Advances in Geometry and Reflectance Acquisition

M. Weinmann[1] and F. Langguth[2] and M. Goesele[2] and R. Klein[1]

[1]University of Bonn, Germany
[2]TU Darmstadt, Germany

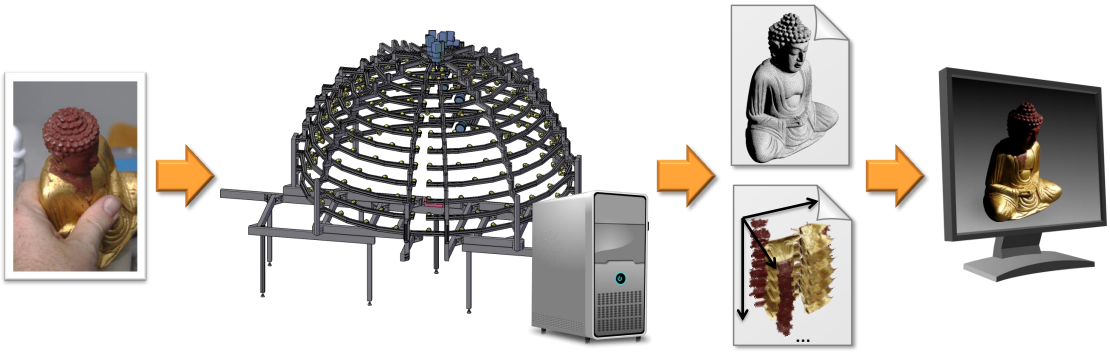

Figure 1: The acquisition pipeline: An object of interest is acquired with specific acquisition devices. The digitized geometry (exemplarily illustrated as a triangle mesh) and appearance (exemplarily illustrated as a BTF, i.e. a stack of surface textures for different configurations of view and light directions) enable a subsequent visualization of the digitized object.

Abstract

This tutorial is focused on acquisition methods for geometry and reflectance as well as strategies towards an efficient acquisition pipeline to fulfill the demands of industry with respect to mass digitization of 3D contents. We provide a thorough overview of the standard methods for the acquisition of both geometry and reflectance of surfaces with different types of reflectance behavior ranging from diffuse over opaque to specular surfaces or even translucent and transparent surfaces as well as the necessary preliminaries of material appearance and setup calibration. As standard acquisition techniques are only well-suited for a limited range of surface materials, we will also discuss strategies on how an efficient, fully automatic acquisition can still be achieved when no prior information with respect to the surface reflectance behavior is available. In addition, a discussion of strategies regarding an acquisition in the wild, i.e. under uncontrolled conditions, is provided.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation—Digitizing and scanning; Digitization and Image Capture [I.4.1]: I.4.1—Picture/Image GenerationRadiometry; Reflectance; Scanning

Schedule

1. Schedule

Part I - Introduction

- Introduction and Welcome
- Motivation
- The Acquisition Ecosystem
- Course Objectives
- Course Outline

Part II - Preliminaries of Material Appearance

- Basics of Material Appearance
- A Taxonomy of Surface Classes
- Commonly used Reflectance Functions

Part III - Calibration

- Geometric Calibration
- Radiometric Calibration

This work is based on an earlier work:
Advances in Geometry and Reflectance Acquisition
SIGGRAPH Asia 2015 Courses, © ACM, 2015.
http://dx.doi.org/10.1145/2818143.2818165

Part IV - Advances in Geometry Acquisition

- Geometry Acquisition of Rough Surfaces with Diffuse or Near Diffuse Reflectance
- Geometry Acquisition of Glossy Surfaces with Mixed Diffuse and Specular Reflectance
- Geometry Acquisition of Smooth Surfaces with Ideal or Near Ideal Specular Reflectance
- Geometry Acquisition of Surfaces Where Light is Scattered Multiple Times Underneath the Surface
- Geometry Acquisition of Smooth Surfaces with Ideal or Near Ideal Specular Refraction

Part V - Mesh Reconstruction

- Graph Cut Based Methods
- Poisson Reconstruction
- Smooth Signed Distance Surface Reconstruction
- Floating Scale Surface Reconstruction

Part VI - Advances in Reflectance Acquisition

- Texture Acquisition
- BRDFs and Their Acquisition
- SVBRDFs and Their Acquisition
- BTFs and Their Acquisition
- BSSRDFs and Their Acquisition
- Light Fields and Reflectance Fields and Their Acquisition

Part VII - Novel Trends in Acquisition

- Increasing the Efficiency of Automated Acquisition Pipelines
- Acquisition *in the Wild*

Part VIII - Applications and Conclusions

- Final Remarks and Open Challenges
- Discussion

2. Introduction

The rich information perceived via the senses of the human perceptual system such as sight, hearing, taste, smell, touch and balance greatly supports us in exploring our environment and, combined with our gained experience, allows us to infer insights regarding daily life tasks such as how we have to interact with the content of the surrounding environment. Among these senses, sight is probably the most important one for these interactions as it allows a touchless close-range and far-range perception of our environment, whereas the other senses allow a more limited perception and interaction. The content we visually perceive in a scene is characterized by the presence and arrangement of objects, their shapes, as well as their attached colors and textures. However, observed colors and textures do not appear arbitrarily, but result from a material-specific surface reflectance behavior based on characteristic material traits and are strongly depending on the interplay of surface geometry, material properties and illumination characteristics. From the gained impression of the observed objects and materials, further insights regarding physical and functional properties such as their deformability, fragility, density, weight, value, thermal conductivity or toxicity can be derived. Indeed, based on visual perception, we not only get impressions about a characteristic look but also an accompanying "feel" for materials.

However, we do not only encounter physical objects and materials in everyday life, but also more and more digital objects or materials. With the steady increasing capacity of computing hardware, the creation of digital content has gained an enormous attention in the recent decades. While the numerous traditional applications focused on the demands in the entertainment industry with respect to video games and movies, digital content has also become an ubiquitous component for a much wider range of applications such as virtual prototyping, advertisement or digital preservation of objects for cultural heritage. In addition, the upcoming interest in virtual reality applications or augmented reality applications further pushes the need for digital 3D content. While content creation can be performed manually by designers in a time-consuming process, there is often no alternative to an accurate digitization of physical objects including both geometry and reflectance characteristics where even the finest details of surface geometry and surface materials should be accurately captured so that they can be reproduced in virtual scenarios.

2.1. Material Acquisition in Industry and Object Digitization in Cultural Heritage

For many industrial applications such as movies, games, advertisement or visual prototyping, there is a need for digitized materials to represent object surfaces of any kind. These digitized materials should transport the characteristic "look" and "feel" for the individual real materials so that the resulting digitized objects such as e.g. cloth, furniture or cars appear as realistic as possible to the observer and increase the realism of the scenarios shown in video games, movies or marketing applications. In order to obtain accurate digitized surrogates of physical materials, both the surface profile and the surface reflectance behavior have to be acquired appropriately.

While these aforementioned examples typically focus on the acquisition of flat material samples, several applications also require objects to be digitized completely with both their 3D shape and their reflectance behavior. For instance, the digital preservation of objects has a great importance in cultural heritage. The typical acquisition scenario for cultural heritage objects is shown in Figure 1. Both the geometry and the reflectance characteristics are acquired with certain acquisition techniques such as e.g. highly optimized devices as discussed in [SSW*14]. While the geometry is often stored as a point cloud, a mesh or in a volumetric representation, the reflectance is stored in terms of different reflectance functions that describe the material appearance depending on a multitude of involved parameters.

2.2. The Acquisition Ecosystem

The basic components of the acquisition ecosystem as depicted in Figure 2 are given by (i) hardware components that are required for the acquisition, (ii) techniques for the radiometric and geometric calibration of the involved components, (iii) geometry acquisition

techniques, (iv) reflectance acquisition techniques, (v) assistance systems that might increase the efficiency of the acquisition, such as e.g. a prior material recognition step to select adequate acquisition techniques for the object of interest, and (vi) visualization techniques to depict the digitized models.

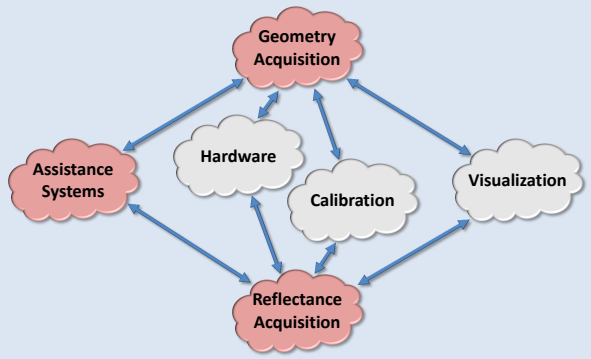

Figure 2: The acquisition ecosystem. In this course, we will discuss the advances with respect to geometry and reflectance acquisition techniques, and we will present techniques that might increase the efficiency of the automatic acquisition pipeline.

The required *hardware components* include a large variety of different sensor types such as cameras or other detectors, components for actively illuminating objects that are required for several of the methods to acquire geometry and reflectance, components that might be required to build the respective acquisition devices such as turntables or other components required to build a gantry where the involved components can be mounted.

Calibration techniques are required to allow an acquisition of objects from multiple viewpoints and under different illumination conditions. Therefore, not only the arrangement of the involved components has to be determined, but also the characteristics of e.g. the illuminants or the sensors have to be considered. This allows measurements taken from different viewpoints or taken under different illuminations to be combined in order to obtain the characteristics of surface reflectance behavior of the acquired objects.

In order to obtain faithful reconstructions of the appearance for the objects to be digitized, *geometry acquisition techniques* have to be applied that allow to acquire the surface geometry with its details adequately. As it is demonstrated by the wide range of diverse approaches presented in literature, there is no technique capable of handling the full range of materials that objects might consist of. Instead, the numerous developed techniques typically focus on only a limited fraction of the materials. Some methods are only applicable to acquire objects with diffuse surface reflectance, some techniques are also applicable if there is an additional specular component in the surface reflectance behavior as seen in many glossy materials. Furthermore, several techniques are directly designed for mirroring objects, for translucent objects or for transparent objects. After applying the individual acquisition techniques, the measured object geometry is typically stored in the form of point clouds, meshes or volumetric representations.

Similar to the geometry acquisition techniques, *reflectance acquisition techniques* that have been developed to capture the surface reflectance characteristics are also only applicable for a limited range of materials. The complexity of the surface reflectance behavior determines the number of parameters that are involved and, hence, have to be measured. These measurements allow surface reflectance to be stored in terms of certain reflectance models that are capable of representing certain materials. Different acquisition devices have been developed for acquiring the respective reflectance models that have been discussed in literature so far. As the acquisition times increase significantly with each additional parameter that is measured, the user typically selects the reflectance model according to his expertise so that only the required parameters are measured.

In many cases, especially in the industrial scenarios, the acquisition process should be as efficient as possible to allow for a highly efficient mass digitization of 3D contents. Therefore, *assistance systems* are required that allow a highly efficient automatic acquisition process. In particular, only the really necessary data should be measured in order to reduce the acquisition time as much as possible.

Finally, the digitized geometry and reflectance properties of the respective objects need to be visualized again which leads to a need for adequate *visualization techniques*.

2.3. Course Objectives

The highly accurate digitization of materials or objects has represented one of the main research topics in computer graphics and computer vision for decades. In particular, there is a demand for an extremely high visual quality of digitized materials or objects. This means that not only highly detailed and hole-free 3D surface geometries, that even include fine surface details such as scratches or engravings, should be captured, but also reflectance characteristics have to be accurately acquired, ideally directly on the true surface geometry, to achieve photo-realistic visual reproductions of real-world objects. This allows fine material characteristics such as e.g. the weave pattern of cloth or the grain of leather to be accurately reproduced in a virtual setting. Automatic reflectance acquisition devices such as gonioreflectometers or camera arrays are capable of taking images of an object under a huge multitude of varying viewing and illumination conditions. In order to reconstruct the underlying surface geometry, these setups are typically equipped with laser scanners or structured light systems as well.

However, accurately capturing the 3D shape of objects and optical properties of materials is especially challenging for objects made of materials with complex reflectance behavior. As a result, there is still a need for accurate acquisition techniques which are suited to handle surfaces with complex reflectance behavior such as mirroring surfaces, translucency, transparency, etc.. Unfortunately, current state-of-the-art acquisition techniques are designed for only a limited range of surface reflectance and the user typically selects appropriate ones based on his experience (see Figure 3). This represents the typical acquisition scenario with respect to cultural heritage artifacts. In a similar way, the user has to select the respective industrial work flow depending on the material of the object in the scope of many industrial applications. In both cases, the ultimate

prerequisite is formed by the existence of acquisition techniques appropriate for the different surface materials that might occur. For objects with heterogeneous surface reflectance behavior due to different surface materials, such as diffuse and mirroring parts of the surface, the acquisition expert has to select adequate techniques for the different material types and to merge the resulting reconstructions (see Figure 4).

Furthermore, the demand for creating large quantities of digital contents also pushes the focus towards fully automatic and highly efficient solutions that allow masses of objects to be acquired as fast as possible. In addition, handling the individual objects appropriately in a supply chain using robotized control systems requires work processes to be adapted to their respective surface materials. Therefore, as-efficient-as-possible solutions have become a crucial prerequisite in many industrial applications. Instead of the above-mentioned traditional user-guided acquisition schemes, the presence of the individual, occurring surface materials should guide the acquisition process if there is no prior knowledge about the object surface available. Hence, automatically recognizing the surface materials represents a key component for automatic acquisition pipelines as it allows making decisions such as reasoning about fragility, deformability, weight, etc., which, in turn, naturally guide the interaction of humans with objects in daily life.

Instead of naïvely processing each region several times using different methods as illustrated in Figure 5, only those methods should be selected from a pool of material-specific acquisition techniques for which at least a certain part of the object surface exhibits the corresponding reflectance behavior assumed by these methods (see Figure 6). For instance, in order to acquire the surface geometry of a heterogeneous object with both diffuse and mirroring surface parts, only a reliable shape acquisition technique for diffuse objects and a respective method for mirroring objects should be involved. Furthermore, based on the material properties, there is also the possibility to automatically detect cases where none of the available reconstruction techniques is appropriate. In an initial stage before the actual acquisition, it is therefore desirable to reliably recognize the present surface materials. Subsequently, these recognized surface materials can be used to guide the acquisition process.

In the scope of this course, we focus on

- advances in geometry acquisition,
- advances in reflectance acquisition, and
- concepts towards a more efficient automatic acquisition process,

which cover techniques belonging to three components of the aforementioned acquisition ecosystem as shown in Figure 2 that are most crucial for an accurate acquisition process.

2.4. Course Outline

The course is organized in six parts. After this introduction with the basic motivation and the main objectives of this course in Section 2, we discuss the preliminaries of material appearance that are relevant for an adequate acquisition process in Section 3. Furthermore, we shortly discuss fundamentals with respect to the calibration of individual setup components in Section 4. This is followed by a detailed review of the advances in geometry acquisition (see Section 5), mesh reconstruction (see Section 6) and in reflectance acquisition (see Section 7). Finally, Section 8 concludes the course with a survey on different application scenarios and novel trends in acquisition.

3. Preliminaries of Material Appearance

In order to analyze material appearance, it is essential to have a closer look at material properties and how these can be visually derived from the image content. In this context, a key observation is that the visual complexity of surface appearance is characterized by the complex interplay of surface material, surface geometry and illumination. For this reason, human perception can only observe material appearance depending on all of the involved modalities of material properties, surface geometry and illumination conditions determined by the environment. Similarly, standard acquisition devices are also only capable of capturing the coupling of the respective modalities, which consequently also influence the results of image analysis such as extracted feature descriptors. Directly separating these modalities would require a-priori information regarding a subset of these modalities and, hence, turns out to be a chicken-and-egg problem.

In the following, we will first give a brief overview on the characteristics of material appearance and discuss the dependency of material appearance with respect to scale, illumination and surface geometry (Section 3.1). Subsequently, we discuss a taxonomy of surface classes with respect to the characteristics of light transport induced by surface reflectance properties which is relevant for 3D geometry acquisition (Section 3.2). This is followed by an overview of commonly used reflectance models (Section 3.3).

3.1. Basics of Material Appearance

When looking at the materials of objects present in our daily life, we may easily get a first impression of the complexity of visual material appearance. Some exemplary objects or material samples are depicted in Figure 7. While some of the respective materials are flat, others have a characteristic relief structure. While some of them have only a single color, others are colorful. On some objects, we observe specular highlights or even see reflections of the environment, while other objects appear matte. Of course, there are many more of such examples in daily life. In order to understand the key effects influencing material appearance, we first may have a closer look at the underlying physical effects that characterize material appearance.

Before we discuss the physical background with respect to light exchange at material surfaces in Section 3.1.2, we first briefly review some basics of radiometry in the following Section 3.1.1.

3.1.1. Radiometric Quantities

The term *light* can be seen as a synonym for electromagnetic radiation. Due to the characteristics of light propagation, light is often represented by straight lines and the respective physical relations based on this ray representation are denoted as ray optics. Based on the mathematical definition, a ray $\mathbf{r}$ can be parameterized as a

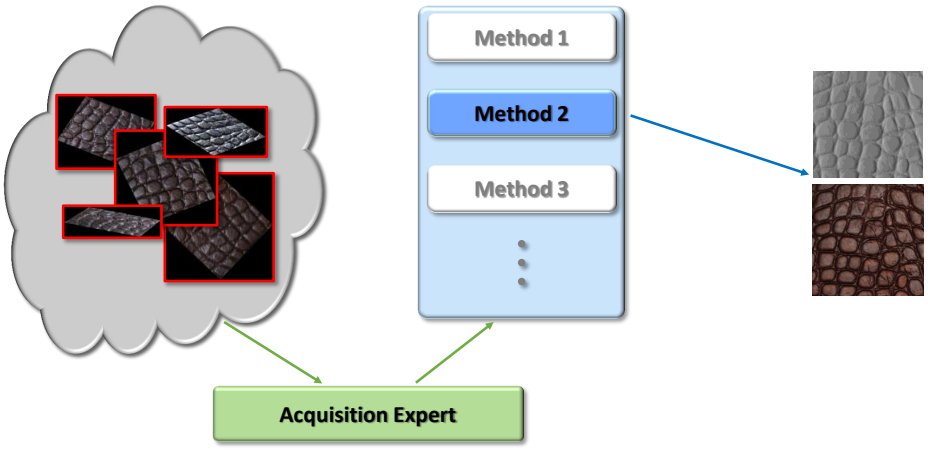

Figure 3: User-guided acquisition process: An acquisition expert selects appropriate acquisition techniques based on his experience.

