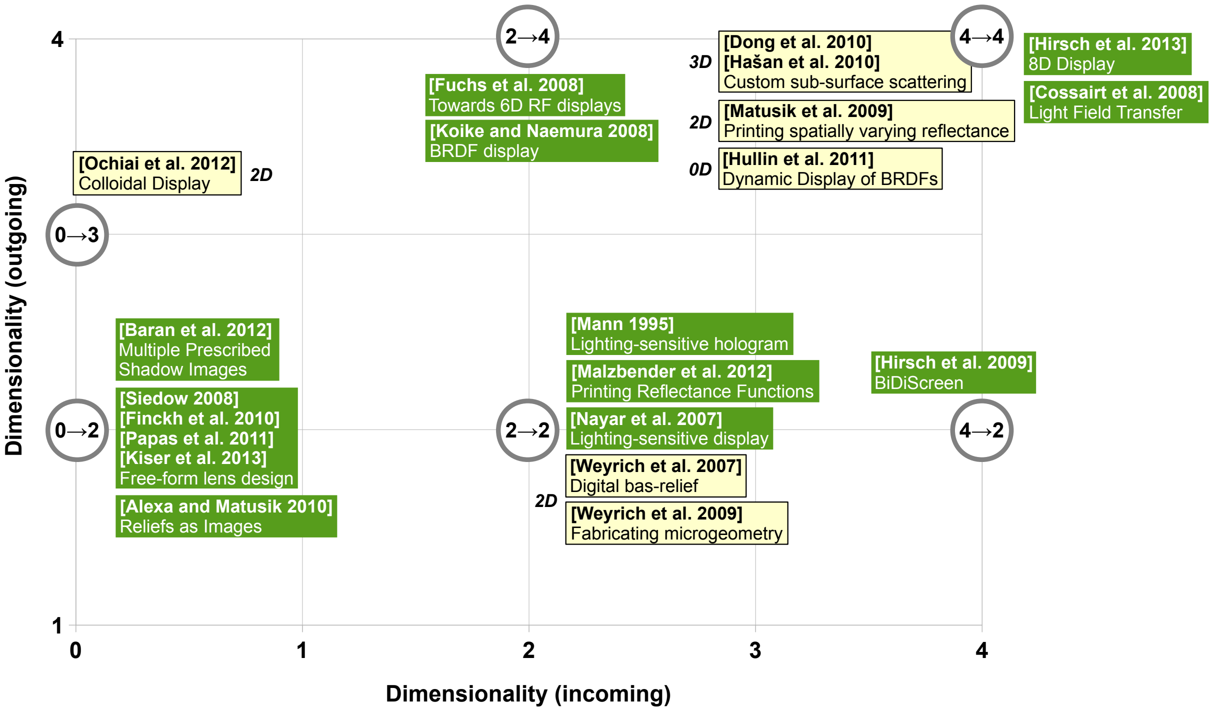

Figure 4: *Map of display / fabrication techniques according to the dimensionality of their incoming and outgoing light fields (interaction dimensionality).. Technically, many implementations offer full control over of the light transport tensor (marked in green); others only have a reduced parameter dimensionality (annotated in italics). We note that a full parameterization may not always be necessary for a convincing viewing experience.*

of surface effects such as spatially varying normals and reflectance distributions. In principle, this approach embodies a reflective 2→2 reflectance field and can therefore produce plausibly lit images for a fixed observer position and distant illumination. A technological challenge and limitation stems from the fact that the printed layer hovers above the mirrors, leading to unwanted shadowing as light traverses the layer in different locations before and after reflection.

The devices discussed so far only showed 2D images and could therefore not reproduce stereoscopic effects such view-dependent highlights. Koike et al.'s "BRDF display" [2008] extends Nayar et al.'s display by adding a layer of lenslets to the LCD panel, obtaining a 2→4 display that is autostereoscopic and reacts to incident light in real time. The practical challenge of color moiré is solved by using field-sequential color from red, green and blue backlight LEDs behind a monochrome LCD.

Holographic techniques can be used to store and reproduce lighting-dependent images by purely optical means, as demonstrated by Mann [1995]. Due to the vast body of work on the topic of holography, we will not discuss it all here, referring the reader to Section 4.4 instead.

Fuchs et al. [2008] use a purely passive approach to display 2→2 and 2→4 reflectance fields in transmission. Using an elaborate arrangement of lenslets, they flatten the 4D or 6D light transport tensor into a plane where it is modulated by an attenuating transparent layer. Inherently to the integral imaging approach, this display concept is very limited in its resolution; the 6D display demonstrated has 7×7 macropixels that map an incident space of 6×5 illumination directions to a small set of outgoing radiant directions. The paper discusses in great detail the challenges in terms of resolution and alignment that result from this extreme example of multiplexing.

The BiDiScreen [Hirsch et al. 2009] combines a mask-based (heterodyned) light field camera and 2D display into a device technically capable of displaying 4→2 reflectance fields. The main suggested use is the exploration of novel kinds of user input, for instance through gestures. Since only the light-capturing mechanism features a multiplexing technique, the displayed image can use the full resolution of the display panel.

Finally, the work on light field transfer by Cossairt et al. [2008] combines 4D light field imaging with 4D display, creating an interface that allows virtual and real scenes to exchange global illumination. For the first time, this device

Figure 5: *This type of split-flap display is probably best known from airport information panels. By replacing the numbers with a set of material samples, one would obtain a trivial two-pixel BRDF display capable of full 4→4 interaction, but with 0D parameterization.*

embodies the most general type of reflectance field (4→4) in the real world. The underlying optical mechanism follows the integral imaging approach and uses a lenslet array that is shared between a camera and a projector. Besides resolution issues, the challenges are mostly computational and stem from the need for real-time 8D rendering. By approximating the incident light field with an array of projective light sources, the rendering technique is made suitable for efficient GPU execution.

Technically very close to Cossairt et al.'s approach is the 8D display by Hirsch et al. [2013]. The authors set their work apart by proposing a viewer-centered usage scenario and a simplified implementation using bidirectional LCD technologies that may become available in the future.

3.2. Local (Surface-Based) Representations

As with the global representations, we clearly see the development from a single 2D microfacet distribution to the most general surface-based appearance description that includes non-local light transport (sub-surface scattering). Due to physical limitations, only a very small subset of such high-dimensional distributions can actually be fabricated. The basis materials used always impose a limited gamut within the full space of imaginable appearances.

The first effort to computationally design a surface geometry to achieve a certain non-trivial target reflectance distribution was presented by Weyrich et al. [2009b]. The surface is divided into an array of facets which is then optimized according to a set of objective functions. Besides the target distribution of normal orientations, the objective includes several energy terms to reduce visual occlusion and enable fabrication. The resulting geometry is then milled into an aluminum block, producing a 2D probability density function of normals and the resulting 4D bidirectional reflectance distribution function. The authors demonstrate a set of such distributions and propose further research directions such as spatial variation and texturing of actual objects

In an attempt to generate such normal distributions in a re-programmable way, Hullin et al. dynamically alter the geometry of a reflective surface (water) by exciting waves on it [2011b]. Although the manipulation of liquid surfaces underlies complex physical constraints, the authors derive an analytical model for the BRDFs that result from the superposition of sinusoidal waves. A single-pixel prototype demonstrates the validity of the model, but also the very limited gamut of BRDFs that can be produced using this approach: only axis-aligned elliptical Gaussian distributions of normals on a fixed basis material. Due to the choice of a dynamic surface, this approach results in a higher effective interaction dimensionality than the method by Weyrich et al. [2009b]: thanks to the temporal multiplexing, each point on the surface assumes the target normal distribution over time, enabling plausible reflection even of spatially varying illumination patterns.

Hersch et al. introduce a framework that integrates full-color printing with spatially varying metallic effects, also modeling effects from layering and halftoning [2003]. In extension of this work, Matusik et al. [2009] employ a set of inks and foils with particular reflectance and transmittance properties in order to print spatially varying reflectance distributions. At the core of both approaches is a careful characterization of the optical properties of the basis materials, their linear combination through halftoning, and the (nonlinear) alteration of angular lobes by applying diffusing layers on top of specular ones, allowing for the generation of target appearance within a continuous gamut spanned by the materials.