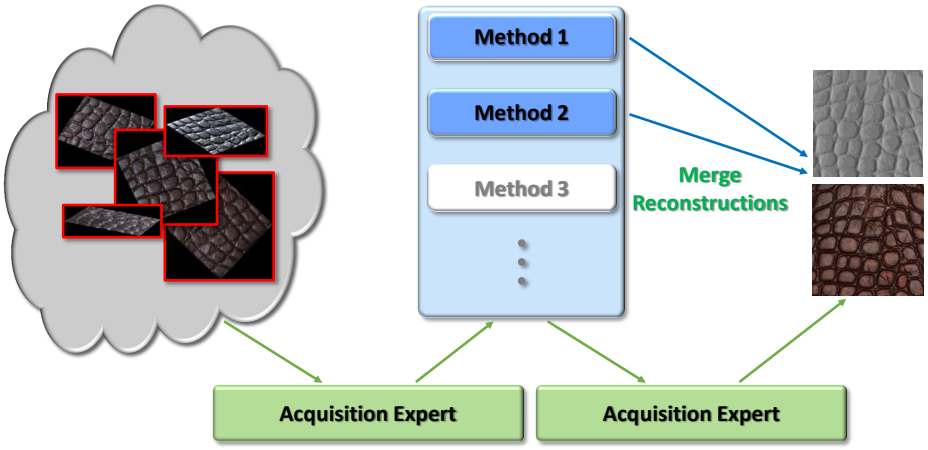

Figure 4: User-guided acquisition process with a manual merging step: An acquisition expert selects appropriate acquisition techniques based on his experience and manually merges the results obtained from the individual techniques to a final reconstruction result.

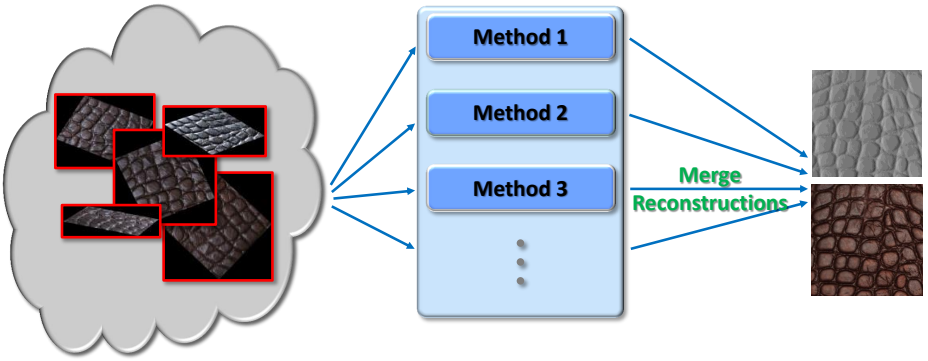

Figure 5: Naïve automatic acquisition process: The object of interest is measured using all the available acquisition techniques. Subsequently, the individual reconstructions have to be merged in order to get an adequate reconstruction.

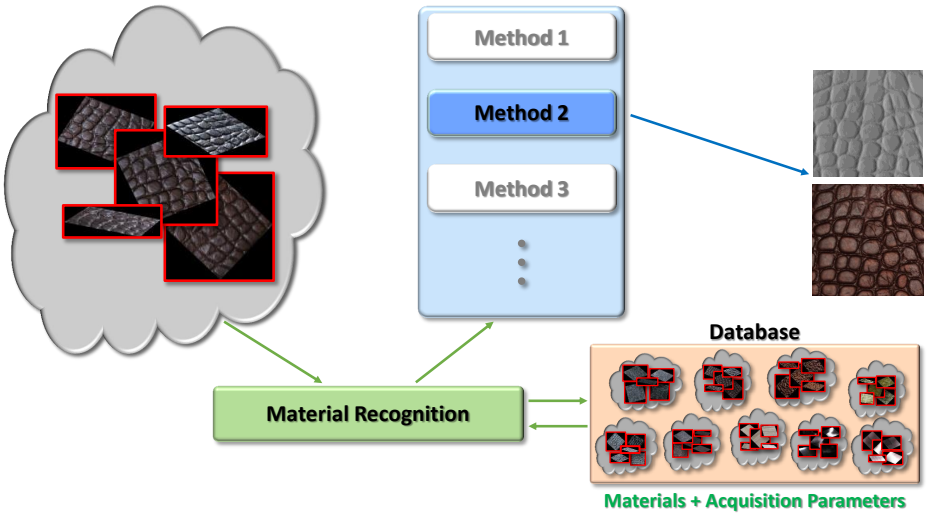

Figure 6: Example of a more efficient automatic acquisition process: Based on a prior material recognition step, the respective annotations of the most similar material in a database can be used to guide the acquisition process. Consequently, only the techniques that are adequate regarding the reconstruction of the particular material are involved which leads to a significant increase in efficiency.

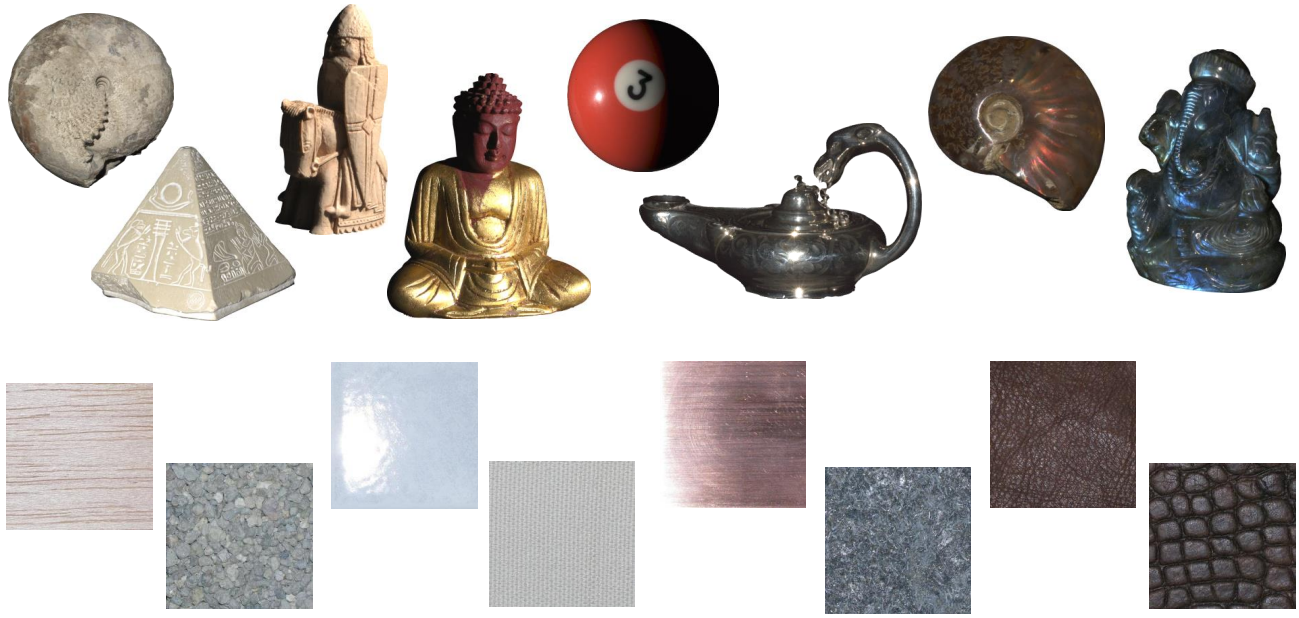

Figure 7: Illustration of exemplary objects and material samples with different appearance characteristics.

mapping $\mathbb{R}^+ \rightarrow \mathbb{R}^3$ using the ray origin $\mathbf{o} \in \mathbb{R}^3$ and a ray direction $\mathbf{d} = (d_1, d_2, d_3)^T$

$$\begin{aligned} \mathbf{r} : \mathbb{R}^+ &\rightarrow \mathbb{R}^3 \\ s &\mapsto \mathbf{o} + s\mathbf{d}, \end{aligned} \tag{1}$$

where $s \in \mathbb{R}$ represents a scalar value.

Light might be monochromatic, i.e. it contains radiation of only one single wavelength which is valid mostly for lasers, or polychromatic, i.e. it contains radiation of multiple wavelengths. The usual case of polychromatic light is typically specified in terms of a spectral power distribution $\lambda \in \mathbb{R}^+ \mapsto L(\lambda) \in \mathbb{R}^+$, that describes the amount of electromagnetic radiative power L at every individual wavelength λ. The portion of the light spectrum that is visible to the human eye covers wavelengths λ between 380nm (violet) and 780 nm (red). This visible spectrum is adjoined by the *ultraviolet*

(UV) spectrum for wavelengths below 380nm and by the *infrared* (IR) spectrum for wavelengths above 780nm. In the scope of this course, we will focus on the tristimulus values red (R), green (G) and blue (B). These tristimulus values can be derived from a spectral power distribution based on inner products with suitable color matching functions such as e.g. CIE RGB [SG31].

Light can be quantified in terms of the *radiant flux* or *radiant power* Φ ([J s^{-1}] or [W]) which is given by the flow of *radiant energy* Q [J] per unit time dt, i.e. $\Phi = \frac{dQ}{dt}$. If more radiant flux is passing a region with fixed extent, the region will appear brighter. Furthermore, the definition of radiant flux Φ does also not depend on the direction of the light flow. For diffuse surfaces, the light flux flows through or from a surface area uniformly into all directions, and, hence, the area looks equally bright under different viewing directions. Otherwise, the surface brightness might also vary for different viewing angles if the material shows a preferred direction of reflection as given for e.g. glossy or specular materials.

In order to provide a better understanding of the light flow and its direction, we have to consider further radiometric quantities. The differential amount of flux that is arriving at the surface per unit area dA is defined as *irradiance* $E = \frac{d\Phi}{dA}$ [W m^{-2}]. Similarly, the *exitance* is also defined by $E = \frac{d\Phi}{dA}$, however, it represents the differential flux leaving the surface per unit area dA. For surfaces with inhomogeneous flux, the irradiance or exitance depends on the local surface point $\mathbf{x}$, i.e. $E = E(\mathbf{x})$. In contrast, for surfaces with homogeneous flux distribution, the irradiance can be represented by the total flux per surface area.

Furthermore, the flux per differential solid angle dω is defined as *intensity* $I = \frac{d\Phi}{d\omega}$ [W sr^{-1}]. The unit of the solid angle is *Steradian* [sr]. If d$\omega = 1$sr, the area of a sphere with radius r inside the cone of the solid angle has the value r^2. As the surface of a sphere with radius r is given by $4\pi r^2$, the sphere of directions around a point has a solid angle of 4πsr. Consequently, the intensity of a point light source that emits a radiant flux of ΦW uniformly into all directions on the unit sphere can be computed as $\frac{\Phi}{4\pi}$W sr^{-1}.

Finally, the *radiance* L denotes the differential amount of radiant flux dΦ per projected differential area dA and per unit solid angle dω, i.e. $L = \frac{d\Phi}{dA\, d\omega}$ [W m^{-2} sr^{-1}]. Projected area means that the area has to be projected along the direction of the light flow. Please note that the same projected areas might correspond to differently sized actual surface areas depending on the local surface orientation. This needs to be accounted for and, hence, the radiance can be expressed by

$$L(\omega) = \frac{d^2\Phi}{dA_\perp\, d\omega} = \frac{d^2\Phi}{dA\cos\theta\, d\omega}[\mathrm{W}\,m^{-2}\,\mathrm{sr}^{-1}]. \quad (2)$$

The inclination angle θ is defined between the local surface normal $\mathbf{n}$ and the direction of the light flow. Integrating the radiance arriving from the upper hemisphere Ω with respect to the local surface normal gives the irradiance

$$E = \frac{d\Phi}{dA} = \int_\Omega L\cos(\theta)\, d\omega. \quad (3)$$

Under the assumption of vacuum, the radiance remains constant along a ray [Gla95]. In the simplified model of ray optics, this assumption is often used. When a medium such as air or water is considered, the light propagation can be modeled by using respective interaction events that change the power of the light ray.

Furthermore, light might be polarized. *Polarization* describes the orientation of the electromagnetic wave perpendicular to the propagation direction in space. In some acquisition techniques, the properties of polarization are used by e.g. employing polarization filters to separate the direct and the global components.

3.1.2. Light Interaction at Surfaces

As already discussed before, material appearance is determined by the complex interplay of light, surface geometry and material properties of the surface. Considering the general case where the incoming radiant flux arrives at the surface at position $\mathbf{x}_i$ at the time t_i with the wavelength λ_i, the flux might enter the material, travel through the material and exit the material at position $\mathbf{x}_r$ and the time t_r with the possibly changed wavelength λ_r. The direction of the incoming flux (θ_i, φ_i) and the direction of the exiting flux (θ_r, φ_r) are usually represented based on using local coordinate frames that depend on the individual surface points $\mathbf{x}_i$ and $\mathbf{x}_r$ that are parameterized over the surface and, hence, parameterized in 2D. Figure 8 illustrates this process. Depending on the material type, this rather general, 12-dimensional model might be significantly simplified. Typical assumptions made in the great majority of publications are that the light transport at the surface happens in an infinitesimal period (i.e. $t_i = t_r$), that there is no time dependency of the reflectance behavior (i.e. $t_0 = t_i = t_r$), that the wavelength remains unchanged (i.e. $\lambda_i = \lambda_r$) and that the incoming flux is completely reflected at the surface (i.e. $\mathbf{x}_i = \mathbf{x}_r$).

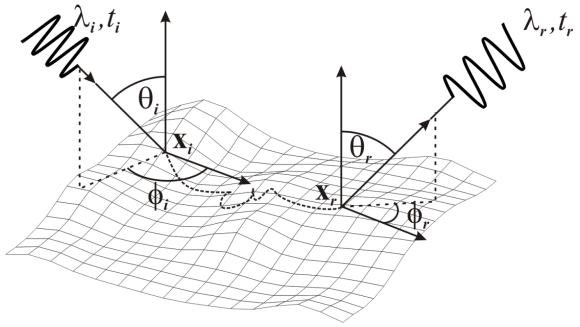

Figure 8: Light exchange at the material surface: The incoming radiant flux hits the material surface at position $\mathbf{x}_i$ and the time t_i with the wavelength λ_i, travels through the material and exits the material at position $\mathbf{x}_r$, the time t_r with the wavelength λ_r. The incoming direction (θ_i, φ_i) and the outgoing direction (θ_r, φ_r) of the flux can be formulated using local coordinate frames, where θ_i and θ_r denote the inclination angles of the incoming and outgoing light and φ_i and φ_r denote the azimuthal angles of the incoming and outgoing light respectively. Image taken from [MMS*04].

In this context, it is furthermore inevitable to take into account that material appearance is a scale-dependent phenomenon. On a *microscopic scale*, i.e. the scale of atoms and molecules, the interactions of photons with the atoms or molecules of a particular material have been analyzed in the domain of quantum optics. Clearly,

these structures cannot directly be observed by the human visual system. Nevertheless, they significantly contribute to material appearance. In particular, they determine the appearance of all materials, e.g. metals, paper, plastics, etc.. On a slightly coarser scale, studies in the field of wave optics have considered the interactions of light with small structures with a size of approximately the wavelength of light to describe effects such as polarization or diffraction. Furthermore, material appearance is also characterized by effects of light exchange happening on a *mesoscopic scale* at fine details in surface geometry such as scratches, engravings, weave-patterns of textiles or embossing of leathers. Such surface structures cause effects like interreflections, self-occlusions or self-shadowing (see Figure 9). While the effects on these aforementioned scales obviously represent the material characteristics and determine the material appearance, the 3D geometry of the object with the respective, considered material also influences the material appearance significantly. Considering this *macroscopic scale*, regular structures such as present in e.g. woven cloth, brushed metal or surface textures of certain objects might appear distorted in the image because of the dependency on the object geometry.

Unfortunately, the consideration of these scales suffices only for a close distance between the surface material and the human observer. For an increasing distance, the effects of light exchange at fine surface details such as scratches, engravings, weave-patterns or embossing will become less visible and finally not be perceivable as mesostructures anymore. Hence, they might be treated as irregularities in a different kind of microscopic scale. To give a further example, shininess of specular objects or translucency might also depend on the distance between object and observer.

When considering a highly specular surface with a rough surface profile from a close range, the resolution of the human visual system is sufficient to perceive the many surface patches with different surface normals, and the material will appear specular. With an increasing distance to the surface, the resolution of the visual system will become insufficient to perceive the appearance of all the individual surface patches with different orientations separately and, instead, perceive a superposition of the appearances of several of these patches. This will lead to a transition from specular to diffuse appearance perception. In contrast, for flat, highly specular surfaces, the surface will also appear highly specular with an increasing distance. In a similar way, the appearance of translucent objects with a rough surface profile is characterized by subsurface scattering effects when viewed from a close range. For an increasing distance, such objects might be perceived as opaque, if only the superposition of the appearances of the individual patches with the subsurface scattering effects is perceived by the visual system. This clearly indicates that the definition of scale is of dynamic nature. Therefore, material appearance involves a multitude of scales $\ldots \subset D_{i-1} \subset D_i \subset D_{i+1} \subset \ldots$ ranging from an atomic scale to the intergalactic scale [Kaj85, Sch97, MMS*04].

3.2. A Taxonomy of Surface Classes

As, to the best of our knowledge, none of the available acquisition techniques is capable of handling arbitrary materials, the idea of grouping the individual materials according to the acquisition principles applicable to the individual material groups becomes immediately evident. It is easy to imagine that these groups strongly rely on a clustering of the materials with respect to the complexity of their visual appearance.

In this spirit, the authors of [IKL*10] consider a taxonomy of object classes based on increasing complexity in light transport due to different types of interaction of the light with matter. In particular, nine classes have been identified for categorizing the properties of surface reflectance behavior:

- rough surfaces with diffuse or near diffuse reflectance
- glossy surfaces with mixed diffuse and specular reflectance
- smooth surfaces with ideal or near ideal specular reflectance
- surfaces where light is scattered multiple times underneath the surface
- smooth surfaces with ideal or near ideal specular refraction
- volumes with light emission or absorption
- volumes where a light ray is scattered only a single time
- volumes where a light ray is scattered multiple times
- mixed scenes containing several of the other types

In the scope of this course, we focus on analyzing the acquisition of solid objects and, hence, do not further discuss acquisition techniques designed to handle volumetric phenomena such as fog or fire. For this reason, we rather give a brief survey on the main trends in shape acquisition for the remaining classes according to the survey in [IKL*10] with extensions to more recently published techniques that improved the state-of-the-art.

3.2.1. Rough Surfaces with Diffuse or Near Diffuse Reflectance

The direction-independent fraction of the light reflected at a surface is denoted as diffuse reflection (see Figure 10a). For diffuse surfaces that are also referred to as materials with Lambertian reflectance behavior, the incoming light is uniformly reflected into the full hemisphere with respect to the local surface normal of the object geometry. This means that the surface can be perceived in a view-independent way, i.e. the surface appears identically from any different direction. In this case, the amount of the reflected light only depends on the direction of the incoming light L_i. This can be seen when looking at the respective reflectance model for diffuse surfaces, which describes the amount of reflected light L_{diffuse} according to

$$L_{\text{diffuse}} = L_i \, k_{\text{diffuse}} \cos\theta. \tag{4}$$

The inclination angle θ is measured between the local surface normal vector $\mathbf{n}$ and the direction of the incoming light. Furthermore, the parameter k_{diffuse} denotes a constant term that describes the degree of diffusivity of the surface. Consequently, the largest fraction of the incoming light is reflected if the light direction of the incoming light is perpendicular to the surface.

As both the direction $\mathbf{d}_i$ of the incoming light and the local surface normal $\mathbf{n}$ are typically normalized to unit length, the relation

$$\cos\theta = \mathbf{n}\,\mathbf{d}_i \tag{5}$$

holds, and the diffuse reflection can be written as

$$L_{\text{diffuse}} = L_i \, k_{\text{diffuse}} \, (\mathbf{n}\,\mathbf{d}_i). \tag{6}$$

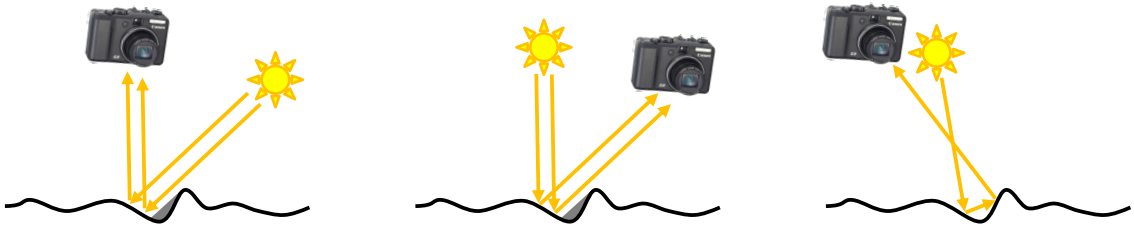

Figure 9: Mesoscopic effects of surface reflectance: Self-shadowing as illustrated with the gray region occurs at surface parts that are not directly illuminated (left), self-occlusions as illustrated with the gray region occurs if surface parts cannot directly be seen in the image taken by a camera (middle) and interreflections occur due to indirect observations where the light is reflected at the surface several times (right).

Due to the view-independent observation of the surface appearance, the geometry of such diffuse objects can be typically acquired in a rather easy way and a multitude of respective acquisition techniques for such surfaces have been developed.

3.2.2. Glossy Surfaces with Mixed Diffuse and Specular Reflectance

For glossy surfaces, there is a (non-ideal) specular reflectance component in addition to a diffuse reflectance component

$$L_{\text{glossy}} = L_{\text{diffuse}} + L_{\text{specular}}. \tag{7}$$

This specular reflectance component depends directionally on the reflected light as illustrated in Figure 10b, i.e. more light is reflected into preferred directions. As a consequence, the amount of reflected light depends on the direction of the incoming light $\mathbf{d}_i$ and the direction of the outgoing light $\mathbf{d}_o$. Hence, the observed material appearance is view-dependent as seen e.g. in the perception of specular highlights on glossy surfaces.

3.2.3. Smooth Surfaces with Ideal or Near Ideal Specular Reflectance

For (near) ideal specular surfaces, the incoming light L_i is (almost) completely reflected into the direction

$$\mathbf{d}_{o,\text{ideal reflection}} = 2\,\mathbf{n}(\mathbf{n}\,\mathbf{d}_i) - \mathbf{d}_i \tag{8}$$

as illustrated in Figure 10c. Here, the angles between the direction of the incoming light $\mathbf{d}_i$ and the local surface normal $\mathbf{n}$ and between the direction of the reflected light $\mathbf{d}_o$ and the local surface normal $\mathbf{n}$ have the same value θ.

Materials with (near) ideal specular reflectance characteristics typically do not have an own characteristic appearance themselves but rather reflect the surrounding environment in a view-dependent manner. Therefore, the geometry reconstruction for objects with such a reflectance behavior is rather challenging.

3.2.4. Surfaces Where Light is Scattered Multiple Times Underneath the Surface

The challenge in accurately acquiring shape and reflectance of translucent objects such as skin, plastics, wax or several types of minerals arises from the light transport within the object. In particular, the incoming light enters the material and travels through the material where it is scattered as illustrated in Figure 10d. This means that a certain fraction of the incoming light leaves the material at surface points that do not coincide with the point where the incoming light has hit the surface. This makes the acquisition of translucent materials rather challenging. When e.g. actively illuminating such translucent objects with a coded structured light pattern, these non-local subsurface scattering effects induce a blurring of the observed pattern and, hence, make a triangulation-based reconstruction from the decoded correspondences rather unreliable.

3.2.5. Smooth Surfaces with Ideal or Near Ideal Specular Refraction

For surfaces of this class, the incoming light L_i which has traveled through a first material completely enters a second material instead of being reflected at the surface (see Figure 10e). At the transition, the refractive indices of the neighboring materials have to be considered, as these material-specific characteristics influence the direction into which the light is refracted. The refractive index of a particular material describes the speed of light within this material. While the direction of the refracted light is determined by the refractive indices η_1 and η_2 of the involved materials following *Snell's law* that is given by

$$\eta_1 \sin\theta_1 = \eta_2 \sin\theta_2, \tag{9}$$

the amount of light that is refracted inside the second material only depends on the properties of the second material. This phenomenon can be observed when looking at the stem of water reed. Where the stem exits the water, there seems to be a kink in it. Acquiring the geometry and reflectance of refractive objects is challenging as such objects might also exhibit inhomogeneous reflectance characteristics induced e.g. by a spatially varying refractive index or by inclusions of Lambertian or opaque material components as given in many minerals.

3.3. Commonly Used Reflectance Models

For an adequate acquisition of surface reflectance, the complexity of visual surface reflectance has to be considered in a similar way as in the context of geometry acquisition. The categorization of materials as discussed in Section 3.2 indicates that visual material appearance is characterized by different phenomena of light exchange with a particular object surface of interest, which might

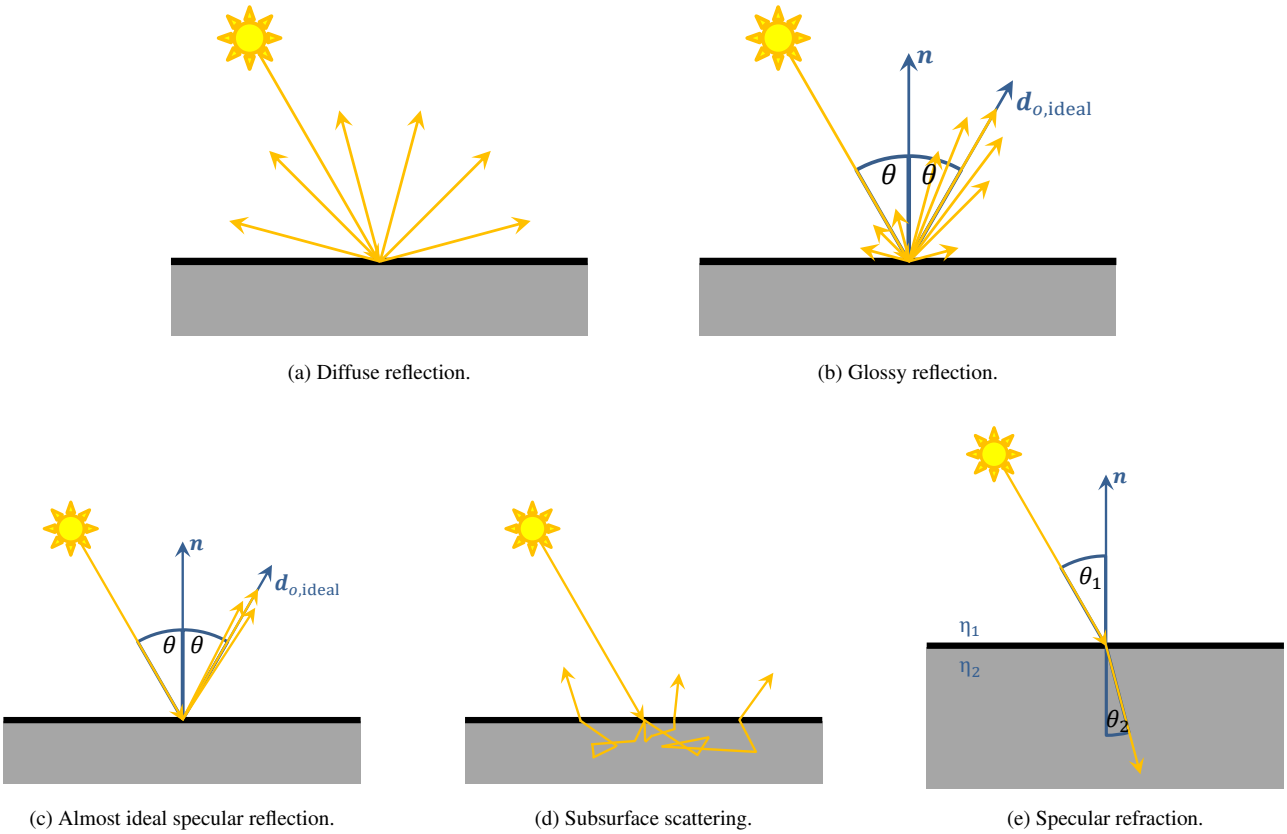

(a) Diffuse reflection.

(b) Glossy reflection.

(c) Almost ideal specular reflection.

(d) Subsurface scattering.

(e) Specular refraction.

Figure 10: Illustration of different types of surface reflectance behavior for incoming light.

also be explored when focusing on reflectance acquisition. In particular, diffuse and specular components as well as potentially occurring subsurface scattering or refraction characteristics have to be considered in the reflectance models, and the respective reflectance acquisition is typically designed according to the assumed underlying model. So far, many different models have been proposed in the literature to model surface reflectance behavior, each focusing on accurately representing a certain subset of the possible materials. However, efficiently modeling surface reflectance behavior is also coupled with the use of an adequate model, which should have as few parameters as possible to still enable a faithful depiction of the material in a synthetic image within an acceptable acquisition time. Therefore, modeling e.g. the surface reflectance behavior of a diffuse object, where the incoming light is reflected uniformly into a hemisphere on the local surface patch, requires considering different material characteristics than modeling surface reflectance of mirrors, which is determined by an almost ideal direct reflection of the incoming light. Similarly, modeling reflectance behavior for materials with both diffuse and specular components or translucent and transparent materials requires considering the respectively relevant characteristics of the individual materials. In this regard, reflectance acquisition strongly depends on the representation used to model the reflectance of a particular material, as some parameters might not have to be measured. For e.g. a diffuse material, there is no need to capture the specular characteristics, which allows to reduce the hardware usage and, hence, to speed up the measurement process.

A rather general way to model surface reflectance can be seen in using a function

$$\rho(\mathbf{x}_i, \theta_i, \varphi_i, t_i, \lambda_i, \mathbf{x}_r, \theta_r, \varphi_r, t_r, \lambda_r), \tag{10}$$

which captures reflectance depending on the twelve parameters of incoming light direction (θ_i, φ_i), the position $\mathbf{x}_i$ where the light hits the surface at the time t_i with the wavelength λ_i, the position $\mathbf{x}_r$ where the light exits the surface with the outgoing light direction (θ_r, φ_r) at the time t_r with the wavelength λ_r (see Section 3.1). Typically, the time and the wavelength characteristics are omitted for simplicity, i.e. it is assumed that $t_i = t_r$ and $\lambda_i = \lambda_r$.

The plenoptic function $P(\mathbf{X}, \theta, \varphi)$ has been introduced in [AB91] as a function that describes the radiance sent from a certain scene point $\mathbf{X} \in \mathbb{R}^3$ into the direction (θ, φ). Assuming that an arbitrarily complex shaped surface S with the bounding volume V is embedded in the considered scene volume, the radiance values observed at points $\mathbf{x} \in \partial V \subset \mathbb{R}^2$ on the surface from a viewpoint outside V with viewing directions (θ_r, φ_r) can be used to represent the appearance of the object surface S for a given, static illumination [GGSC96, LH96]. As the surface points can be parame-

terized over the surface S, it is possible to use a four-dimensional function $\rho_{\mathrm{SLF},r}(\mathbf{x},\theta_r,\varphi_r)$, the radiant surface light field, instead of the five-dimensional plenoptic function $P(\mathbf{X},\theta_r,\varphi_r)$. If the viewpoint is inside of the volume V, the incident surface light field $\rho_{\mathrm{SLF},i}(\mathbf{x},\theta_i,\varphi_i)$ can be observed [LH96]. Consequently, a radiant light field $\rho_{\mathrm{LF},r}(\mathbf{x}_r,\theta_r,\varphi_r)$ observed at particular scene points $\mathbf{x}_r$ from the viewing angles (θ_r and φ_r) can be interpreted as a general response of the arbitrary complex scene to an incoming light field $\rho_{\mathrm{LF},i}(\mathbf{x}_i,\theta_i,\varphi_i)$ [LH96]. From the definition of surface light fields it becomes obvious that they can only be used to describe static scenes without variations in illumination, scene geometry and surface materials. The term *reflectance field* [DHT*00] describes the dependency of the radiant light field $\rho_{\mathrm{LF},r}(\mathbf{x}_r,\theta_r,\varphi_r)$ on the incident light field $\rho_{\mathrm{LF},i}(\mathbf{x}_i,\theta_i,\varphi_i)$ and can be formulated as an eight-dimensional function

$$\rho_{\mathrm{RF}}(\rho_{\mathrm{LF},i},\rho_{\mathrm{LF},r}) = \rho_{\mathrm{RF}}(\mathbf{x}_i,\theta_i,\varphi_i,\mathbf{x}_r,\theta_r,\varphi_r), \tag{11}$$

i.e. a reflectance field describes the exitant radiance depending on the possibly occurring incident illuminations. Typically, reflectance fields are defined on convex surfaces that surround the respective object and it is assumed that the viewpoint is outside this bounding volume and that the light is also coming from outside. This allows to use the reflectance field to represent material appearance under arbitrary new viewpoints and illumination conditions by sampling the outgoing light fields under a set of basis incident light fields. The linearity of light transport allows the definition of new illumination conditions in terms of a linear combination of the illumination basis. The definition (11) is closely related to the definition of the *bidirectional scattering-surface reflectance distribution function (BSSRDF)* [NRH*77] given by

$$\rho_{\mathrm{BSSRDF}}(\mathbf{x}_i,\theta_i,\varphi_i,\mathbf{x}_r,\theta_r,\varphi_r) \tag{12}$$

and even would be identical if the true surface is used.

Assuming that the surface reflectance is defined on the object surface similar as for BSSRDFs but assuming $\mathbf{x}_i = \mathbf{x}_r = \mathbf{x}$, i.e. that light is not scattered inside the material, the inhomogeneous reflectance behavior can be explained by six-dimensional *spatially-varying bidirectional reflectance distribution functions (SVBRDFs)*

$$\rho_{\mathrm{SVBRDF}}(\mathbf{x},\theta_i,\varphi_i,\theta_r,\varphi_r). \tag{13}$$

In contrast, another six-dimensional representation is given by *bidirectional texture functions (BTFs)* which can be obtained by assuming far-field illumination where the light sources are infinitely far away. This means that the incident radiance is the same for all surface points, i.e. $\rho_{LF,i}(\mathbf{x}_i,\theta_i,\varphi_i) = \rho_{LF,i}(\theta_i,\varphi_i)$. As a result, the definition of the BTF is given by

$$\rho_{\mathrm{BTF}}(\theta_i,\varphi_i,\mathbf{x}_r,\theta_r,\varphi_r). \tag{14}$$

In comparison to SVBRDFs, BTFs allow to capture local subsurface scattering characteristics as well as mesoscopic effects such as interreflections, self-masking or self-occlusions.

In contrast, when assuming homogeneous reflectance behavior, the BSSRDF can be relaxed to the *bidirectional subsurface scattering distribution function (BSSDF)*

$$\rho_{\mathrm{BSSDF}}(\theta_i,\varphi_i,\mathbf{x}_r-\mathbf{x}_i,\theta_r,\varphi_r), \tag{15}$$

which still is capable of modeling subsurface-scattering effects. Additionally assuming non-subsurface scattering reflectance, the *bidirectional reflectance distribution function (BRDF)*

$$\rho_{\mathrm{BRDF}}(\theta_i,\varphi_i,\theta_r,\varphi_r) \tag{16}$$

models the reflectance behavior under the remaining four parameters. Considering BTFs or SVBRDFs and additionally assuming homogeneous surface reflectance also leads to the BRDF model. Assuming non-anisotropic reflectance, isotropic BRDF models represent a further simplified reflectance model. In addition, the diffuse surface reflectance function represents a further four-dimensional representation depending on the parameters $\theta_i,\varphi_i,\theta_r$ and φ_r to model diffuse surfaces. Further non-material dependent simplifications of BTFs, as illustrated in Figure 11, are the restriction to a fixed lighting or a fixed view resulting in four-dimensional *surface light fields (SLFs)*

$$\rho_{\mathrm{SLF}}(\mathbf{x},\theta_r,\varphi_r) \tag{17}$$

or *surface reflectance fields (SRFs)*

$$\rho_{\mathrm{SRF}}(\mathbf{x},\theta_i,\varphi_i) \tag{18}$$

respectively. In the case of diffuse reflectance, this can be further simplified to two-dimensional *texture maps* or *bump maps*

$$\rho_{\text{Texture Map/Bump Map}}(\mathbf{x}). \tag{19}$$

This hierarchy of commonly used reflectance functions according to [MMS*04, DLG13] is shown in Figure 11. Depending on the parameters of the respective material model, the acquisition device has to be designed in a way that material appearance can be captured under the involved parameter configurations. Further details regarding the individual reflectance models and their acquisition are discussed in Section 7.

4. Calibration

While the course mainly focuses on the review of methods for geometry and reflectance acquisition, we also provide a short discussion of fundamental properties that have to be considered for the calibration of the respective setups as calibration is an important aspect of the acquisition ecosystem. If several measurements of the geometry in a scene or the corresponding reflectance behavior have been performed, these measurements need to be brought into some kind of reference system. Unfortunately, the discussion of all the details that have to be considered and the individual techniques that have been proposed in literature so far would be too comprehensive to include them into this course. For this reason, we will only provide brief discussions in this section in order to establish an awareness regarding why individual calibration techniques are needed and some rather general ideas on how a calibration of individual setup components can be performed. For more detailed surveys on calibration, we refer to the literature.