In parallel efforts, Hašan et al. [2010] and Dong et al. [2010] explore the inclusion of custom sub-surface scattering into the fabrication process. On the theoretical side, both exploit the insight that light transport between two surface points can be closely approximated as the product of the diffuse scattering profiles of those points [Song et al. 2009]. Thus, out of all imaginable (8-dimensional) BSSRDFs, a 4-dimensional subspace is selected, greatly reducing the complexity of the mathematical problems that need to be solved. On the technological side, both approaches are executed differently but based on the similar idea: they combine the characteristic optical properties of a small set of basis materials in order to span a gamut of scattering functions. In the case of Hašan et al. [2010], layers from three different UV-cured resins are 3D printed to mix their different transmission, reflectance and absorption properties. A forward model for the scattering profile of stacks of layers is used to formulate a discrete optimization problem: given a (homogeneous) target profile, find the stack of layers that best approximates it. After solving this inverse problem for a set of surface points, 3D geometry for the entire object is computed by interpolating between the stacks. Dong et al. [2010] access a larger set of basis materials by employing both 3D printing and milling processes to fabricate a stack of layers that are then assembled as a separate step. Instead of optimizing the layer stack for a sparse set of surface points, they apply an inverse diffusion optimization to solve for a dense map of layer thicknesses.

For flat objects, a printed transparency layered on top of the fabricated layers enables full-color output.

3.3. Hybrid and non-appearance displays

In this section, we discuss a few examples where the displayed content is not material appearance in the strict sense. Nevertheless, this work shares technological aspects with other fabrication approaches discussed above.

Weyrich et al. propose a technique that allows for the generation of bas-relief geometry based on photometric normals, i.e., a surface that reproduces the shading behavior of a given scene [2007]. Alexa and Matusik generalize this idea by computing reliefs that, when lit from different directions, shade into a pair of distinct grayscale images [Alexa and Matusik 2010]. The authors note that not all possible images can be interpreted as shaded smooth height fields. Their solution is the use of non-smooth surface primitives (pyramids) that partly evade this limitation by introducing additional degrees of freedom.

A significant amount of effort has gone into the design of free-form reflectors and lenses, a research area driven mainly by the design of automotive headlamps, but with roots that reach back as far as the ancient Chinese "magic mirrors" [Berry 2006]. Here, the goal is to shape the light upon reflection or refraction so as to produce certain target caustic images on a projection surface. Compared to conventional image projection, this approach has the advantage that it is in principle lossless since it does not attenuate but only re-distribute the light that falls through the device. Unlike holographic light shaping, the geometric structures are macroscopic and do not require detail on the scale of a wavelength.

Researchers in industrial mathematics have approached the problem of full 2D caustic display using free-form lenses [Siedow 2008]. Berry [2006] shows that for small phase modulations, the intensity of the caustic image is directly related to the *Laplacian image*, which represents the size of the phase modulation.

Within the graphics community, very similar problems appear to have been rediscovered several times. Patow et al. investigated the inverse problem of reflector design exhaustively and collected their insights in several publications including a survey article [2005], but came to the conclusion that to that date, no mathematical optimization scheme had "proven to be the best". Finckh et al. [2010] proposed to apply a fairly recent stochastic optimization approach in order to determine a C^2 continuous B-Spline surface that would result in a given caustic pattern, and verified the result in simulation. Papas et al. [2011] are the first to demonstrate a working prototype object by milling down a slab of refractive material. According to the authors, the continuous geometries obtained by Finckh et al. could not be manufactured using the technique at hand because of resolution limitations. Similar to Weyrich et al. [2009b], their own surface optimization scheme is based on a discretization of the surface into square patches, each of which builds a plano-convex lenslet that casts a Gaussian caustic. The target images is decomposed into a mixture of Gaussians that are then assigned to lenslets with additional continuity and smoothness constraints to enable manufacturing of the resulting surface. Kiser et al. [2013] are the first to demonstrate actual manufactured continuous surfaces that produce high-quality target caustics without prior discretization into patches although technical details of their approach are not available a the time of writing of this report. The visual quality achieved appears to be on par with the simulations by Siedow [2008].

A recent work tackles the problem of encoding multiple target images in a single caustic light field [Tandianus et al. 2012]. By lifting the constraint of continuity between surface cells and employing an expensive stochastic optimization approach, a hypothetical refractive surface is obtained that would generate multiple distinct target patterns at different distances. Due to the lack of optical manipulators between the image planes, the authors find that the image quality degrades quickly as light distributions need to satisfy two or more target images.

The use of computer-controlled sound waves has also been used in a multi-layer display that can dynamically change the scattering function of its projection surfaces in a way similar to Hullin et al. [2011b]. Ochiai et al. used a mixture of colloids, soap and milk to produce films that are stable over a time frame of about 5 mins [Ochiai et al. 2012]. They use ultra-sonic vibrations to deform the surface which can be changed from transparent to scattering. The resulting surface can be used as a switchable rear-projection screen though the authors do not provide a quantitative link of their device to computational reflectance or scattering models.

An interesting approach that combines some of the above ideas with active illumination is the work on shading illusion [Amano 2012]. Here, a 2D-printed image encodes geometry information (a normal map) in the color space. A camera-projector system shades the normal map in accordance with virtual light sources, adding shading to the viewing experience. Unlike prior work in augmented reality literature [Raskar et al. 2001], the camera-projector system does not require prior information on the scene geometry. However, the authors do not investigate the use of more sophisticated material appearance and restrict themselves to diffuse (Lambertian) shading.

4. Working Principles

While traditional appearance research in computer graphics focuses on measurements, representations, and rendering, and considers aspects such as physical and perceptual realism, memory requirements and computational efficiency, work on appearance *fabrication* additionally has to develop novel physical means of implementing a desired effect. Doing so

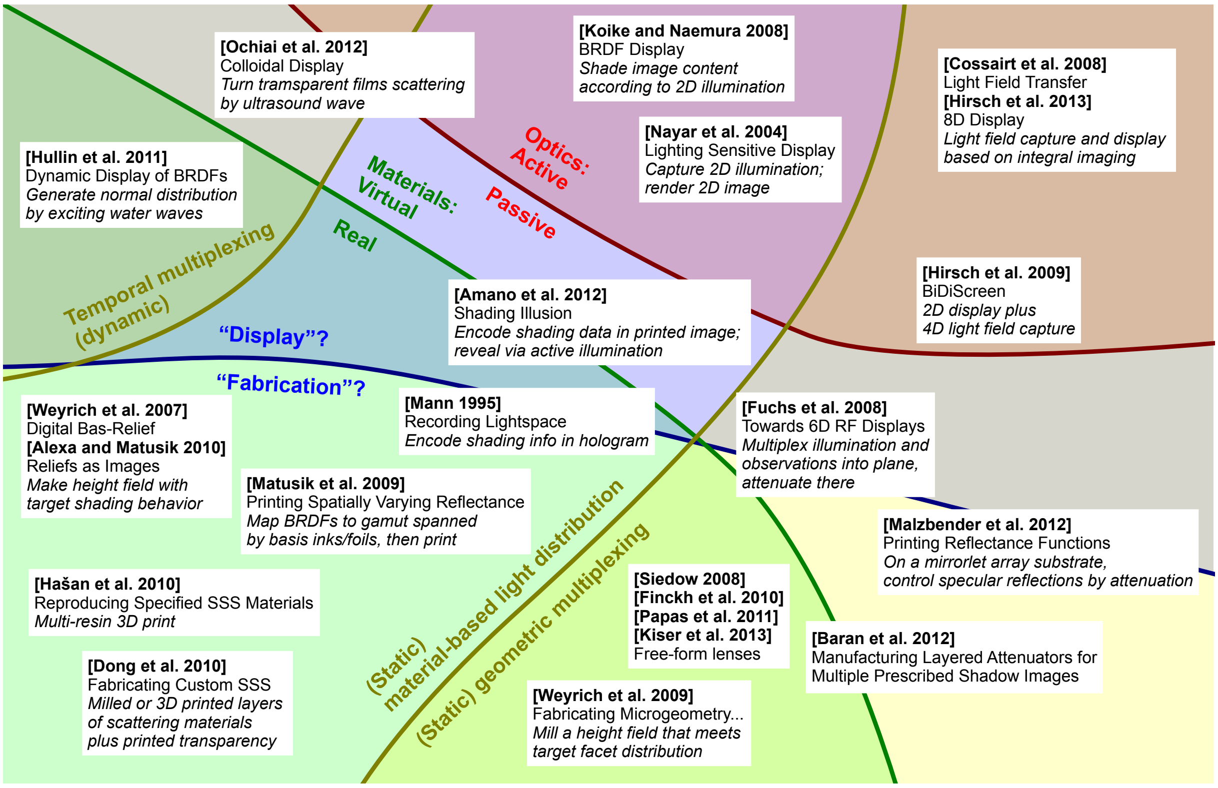

Figure 6: *The design space of appearance fabrication and display to date.*

requires identifying suitable materials and manufacturing techniques, but also addressing physical constraints that arise from this choice. In the remainder, we define the distinction between optically passive and active devices (Section 4.1). We then discuss commonly used methods of fabrications and their characteristics (Section 4.2), and identify reoccurring working principles in appearance fabrication (Sections 4.3–4.5).

4.1. Passive and Active Devices

The distinction between active and passive devices is arguably the most defining design choice for the fabrication and display of appearance, since it fundamentally affects the interaction of the device with light.

Optically passive approaches are systems where no electronic parts obstruct the light path. Instead, they use light that is present in the environment for display. This allows such devices not only to react instantaneously to changes in illumination, but also to reproduce arbitrary brightnesses up to the threshold of thermal destruction. Passive systems are always based on explicitly constructed hardware that may be cheaper than a system consisting of a sensor, a computer and a display. As such, they often do not rely on external power supply and are therefore freely mobile. However, a purely passive approaches always brings with itself limitations and construction challenges.

Active approaches, on the other hand, have a light path that is intercepted by an image sensor and a display or projection device, connected via an optional processing stage. These devices require the presence of a computer platform in some way, and they underly the technical limitations of all components involved in measuring and processing the input and generating an output (Section 4.5).

In practice, very similar devices can occur in active and passive implementations. Reflective LCD panels, for instance, actually change their reflectance with respect to surrounding light, and are hence optically passive devices under our definition. LCDs with backlight, in contrast, are active. The distinction only really makes sense when the modality of display is appearance, i.e., when the way in which the scene reacts to surrounding light is an integral part of the displayed content.

4.2. Fabrication Techniques

Appearance fabrication highlights on readily available means of computer-controlled physical output; this section summarises the set of techniques used to date. Any design creating custom appearance has to consider their optical characteristics, but also include more practical factors, such as costs, ease of fabrication, mechanical constraints (statics), and effects specific to the tooling or deposition technique used.

2D Printing

A trivial means of generating spatially-varying diffuse reflectance is traditional (2D) printing. Key parameters are the print resolution and the printer gamut defined by the substrate and the set of inks or toners employed. Extending this concept, Matusik et al. [2009] used a regular printer, combining inks of basis BRDFs to output spatially-varying reflectance variations. 2D prints are also used by various designs to spatially modulate light (coloured slides, patially coded masks, etc.), as we discuss in Section 4.3.

Additive Manufacturing

Recent years have seen a rise in additive rapid prototyping technology. The general concept is to build up 3D objects from a small number of basis materials by adding the material layer by layer. Popular techniques are: extrusion (fused deposition of material from a computer-controlled nozzle), granular (sintering or melting of granulated base material through local heating by a laser), and light-polymerized 3D printing (a photopolymer "resin" cured by controlled UV illumination). All methods achieve similar resolutions (typically between 15 and 100 μm layer thickness) and mainly differ in the material selection, which affects optical and mechanical parameters.

Any physical output has to respect mechanical limitations of the material to ensure structural integrity [Stava et al. 2012]. For appearance fabrication, however, optical properties play a critical role as well. For instance, most 3D printing materials exhibit noticeable subsurface scattering, which leads to subtle colour bleeding, blurring the spatially-varying albedo on any 3D printed surface. Normally treated as an undesired artifact to be compensated for [Cignoni et al. 2008], Hašan et al. [2010] and Dong et al. [2010], for instance, learnt to exploit this effect by creating material combinations that approximate desired subsurface scattering kernels. Also, the limited number of materials constraints the gamut of optical properties.

A very different additive manufacturing technique has been presented by Holroyd et al. [2011] who create a three-dimensional object replica by stacking multiple 2D-printed layers of transmissive slides between acrylic glass panes. The design comes with a number of visual artifacts due to its physical limitations (including parallax discontinuities due to the coarse layer spacing and gradual light attenuation within the plane assembly), so the authors present an optimization framework to compensate for them. Their work enters the continuum between material appearance fabrication and work that explores the creation of 3D objects using non-traditional processes, such as [Mitani and Suzuki 2004, Igarashi et al. 2008, Hildebrand et al. 2012, Skouras et al. 2012]. Although outside the scope of this report, such work is similar to many appearance fabrication approaches, as it optimizes the output design to address limitations and properties of the output medium to make a target shape fabricable.

Subtractive Manufacturing

In contrast to 3D printing, which incrementally builds up an object, subtractive manufacturing techniques use a sharp cutting tool to gradually remove material from an unmachined part until it has a desired shape. The most versatile subtractive process is *milling*. It is amenable to a large variety of materials, including metals and optical-grade plastics, which makes it very attractive for reflectors and refractors in appearance fabrication.

Similar to 3D printing, however, computer numerical control (CNC) mills exhibit a finite resolution: stepper motors control the individual axes of the tool movement. This leads to quantization, in particular to step artifacts at shallow slopes, and has to be taken into account in any milled lens design [Papas et al. 2011].

Another source of artifacts specific to milling are the tool paths, which leave tiny grooves that introduce spurious (anisotropic) reflections that require development of optimal milling parameters (not always discussed in the literature but reported, e.g., in [Weyrich et al. 2009b, Papas et al. 2011]). Furthermore, the finite shape of the drill bit limits, for instance, the precision at which concave edges can be milled, which may require additional consideration when optimizing the output geometry [Weyrich et al. 2009b].

Optical Components

In addition to custom-fabricated optical elements, some works make use of off-the-shelf lenses, mirrors and lenslet arrays. Active techniques further employ LCD panels, projectors and lasers. Precision and artifacts inherent to these elements are well understood in the computer vision and computer graphics community, so we omit their discussion here.

4.3. Redirection and Modulation of Light

We recall that the macroscopic appearance of materials stems from their optical properties, i.e., their interaction with light, on a microscopic scale. The physical principles behind these interactions are the redirection of light due to reflection or refraction, modulation of the spectral composition of the incoming light and diffraction. Fabrication, i.e., the reproduction of these properties is made difficult by the fact that the accuracy and resolution of current techniques (see previous section) cannot reproduce the microscopic structures present in real materials. Nevertheless, researchers have attempted to reproduce the statistical behavior of materials interacting with light by exploiting the same principles of light redirection, modulation, and diffraction. This section concentrates on light redirection by means of reflection or refraction and on light modulation, whereas Section 4.4 focuses on diffraction in the form of holographic reproduction.

Light Redirection Only

Among the techniques introduced in Section 3, we identify a family of approaches that concentrates on pure redirection of light by computationally generating geometric structures that yield predefined reflection or refraction patterns.

This principle has been implemented in various modes of operation. Reflective [Weyrich et al. 2009b] and refractive media [Kiser et al. 2013] are not fundamentally different as far as the image formation model is concerned, but allow different application scenarios. A surface geometry that is assembled from discrete patches (facets) [Papas et al. 2011] facilitates the generation of target distribtions, e.g., of surface normals or curvatures. However, it can lead to challenging discrete optimization problems to fulfill additional constraints for manufacturability that are naturally met by smooth representations such as B-splines. Simple normal distributions can be synthesized by shaping the statistics of the surface even when its exact displacement field is impossible to control [Hullin et al. 2011b, Ochiai et al. 2012].

Light Modulation Only

Recently, a group of researchers realized that stacks of light modulating layers can be used to produce an image that varies with the viewing direction resulting in a light field display [Wetzstein et al. 2011]. The layer transparencies are computed by tomographic means from a set of predefined light field views. Modulating the polarization state of light yields a programmable display with superior light throughput [Lanman et al. 2011]. By combining the stacked modulation layers with temporal multiplexing, the researchers demonstrated that higher fidelity of the light field reproduction can be achieved [Wetzstein et al. 2012]. In the former work, a uniform backlight is assumed as illumination source. However, the concept of stacked transparencies can be applied to varying illumination as well [Baran et al. 2012]. In this case, predefined environment illumination and specified images that are to be cast as shadows are used as constraints on the tomographic reconstruction. The authors' results indicate that for sufficiently localized light sources, the desired effect can be achieved.