The relative positions and orientations of the individual components with respect to the object surface as well as device-specific characteristics are important for an accurate acquisition of geometry and reflectance. As a consequence, there is a need for geometric calibration (see Section 4.1) that allows to obtain a geometric relationship between the involved illuminants, imaging sensors and the

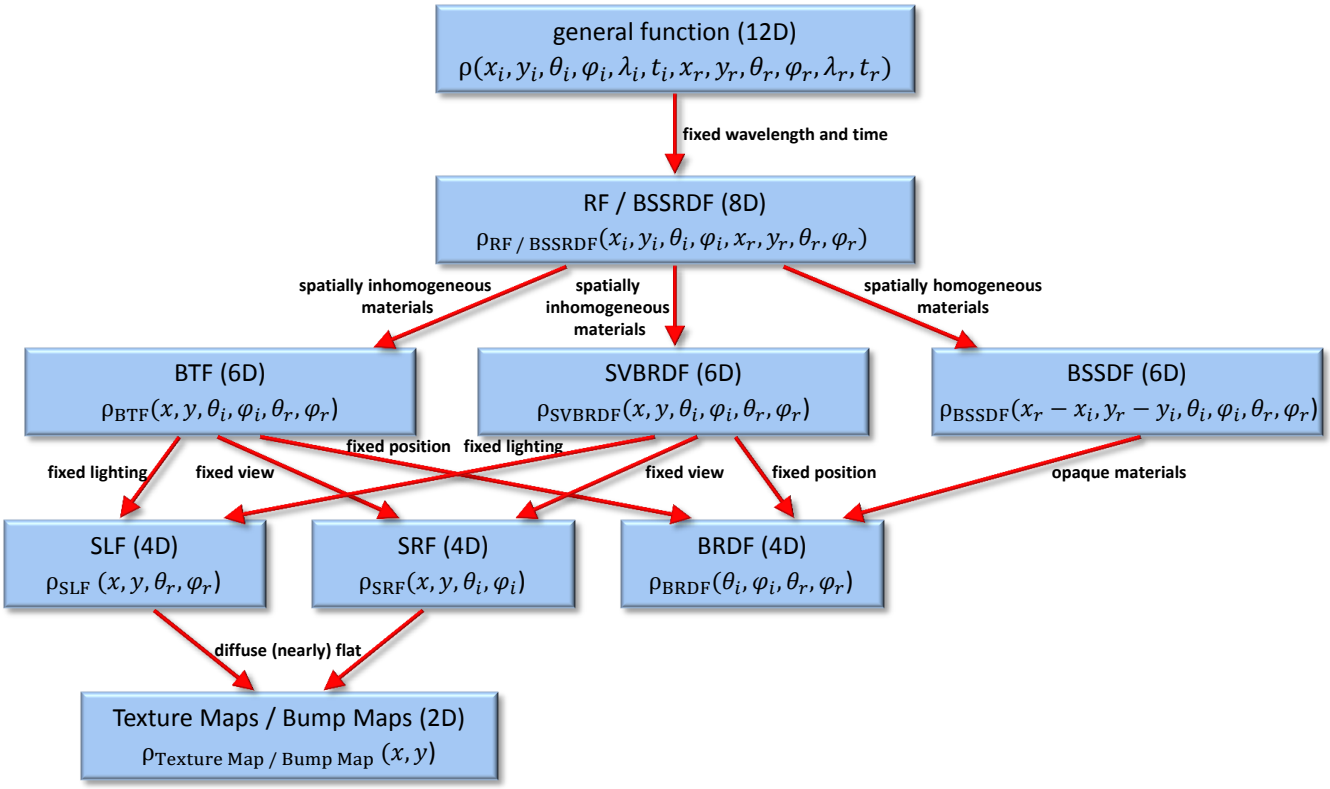

Figure 11: A hierarchy of reflectance functions according to [MMS*04] and [DLG13].

object surface during the acquisition. Furthermore, it is usually not sufficient to only consider the position and orientation of components as e.g. the mapping of a 3D scene onto a 2D camera image involves a projection that depends on further camera parameters such as the focal length or the principal point. In addition, the radiometric calibration (see Section 4.2) of the setup allows to take the radiometric properties of the imaging sensors and illuminants into account. This is particularly important for appearance acquisition where the measured reflectance values have to be brought into some kind of reference system for radiance as well.

4.1. Geometric Calibration

The geometric calibration aims at the specification of relative positions and orientations of the individual components involved in the setup by either using certain known calibration objects or possibly also objects with unknown shape. The observations of the object by the imaging sensor at possibly various positions need to be geometrically registered with respect to each other and also regarding the positions of the light source. Often it is not practical to have an adequate number of sensors to densely acquire certain view-light configurations during reflectance acquisition. For this reason, objects are often placed on turntables to simulate the effect of having many sensors. We will therefore give a short overview on camera calibration (see Section 4.1.1), light source calibration (see Section 4.1.2) and turntable calibration (see Section 4.1.3).

4.1.1. Camera Calibration

In order to infer information regarding the shape or reflectance behavior of objects, the relationship between the image domains of the cameras involved in the respective setup and the 3D scene have to be taken into account. In other words, there is a need for a mathematical model that describes how the scene content is imaged by the sensor. This model is often denoted as image formation model as it describes the 3D-to-2D mapping of the scene content into the 2D image plane. After a short discussion of camera models, we provide a short analysis of widely applied strategies for camera calibration. For a more detailed discussion of the corresponding aspects we refer to [HZ04].

Most acquisition techniques are based on the pinhole camera model to describe the projective mapping of the 3D scene content into the 2D image domain. Thus, the image formation can be modeled in terms of a simple matrix multiplication using homogeneous coordinates. Following the principle of a pinhole camera, the relationship between a 3D point $\mathbf{X} \in \mathbb{R}^3$ with homogeneous coordinates $\mathbf{X}_h$ on the surface of an object in the scene and its observed location $\mathbf{x} \in \mathbb{R}^2$ with homogeneous coordinates $\mathbf{x}_h$ in the image can be formulated according to

$$\mathbf{x}_h = \underbrace{\mathbf{K}\,[\mathbf{R}\,\mathbf{t}]}_{\mathbf{P}}\,\mathbf{X}_h. \tag{20}$$

The projection matrix $\mathbf{P} \in \mathbb{R}^{3\times 4}$ of the camera describes the intrin-

sic camera parameters that are contained in the calibration matrix $\mathbf{K} \in \mathbb{R}^{3\times3}$ as well as the extrinsic camera parameters given by rotation parameters in a rotation matrix $\mathbf{R} \in \mathbb{R}^{3\times3}$ with three degrees of freedom and three translation parameters in a vector $\mathbf{t} \in \mathbb{R}^3$. The intrinsic parameters include parameters f_x and f_y that describe the focal length in x-direction and y-direction, the parameter s that describes the skewness of the image pixels and the 2D location $(p_x, p_y)^T$ of the principal point where the optical axis intersects the image plane. These intrinsic parameters determine the calibration matrix $\mathbf{K}$ according to

$$\mathbf{K} = \begin{bmatrix} f_x & s & p_x \\ 0 & f_y & p_y \\ 0 & 0 & 1 \end{bmatrix}. \tag{21}$$

Obviously, the projection matrix $\mathbf{P}$ has a total of eleven degrees of freedom which have to be estimated from observations. This means that at least six pairs of known 3D points $\mathbf{X}_i \in \mathbb{R}^3$ (in general positions) and their corresponding 2D projections $\mathbf{x}_i \in \mathbb{R}^2$ in the camera image are required for the camera calibration for the case of a pinhole camera, i.e. to estimate the camera parameters [HZ04].

However, while such a rather simple camera model approximately already models the image formation process of real cameras rather well, there are some more effects that often need to be handled as well. For instance, the lens-based optics of many cameras might cause non-linear distortions so that a 3D point $\mathbf{X} \in \mathbb{R}^3$ is observed at the 2D image location $\tilde{\mathbf{x}} \in \mathbb{R}^2$ instead of the position $\mathbf{x} \in \mathbb{R}^2$. As a result, distortion models need to be considered as well which typically model radial and tangential distortions. If the radial distortion coefficients are given by k_1, k_2, k_3, k_4, k_5 and k_6 and the tangential distortion parameters are specified by p_1 and p_2, the basic pinhole model can be extended by simply adding some intermediate calculations. In a first step, the 3D point $\mathbf{X}$ is transformed to a point $\mathbf{X}^C$ in the camera coordinate system using the extrinsic camera parameters according to

$$\mathbf{X}_h^C = [\mathbf{R}\,\mathbf{t}]\,\mathbf{X}_h \tag{22}$$

and the components $x' = \frac{x}{z}$ and $y' = \frac{y}{z}$ are extracted from $\mathbf{X}_h^C = (x,y,z)^T$. After this transformation into the camera coordinate system, an additional intermediate step is applied that models the distortion due to the lens characteristics. Applying the distortion model yields

$$x'' = x' \frac{1+k_1r^2+k_2r^4+k_3r^6}{1+k_4r^2+k_5r^4+k_6r^6} + 2p_1x'y' + p_2(r^2+2x'^2) \tag{23}$$

$$y'' = x' \frac{1+k_1r^2+k_2r^4+k_3r^6}{1+k_4r^2+k_5r^4+k_6r^6} + p_1(r^2+2y'^2) + 2p_2x'y' \tag{24}$$

where $r^2 = x'^2 + y'^2$. After the application of the distortion model, the image locations $\mathbf{x} = (u,v)^T$ of the respective 3D point can be obtained by applying the calibration matrix $\mathbf{K}$ according to

$$\begin{bmatrix} u \\ v \\ 1 \end{bmatrix} = \mathbf{K} \begin{bmatrix} x'' \\ y'' \\ 1 \end{bmatrix} = \begin{bmatrix} f_x x'' + s y'' + p_x \\ f_y y'' + p_y \\ 1 \end{bmatrix}. \tag{25}$$

When distortion parameters are assumed to be relevant as well, they can be estimated based on a sufficiently large number of observations of known 3D positions in the camera image.

As the accuracy of the acquisition is also determined by the accuracy in estimating the camera parameters, the latter ones should typically be reasonably accurate. For this reason, calibration patterns with specified 3D locations of the specific calibration features that might be given by e.g. corners of a checkerboard pattern, centers of circles arranged in a certain structure or corners of QR markers are typically moved in front of the individual cameras and the calibration can be obtained by acquiring a sufficient number of images. If no such calibration target with a known arrangement of reliably specified 3D positions is available, observations of either passively generated point correspondences obtained e.g. based on SIFT features [Low04] or SURF features [BETvG08] or actively generated point correspondences generated by e.g. using structured light patterns in the images from multiple cameras can be used for camera calibration via techniques such as bundle adjustment (typically used in photogrammetry), structure-from-motion (typically used in computer vision) and simultaneous-localization-and-mapping (typically used in robotics). For this purpose, image correspondences can be established by extracting and matching characteristic feature points such as SIFT features [Low04] or SURF features [BETvG08] in different images or by actively projecting patterns into the scene and matching the corresponding encodings as e.g. applied in [WSRK11]. These related techniques are similar in the aspect that they jointly optimize the camera parameters and the 3D position of points based on corresponding observations of the same points in the images acquired by cameras at several positions.

Sparse Bundle Adjustment (SBA) [LA09] is based on considering a set of observations $\mathbf{x}_{i,j}$ of i unknown 3D points $\mathbf{X}_i$ in the images acquired by j cameras with unknown projection matrices $\mathbf{P}_j$ and minimizing the sum of reprojection errors

$$\sum_i \sum_j \left\| \mathbf{x}_{i,j} - \mathcal{P}_j(\mathbf{X}_i) \right\|. \tag{26}$$

This formulation is typically solved by using a non-linear Levenberg-Marquardt optimization [Lev44].

In some setups, cameras are remotely positioned with robotized arms. Therefore, the position might theoretically be specified directly. However, the positioning systems are typically not accurate enough to allow for a subpixel-accurate registration of images taken from different views.

4.1.2. Light Source Calibration

While light sources are particularly used in appearance acquisition, they also play an important role in several geometry acquisition techniques such as photometric stereo techniques, shape-from-specularity techniques or techniques based on Helmholtz reciprocity that will be described later. If light sources are mounted on robotized positioning arms, their position might be specified directly. However, the positioning systems often do not offer the possibility for a precise specification of a light source position and orientation, so that typically a separate estimation of these parameters has to be performed. Several setups such as the ones in [CGS06, WRO*12, SSWK13] rely on an estimation of light source positions based on highlight observations on mirroring calibration spheres with typically known position and radii. The key

idea is given by tracing rays originating from the perspective center of a calibrated camera through the image positions of the observed highlights into the scene where they are reflected on the known sphere geometries. Reflected rays obtained from highlights observed from multiple spheres allow the reconstruction of the light source position which is given by their intersection (see Figure 12). Optionally, a further non-linear optimization that simultaneously refines both the light source positions and the positions of the calibration spheres can be performed to reduce the re-projection error of the observed reflections [SSW*14]. A similar technique has been presented in [AFG13], where the position of the calibration sphere is not refined during the optimization. If flashes of cameras serve as light sources, the light source positions can be estimated based on offsets of the flashes with respect to the perspective centers of calibrated cameras as applied in [MMS*04, SWRK11].

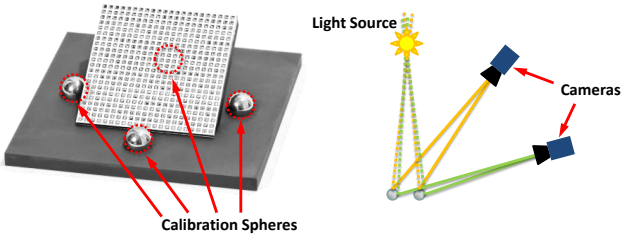

Figure 12: (left) Depiction of a calibration target used for camera calibration and light source calibration [SSW*14]; (right) Estimation of light source position based on calibration spheres with known position and known radius: The light source is observed as highlights on the mirroring calibration spheres. Rays are back-projected through the corresponding image positions of the individual highlights and reflected at the surface of the spheres. The intersection of these reflected rays defines the light source position.

4.1.3. Turntable Calibration

Turntables have been used in a multitude of acquisition setups to simulate the presence of further light sources or imaging sensors. Rather accurate estimates for the turntable center and the turntable axis have to be determined which can be achieved by rotating calibration targets. The targets can be observed by cameras, that can be calibrated based on such targets as well, and the corresponding observations allow to get the parameters of the turntable. Some high-quality turntables also directly offer the possibility to accurately specify the rotation angle.

4.2. Radiometric Calibration

In addition to the geometric calibration of the acquisition setup, characteristics of the light emitted by the light sources as well as the sensor characteristics that transform the received light into certain pixel values on the sensor have to be taken into account when the observed pixel values are directly involved in either geometry or reflectance reconstruction.

When using continuous light sources such as halogen lamps or LEDs, this correction can be performed by first subtracting dark frames, i.e. images acquired under completely darkened lab conditions, from all images to remove the effect of hot pixels or sensor bias and subsequently inverting the response curve of the respective cameras to obtain energy values from the observed pixel values. The inverse camera response curve can be determined based on observations of a white-standard such as SphereOptics Zenith UltraWhite [Sph] or Labsphere Spectraleon [lab] or a color-target such as the X-Rite ColorChecker Passport under different exposure times or different wavelengths. The influence of spatially varying illumination effects such as vignetting, chromatic aberrations or distance fall-off is typically canceled out when computing the reflectance values. The use of flash light sources as in [MMS*04, SWRK11] further complicates the radiometric correction as light color and light intensity typically vary per discharge. For this reason, it is not possible to perform an initial calibration step but, instead, a separate calibration for each flash discharge is required.

5. Advances in Geometry Acquisition

There has been a significant technological progress regarding the development of geometry acquisition methods for decades and many acquisition techniques have been presented. However, it turned out that many of the conventional techniques do not allow a reconstruction of the surface geometry of objects with arbitrary surface reflectance but only a geometry reconstruction for objects with a rather simple diffuse surface reflectance behavior. The complexity of surface reflectance behavior has led to the development of a huge range of diverse material-specific acquisition techniques, each tailored to only a very limited range of materials. Most of the developed methods follow the categorization according to the classes considered in [IKL*10] as already discussed in Section 3. As we will focus on analyzing solid objects in the scope of this course, we do not further discuss acquisition techniques designed to handle volumetric phenomena such as fog or fire and focus on discussing approaches for the remaining surface types.

Before we review the main trends that have been introduced in order to approach the 3D geometry acquisition objects with different surface materials (Sections 5.2- 5.6), we will provide an overview of standard principles (Section 5.1) that are used in many of the geometry acquisition techniques.

5.1. Basic Geometry Acquisition Principles

The acquisition of the 3D surface geometry of objects can be approached based on different principles. The key differences of the individual methods include whether methods are active or passive, contact-based or non-contact-based and optical or non-optical. In Figure 13 and Figure 14, we provide a categorization of the individual techniques according to these attributes.

Active methods rely on the direct interaction with the object or scene of interest which might be achieved e.g. in terms of a projection of a certain type of energy on the object surface that can be observed by the components of the respective acquisition setup or in terms of contact-based feelers that produce a characteristic signal when touching the surface [Cur97]. In contrast, no such active

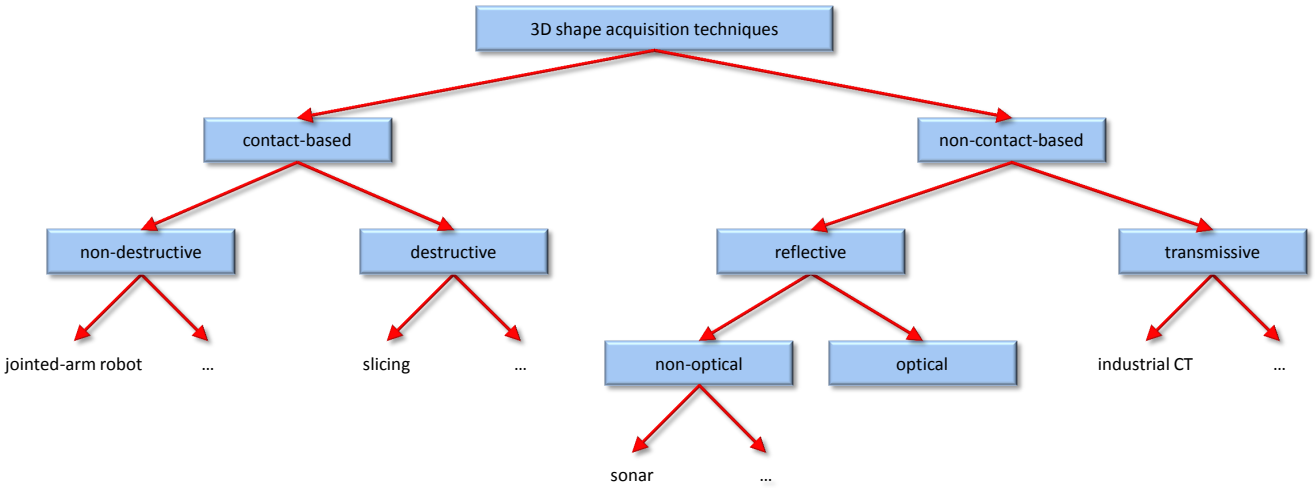

Figure 13: A general taxonomy of principles for geometry acquisition.

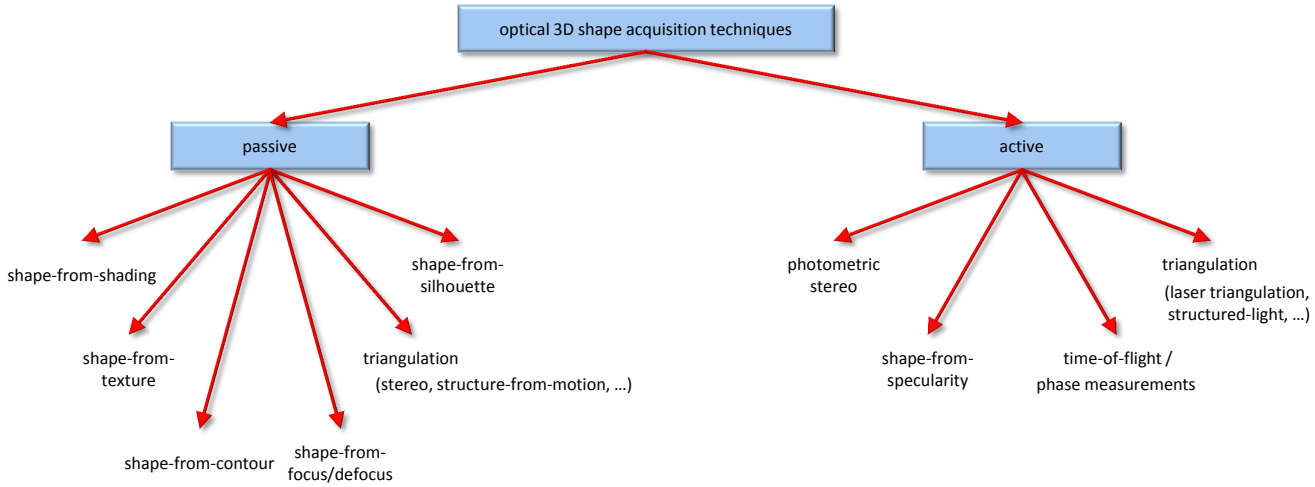

Figure 14: A taxonomy of principles for optical geometry acquisition.

manipulation is involved in *passive methods* which typically focus on the analysis of characteristics of optical object appearance as observed under rather fixed and distant illumination.