Combined Redirection and Modulation

The combination of the redirection and modulation of light leads to yet another way of achieving desired light distributions: the principle of spatial multiplexing. Out of the examples discussed in Section 3, this idea was implemented by two approaches that combine integral optics with printing. Fuchs et al. [2008] describe a set of optical designs that are operated in through-mode, i.e., they are illuminated from one side and viewed from the other. An arrangement of lenses maps pairs of illumination and viewing rays to distinct points in a two-dimensional plane, where a printed transparency serves as the modulation layer. Malzbender et al. proposed a reflection-mode redirection and modulation approach as a means of printing reflectance [2012]. On a special type of substrate consisting of small spherical mirrors, a patterned light-blocking layer is super-positioned. This results in the blocking of certain light/view pairs, mimicking a specular reflection function.

Common to both approaches is that their light throughput is limited by the maximum angular frequency content and dynamic range of the displayed scene. A pathological worst-case usage scenario would be the "display" of a mirror where each incident light ray is reflected in exactly one mirror direction. A multiplexing setup that provides light paths for all possible combinations of incoming to outgoing rays then needs to attenuate those paths that carry no light, i.e., almost all of them. Hence, the mirror will reflect only the small fraction of the incoming light that is assigned by the optical mapping to its corresponding mirror direction.

The devices discussed so far were all optically passive. However, the spatial multiplexing principle is more commonly associated with sensor technology and light field recording and projection devices that have been used to enable arbitrary "computational modulation". This approach enables virtual objects lit by real light sources [Hirsch et al. 2009, Hirsch et al. 2013] and even global illumination exchange between real and virtual scene parts [Cossairt et al. 2008]. They circumvent the light throughput problem to some extent by replacing the light ray assignment by computation. We will discuss their practical challenges and limitations separately in Section 4.5.

4.4. Holography

The term "holography" in popular culture is probably the one that is most closely connected to the techniques and goals discussed in this report. Science fiction concepts such as Star Trek's *holographic screens*, the *replicator*, and the *holodeck* in which a complete tangible world indistinguishable from the real one can be created and arbitrarily changed by a computer program has influenced popular and scientific imagination alike. While these futuristic visions are far from realized, holography plays an important role in the life-like representation of usually static objects under usually fixed illumination.

The holographic principle was invented by Dennis Gabor with the goal of improving the imaging capabilities of electron microscopes. He demonstrated the feasibility of his idea in an optical setup [Gabor 1948] and showed that an intensity image of two interfering wave fields, an object

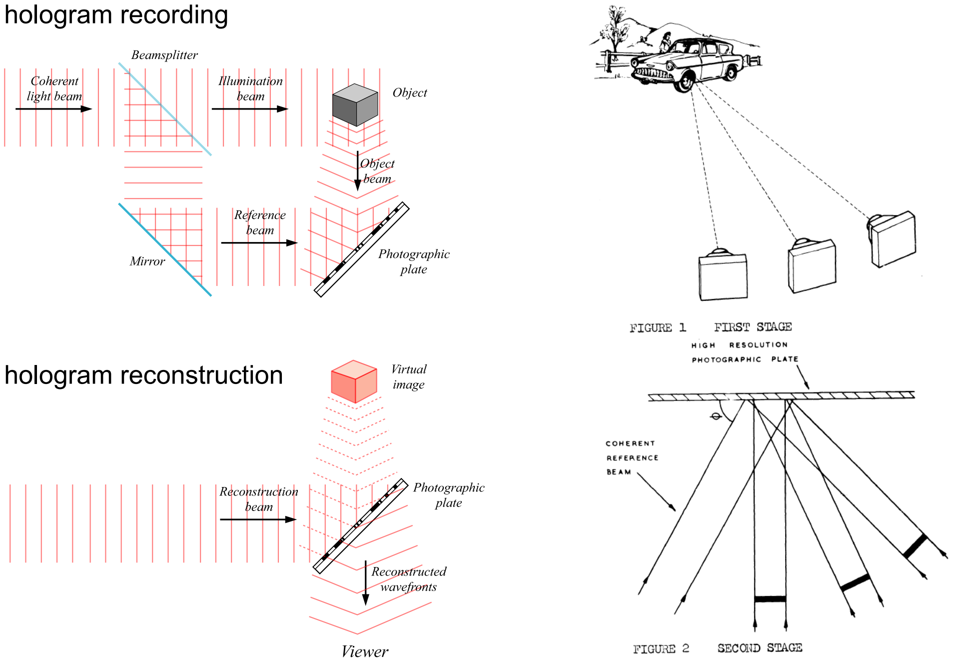

Figure 7: Left: *Principle of holographic recording and reconstruction (illustration adapted from Bob Mellish/Wikipedia).* Right: *Operation principle of holographic stereograms (reproduced from [Redman 1968]). On the top, a light field is recorded with a standard camera. At the bottom, the individual frames are stored in the angular components of a hologram.*

wave and an illuminating reference wave, could reproduce the amplitude and phase information of a complex-valued wave field when illuminated by exactly the same reference wave, see Fig. 7 (left). He realized that this would store three-dimensional information, or full light field information [Levoy and Hanrahan 1996] as we would say today, about the object in a single intensity image. The practical utility of his invention for three-dimensional imaging was, however, limited due to the inavailability of a coherent monochromatic point light source. With the advent of the laser, however, the situation changed and holography became a feasible option. In a landmark paper [Leith and Upatnieks 1962], Leith and Upatnieks showed that an out-of-focus twin image that had diminished the contrast in Gabor's original design could be removed by an off-axis configuration of the reference wave with respect to the hologram plane and the direction of the object wave, enabling the practical application of the concept. The necessity for monochromatic light sources for hologram viewing was lifted by Denisyuk in 1962, enabling the viewing of reflection holograms under normal incandescent illumination [Denisyuk 1963]. Stephen Benton invented white light transmission holograms [Benton and Bove 2007]. An interesting 2003 interview [Johnston 2003] gives his perspective on the development of holography, especially in the early days. A hologram can provide horizontal parallax only (HOP) or full parallax. In our interpretation, a typical full-parallax hologram stores a 4D view space under 0D illumination.

The holographic principle has been extended in many ways. However, the most relevant for the purposes of this report is the development of the *angle multiplexed hologram*. We prefer this term over the more commonly used “holographic stereogram” for reasons that will become clear soon. Angle multiplexed holograms, invented in 1968 by Redman, are very similar to light fields stored in a hologram, in fact a light field recording architecture was used in his experiments [Redman 1968]. The principle of operation is depicted in Fig. 7 (right). The scene itself is replaced by a discrete set of conventional photographs which are then stored in the angular components of a hologram. This step decouples the delicate optical laser recording setup from the scene that is being depicted and enables the hologram generation of outdoor scenes and people. The replacement of the scene depth by the one of the photograph that substitutes it, in conjunction with the discrete number of view points, makes this type of hologram very similar to a light field. The simple superposition of holograms from different directions resulted in contrast loss of the resulting hologram and a modified architecture involving a slit aperture was invented to circumvent this problem [DeBitetto 1969]. A good overview of the developments until the early eighties is given in [Benton 1982]. It was obvious by then that not only perspective views of a scene, but also animations or different x-ray images could be stored in an angle-multiplexed hologram, the latter enabling a three-dimensional view of a patient's internal body structure.

By the early 90's it was reported that 5000 angular views with an angular resolution of 0.01 degrees could be stored in a single holographic medium [Mok 1993].

The DeBitetto recording arrangement was further refined to enable hologram printing [Klug et al. 1993], a technology that was later commercialized [Klug et al. 2001] and that is being offered by Zebra Imaging today. The commercially printed holograms offer full-parallax with 800 × 800 angular views over a 90 degrees viewing zone (personal communication Michael Klug, Zebra Imaging, 2012) and can be produced with a large size.