In the scope of our survey, we start with a more thorough discussion of principles for passive geometry acquisition (see Section 5.1.1) which is followed by a discussion of principles for active geometry acquisition (see Section 5.1.2).

5.1.1. Principles for Passive Geometry Acquisition

As already discussed above, the principles for passive shape acquisition typically rely on the observation and analysis of the optical appearance of objects under fixed illumination without systematic manipulations.

5.1.1.1. Shape-from-Shading Techniques

As discussed in Section 3, the observed intensity information can be seen as a result of the complex interplay of material properties, geometric surface properties and illumination conditions. However, the 2D information given in an image acquired from a single viewpoint is not sufficient to reconstruct the 3D shape of the object surface which leads to an ill-posed nature of the considered problem. In order to approach the reconstruction of the surface geometry, this means that some kind of underlying prior information or assumptions have to be included to obtain additional constraints when solving for the surface.

Conventional *shape-from-shading techniques* [Hor70] rely on analyzing the observed intensity information obtained under a single viewpoint with respect to the object in a controlled environment

with controlled illumination by a point light source, where the illumination characteristics such as strength of the illumination as well as its origin (typically distant illumination, i.e. the light source position is typically significantly farer away than the extents of the considered object surfaces) and direction are assumed to be known as a result from a prior calibration procedure. However, despite the rather controlled scenario, the problem of surface reconstruction is still ill-posed which imposes the need for even more constraints. In particular, the assumption of diffuse surface reflectance following Lambert's law allows to derive additional constraints that encapsulate the relation between the intensity $I(x,y)$ observed in the 2D image domain of the involved camera, the known direction $\mathbf{l}$ of the incoming illumination, the local surface normal $\mathbf{n}$ that has to be determined, the known intensity r of the illumination and the known surface albedo $\alpha(x,y)$ according to

$$I(x,y) = \alpha(x,y)\, r\, \mathbf{l}\, \mathbf{n}. \tag{27}$$

Additional widely used assumptions include an orthographic image projection and the absence of shadows and interreflections. The resulting equation system includes a single constraint per pixel and, hence, is still under-constrained as there are two unknown parameters given by the components $p = \frac{dz}{dx}$ and $q = \frac{dz}{dy}$ of the normal directions per local surface normal. For this reason, further typical assumptions focus on a constant surface albedo $\alpha(x,y) = \alpha$ and the normalization of the z-component that results in representations of the form $\mathbf{n} = [-p,-q,1]^T$ and constraints on the respective gradients for e.g. enforcing smoothness or using the self-shadow boundary [IH81]. A more detailed discussion regarding different constraints that might be used for shape-from-shading as well as a performance evaluation of different respective techniques is given in the surveys in [ZTCS99, DFS08]. Based on the surface normals, the surface geometry can be derived via normal field integration techniques. Most of the shape-from-shading techniques focus on geometry reconstruction from a single image, and, hence, only a 2.5D height map can be derived.

5.1.1.2. Shape-from-Texture Techniques

Several materials exhibit a rather regular surface pattern that might be used for surface reconstruction. Examples might be given by fabrics where the regular structure is typically the result of the weaving pattern. In particular, *shape-from-texture techniques* focus on the reconstruction of surface geometry by assuming that objects have such a regular statistical or geometric surface pattern. For curved surface parts, this regular surface pattern appears distorted within the image. The reason for this is that the non-regularity of the surface pattern as observed e.g. in anisotropic statistics of edge orientations for isotropic textured surface patterns results from the projection of curved surface patches into the 2D image domain. In turn, this means that information regarding the surface geometry might be derived based on the distortion of the respective surface pattern. A statistical approach for recovering surface shape and orientation from texture has been proposed in [Wit81]. This approach is based on the assumption that texture does not mimic projective effects, and it focuses on attributing as much of the variation in the image as possible to the projection, while the texture itself is made as uniform as possible. Thereby, the surface orientation which best explains the data in this sense represents the best estimate for the surface orientation. Considering that texture consists of small elements which are referred to as texels [Gib50], an approach based on the assumption that such texels are uniformly distributed has been presented in [Alo88] as well as an extension only relying on the partial boundaries of the texels instead of a detection and identification of the texels themselves.

5.1.1.3. Shape-from-Contour Techniques

For several objects, the contours of structures given in the surface texture might also be used to derive the surface geometry as done by *shape-from-contour techniques*. The distortion of these contours depends on the surface orientation. If an object is e.g. covered by a known pattern with circular structures that might appear as ellipses in the acquired image depending on the local surface curvature, the respective rotations of the individual circles can be derived. A more general approach relies on using the tangent of the contour to estimate surface orientation based on a statistical model of the scene parameters [Wit80]. In addition, several more sophisticated approaches have been proposed such as the maximization of the ratio of the area and the square of the perimeter [BY84]. Similarly, symmetry characteristics that are not preserved after projecting the object into the image domain might also serve to infer the surface geometry by searching for the tilt and slant angles that maximize the restored symmetry.

5.1.1.4. Shape-from-Focus / Shape-from-Defocus Techniques

A different strategy for geometry acquisition relies on the fact that cameras have a limited depth-of-field as only points at a particular distance will be imaged with a sharp projection in the image plane [MVGV10]. As illustrated in Figure 15, for a thin lens, the distance d_p between a scene point and the lens, the distance d_i between the lens and the image plane, and the focal length f of the lens are related via the Gaussian lens formula according to

$$\frac{1}{f} = \frac{1}{d_p} + \frac{1}{d_i}. \tag{28}$$

Thus, a single scene point is projected onto a single point on the image plane, and a clear or *focused* image is formed on the image plane if the sensor plane coincides with the image plane [NN94]. However, if the sensor plane does not coincide with the image plane, but is displaced, a blurred or *defocused* image is formed on the image plane. Since objects with different distances d_p to the lens are focused at different distances d_i, only a specific part of a respective scene appears sharp in the camera image.

Based on these considerations, *shape-from-focus techniques* [NN94] rely on estimating the scene geometry by actively changing the optics of the involved camera (e.g. by displacing the sensor with respect to the image plane, by moving the lens or by changing the distance of the object with respect to the camera) until the respective scene point of interest is focused. Typically, numerous images are taken from the same viewpoint by varying the distance between image plane and sensor plane in small increments and evaluating a focus measure (e.g. as presented in [Kro87, DW88, NS92, NN94, XS97]) for each configuration. For

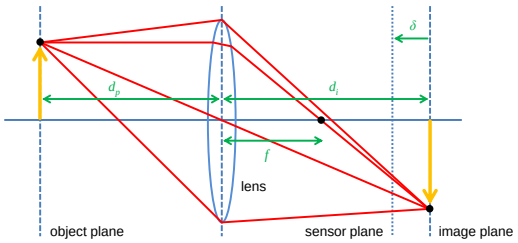

Figure 15: Formation of focused and defocused images.

each scene point, the respective maximum of the focus measure is used to estimate its distance from the camera [NWN96].

In contrast to shape-from-focus techniques, *shape-from-defocus techniques* [NWN96] typically rely on considering two images taken from the same viewpoint and estimating the relative blurring due to a change in the optics of the involved camera in order to acquire the geometry of an object. In this context, a relative blurring can principally be estimated by frequency analysis since, in the frequency domain, blurring may be considered as low-pass filtering of scene texture [NWN96]. Thus, the relative blurring in two images may be derived by applying narrow-band filters (e.g. Gabor filters) to isolate more or less single frequencies and subsequently estimating the relative attenuation of the isolated frequencies. Due to the consideration of only two images, geometry acquisition via shape-from-defocus techniques generally tends to be faster than geometry acquisition via shape-from-focus techniques, but it also tends to be less accurate since the depth resolution is typically significantly lower than the resolution perpendicular to the view direction [STD09]. Furthermore, it becomes obvious that both shape-from-focus techniques and shape-from-defocus techniques share a major weakness with stereo techniques and structure-from-motion techniques which consists in the assumption of well-textured object surfaces. In order to acquire the geometry of textureless object surfaces with respective techniques, an active geometry acquisition via the projection of a light/illumination pattern on the scene may be applied to obtain a synthetic texture.

5.1.1.5. Shape from Passive Triangulation

The underlying principle of *passive triangulation* is based on the assumption that certain characteristic features of the scene geometry are observed simultaneously in several images taken from different viewpoints. In standard *stereo techniques*, the object or scene of interest is observed simultaneously from two viewpoints as illustrated in Figure 16. In more detail, a particular point $\mathbf{X}$ on the scene geometry is projected onto individual image locations $\mathbf{x}_1$ and $\mathbf{x}_2$ in the image planes I_1 and I_2 of two typically synchronized cameras with projective centers $\mathbf{C}_1$ and $\mathbf{C}_2$. In turn, if the cameras are calibrated, i.e. if their positions and orientations as well as the internal parameters are known, the point $\mathbf{X}$ on the object surface can be reconstructed as the intersection of the backprojected rays $\mathbf{r}_1$ and $\mathbf{r}_2$ that pass through the projective center of the respective camera and the image location in the corresponding image plane. As the points $\mathbf{C}_1$, $\mathbf{C}_2$ and $\mathbf{X}$ form a triangle, this process is named (passive) triangulation. Obviously, many points $\mathbf{X}_i$ on the object surface have to be reconstructed to obtain a rather dense and accurate reconstruction of the acquired object surface. A comparison and a performance evaluation of stereo techniques is given in [SS02].

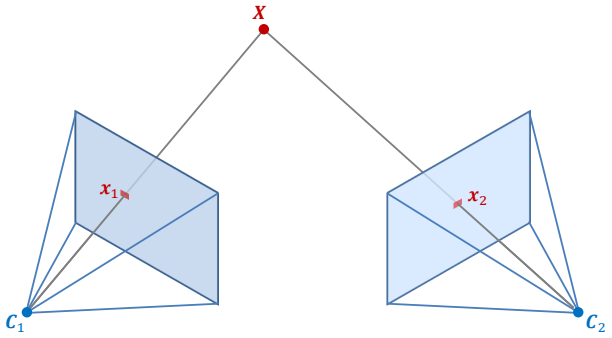

Figure 16: Principle of triangulation: A point $\mathbf{X}$ on the object surface can be reconstructed based on the intersection of the backprojected rays of two corresponding image points $\mathbf{x}_1$ and $\mathbf{x}_2$.

For static scenes, a rather similar approach termed as *structure-from-motion* can be followed to reconstruct the scene by replacing the two synchronized cameras used in stereo techniques with a single moving camera. Generally, structure-from-motion techniques rely on the use of point tracks, i.e. corresponding image locations between consecutive images, in order to simultaneously recover both the 3D structure of a scene and the camera pose [Sze11].

Both of the aforementioned approaches can be extended by using more images for the reconstruction of 3D points on the object surface. The extension of classical stereo techniques towards geometry reconstruction from a collection of images taken from known viewpoints results in *multi-view stereo techniques* for which a survey is given in e.g. [SCD*06]. One of the most general structure-from-motion techniques is represented by *bundle adjustment* [TMHF00] which relies on the consideration of bundles of rays connecting the projective centers to 3D scene points and an adjustment via an iterative minimization of the reprojection error and thus allows a simultaneous recovery of the 3D structure of a scene, the camera pose for each image and even the intrinsic camera parameters. Furthermore, the use of more images typically not only allows for a more complete reconstruction of the object surface but also for including self-calibration techniques so that the requirement of having pre-calibrated cameras can be omitted.

The main challenge of all these passive triangulation-based methods is given by the need for an establishment of reliable point correspondences between the available information given in the different images acquired under different viewing conditions. Typically, characteristic feature points are detected based on finding distinctive locations within the image domains via an analysis of the local texture, and they are described with vector-based representations that encode the characteristics of the respective local texture gradient distributions and allow an efficient comparison. Due to the local nature of such features, they are commonly referred to as *local*

features. As such features facilitate a variety of applications, feature extraction and matching in this regard has been intensively studied in literature for decades. Particularly SIFT features [Low04] and SURF features [BETvG08] belong to the standard methods for the detection and matching of common image contents. Respective surveys on local features can be found in [TM08, Wei13]. However, since such local features analyze the local statistics of intensity gradients for the extraction of distinctive feature points and the computation of local descriptors, they can only be applied on textured surfaces. In contrast, there is no distinctive local texture gradient information available for textureless surfaces that would support the extraction and matching of reliable feature correspondences. As a consequence, a robust matching of features is often difficult.

While triangulation-based methods are capable of accurately reconstructing the low-frequency shape of the absolute 3D surface geometry, their limitations can be identified in the high-frequency noise in the reconstructed depth values as a result of inaccuracies in the 2D positions of the extracted point correspondences and inaccuracies in camera calibration. Additionally, these methods are not suitable for capturing objects with highly complex reflectance behavior including effects such as specular highlights, transparency, translucency and interreflections as they are – by design – tailored to surfaces exhibiting Lambertian reflectance behavior. To be more specific, multi-view stereo methods are typically based on the assumption that the appearance of a certain point on the object surface does not change under variations of viewing conditions, i.e. the emitted radiance is independent from the view direction. However, this assumption is only valid for diffuse surfaces.

5.1.1.6. Shape-from-Silhouette Techniques

Using silhouette information in order to reconstruct the surface geometry has been approached with shape-from-silhouette techniques as used in e.g. [Lau94, GGSC96, FKIS02, MBK05]. Such approaches rely on obtaining accurate silhouette information by a prior segmentation of the respective images into foreground regions that contain the object and background regions that include the remaining scene content. Subsequently, the observed silhouette information is projected into the volume of interest as illustrated in Figure 17. This can be achieved by using a volumetric, voxel-based representation where the voxels that are projected into the foreground regions within the images of all individual views are considered as occupied and the remaining voxels are set to empty. As a result, the respective 3D surface is represented by the isosurface between the occupied voxels and the empty voxels. However, reliably extracting silhouettes often represents a rather challenging task due to effects such as shadows and occlusions, and, additionally, the properties of the surface materials have to be considered. Moreover, only the shape of objects with a rather simple surface geometry can be accurately acquired. Therefore, this kind of technique is rather limited and typically has to be combined with additional normal information as in e.g. [CLL07, Dai09] or multi-view stereo consistency as in e.g. [ES04, CK11] to allow for the acquisition of concavities in the surface geometry of the considered objects. In general, shape-from-silhouette techniques can neither be categorized according to active or passive methods or directly be categorized based on the taxonomy of approaches for the different surface types as the required information might be obtained in different ways for objects made of different materials. However, as the large majority of the corresponding techniques relies on extracting the silhouettes without any active scene manipulation, we list shape-from-silhouette approaches under the passive techniques.

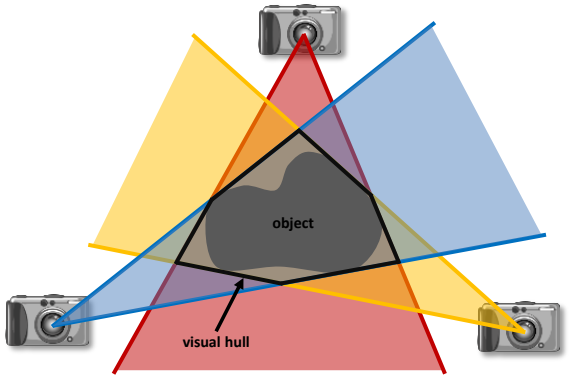

Figure 17: Principle of shape-from-silhouette: Silhouettes are extracted in the individual images and projected into the volume. The respective cones are intersected to obtain the visual hull of the object.

5.1.2. Principles for Active Geometry Acquisition

In contrast to passive methods, the principle of active acquisition techniques relies on the interaction with the scene content based on the acquisition hardware. For instance, the scene illumination might be manipulated by actively projecting light patterns that might facilitate the establishment of point correspondences or contact-based feelers might be used to reconstruct the surface geometry based on haptics.

While contact-based 3D geometry acquisition techniques are capable of reconstructing the 3D geometry of almost all solid objects by e.g. using feelers attached to manipulable arms, such approaches are typically considered as impractical. The pointwise surface acquisition results in long acquisition times as dense measurements of the object surface are desirable to accurately reconstruct the surface details. Furthermore, several objects made of e.g. fabrics or clay cannot be acquired contact-based as the pressure of contact might deform or even break them. In particular, the invaluable and typically sensitive cultural heritage artifacts have to be handled as careful as possible which is contradictory to a contact-based acquisition. For this reason, we will skip the discussion of such contact-based approaches in our review.

5.1.2.1. Photometric Stereo

The principle of *photometric stereo* as introduced in [Woo80] is closely related to the principle of shape-from-shading (see Section 5.1.1). However, while shape-from-shading techniques typically rely on a surface reconstruction from a single image with several additional constraints, the key idea of photometric stereo is the

acquisition of several images depicting an object from the same viewpoint under N varying illumination directions $\mathbf{l}_i \in \mathbb{R}^{3\times 1}$ with the corresponding illumination strengths $r_i \in \mathbb{R}$ and $i = 1,\ldots,N$. The resulting equation system is based on Lambert's law which relates the intensity $I_i(x,y)$ observed at pixel position (x,y) in the i-th image acquired under illumination direction $\mathbf{l}_i$ and illumination strength r_i to the surface albedo $\alpha(x,y)$ and the surface normal $\mathbf{n}(x,y)$ according to

$$I_i(x,y) = \alpha(x,y)\, r_i\, \mathbf{l}_i\, \mathbf{n}(x,y). \tag{29}$$

The constraints obtained for all acquired images can be combined in an equation system

$$\underbrace{\begin{bmatrix} I_1(x,y) \\ \vdots \\ I_N(x,y) \end{bmatrix}}_{\mathbf{I}(x,y)} = \alpha(x,y) \underbrace{\begin{bmatrix} r_1\, \mathbf{l}_1^T \\ \vdots \\ r_N\, \mathbf{l}_N^T \end{bmatrix}}_{\mathbf{L}} \mathbf{n}(x,y) \tag{30}$$

where the rows of $\mathbf{L}$ need to be linearly independent so that the inverse of $\mathbf{L}$ can be calculated. To solve this equation system, $N \geq 3$ images need to be acquired under varying illumination conditions. This allows to obtain a solution for the surface normals and the surface albedo by e.g. least squares estimation. Representing the surface normals as $\mathbf{n}(x,y) = [p(x,y), q(x,y), 1]^T$ with the respective surface gradients $p(x,y) = \frac{\mathrm{d}z(x,y)}{\mathrm{d}x}$ and $q(x,y) = \frac{\mathrm{d}z(x,y)}{\mathrm{d}y}$ allows to apply surface normal integration schemes.