An interesting application of holograms in the context of this article was the storage of views with changing illumination in a single hologram [Mann 1995]. Mann describes what in computer graphics is called a reflectance field [Debevec et al. 2000] today and calls it the light space. He describes attempts of super-positioning several holograms with different directions of the reference beam within the same hologram and reports success in storing the reflectance field for few point light source positions. However, loss of contrast and distortions of the individual holograms for different reference beam directions prevented him from recording a full reflectance field. He resorted to keeping the viewpoint static and recording the illumination response of a scene. He therefore created a 2D view space with 2D illumination variation.

Other extensions of the holographic principle include the development of temporally varying holograms by means of a spatial light modulator (SLM), an early example being the MIT holo-video system [Benton 1991, Hilaire et al. 1992]. These devices were driven by computer-generated imagery and the holographic fringe patterns to be shown on the SLM where also generated by a computer. The subject is known as computer generated holography. An early overview is given by Dallas [1980]. The computational power required to compute the interference patterns still required a supercomputer with 16K processors in 1993 to achieve update times of 1 frame per second [Lucente 1993]. However, specialized graphics processors (SGI Onyx) where shown to accomplish the same task [Lucente and Galyean 1995] with less hardware effort. In the 90's, computer-generated holography was restricted to horizontal-parallax-only holograms. Today GPU's can synthesize holographic patterns for a few thousand object points in real-time [Ahrenberg et al. 2006] for the full-parallax case. However, the need for computer graphic primitives other than points is not fully accomplished yet. An initial attempt at holographic triangle primitives, offering occlusion, is described in [Ahrenberg et al. 2008]. Alternatively, the direct conversion of light fields into holograms might turn out to be beneficial. To generate a real hologram, as opposed to an angular multiplexed one, however, depth estimation is required [Ziegler et al. 2007].

4.5. Camera-Display Systems

All active approaches introduced so far consist of a light sensing element (camera), a processing step (rendering) and an output mechanism (display). In this section, we will discuss the technical limitations that govern these stages and therefore need to be considered when implementing a display system for material appearance. Our focus is not so much on general limitations of traditional cameras and displays with respect to standard 2D imaging, but rather on aspects that are specific to our application case.

Cameras

In standard digital photography, image sensors have traditionally been marketed by their pixel count, although consumers are becoming increasingly aware that megapixels are not the only valid measure of camera performance, and that small pixels produce noisy pictures in low-light conditions. Many instances of appearance display, however, are based on the capture of higher-dimensional light fields that are multiplexed into the sensor plane, trading angular against spatial resolution. Consequently, pixel count, filtering and alignment issues become much more important than in traditional imaging. A lenslet array on top of a sensor with a Bayer-like color filter arrangement, for instance, could easily create spatial frequencies high enough to challenge even the best demosaicing algorithms. In a system where the main task of the camera is to capture environment light, special attention needs to be paid to its dynamic range. Underexposure and photon count or quantization noise, in this scenario, may be hardly noticeable if the measured light field is used to computationally relight a given scene, effectively averaging many spatio-angular samples in each surface point. Clipping of strong directional light sources, on the other hand, may quickly degrade the quality of the outcome as it causes significant amounts of radiant power to be missing in the measurement. The choice of hardware and the settings with which to operate it need to take these aspects into account.

Displays

Throughout this paper, we use the concept of "display" in a more general sense than usual. Just for this section, let us revert to the traditional notion of displays as devices that emit two-dimensional images based on computerized data.

On the output side of our setup, some requirements are very similar to the ones discussed above for cameras. A high dynamic range is obviously a key ingredient to the high-fidelity reproduction of real-world scenes. If (integral) light field display is the goal, a high resolution is desirable, although first efforts have demonstrated the use of layered LCD panels for light field display without additional lenslets and the associated loss of resolution, but introducing new kinds of artifacts [Lanman et al. 2010]. The combination of LCD panels with other optics or their stacking in multiple layers often causes moiré patterns as a result of interference

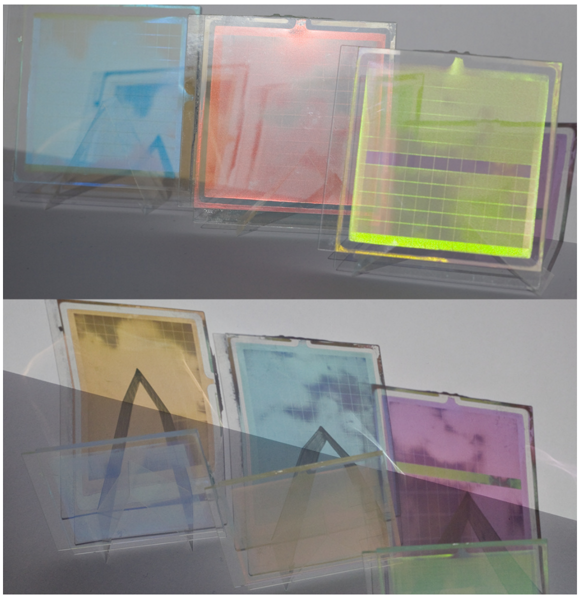

Figure 8: *Three lab prototypes of cholesteric LCD panels. Acting like switchable dichroic mirrors, each of these displays selectively reflects one primary color with a glossy angular distribution (top) while transmitting the rest of the spectrum (bottom). Photos reproduced from [Hullin 2010].*

between the pixel grids. Some of the prototype devices discussed in this report attempt to reduce the worst of these effects by using monochrome displays with color backlights in field-sequential operation.

At a high level, all available display technologies fall within three categories: displays that emit light locally at the pixel level; displays that transmit light from a light source spanning all or many pixels; and reflective displays. Each class has its own set of technologies, which in turn each have benefits and shortcomings.

A particularly interesting direction to consider in our particular context are reflective digital displays. While less dominant in the market, they probably make up the field with the most diverse set of technologies. The best-known working principle is electrophoretic reflective technology (branded as E-ink [Comiskey et al. 1998]) based on transparent microcapsules containing positively charged white pigments and negatively charged dark pigments that are oriented by applying an external electric field. For traditional display purposes, this technology is very attractive in terms of power consumption, the readability in bright environments and the steady-state of the system, even when no power is present at the electrodes. Its disadvantages include low reflectivity, slow switching speeds, limited gray scale and color capabilities. It remains to be seen whether approaches such as Malzbender et al.'s [2012] or direct display of spatially varying BRDFs [Matusik et al. 2009] could be implemented and made re-programmable using this sort of technology.

Other reflective technologies that are in stages closer to prototypes than commercially viable products are based on electrowetting, or use special materials with electrochromic, electrokinetic, electrofluidic or thermochromic properties. For an exhaustive overview we point to Heikenfeld et al. [2011]. Another interesting optical property of light, frustrated total internal reflection, is exploited by Mossman et al. to achieve reflective color displays [2001]. Notably, some of these techniques do not only modulate the color of light, but also its angular distribution (Figure 8), which would make them potential candidates for passive reflectance displays as well.

Finally, one of the most common technologies in video projection is based on MEMS binary mirrors that either direct light towards the screen via a lens or onto a light trap with heat sink. Very fast switching speeds allow the creation of grayscale and color in a field-sequential manner [Armitage et al. 2006]. An extension of this technology to include dual-axis analog positioning of the mirrors could potentially allow for the temporally multiplexed generation of multifacet distributions in extension of Weyrich et al. [2009b] and Hullin et al. [2011b].

5. Discussion

We hope that by now, the reader has realized that the distinction between "fabricated" things and "displays' is mainly a philosophical one. On the one hand, one could argue that it is the most obvious decision between an optically active or passive approach that separates one from the other: if it does not draw electrical power or emit light, we call it fabrication. A closer look, however, reveals that this definition is insufficient. Electronic paper, for instance, actually changes its reflectance with respect to surrounding light, and is hence optically passive. Nevertheless, although such devices lack a backlight, they are commonly considered displays. Further, the term "display" has also been used in a more general sense, for devices that are more or less flat and convey a custom appearance. Devices like Fuchs et al.'s reflectance field display [2008] or the printed reflectance functions by [Malzbender et al. 2012] would fall under this category, even though they involve significant amounts of fabrication and no active components: their display character is motivated by their programmability (during manufacturing or dynamically) to create the illusion of an arbitrary scene or object.

Ultimately, the boundary between appearance fabrication and displays remains ill-defined. We hope that our choice to unite both in a common framework is a constructive contribution, as imperfect as it is bound to be.

As discussed in Section 4.2, appearance fabrication is also closely related to recent work in computer graphics that aims at creating custom shapes and physical properties using innovative fabrication processes. Such work shares many underlying concepts, for instance the requirement of maintaining fabricability for a given medium. At this point, however, we believe that more work in this area will be required before emerging concepts crystalize to the extent allowing for a state-of-the-art report.

Acknowledgements

This work was supported by the German Research Foundation (DFG) through Emmy Noether fellowship IH 114/1-1.

References

[Adelson and Bergen 1991] Adelson, E. H., and Bergen, J. R. 1991. The plenoptic function and the elements of early vision. In *Computational models of visual processing*, M. S. Landy and J. A. Movshon, Eds. MIT Press, ch. 1, 3–20.

[Ahrenberg et al. 2006] Ahrenberg, L., Benzie, P., Magnor, M., and Watson, J. 2006. Computer Generated Holography using Parallel Commodity Graphics Hardware. *Optics Express 14*, 7636–7641.

[Ahrenberg et al. 2008] Ahrenberg, L., Benzie, P., Magnor, M., and Watson, J. 2008. Computer Generated Holograms from Three-Dimensional Meshes Using an Analytical Light Transport Model. *Applied Optics 47*, 10, 1567–1574.

[Alexa and Matusik 2010] Alexa, M., and Matusik, W. 2010. Reliefs as images. *ACM Transactions on Graphics (Proc. SIGGRAPH) 29*, 4 (July), 60:1–60:7.

[Amano 2012] Amano, T. 2012. Shading illusion: A novel way for 3-D representation on paper media. In *CVPR Workshops*, 1–6.

[Armitage et al. 2006] Armitage, D., Underwood, I., and Wu, S. 2006. *Introduction to microdisplays*. Wiley SID series in display technology. Wiley.

[Baran et al. 2012] Baran, I., Keller, P., Bradley, D., Coros, S., Jarosz, W., Nowrouzezahrai, D., and Gross, M. 2012. Manufacturing layered attenuators for multiple prescribed shadow images. *Computer Graphics Forum (Proc. Eurographics) 31*, 2 (May), 603–610.

[Benton and Bove 2007] Benton, S. A., and Bove, V. M. 2007. *White-Light Transmission "Rainbow" Holography*. John Wiley & Sons, Inc., 145–157.

[Benton 1982] Benton, S. A. 1982. Survey of Holographic Stereograms. In *Proceedings of SPIE Vol. 367 Processing and Display of Three-Dimensional Data*, 15–19.

[Benton 1991] Benton, S. A. 1991. Experiments in Holographic Video Imaging. In *Proceedings of SPIE Vol. IS#08*, 247–267.

[Berry 2006] Berry, M. V. 2006. Oriental magic mirrors and the laplacian image. *European Journal of Physics 27*, 1, 109.

[Bickel et al. 2010] Bickel, B., Bächer, M., Otaduy, M. A., Lee, H. R., Pfister, H., Gross, M., and Matusik, W. 2010. Design and fabrication of materials with desired deformation behavior. *ACM Transactions on Graphics (Proc. SIGGRAPH) 29*, 4 (July), 63:1–63:10.

[Cignoni et al. 2008] Cignoni, P., Gobbetti, E., Pintus, R., and Scopigno, R. 2008. Color enhancement for rapid prototyping. In *Proceedings of International Symposium on Virtual Reality, Archeology and Cultural Heritage (VAST)*, Eurographics, 9–16.

[Comiskey et al. 1998] Comiskey, B., Albert, J., Yoshizawa, H., and Jacobson, J. 1998. An electrophoretic ink for all-printed reflective electronic displays. *Nature 394*, 6690, 253–255.

[Cossairt et al. 2008] Cossairt, O., Nayar, S. K., and Ramamoorthi, R. 2008. Light field transfer: Global illumination between real and synthetic objects. *ACM Transactions on Graphics (Proc. SIGGRAPH) 27*, 3 (Aug.), 57:1–57:6.

[Dallas 1980] Dallas, W. J. 1980. Computer-Generated Holograms. In *The Computer in Optical Research*, vol. 41 of *Topics in Applied Physics*. Springer Berlin, 291–366.

[Dana et al. 1999] Dana, K. J., van Ginneken, B., Nayar, S. K., and Koenderink, J. J. 1999. Reflectance and texture of real-world surfaces. *ACM Transactions on Graphics 18*, 1 (Jan.), 1–34.

[Debevec et al. 2000] Debevec, P., Hawkins, T., Tchou, C., Duiker, H.-P., Sarokin, W., and Sagar, M. 2000. Acquiring the Reflectance Field of a Human Face. In *Proceedings of SIGGRAPH*, 145–156.

[DeBitetto 1969] DeBitetto, D. 1969. Holographic Panoramic Stereograms Synthesized from White Light Recordings. *Applied Optics 8*, 8 (Aug.), 1740–1741.

[Denisyuk 1963] Denisyuk, Y. N. 1963. On the reproduction of the optical properties of an object by the wave field of its own scattered radiation. *Optics Spectroscopy 18*, 365–368.

[Dobrowolski 1973] Dobrowolski, J. 1973. Subtractive method of optical thin-film interference filter design. *Applied Optics 12*, 8, 1885–1893.

[Dong et al. 2010] Dong, Y., Wang, J., Pellacini, F., Tong, X., and Guo, B. 2010. Fabricating spatially-varying subsurface scattering. *ACM Transactions on Graphics (Proc. SIGGRAPH) 29*, 4 (July), 62:1–62:10.

[Dorsey et al. 2008] Dorsey, J., Rushmeier, H., and Sillion, F. 2008. *Digital Modeling of Material Appearance*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.

[Finckh et al. 2010] Finckh, M., Dammertz, H., and Lensch, H. P. 2010. Geometry construction from caustic images. In *Proc. European Conference on Computer Vision (ECCV)*, Springer, vol. 6315 of *Lecture Notes in Computer Science*, 464–477.

[Fuchs et al. 2008] Fuchs, M., Raskar, R., Seidel, H.-P., and Lensch, H. P. A. 2008. Towards passive 6D reflectance field displays. *ACM Transactions on Graphics (Proc. SIGGRAPH) 27*, 3 (Aug.), 58:1–58:8.

[Fuchs 2008] Fuchs, M. 2008. *Advanced methods for relightable scene representations in image space*. PhD thesis, Universität des Saarlandes, Saarbrücken, Germany.

[Gabor 1948] Gabor, D. 1948. A New Microscopic Principle.

Nature 161, 4098 (May), 777–778.

[Hašan et al. 2010] Hašan, M., Fuchs, M., Matusik, W., Pfister, H., and Rusinkiewicz, S. 2010. Physical reproduction of materials with specified subsurface scattering. *ACM Transactions on Graphics (Proc. SIGGRAPH) 29*, 4 (July), 61:1–61:10.

[Heikenfeld et al. 2011] Heikenfeld, J., Drzaic, P., Yeo, J.-S., and Koch, T. 2011. Review paper: A critical review of the present and future prospects for electronic paper. *Journal of the Society for Information Display 19*, 2, 129–156.

[Hersch et al. 2003] Hersch, R. D., Collaud, F., and Emmel, P. 2003. Reproducing color images with embedded metallic patterns. *ACM Transactions on Graphics (Proc. SIGGRAPH) 22*, 3 (July), 427–434.

[Hilaire et al. 1992] Hilaire, P. S., Benton, S. A., Lucente, M., and Hubel, P. M. 1992. Color Images with the MIT Holographic Video Display. In *Proceedings of SPIE Vol. 1667 Practical Holography VI*, 73–84.

[Hildebrand et al. 2012] Hildebrand, K., Bickel, B., and Alexa, M. 2012. crdbrd: Shape fabrication by sliding planar slices. *Computer Graphics Forum (Proc. Eurographics) 31*, 2 (May), 583–592.

[Hirsch et al. 2009] Hirsch, M., Lanman, D., Holtzman, H., and Raskar, R. 2009. BiDi screen: A thin, depth-sensing LCD for 3D interaction using light fields. *ACM Transactions on Graphics (Proc. SIGGRAPH ASIA) 28*, 5 (Dec.), 159:1–159:9.