Several techniques such as the ones proposed in e.g. [BJK07, WGS*11, WLDW11] consider extending photometric stereo towards general unknown illumination. Furthermore, photometric stereo has been explored in multi-view setups (e.g. [EVC08, BAG12]). The approach presented in [BAG12] is based on using normal consistency by determining a maximal set of inliers per voxel on which regular photometric stereo is applied in a multi-view approach. While producing good reconstructions on synthetic data, the estimated surface consistency tends to being localized inaccurately for real-world data due to the lack of a per-voxel normalization. Furthermore, multi-view normal field integration approaches as proposed in [CLL07, Dai09] have been considered in the context of photometric stereo. These techniques overcome the problem of obtaining only 2.5D reconstructions of partial surfaces in the single-view case. In [CLL07], an initial visual hull reconstruction is followed by an iterative surface evolution based on level sets in a variational formulation. As no global optimization is performed, the surface evolution is sensitive to the initial visual hull. In contrast, the technique proposed in [Dai09] is based on a Markov Random Field (MRF) energy function where the surface is computed via min-cut to find a global minimum. This is followed by a smoothing step similar to the one applied in [CLL07]. A surface orientation constraint has been included in the energy functional which enforces the reconstructed geometry to agree with the observed surface normals. Both techniques employ additional silhouette information which, however, are rather difficult to determine. Furthermore, several other approaches focus on improving photometric stereo by considering normal hypotheses in a volumetric representation and the surface is assumed to pass through voxels with a high consistency of the hypothesized normals. In e.g. [CJ82, MC09], classical single-view photometric stereo has been improved by selecting only hypotheses which agree with the underlying model assumptions. Generating per-pixel normal hypotheses for varying lighting directions has also been used in [HMI10] where consistency is obtained by considering monotonicity, visibility, and isotropy properties. Therefore, the approach can handle both diffuse and specular surfaces. The multi-view normal integration framework presented in [WORK13] relies on the generation of normal hypotheses in the considered volume. Several normal hypotheses are obtained per voxel using photometric stereo and both local surface normals and local consistencies can be obtained based on a non-parametric clustering of normals. Both the local normal estimates and the consistencies are used as input of a continuous min-cut based optimization with a final refinement.

5.1.2.2. Shape from Time-of-Flight / Phase Measurements

A different strategy for active geometry acquisition is followed by scanning devices which illuminate a scene with light modulated by an approximative sinusoid of a specific frequency f_m). The emitted laser light is (partially) reflected at an object surface and a certain amount of the laser light returns to the receiver unit of the scanning device. Focusing on different properties, respective techniques can further be categorized with respect to the laser type, the modulation technique, the measurement principle, the detection technique or the arrangement of emitting and receiving component of the involved scanning device [SJ09]. In this regard, however, the measurement principle may be identified as one of the most important aspects as it may generally rely on different signal properties such as amplitude, frequency, polarization, time or phase.

The most prominent technique for active geometry acquisition is based on the *time-of-flight principle*, where the key idea consists in the measurement of the time that has elapsed between the timestep t_e when the laser light has been emitted and the timestep t_r when the laser light has reached the sensor. This allows to compute the distance the laser light has traveled according to

$$d_{\text{full light path}} = v_{\text{light}}\,(t_r - t_e), \tag{31}$$

where v_{light} denotes the speed of light. Accordingly, the distance d to the respective 3D scene point is represented by

$$d = \frac{d_{\text{full light path}}}{2} = \frac{v_{\text{light}}(t_r - t_e)}{2}. \tag{32}$$

As many scanning devices emit the laser light periodically with a certain modulation frequency f_m, the measured distance d is limited to the interval between 0m and a maximum value of

$$d_{\max} = \frac{v_{\text{light}}}{2 f_m}, \tag{33}$$

where $d_{\max}$ represents the non-ambiguous range.

In contrast, it is also possible to measure the phase shift $\Delta\phi$ between emitted and received signal instead of the respective time-of-flight. In this case, the modulation is typically a nearly sinusoidal signal. This special case is known as amplitude modulated continuous wave (AMCW) lidar and allows to derive the distance d to the respective 3D scene point according to

$$d = \frac{v_{\text{light}}}{2 f_m} \frac{\Delta\phi}{2\pi}, \tag{34}$$

where the derived distance d is also non-ambiguous for the interval between 0m and $d_{\max}$.

For distances beyond $d_{\max}$, the obtained distance d is ambiguous. In the case of phase measurements, the measured phase shift $\Delta\phi$ is only a wrapped phase and, in order to recover the unwrapped phase, multiples of 2π have to be added according to

$$\phi = 2\pi k + \Delta\phi \tag{35}$$

with $k = 0, 1, 2, \ldots$, which results in the correct distance

$$d_{\text{corr}} = d_{\max} k + d. \tag{36}$$

For a respective disambiguation by finding the suitable value k, various unwrapping procedures have been proposed [Jut09, DHB10, Jut12].

However, except for the case where the direction of the emitted laser light is perpendicular to the object surface (i.e. $\theta_i = 0$), a respective geometry reconstruction faces problems at shiny surfaces exhibiting glossy reflection, since only a small amount of the laser light reaches the receiver unit again, and at those surfaces exhibiting specular reflection, where no laser light will reach the receiver unit again.

In order to acquire the geometry of object surfaces instead of single 3D scene points, the involved sensor is typically mounted on a rotating platform which allows a successive scanning of the scene with respect to a given scan grid. While a rotation around a single rotation axis is for instance applied for line laser scanners which provide measurements on a 1D scan grid, standard terrestrial laser scanners allow a rotation in both horizontal and vertical direction and thus provide measurements on a 2D scan grid. As a consequence of such a sequential scanning of a scene, however, respective acquisition systems are not suited for the acquisition of dynamic scenes.

Instead of a successive scanning of points on a specific scan grid, modern range cameras use a sensor array in order to simultaneously acquire range information for all points on the considered scan grid and thus estimate the scene geometry in a single shot. Most of these range cameras are even able to deliver depth maps at video rates [KBKL10, DZC12, RS13, WLH*14, FSK*14] and hence they are also suited for the acquisition of dynamic scenes.

Another problem more specific to AMCW lidar systems arises from global illumination effects. These cause the light to reach points in the scene on multiple paths of different lengths. This is known as multi-path interference and can cause strong systematic distortions in range measurements. Several investigations use measurements at multiple modulation frequencies to separate paths of different lengths thus mitigating multi-path interference [DGC*11, FSK*14, BFI*14, PKHK15].

5.1.2.3. Shape from Active Triangulation

Active geometry acquisition can also be performed similar to passive geometry acquisition via triangulation. Respective techniques are inspired by stereovision which relies on imaging the scene from two points of view and subsequently deriving correspondences between the different images in order to triangulate the 3D position of a respective scene point. However, for typical scenes, difficulties in finding correspondences may arise, e.g. due to weakly textured object surfaces. Such stereo techniques can directly be transferred to active stereo techniques by replacing one of the cameras by a projection device as illumination source. In the easiest case, the projection device consists of a laser and projects a spot on the object surface [MVGV10]. The laser spot on the object surface, in turn, can easily be detected in the image taken by the camera and, in case the relative position and orientation between laser and camera are known, the laser beam origin, the projective center of the camera and the laser spot on the object surface form a triangle from which the 3D coordinates of the laser spot on the object surface can be derived [HW11]. To acquire the shape of an object, the laser is directed at different points on the object surface and a photograph is taken for each case in order to successively derive 3D points on the object surface. Instead of using such a single-point *laser triangulator*, it is also possible to equip the laser with a cylindrical lens in order to expand the laser beam along one direction and, hence, obtain laser stripes [STD09, MVGV10]. As a result, the intersection between the object surface and the plane containing the laser rays yields a curve, and the respective 3D points on this curve can again be determined via triangulation.

To avoid involving a mechanical apparatus, a more complicated and two-dimensional pattern is required, where some kind of code is included into the pattern to allow a distinction of different parts of the pattern and thus be able to reliably derive point correspondences [MVGV10]. In this regard, it has proven to be feasible to replace one of the cameras of a typical stereo setup by a device (e.g. an LCD video projector or a slide projector) that projects a coded *structured light* pattern and thus manipulates the illumination of the scene (see Figure 18). Thereby, the coded structured light pattern has a specific structure. When imaging the scene with the pattern reflected at the respective surfaces with a single camera or a set of cameras, pixels in the respective camera images, where the patterns have been observed, can be assigned a code-word and, hence, different pixels can easily be distinguished by means of a local coding strategy. As a consequence, the correspondence problem can be solved without taking into account geometrical constraints and, hence, active geometry acquisition based on the use of coded structured light can be considered as a reliable strategy for recovering object surfaces [SPB04].

Typically, the pattern design focuses on obtaining code-words which are effectively decodable in case of a non-ideal pattern projection or acquisition process and, accordingly, techniques relying on the use of structured light patterns can further be categorized with respect to the applied pattern codification strategy [SPB04, DZC12]:

- *Direct codification strategies* rely on the use of a pattern which typically involves a large range of either gray or color values. Thus, the code-word associated to a pixel is simply represented by the respective gray or color value of the pattern at this pixel. While respective techniques are easy to implement, they also allow a geometry acquisition in case of dynamic scenes due to the use of a single pattern. However, such techniques tend to be extremely sensitive to color or gray-level distortion due to scene

color distribution, reflectivity properties of scene objects and external illumination.

- *Time-multiplexing codification strategies* rely on the use of a temporal coding, where a sequence of structured light patterns is projected onto the scene. Thus, the code-word associated to a pixel is represented by the sequence of illumination values across the projected patterns. In this context, the respective patterns are often based on binary codes (e.g. represented in the form of stripe patterns or Gray codes) which can robustly be detected. However, due to the sequence of structured light patterns involved in the acquisition of the scene geometry, respective techniques are not suited for geometry acquisition in case of dynamic scenes.
- *Spatial neighborhood codification strategies* rely on the use of a unique pattern. The code-word associated to a pixel is derived from the spatial pattern distribution within the local neighborhood of the considered pixel. Thus, code-words of neighboring pixels share information and therefore provide an interdependent coding. Respective techniques allow a geometry acquisition in case of dynamic scenes due to the use of a single pattern and are relatively robust. However, the choice of the local neighborhood is crucial.

As already discussed for passive triangulation-based techniques, reconstructions obtained from triangulation-based methods typically contain high-frequency noise in the reconstructed depth values as a result of inaccuracies in the 2D positions of the extracted point correspondences and inaccuracies in camera calibration. Nevertheless, such active triangulation-based laser scanner systems or structured light systems reach a remarkable scanning accuracy for diffuse surfaces and can also be used for surfaces with an additional specular component as long as the material still shows a sufficient surface albedo. However, these techniques are not robust with respect to effects such as interreflections, subsurface scattering or highly specular surface reflectance. Multi-view observations might provide a certain robustness with respect to these effects. However, if the surface is almost ideally mirroring and does not show a sufficient albedo, the observed pattern information becomes completely view-dependent and cannot be used to establish correspondences across images taken from different views. In a similar way, the correspondences cannot be established in case of transparent or translucent surfaces.

When using projector-camera systems for 3D reconstruction, one key observation can be identified in the fact that the resolution of currently available projectors is significantly lower than the one available in standard cameras. As a result, the footprint obtained by projecting a unique projector pixel extent onto the object surface usually covers several pixels in the camera which observes the scene. As active illumination relies on the unique encoding and decoding of the per-pixel projector illumination, this means that the resolution of the surface points obtained during the reconstruction process is limited by the projector resolution. In order to obtain a denser and more accurate reconstruction of the surface with its fine surface details, it is therefore mandatory to overcome this limitation.

Several approaches have focused on this issue. In [RSGS10], a 1D mechanical lens-shifter extension has been used in front of the projector which allows a fine shifting of the projected pattern in steps that are significantly smaller than the pixel size. This strategy allows remarkably accurate reconstructions of the object surfaces at high resolution. Other techniques have focused on increasing the resolution by the use of setups with several cameras and projectors. In [AX08, AX10] a self-calibrating, multi-view acquisition technique based on structured light and photometric methods has been proposed, where either a single projector is placed at several positions to simulate the availability of multiple projectors or multiple projectors are involved respectively as well as multiple cameras. This technique exploits the principle that digital projectors can be simultaneously used as either active light sources or as virtual cameras, which is particularly useful if the number of involved cameras is rather small. An iterative bundle-adjustment is applied in combination with an outlier rejection.

In [WSRK11], a setup with a multitude of cameras and projectors mounted on a hemispherical gantry above the object to be scanned has been used. This allows to perform the acquisition of the object geometry without moving either the object or parts of the acquisition setup and, consequently, there is no need to register several independent measurements. To overcome the resolution limitations of the individual projectors, a novel super-resolution scheme has been introduced. Sequentially illuminating the object using projectors at a sufficient number of different locations to cover the whole object surface with several structured light patterns allows to exploit the overlappings of the patterns of projectors at different positions. Therefore, much smaller regions on the surface can be uniquely identified (see Figure 19) which results in significantly denser and more accurate point clouds. Similar to the techniques in [AX08, AX10], that do not utilize the overlapping of different projected patterns on the object surface, an iterative bundle adjustment for the refinement of the camera calibration has been applied. However, the calibration of the projectors has been omitted, since there is a sufficient number of cameras for establishing correspondences within the captured images. In addition, high dynamic range imaging has been used to relax the assumptions regarding the reflectance behavior of the objects to be scanned, and even complicated objects exhibiting strong specularities can be scanned.

5.1.2.4. Shape-from-Specularity Techniques

Shape-from-specularity (see Figure 20) approaches are rather often used for the acquisition of highly specular materials. Such techniques rely on directly using the information revealed by specular highlights observed on the object surface. Based on the law of reflection discussed in Section 3.2.3, specularities can be observed if the angle between the light direction and the local surface normal **n** and the angle between the view direction and the local surface normal **n** are identical. Given both the viewing direction and the light direction, the observation of specular highlights allows to directly infer corresponding local surface normals. In order to get an accurate reconstruction of the object surface, normals have to be densely. For this reason, it is crucial to obtain dense observations of specular highlights on the object surface which can be obtained by varying the viewing conditions [ZGB89], varying the illumination conditions [CGS06], using extended light sources [Ike81] or using arrays of light sources [SWN88]. Shape-from-specularity techniques rely on the presence of strong highlights. If materials

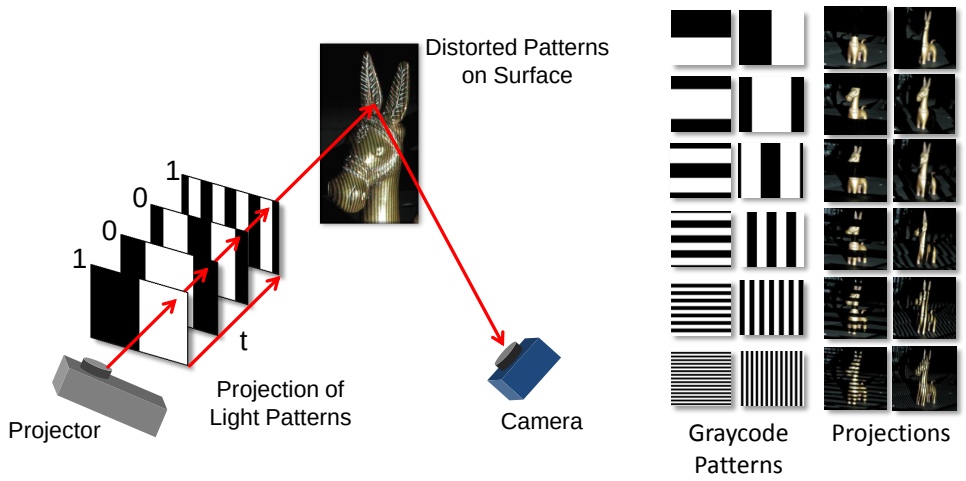

Figure 18: Illustration of active triangluation based on projections of structured light patterns.

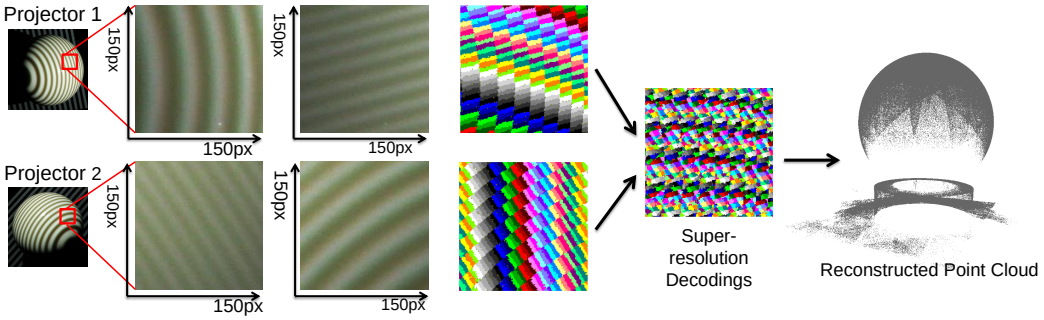

Figure 19: Illustration of the effect of multi-projector super-resolution on structured light decoding.

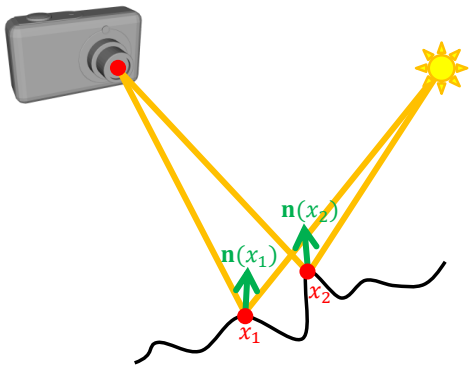

Figure 20: Principle of shape-from-specularity approaches: Specular highlights can be observed on the object surface of a specular material, if the angle between the light direction and the local surface normal and the angle between the view direction and the local surface normal are identical.

are only glossy and only show blurred highlights, the center of the highlight cannot easily be determined, and, hence, surface normals cannot be reliably estimated.

5.1.2.5. Geometry Estimation Based on Prior Knowledge from Databases

In recent years, the advances in the field of machine learning have led to the development of geometry estimation techniques that rather make use of an a-priori collected database of registered images and scene geometries than using standard cues for geometry estimation such as photoconsistency, specularities, etc.. These techniques usually involve a training based on pre-registered RGB-D images as provided by the Make3D dataset [SSN09] or the NYU datasets [SF11, SHKF12] where a depth image is provided per RGB image for different outdoor and indoor scenes. After the training phase, these frameworks allow the inference of coarse depth information for novel RGB images. Several of these approaches are based on using Markov Random Field (MRF) formulations for depth estimation [LGK10, SCN06, SSN09]. Although these methods typically only use approximate techniques for MRF learning and inference, they still require several seconds

for the depth prediction for a new image. In a similar spirit, a discrete-continuous Conditional Random Field (CRF) is used in [LSH14] where the relations between adjacent local regions are taken into account.

Other approaches exploit the opportunities offered by Convolutional Neural Networks (CNNs), which represent one of the currently most powerful techniques for various tasks including e.g. classification and synthesis. The most accurate techniques of this kind differ in the way the neural networks are used. In [EPF14], two CNNs are combined where a coarse-scale network is used to predict depth information at a global level and a subsequent fine-scale network is used to refine the coarse depth information within local regions, and, hence, better accounts for objects or wall edges. This approach has been extended in [EF15] where different scales are considered by sequentially applying three CNNs in order to address the derivation of multi-scale depth information with an additional estimation of surface normals and a semantic labeling. The coarse-scale depth prediction provided by the first network based on the entire image is refined by predictions at a mid-level resolution of the second network and by even more fine-scaled predictions of the third network. A major drawback of these strategies is the fact that the learning of depth information shows some kind of overfitting with respect to specific scene layouts and, hence, a huge database of millions of RGB-D images is required to include as much scene layouts as possible into the training data.

In [LSL15, LSLR15] one CNN is used to derive coarse depth information for the centroids of local regions extracted from the image over which the depth is assumed to be constant and one network is used to account for similarities of neighboring patches to refine the depth estimates. The networks are combined using a CRF formulation. As no spatial information of the local patches is considered for the coarse depth estimation, the technique offers the advantage of translation invariance and, as a consequence, less training data is required.

5.1.2.6. Combinations of Different Methods

The individual strengths and weaknesses of triangulation-based techniques and photometric techniques are complementary. While triangulation-based techniques offer a remarkable accuracy regarding the reconstruction of the low-frequency object geometry, they typically suffer from noise due to inaccuracies in the image positions of the correspondences. In contrast, photometric techniques typically allow an accurate reconstruction of high-frequency details of the surface geometry such as scratches or engravings and have problems in accurately reconstructing the low-frequency geometry. Consequently, a combination of both types of acquisition techniques may allow to overcome the drawbacks of the individual techniques and, hence, to obtain a precise acquisition of both the low-frequency shape and surface details. However, the adequate combination of both types of techniques has been proven to be rather challenging as e.g. their simple combination is not sufficient when the low-frequency geometry information is not sufficiently accurate and leads to incorrect parallax and occlusion effects at grazing angles [NRDR05].

Much effort has been spent on the investigation of strategies regarding an adequate combination of techniques which mainly differ in the individually involved acquisition techniques and the respective reconstruction frameworks. Many techniques based on deformable surface models to guide an iterative surface evolution can be based on volumetric surface representations [ES04, JCYS04] or polygonal meshes [NRDR05, BCJS06, EVC08, LTBEB10, AX10, DPB10, DP11, WLDW11]. Such techniques evolve the surface based on variational formulations using gradient flows. One of the most popular techniques has been presented in [NRDR05], where a triangulation-based range scanning system is used to acquire low-frequency information and photometric stereo is used to derive normal information. The respective optimization functional includes a term that penalizes positional deviations between optimized positions and measured positions and a term that penalizes deviations of the respective surface normals. The optimization with respect to the depth is performed based on linear least squares optimization and the surface is deformed accordingly. As such surface evolution approaches are typically susceptible to local minima, they rely on a good initialization. Furthermore, topological changes of the surface geometry such as self-intersecting regions in the mesh during the iterative optimization are usually problematic for mesh-based techniques and hybrid approaches such as the one in [YY11] might have to be used to circumvent this problem. In [RK09], a combination of multi-view stereo and photometric stereo has been used to directly reconstruct surface heightfields and a SVBRDF is recovered as well based on a non-linear optimization. While this approach also takes interreflections into account, it is limited to planar samples and not applicable to objects with a complex surface geometry.

Other techniques use global optimization frameworks such as graph-cuts [SP05, HK06, YAC06, VHTC07, LBN08, HMJI09] or convex optimization [KPC10], where the cost function is globally defined on the complete volume. This allows to avoid the problem of local minima. However, obtaining accurate geometry reconstructions with discrete graph-cut techniques requires high-resolution volumetric grids. In addition, a rather high grid connectivity has to be used to alleviate metrification errors. Taking these aspects into account leads huge memory requirements. Strategies to reduce the computational burden include the use of adaptive grid structures as applied in e.g. [WRO*12] and continuous formulations of the optimization functionals [CK11, KKH*11, YBT10, WRO*12]. In particular, the approach in [WRO*12] is based on a continuous min-cut formulation [Str83, Str10] defined on an octree structure, where a structured light based surface consistency measure and normals derived based on Helmholtz stereopsis [ZBK02] are used as input. A final refinement is carried out according to the measured normal information in order to adjust the surface in a narrow band around the surface estimated using the continuous min-cut framework. This allows to obtain highly accurate reconstructions.

5.2. Geometry Acquisition of Rough Surfaces with Diffuse or Near Diffuse Reflectance

For diffuse surfaces (see Section 3.2.1), the incoming light is uniformly reflected into the full hemisphere with respect to the local surface normal of the object geometry. This means that the surface can be perceived in a view-independent way. Therefore, the geometry of such diffuse objects can typically be acquired in a rather easy way. In particular, most of the standard principles for geom-

etry acquisition discussed in the previous section can be applied to reconstruct the surface geometry of diffuse objects. Multi-view stereo techniques, active triangulation techniques and photometric stereo techniques are among the widely used methods. However, the acquisition of objects with spatially varying albedo is still a challenging problem.

In order to handle objects with non-uniform surface albedos, two different exposure times are used in [SS03] and the one which leads to the largest absolute difference between the two illuminations is selected. For the same reason, the idea of projecting multiple Gray code patterns at different illumination intensities and forming high dynamic scale radiance maps in order to decrease the susceptibility to misclassification caused by the reflectance properties of the considered surfaces has been introduced in [SL00]. Though these approaches address the problem of over- and underexposure, they do not compensate for the non-linearity of the response curve of the camera. In contrast, high dynamic range imaging has been applied in [WSRK11] to overcome this problem. An alternative approach has been proposed in [KPDVG05] where the projector intensities are adapted locally so that the dynamic range of the illuminated scene is reduced in order to avoid over- or underexposure. However, this method suffers from a decreased contrast in the adapted projector pixels. This problem does not occur when, instead of adapting the dynamic range of the scene, the scene is completely captured by taking a sequence of images with different exposure times.

5.3. Geometry Acquisition of Glossy Surfaces with Mixed Diffuse and Specular Reflectance

Considering materials with both diffuse and specular components (see Section 3.2.2) makes the 3D shape acquisition more complex as the perception of glossy highlights is a view-dependent phenomenon. The methods that have been proposed in this context are based on different principles and are described in the following sections.

5.3.1. Structured Light Techniques

Structured light techniques are also applicable for materials with both a diffuse and a specular reflectance component as long as the surface material shows a sufficient albedo under different viewing conditions. In these cases, the projected patterns can still be reliably decoded in the acquired images. Some techniques focus on further increasing the robustness with respect to the range of materials that can be coped with by using pattern projections at different intensity levels [SL00] or by capturing images for different exposure times [SS03]. In [WSRK11], it has been demonstrated that even the shape of moderately specular objects can be accurately acquired with a structured light system – which is also based on high dynamic range imaging – as long as the material still shows a sufficient surface albedo (see Figure 21). An example is shown in Figure 22. Limitations of structured light approaches, as mentioned in the previous section, are the lacking robustness with respect to effects such as interreflections, subsurface scattering, or highly specular surface reflectance. To some degree, using multi-view observations might improve the robustness with respect to these effects. However, for almost perfectly specular surfaces with a low albedo, the pattern information observed by the cameras is view-dependent. This makes the establishment of correspondences across images obtained under different view conditions based on the observed pattern information impossible.

5.3.2. Extensions of Photometric Stereo

With the objective of also acquiring such objects with more complex surface reflectance behavior, several methods considered extending classical photometric stereo. Therefore, violations of the underlying assumption of Lambertian reflectance due to specularities have to be considered as well as shadows. One way of relaxing the assumption of diffuse surface reflectance behavior made by conventional photometric stereo techniques is given by the replacement of the underlying diffuse model with a more general model that includes both a diffuse and a specular component [Ike81]. Such a model with a diffuse and an additional specular component has been used in [NWSS90], where multiple extended light sources are used to illuminate the object. The extended light sources allow to widen the specular lobe and to balance the intensities of diffuse and specular reflections [HW11]. A combination of a diffuse component and a specular component has also been used in [Geo03]. However, the specular component has been modeled based on the Torrance-Sparrow model [TS67]. As rather small angular deviations of the view direction and the light direction have been assumed, Fresnel reflectivity can be neglected as well as shadowing effects. The estimation of surface albedo, light sources and BRDF model parameters have been determined iteratively based on a linear least-squares optimization. In e.g. [HS05,GCHS05,RK09], spatially-varying BRDFs have been considered to broaden the range of materials that can be handled. Based on the assumption that the surface reflectance behavior can be modeled based on only a few base materials and based on the use of an isotropic Ward model [War92], both surface normals and the parameters of the reflectance model have been estimated per spatial point on the object surface. However, several effects such as e.g. interreflections and shadows are not taken into account in the techniques presented in [HS05,GCHS05]. The extension of photometric stereo proposed in [CJ08] is based on a normal estimation using shadow boundaries and the estimation of a BRDF where the Ward model [War92] is used. The obtained BRDF parameters are then clustered to find the material types and discard noise and outlier values. Once the specular parameters are known, their method proceeds with estimating normals and surface albedo. Limitations of this method can be seen in the use of a Ward BRDF model, which is only valid for a rather small range of materials. Therefore, this technique is only capable of handling isotropic materials with a single lobe. Materials for which more than a single lobe has to be fitted as well as the huge range of materials which exhibit mesoscopic effects cannot be handled adequately. Furthermore, only a 2.5D reconstruction from a single view is performed based on an orthographic projection model and normal field integration techniques are susceptible to errors in the estimated normals that are accumulated during the surface reconstruction. In [ZMLC10], photometric stereo has been extended by representing specular reflection with a set of specular basis functions with different roughness values. Therefore, specular objects can be handled as well by their photometric stereo technique and the final reconstruction consists of information regarding surface shape and also reflectance properties. However, the range of

Figure 21: Example reconstructions of glossy objects using structured light techniques.

materials that can be handled is limited to isotropic materials that can be represented by the Ward model with the specular basis functions. Furthermore, only a 2.5D reconstruction from a single view is considered and an orthographic camera projection is assumed.

5.3.3. Techniques Based on the Separation of the Diffuse Component and the Specular Component

Other techniques approach the reconstruction of objects with a more complex reflectance behavior beyond diffuse reflection by the separation of the diffuse and the specular component. A survey on respective techniques is provided in [HW11]. As discussed in [TI05, HW11], specular reflections show a larger degree of polarization in comparison to diffuse reflections, the intensity distributions of diffuse reflections follow Lambert's law whereas specular reflections follow the Torrance-Sparrow model [TS67] or the Beckmann-Spizzichino model [BS87], and specular reflections are largely independent of the spectral reflectance behavior of the object surface in the visual spectrum in contrast to diffuse reflections. Polarization has been exploited in [Wol89, WB91] to identify the specular component. In [NFB97], additional color information has been used to allow the separation of both components based on a single image. Further investigations such as the ones in [TI03, TI05] have focused on the separation of the diffuse component and the global component in a single image using color. This has been achieved by an initial normalization of the illumination color of the respective image to obtain white specular components and a subsequent per-pixel analysis of the maximum chromaticity over the intensity space. In [TS91], this approach has been extended so that the illumination constraint can be relaxed. The illumination directions have been estimated based on observed highlights and the assumption of mirroring reflectance. Similar light directions have been merged to obtain an approximation of the illumination environment. This allows to perform the separation based on an analysis of the consistency of the observed intensity values with respect to the illumination conditions. For non-consistent pixel values, an extrapolation of their neighborhood is used and, finally, both components can be separated using the technique presented in [TI03, TI05]. Furthermore, removing specular highlights has been exploited in [TLQS03] to allow the separation of the diffuse and the specular component in a single image. An illumination-constrained inpainting has been used to estimate the diffuse color. In contrast, the separation technique for dichromatic surfaces presented in [MZKB05] is based on an analysis of a set of at least three input images regarding specular highlights. The latter have been removed and a photometric stereo technique has been applied on the diffuse component.

5.3.4. Techniques Based on Helmholtz Reciprocity

In order to handle the even larger range of opaque materials, the approaches presented in [ZBK02, ZHK*03] exploit the Helmholtz reciprocity for surface normal estimation in a BRDF-invariant manner. At each point, the corresponding normal information is estimated via the Helmholtz principle [ZBK02]. This principle is based on the reciprocity of the BRDF, i.e. $\rho(\mathbf{v},\mathbf{l}) = \rho(\mathbf{l},\mathbf{v})$, where $\mathbf{l} = (\theta_i, \varphi_i)$ denotes the light direction and $\mathbf{v} = (\theta_r, \varphi_r)$ the view direction, and can be utilized wherever several image pairs are available in which the position of the light source and the camera have been exactly exchanged [ZBK02]. In [DPB10], Helmholtz normals have been used in a multi-view setting. Furthermore, a combination of structured light consistency and Helmholtz normals has also been used in the multi-view reconstruction in [WRO*12] (see Figure 22).

5.4. Geometry Acquisition of Smooth Surfaces with Ideal or Near Ideal Specular Reflectance

The challenge in reconstructing highly specular surfaces arises from the fact that such surfaces typically do not have an own characteristic appearance but rather reflect the surrounding environment in a view-dependent manner (see Section 3.2.3). Recent surveys on techniques for surface reconstruction of specular objects have been presented in [IKL*10, BW10]. Furthermore, a more theoretical discussion is given in [KS08]. Approaches based on different principles are discussed in the following sections.

5.4.1. Shape-from-Specular-Flow Techniques

Some of the methods, such as shape-from-specular-flow techniques [RB06, AVBSZ07], are based on considering the movement of environment features which are mirrored on the surface of specular objects. Typically, a known motion of the mirroring object, its environment or the cameras is assumed. Unfortunately, interreflections might cause a single environment feature to be observed several times on the mirroring surface which makes the estimation of dense

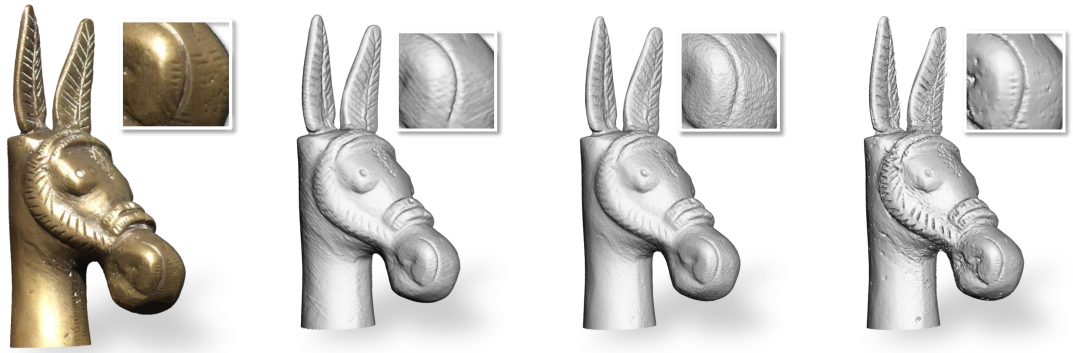

(a) Photo (b) Laser scanning (c) [WSRK11] (d) [WRO*12]

Figure 22: A comparison of geometry reconstruction techniques: photo of an object, reconstruction obtained using a highly accurate laser scanner, reconstruction obtained using the super-resolution structured light system presented in [WSRK11] and reconstruction obtained by using both structured light information and normal information in the efficient framework of [WRO*12]. The additional integration of normal information significantly improves the accuracy of the reconstructed geometry.

optical flow highly non-trivial. In addition, such methods usually rely on the assumption of a distant environment and do not consider a more complex scene geometry. Instead of considering dense correspondences, the approach in [SVTA10] is based on using sparse reflectance correspondences to locally approximate specular surfaces using quadrics.

5.4.2. Shape-from-Specularity and Shape-from-Reflection Techniques

As discussed in Section 5.1.2.4, shape acquisition of mirroring objects has been approached by utilizing the information revealed by specular highlights or specular reflections which are observed on the mirroring object surface due to specular reflection in controlled environments. However, densely sampled observations of specular highlights or specular reflections on the mirroring surface are required to obtain an accurate reconstruction of the surface geometry. Using specular highlights, this can be achieved by using a moving camera [ZGB89], moving the light source [CGS06], using extended light sources [Ike81] or sequentially switching on individual elements of a grid of light sources [SWN88].

As, in the case of using point light sources, the number of required images increases linearly with the number of utilized light source positions, taking one photo for each of the utilized light source positions would be impractical for the dense sampling of the light source positions required to obtain an accurate reconstruction. For this reason, several techniques focus on the reduction of the number of required images by performing measurements in parallel. For this purpose, rotating the object and using a circular light source has been proposed in [ZM00]. Furthermore, printed, static or moving calibrated patterns have been used in [BS03, SCP05, LWDC10]. In [Ike81], a planar Lambertian surface is illuminated by linear lamps and the light reflected from the surface illuminates the object as an extended light source. Ideal surface reflectance is assumed and the measured intensities in the acquired images can be related to the integrated irradiance over the whole extent of the light source. Finally, using several reflectance maps, surface normals are computed using a photometric stereo technique [Woo80]. Further methods make use of the simultaneous encoding of multiple light sources. Such encoding schemes have already been investigated in [NWSS90] for light source arrays, and several more recently published approaches build upon this idea by simulating dense illumination arrays using LCD screens and encoding the illumination emitted from the pixels using structured light patterns [TLGS05, FCM*08, NWR08, YIX07, BHB11]. While the projection of structured light patterns using displays can be applied as a parallel variant of shape-from-specularity techniques, they can be used to derive the surface geometry based on shape-from-distortion techniques as well, which typically rely on the observation of a distorted version of a known or unknown pattern by an imaging sensor, where the distortion is induced by the surface geometry [TLGS05]. So far, most of the approaches still rely on assuming far-field illumination or a distant environment. However, these assumptions are typically not fulfilled as the printed patterns or the LCD displays used for pattern projection have to be located rather close to the object surface in order to obtain dense observations of light directions or feature directions on the mirroring surface.

Knowing the view direction that is determined by the camera parameters and the pixel-based observation of a specular highlight, and the 3D position of the light source or feature on the utilized patterns is still not sufficient to infer the surface geometry due to a remaining normal-depth ambiguity. This ambiguity can be overcome in a multi-view setting as e.g. presented in [BS03], where a calibrated pattern is used to produce reflections on the specular surface. Based on a volumetric representation, the law of reflection is used to hypothesize a normal at each voxel. As a result, generally several normal hypotheses are obtained for the different view directions and light source positions or feature positions. This allows to approach surface reconstruction by assuming that the surface passes through the voxels with the most consis-

tent normal hypotheses. The respective normal consistency is computed per voxel according to a normal disparity measure. As voxels with a low consistency are discarded, such principles are often named as *voxel carving techniques*. The technique proposed in [Pak12] is also based on voxel carving but considers the consistency of normal vector maps per voxel that are used to parameterize the normal hypotheses per voxel. Several other works such as e.g. [CJ82, MC09, HMI10, BAG12] also consider techniques based on hypothesizing surface normals in the context of extending photometric stereo techniques (see Section 5.3). For specular surfaces, investigations on surface reconstruction based on the idea of matching hypothesized normals include the approaches proposed in [WI93, BHB11]. While considering overlapping deflectometric measurements obtained from multiple views can be used to reconstruct large mirroring surfaces as shown in [BHB11], self-occlusions are problematic for this approach and a lot of manual work is involved in configuring the individual views.