[Hirsch et al. 2013] Hirsch, M., Izadi, S., Holtzman, H., and Raskar, R. 2013. 8D display: a relightable glasses-free 3D display. In *Proc. of ACM SIGCHI (CHI 2013)*, to appear.

[Holroyd et al. 2011] Holroyd, M., Baran, I., Lawrence, J., and Matusik, W. 2011. Computing and fabricating multilayer models. *ACM Transactions on Graphics (Proc. SIGGRAPH ASIA) 30*, 6 (Dec.), 187:1–187:8.

[Hullin et al. 2008] Hullin, M. B., Fuchs, M., Ajdin, B., Ihrke, I., Seidel, H.-P., and Lensch, H. P. A. 2008. Direct visualization of real-world light transport. In *Proceedings of Vision, Modeling, and Visualization (VMV)*, 363–371.

[Hullin et al. 2011a] Hullin, M. B., Lensch, H. P. A., Raskar, R., Seidel, H.-P., and Ihrke, I. 2011. A dynamic BRDF display. In *SIGGRAPH Emerging Technologies*, ACM, New York, NY, USA, SIGGRAPH '11.

[Hullin et al. 2011b] Hullin, M. B., Lensch, H. P. A., Raskar, R., Seidel, H.-P., and Ihrke, I. 2011. Dynamic display of BRDFs. In *Computer Graphics Forum (Proc. Eurographics)*, Blackwell, Llandudno, UK, O. Deussen and M. Chen, Eds., Eurographics, 475–483.

[Hullin 2010] Hullin, M. 2010. *Reconsidering Light Transport – Acquisition and Display of Real-World Reflectance and Geometry*. PhD thesis, Universität des Saarlandes, Saarbrücken, Germany.

[Igarashi et al. 2008] Igarashi, Y., Igarashi, T., and Suzuki, H. 2008. Knitting a 3D Model. *Computer Graphics Forum (Proc. Pacific Graphics) 27*, 7, 1737–1743.

[Ito et al. 2010] Ito, K., Kikuchi, H., Sakurai, H., Kobayashi, I., Yasunaga, H., Mori, H., Tokuyama, K., Ishikawa, H., Hayasaka, K., and Yanagisawa, H. 2010. Sony RayModeler: 360° autostereoscopic display. In *SIGGRAPH Emerging Technologies*, ACM.

[Johnston 2003] Johnston, S. F., 2003. Interview with Stephen Benton. `http://www.aip.org/history/ohilist/29281.html`. Last accessed: 14 Feb 2013.

[Jones et al. 2007] Jones, A., McDowall, I., Yamada, H., Bolas, M., and Debevec, P. 2007. An interactive 360° light field display. In *SIGGRAPH Emerging Technologies*, ACM.

[Kiser et al. 2013] Kiser, T., Eigensatz, M., Nguyen, M., Bompas, P., and Pauly, M. 2013. Architectural caustics – controlling light with geometry. In *Advances in Architectural Geometry 2012*, L. Hesselgren, S. Sharma, J. Wallner, N. Baldassini, P. Bompas, and J. Raynaud, Eds. Springer Vienna, 91–106.

[Klug et al. 1993] Klug, M. A., Halle, M. W., Lucente, M., and Plesniak, W. J. 1993. A Compact Prototype One-Step Ultragram Printer. In *Proceedings of SPIE Vol. 1914*, 15–24.

[Klug et al. 2001] Klug, M. A., Holzbach, M., and Ferdman, A., 2001. Method and Apparatus for Recording One-Step, Full-Color, Full-Parallax, Holographic Stereograms. US Patent 6,330,088.

[Koike and Naemura 2008] Koike, T., and Naemura, T. 2008. BRDF display: interactive view dependent texture display using integral photography. In *IPT/EDT '08: Proceedings of the 2008 Workshop on Immersive Projection Technologies/Emerging Display Technologies*, ACM, New York, NY, USA, 1–4.

[Lanman et al. 2010] Lanman, D., Hirsch, M., Kim, Y., and Raskar, R. 2010. Content-adaptive parallax barriers: optimizing dual-layer 3D displays using low-rank light field factorization. *ACM Transactions on Graphics (Proc. SIGGRAPH ASIA) 29*, 6 (Dec.), 163:1–163:10.

[Lanman et al. 2011] Lanman, D., Wetzstein, G., Hirsch, M., Heidrich, W., and Raskar, R. 2011. Polarization fields: Dynamic light field display using multi-layer LCDs. *ACM Transactions on Graphics (Proc. SIGGRAPH ASIA) 30*, 6 (Dec.), 186:1–186:10.

[Leith and Upatnieks 1962] Leith, E. N., and Upatnieks, J. 1962. Reconstructed Wavefronts and Communication Theory. *Journal of the Optical Society of America A 52*, 10 (Oct.), 1123–1130.

[Lensch 2003] Lensch, H. P. A. 2003. *Efficient, Image-Based Appearance Acquisition of Real-World Objects*. PhD thesis, Universität des Saarlandes.

[Levoy and Hanrahan 1996] Levoy, M., and Hanrahan, P. 1996. Light Field Rendering. In *Proceedings of SIGGRAPH*, 31–42.

[Lippmann 1908] Lippmann, G. 1908. La Photographie Intégrale. *Comptes Rendus 146*, 446–451.

[Lucente and Galyean 1995] Lucente, M., and Galyean, T. A. 1995. Rendering Interactive Holographic Images. In *Proceedings of SIGGRAPH*, 387–394.

[Lucente 1993] Lucente, M. 1993. Interactive Computation of Holograms Using a Look-up Table. *Journal of Electronic Imaging 2*, 1 (Jan.), 28–34.

[Malzbender et al. 2012] Malzbender, T., Samadani, R., Scher, S., Crume, A., Dunn, D., and Davis, J. 2012. Printing reflectance functions. *ACM Transactions on Graphics 31*, 3 (May), 20:1–20:11.

[Mann 1995] Mann, S. 1995. Recording Lightspace so Shadows and Highlights Vary with Varying Viewing Illumination. *Optics Letters 20*, 24 (Dec), 2538–2540.

[Matusik et al. 2009] Matusik, W., Ajdin, B., Gu, J., Lawrence, J., Lensch, H. P. A., Pellacini, F., and Rusinkiewicz, S. 2009. Printing spatially-varying reflectance. *ACM Transactions on Graphics (Proc. SIGGRAPH ASIA) 28*, 5 (Dec.), 128:1–128:9.

[Mitani and Suzuki 2004] Mitani, J., and Suzuki, H. 2004. Making papercraft toys from meshes using strip-based approximate unfolding. *ACM Transactions on Graphics (Proc. SIGGRAPH) 23*, 3 (Aug.), 259–263.

[Mok 1993] Mok, F. H. 1993. Angle-Multiplexed Storage of 5000 Holograms in Lithium Niobate. *Optics Letters 18*, 11, 915–917.

[Mossman et al. 2001] Mossman, M. A., Kotlicki, A., Whitehead, L. A., Biernath, R. W., and Rao, S. P. 2001. 39.1: New reflective color display technique based on total internal reflection and subtractive color filtering. *SID Symposium Digest of Technical Papers 32*, 1, 1054–1057.

[Nayar et al. 2004] Nayar, S., Belhumeur, P., and Boult, T. 2004. Lighting sensitive display. *ACM Transactions on Graphics 23*, 4, 963–979.

[Nicodemus et al. 1977] Nicodemus, F., Richmond, J., Hsia, J., Ginsberg, I., and Limperis, T. 1977. Geometrical considerations and nomenclature for reflectance. *National Bureau of Standards, Washington, DC. Inst. for Basic Standards*.

[Nimeroff et al. 1994] Nimeroff, J. S., Simoncelli, E., and Dorsey, J. 1994. Efficient re-rendering of naturally illuminated environments. In *Proceedings of Eurographics Workshop on Rendering*, 359–373.

[Ochiai et al. 2012] Ochiai, Y., Oyama, A., and Toyoshima, K. 2012. A colloidal display: membrane screen that combines transparency, BRDF and 3D volume. In *SIGGRAPH Emerging Technologies*, ACM, New York, NY, USA, SIGGRAPH '12, 2:1–2:1.

[Papas et al. 2011] Papas, M., Jarosz, W., Jakob, W., Rusinkiewicz, S., Matusik, W., and Weyrich, T. 2011. Goal-based caustics. *Computer Graphics Forum (Proc. Eurographics) 30*, 2 (June), 503–511.

[Patow and Pueyo 2005] Patow, G., and Pueyo, X. 2005. A survey of inverse surface design from light transport behavior specification. *Computer Graphics Forum 24*, 4, 773–789.

[Raskar et al. 2001] Raskar, R., Welch, G., Low, K.-L., and Bandyopadhyay, D. 2001. Shader lamps: Animating real objects with image-based illumination. In *Proceedings of Eurographics Workshop on Rendering*, Springer-Verlag, London, UK, 89–102.

[Redman 1968] Redman, J. 1968. The Three-Dimensional Reconstruction of People and Outdoor Scenes Using Holographic Multiplexing. *Proceedings of SPIE Seminar-in-Depth on Holography 15*, 117–122.

[Rusinkiewicz and Marschner 2000] Rusinkiewicz, S., and Marschner, S., 2000. Measurement I – BRDFs, Oct. Appearance Models for Computer Graphics and Vision. Script for the course CS448C: Topics in Computer Graphics, Stanford University.

[Siedow 2008] Siedow, N. 2008. *Information sheet on freeform lenses*. Fraunhofer Institute for Industrial Mathematics (ITWM).

[Skouras et al. 2012] Skouras, M., Thomaszewski, B., Bickel, B., and Gross, M. 2012. Computational design of rubber balloons. *Computer Graphics Forum (Proc. Eurographics) 31*, 2 (May), 583–592.

[Song et al. 2009] Song, Y., Tong, X., Pellacini, F., and Peers, P. 2009. Subedit: a representation for editing measured heterogeneous subsurface scattering. *ACM Transactions on Graphics (Proc. SIGGRAPH) 28*, 3 (July), 31:1–31:10.

[Stava et al. 2012] Stava, O., Vanek, J., Benes, B., Carr, N., and Měch, R. 2012. Stress relief: improving structural strength of 3D printable objects. *ACM Transactions on Graphics (Proc. SIGGRAPH) 31*, 4 (July), 48:1–48:11.

[Tandianus et al. 2012] Tandianus, B., Johan, H., and Seah, H. S. 2012. Caustic object construction based on multiple caustic patterns. In *Proceedings of International Conference on Computer Graphics, Visualization and Computer Vision (WSCG)*, vol. 20, 37–46.

[Wada 1992] Wada, O. 1992. Control of fiber form and yarn and fabric structure. *Journal of the Textile Institute 83*, 3, 322–347.

[Wetzstein et al. 2011] Wetzstein, G., Lanman, D., Heidrich, W., and Raskar, R. 2011. Layered 3D: tomographic image synthesis for attenuation-based light field and high dynamic range displays. *ACM Transactions on Graphics (Proc. SIGGRAPH) 30*, 4 (July), 95:1–95:12.

[Wetzstein et al. 2012] Wetzstein, G., Lanman, D., Hirsch, M., and Raskar, R. 2012. Tensor Displays: Compressive Light Field Synthesis using Multilayer Displays with Directional Backlighting. *ACM Transactions on Graphics (Proc. SIGGRAPH) 31*, 4 (July), 80:1–80:11.

[Weyrich et al. 2007] Weyrich, T., Deng, J., Barnes, C., Rusinkiewicz, S., and Finkelstein, A. 2007. Digital bas-relief from 3D scenes. *ACM Transactions on Graphics (Proc. SIGGRAPH) 26*, 3 (July).

[Weyrich et al. 2009a] Weyrich, T., Lawrence, J., Lensch, H. P., Rusinkiewicz, S., and Zickler, T. 2009. Principles of appearance acquisition and representation. *Foundations and Trends in Computer Graphics and Vision 4*, 2 (Oct.), 75–191.

[Weyrich et al. 2009b] Weyrich, T., Peers, P., Matusik, W., and Rusinkiewicz, S. 2009. Fabricating microgeometry for custom surface reflectance. *ACM Transactions on Graphics (Proc. SIGGRAPH) 28*, 3 (July), 32:1–32:6.

[Ziegler et al. 2007] Ziegler, R., Bucheli, S., Ahrenberg, L., Magnor, M., and Gross, M. 2007. A Bidirectional Light Field - Hologram Transform. *Computer Graphics Forum 26*, 3, 435–446.

About the Authors

Matthias B. Hullin is a post-doctoral fellow at the University of British Columbia. Prior to joining the UBC, Hullin worked as a researcher in the graphics department at the MPI Informatik. He received a PhD (summa cum laude) in computer science from Saarland University in 2010, and a Diplom in physics from Kaiserslautern University in 2007. His doctoral dissertation focused on 3D scanning of optically challenging materials as well as acquisition, modeling and display of real-world material appearance, and was awarded the Otto Hahn Medal of the Max Planck Society in 2011. Among his research interests are physical models of light transport, and related forward and inverse problems in computational photography and display.

Ivo Ihrke is head of the research group "Generalized Image Acquisition and Analysis" which is supported by an Emmy-Noether fellowship of the German Research Foundation (DFG). Prior to that he was leading a research group within the Cluster of Excellence "Multimodal Computing and Interaction" at Saarland University. He is an Associate Senior Researcher at the MPI Informatik, and associated with the Max-Planck Center for Visual Computing and Communications. Prior to joining Saarland University he was a postdoctoral research fellow at the University of British Columbia, supported by the Alexander von Humboldt-Foundation. He received a MS degree in Scientific Computing from the Royal Institute of Technology (KTH), Stockholm (2002) and a PhD (summa cum laude) in Computer Science from Saarland University (2007). His main research interest are the modeling of forward and inverse light transport processes and computational algorithms for solving these large scale problems in the context of computational imaging, measurement, and display.

Wolfgang Heidrich holds the Dolby Research Chair in Computer Science at the University of British Columbia. He received a PhD in Computer Science from the University of Erlangen in 1999, and then worked as a Research Associate in the Computer Graphics Group of the Max-Planck-Institute for Computer Science in Saarbrücken, Germany, before joining UBC in 2000. Heidrich's research interests lie at the intersection of computer graphics, computer vision, imaging, and optics. In particular, he has worked on High Dynamic Range imaging and display, image-based modeling, measuring, and rendering, geometry acquisition, GPU-based rendering, and global illumination. Heidrich has written over 100 refereed publications on these subjects and has served on numerous program committees. He was the program co-chair for Graphics Hardware 2002, Graphics Interface 2004, the Eurographics Symposium on Rendering 2006, and ProCams 2011.

Tim Weyrich is an associate professor (Senior Lecturer) in the Virtual Environments and Computer Graphics group in the Department of Computer Science, University College London, and co-founder and associate director of the UCL Centre for Digital Humanities. Prior to coming to UCL, he was a Postdoctoral Teaching Fellow of Princeton University, working in the Princeton Computer Graphics Group, a post he took after having received his PhD from ETH Zurich, Switzerland, in 2006. Weyrich's research interests are appearance modeling and fabrication, point-based graphics, 3D reconstruction, cultural heritage acquisition and digital humanities.

Gerwin Damberg is a PhD student in the Computer Science department at the University of British Columbia (UBC), Canada, and one of the founders of MTT Innovation, a start-up company developing next generation projection technology. Following his electrical engineering degree at the University of Applied Sciences in Karlsruhe, Germany in 2004, he joined BrightSide Technologies, a start-up company out of UBC, as a researcher to develop the first commercially available high dynamic range (HDR) displays, cameras and video codecs. Following the successful acquisition of the company in 2007, Gerwin joined Dolby Laboratories, where he led research groups, working on algorithms derived from perceptual models as well as developing novel display and projection prototypes, both for research studies and commercial applications.

Martin Fuchs holds a junior professor position at the visualization research institute at Stuttgart University (VISUS), Germany, which he assumed in 2011 after postdoctorate studies at Princeton University, USA, which were in part supported by the German Academic Exchange Service (DAAD). He received his Diploma (2004) and PhD (2008, summa cum laude) degrees in computer science from Saarland University / MPI Informatik, Germany, where he initially focused his research interests on appearance measurement and image-space lighting techniques. Since, he contributed to the state of the art in lighting-sensitive passive display technology, appearance synthesis, and computational photography techniques.