The technique presented in [YWT*11] considers normal consistency in a single-view setting by clustering per-pixel normal observations using the k-means algorithm [Ste57, Llo57, Mac67]. In [NWR08], a specularity consistency similar to the one in [BS03] is derived between a set of views in a triangulation-based scheme using a display with Gray codes for illumination. After triangulation, normals are refined for the estimated depth values in a way similar to the iterative scheme proposed in [TLGS05]. In [Pak13], a probabilistic voxel carving technique has been presented that uses an optimization based on loopy belief propagation. However, only synthetic data has been considered and, even in this rather ideal scenario, the reconstruction results are rather inaccurate.

The technique proposed in [WORK13] focuses on dense geometry reconstruction of mirroring objects by introducing a novel, robust, multi-view normal field integration technique. In this setup, a mirroring object is placed on a turntable where it is observed from eleven cameras on a vertical arc and illuminated with a series of structured light patterns by displays next to it. As not all of the projected structured light codes might be reliable decoded because of the dependency of the pattern distortion on the local surface curvature and the projector position or due to non-perfect mirroring surface characteristics, a local adaption of the decoding to the finest still resolvable pattern resolution is used. The corresponding information observed by a camera represents some kind of light map that encodes for each pixel with a structured light observation the origin of the light source on the display. These light maps are used to calculate individual volumetric normal fields for each combination of camera and illumination configuration and different rotation angle of the turntable. However, occlusions, outliers due to interreflections or noise are likely to occur for more complex object geometries that can be typically expected. As a consequence, the resulting normal fields typically contain regions with missing or even unreliable observations. To still allow a reliable acquisition of the object geometry, a non-parametric clustering of normal hypotheses derived for each point in the scene is used to derive both the local surface normal which is most likely and a corresponding local surface consistency estimate. The normal information and the consistency information per point in the considered volume are then used as input to a variational method, where an iterative, continuous min-cut approach [Str83, Str10] is used to reconstruct the surface geometry. Finally, the surface is refined by means of the measured normal in a narrow band similar to the approach in [WRO*12]. A highly accurate acquisition of the full 3D shape of mirroring objects with complex surface geometry can be achieved using this technique as demonstrated in Figure 23. In particular, the involved multi-view normal field integration scheme is probably the first of its kind which enables an accurate 3D reconstruction not only on synthetic data but also on real-world data. The optimization framework based on the continuous min-cut formulation and the normal-based refinement can also be used with other types of normals as input. For instance, a rather diffuse clay object with some glossy paint can be reconstructed based on using normals estimated with a standard photometric stereo technique as demonstrated in [WORK13].

5.4.3. Techniques Based on the Separation of the Diffuse Component and the Specular Component

Furthermore, the technique recently published in [TFG*13] focuses on the acquisition of objects with diffuse or specular surface reflectance behavior by using continuous spherical harmonic illumination. The response of the object to the harmonics can be used to separate the diffuse reflectance component from the specular reflectance component. Unfortunately, this technique is less suitable for the reconstruction of concave objects.

5.5. Geometry Acquisition of Smooth Surfaces with Ideal or Near Ideal Specular Refraction

Reconstructing the 3D shape of refractive objects (see Section 3.2.5) is even more challenging in comparison to the cases mentioned in the previous sections. In general, such objects might exhibit inhomogeneous reflectance characteristics induced e.g. by a spatially varying refractive index or by inclusions of Lambertian or opaque material components. As pointed out in the recent survey given in [IKL*10], research has mainly been spent on solutions relying on certain simplifying assumptions such as homogeneous material characteristics or considering only the reconstruction of a single surface separating the two enclosing media. The authors identify the main approaches for acquiring refractive surfaces being shape-from-distortion techniques, direct ray measurement techniques, reflectance-based techniques, techniques based on inverse ray tracing, tomography-based approaches and direct sampling techniques. In the scope of this section, we therefore group the related approaches according to these principles.

5.5.1. Shape-from-Distortion Techniques

For the simpler case of acquiring a single refractive surface, shape-from-distortion techniques have been successfully applied. While this kind of methods can also be applied for specular surface reconstruction in a simpler form, refractive surface reconstruction requires considering the refractive index in addition to the surface normal in order to analyze the light path. Early work [Mur90, Mur92] has focused on reconstructing water surfaces from a single view. The movement of the water induces the a-priori unknown background pattern placed at the bottom of the liquid to be observed in a distorted way. Assuming an orthographic camera, optical flow [HS81, LK81] and a subsequent integration of the surface

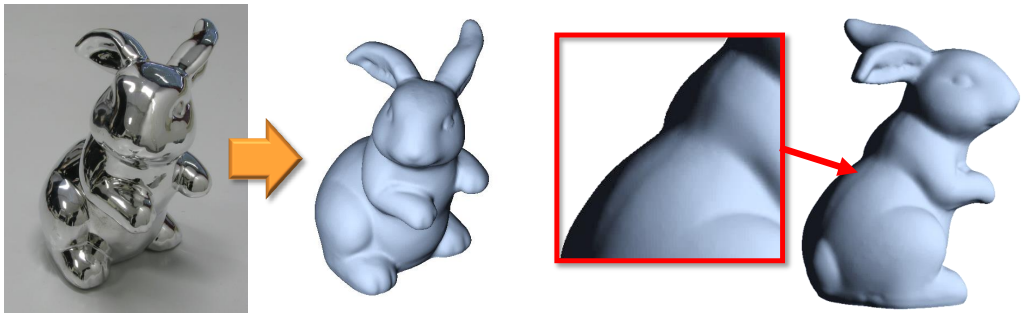

Figure 23: Bunny figurine and reconstructed model [WORK13].

gradient are used to reconstruct the surface up to a certain scale. This seminal work has been extended in [MK05] by using a stereo camera system and a known pattern to estimate the refractive index, per-pixel depth and surface normals. As a further improvement, no average surface model is used in comparison to the approaches in [Mur90, Mur92]. In particular, the considered consistency measure is similar to the normal consistency used in [NWR08] for specular surface reconstruction. Further work has been dedicated to the reconstruction of glass objects. Projecting structured light patterns into the refractive object with a projector and observing the respective distorted patterns in the camera image has been analyzed in [HSKK96]. In [BEN03], an unknown distant background pattern is used in combination with a known parametric model including shape and refractive index. The object of interest is moved in front of a single, static camera and features are tracked over time similar to [Mur90, Mur92]. In [AMKB04], an extension of optical flow has been proposed to track refracted scene features for which the intensity may vary due to the presence of non-ideally transparent surfaces with e.g. additional absorption.

5.5.2. Techniques Based on the Direct Measurement of the Light Path

Furthermore, refractive surfaces have also been reconstructed by directly measuring the light path. For this purpose, calibrated planar patterns in several positions with respect to the object have been used in [KS05, KS08] to measure the light rays. In their theoretical analysis [KS05, KS08], the authors consider the categorization of reconstruction techniques based on ray measurements independently performed for each pixel. The introduced notation $\langle N, M, K \rangle$ contains the relevant information with respect to the number N of views that are required for the reconstruction as well as the number M of points on specular or refractive surfaces that are located on a piecewise linear light path and the number K of calibrated reference points on a ray exitant from the object. The authors discuss that such a reconstruction cannot be performed for more than two intersections of the light ray with specular or refractive surfaces. The number N of views and the number K of calibrated reference points on a ray do not influence this observation. Following this concept, the authors consider a $\langle 3, 2, 2 \rangle$ reconstruction for refractive surfaces. As a result, four surface points with attached normals can be estimated per pixel. While one such pair of point and normal is located at the front surface of the object, the remaining pairs depend on the three differently refracted viewing rays and are located at the back surface.

5.5.3. Reflectance-Based Techniques

A reflectance-based approach has been followed in [MK07]. Dense per-pixel reflectance measurements in a static camera are observed as a result of sequentially illuminating the static object of interest with a light source at varying positions on a regular grid. As a result, a 2D slice of the BRDF is recorded. However, indirect lighting effects influence these measurements. By separating the direct and indirect components of light transport, the authors achieve high-quality reconstructions for depth and normal, even for inhomogeneous refractive objects.

5.5.4. Techniques Based on Inverse Ray Tracing

Other methods rely on the principle of inverse ray tracing. The underlying idea is based on the optimization of the residual of the acquired data and synthetically generated data. In order to reconstruct the surface of time-varying water surfaces, the water has been mixed with a fluorescent dye in [IM06] and a chemiluminescent chemical in [GILM07]. While using UV illumination makes the mixture of water and fluorescent dye self-emissive [IM06], in the case of chemiluminescence [GILM07] a chemical process has to be used for this purpose. Assuming homogeneous emission, both methods use synthetic images for surface fitting via level set optimization. In [WLZ*09], the liquid to be reconstructed is dyed with an opaque white paint. As a result, patterns can be projected onto the liquid and the correspondences observed in the cameras allow the reconstruction of the surface. In addition, a physically-based fluid simulation is used in this approach.

5.5.5. Tomography-Based Techniques

As discussed in [IKL*10], refractive object reconstruction can be performed by making use of certain acquisition strategies. One possibility is to consider sufficiently high wavelengths for the incident illumination as given for x-rays. In this spirit, computer tomography has been used for scanning objects in [KTM*02], and the proposed approach is in principle capable of scanning glass objects.

Furthermore, as mentioned in [IKL*10], a reconstruction of refractive objects is also possible when the refractive indices of the object and the surrounding medium are identical. In [TBH06], glass objects are immersed into a liquid with carefully controlled refractive index. Controlling the refractive index to a value of approximately 1.55 has been achieved by adding chemicals to water. As an ideal transparent object would disappear inside a medium with identical refractive index, the surrounding medium has to be dyed, which can be omitted if the object itself is absorptive [IKL*10].

5.5.6. Direct Sampling Techniques

Furthermore, several methods have been proposed based on a direct sampling. In [HFI*08], fluorescent immersion range scanning has been proposed to reconstruct refractive objects. The objects are placed in a different immersing medium with known refractive index. This liquid has additionally been dyed with a fluorescent chemical. During measurement, this fluorescent liquid causes the utilized laser sheet to be rendered visible while the refractive object to be scanned remains dark. A similar strategy has been explored by observing objects in a different spectrum. Several investigations e.g. focus on shape-from-heating. In [EAM*09], the glass surface is first heated by the incident infrared radiation which is then measured by an infrared camera. This allows to reconstruct the glass surface. Extending this approach, the shape of the hot spots observed by the infrared camera is analyzed in [AEB*12] to derive information regarding local surface orientations. Other works consider structured patterns of infrared radiation [MSSE*10] or the polarization in the infrared domain [MRA*12] for 3D reconstruction. Furthermore, several publications focus on exploiting ultraviolet radiation. In e.g. [RSFM10b,RSFM10a], structured light in the ultraviolet domain has been explored to reconstruct glass surfaces.

5.6. Geometry Acquisition of Surfaces Where Light is Scattered Multiple Times Underneath the Surface

The challenge in acquiring the 3D shape of translucent objects arises from the light transport within the object (see Section 3.2.4). In particular, the incoming light enters the material and travels through the material where it is scattered. When we actively illuminate such translucent objects with a pattern, these non-local subsurface scattering effects induce a blurring of the observed pattern and, hence, make e.g. a triangulation-based reconstruction from the decoded correspondences rather unreliable.

5.6.1. Techniques Based on the Manipulation of Reflectance Properties

In [GLL*04], the authors circumvent these problems arising from the subsurface scattering characteristics by covering the object of interest with a thin, diffuse dust before the actual 3D geometry acquisition via laser scanning is started. Later, this dust can easily be removed again.

5.6.2. Techniques Based on the Separation of Diffuse and Specular Components

Apart from using such tricks, surface geometry of translucent objects can also be acquired by utilizing certain material-specific characteristics of light transport. As many translucent objects also have a strong specular component, shape-from-specularity approaches can be applied where a moving light source is involved and the observed highlights can be used to estimate surface normals which is followed by a normal field integration [CGS06]. In [MHP*07], linearly polarized and circularly polarized spherical gradient illumination patterns are used, and both the diffuse and the specular reflectance is considered to estimate surface normals. The advantage of the circularly spherical patterns can be identified in the fact that they allow the simultaneous estimation of surface normals from different viewpoints. The proposed polarized illumination schemes allow an independent estimation for both diffuse and specular normal maps. The latter have been proven to be adequate for subsurface scattering materials in contrast to the diffuse normal estimates which are affected by the subsurface scattering. More recently, this method has also been used in [GCP*10] with circularly polarized spherical illumination for normal estimation from several viewpoints.

Furthermore, the investigations in [NKGR06] have demonstrated that specular and diffuse components of surface reflection can be separated by phase shifting of high-frequency structured light patterns. This observation has e.g. been explored in [CLFS07], where a phase shifting based structured light approach has been combined with a polarization-based removal of specular highlights at the surface. Based on the fact that global light transport characteristics remove the polarization of light, polarization filters are used in front of both the light source and the camera and multiple scattering effects can be separated from the structured light observations. In subsequent work [CSL08], the same authors remove the dependency on polarization and instead modulate the low-frequency phase shifting patterns to separate direct and global components of light transport. By this modification, the obtained reconstruction quality is further improved in comparison to the technique presented in [CLFS07].

5.6.3. Techniques Based on Structured Light

In [GAVN11], certain structured light patterns tailored to translucent surfaces have been proposed. While the analysis shows that high-frequency patterns are not applicable for translucent objects due to the blurring of the observed pattern, using Gray codes with a certain minimum stripe width following [GG03] shows more robustness on translucent surfaces and allows reliable reconstructions.

5.6.4. Extensions of Photometric Stereo

The recent technique presented in [DMZP14] represents an extension to conventional photometric stereo which enables the simultaneous estimation of both scattering properties and accurate surface normals for planar, homogeneous translucent objects based on observations from at least three different directional illumination configurations based on blind deconvolution. This avoids the problem of blurry normal estimates that would result from an acquisition via conventional photometric stereo.

6. Mesh Reconstruction

Most of the geometry acquisition techniques including the widely used time-of-flight and structured light scanners as well as most of the passive multi-view stereo techniques allow the reconstruction of surfaces in terms of point clouds. As already mentioned in the previous chapter, several e.g. from multiple views separately acquired point clouds can be aligned according to the geometric (extrinsic) calibration of the sensors in order to allow a dense, more complete surface representation of objects or scenes. However, for many purposes, a surface representation in terms of a closed, watertight surface such as a polygon-mesh is desirable. For instance, the accurate depiction of the appearance of an object or scene typically requires the use of closed surface representations on which e.g. reflectance information can be parameterized.

Extracting an accurate, consistent closed surface from point clouds still remains a challenging task and also an active research topic. In particular, the main challenges include the handling of noisy input samples due to inaccuracies of the scanners and their calibration or because of more complex material characteristics that cannot be handled by them. This requires algorithms that are robust enough to avoid overfitting to the noise and make use of redundancy to remove noise from the final surface. As some areas might not be accurately captured because of not being sampled or because of material characteristics, resulting holes in the scans need to be appropriately closed by the reconstruction algorithms. Furthermore, a possibly uneven sampled point cloud needs to be robustly handled and the amount of input data might be very large. The latter requires efficient data structures and optimization techniques are required to keep the computations feasible regarding both time and memory consumption.

In the scope of this section, we will discuss recent approaches for closed surface reconstruction based on two broad categories (see also Figure 24):

- *Global Combinatorial Algorithms:* These methods usually rely on a discretization of space using a tetrahedralization or a voxel based grid or tree. An energy function is then constructed at the discrete points and the domain is partitioned into inside and outside using global graph operations such as min-cut. Recent approaches in this domain include [LPK09, MKG11, WRO*12, WORK13].
- *Implicit Functions:* Another branch of very successful algorithms construct an implicit function that is fitted to the points and extract the isosurface from this function [FKG15]. The Poisson surface reconstruction [KBH06, KH13] transforms the oriented input points into a continuous vector field and reconstructs an indicator function whose gradients match the vector field using a 3D Poisson equation. The method by Calakli and Taubin [CT11] builds a global signed-distance function that is zero at the input points and that has derivatives that correspond to the normals of the points. They also include a regularization based on the Hessian of the function. Fuhrmann and Goesele [FG14] define their implicit function as a sum of weighted basis functions. Each input sample generates a basis function and a weighting function with compact support so that only a small set of these functions needs to be evaluated at any position. The evaluation of the implicit function is performed inside an octree data structure which is then also used for the isosurface extraction.

Figure 24: Comparison of different surface reconstruction methods. *Left*: A global approach by [MKG11]. *Right*: An approach based on a local implicit function [FG14].

6.1. Global Combinatorial Algorithms

Graph-cut based approaches rely on a volumetric interpretation of 3D space. The latter is partitioned into cells based on a regular or adaptive grid structure. A graph-structure is used that links the individual cells that define the nodes of the graph with terminal nodes that represent the labels outside and inside and with the nodes in a local neighborhood following the applied grid structure. The surface to be reconstructed is assumed to pass between the inside and outside nodes.

In [WRO*12], different scalar fields are defined on the continuous volume $\mathbb{R}^3$ to represent surface consistency, outside votes and visibility information. Furthermore, a normal field is derived from the measurements. While in this approach a structured light system is used for the estimation of surface consistency, outside votes and visibility information and the Helmholtz reciprocity [ZBK02] is used for the estimation of a normal field $\mathbf{H}$ that assigns normal information to each point, different techniques could be used to estimate such data. The reconstruction of the object surface δV is formulated as a variational problem, where the object interior $V \subset \mathbb{R}^3$ is obtained by solving

$$\min_V \left\{ -\lambda_1 \underbrace{\int_{\delta V} \langle c\mathbf{H}, \mathbf{n} \rangle \, \mathrm{d}A}_{E_1} + \lambda_2 \underbrace{\int_V \hat{o} \, \mathrm{d}V}_{E_2} + \lambda_3 \underbrace{\int_{\delta V} (\alpha - \hat{c}) \, \mathrm{d}A}_{E_3} \right\}, \tag{37}$$

where λ_1, λ_2 and λ_3 are relative weights of the individual terms and $\alpha > 1$ denotes a constant determining the minimum regularization strength within a completely consistent region. The first term E_1 considers the flux of the vector field $c\mathbf{H}$ through the object surface. This term is minimized by a surface that is perpendicular to the reconstructed Helmholtz normals $\mathbf{H}$ and in regions with a high consistency c. The second term E_2 is used as an outside constraint to penalize regions of large values $\hat{o}$. This prevents the algorithm from short-cutting through concavities. The last term E_3 represents a regularization term and enforces a minimal surface. This penalty is

weighted with the consistency $\hat{c}$ obtained from the structured light. After calculating the estimates for the surface consistency and the normal field an iterative optimization procedure is carried out. After an initialization of the utilized octree at a coarse level, the grid is successively refined according to the local surface consistency estimates in the volume. In a subsequent iterative process, the memory efficient continuous min-cut [Str83, YBT10] is applied for a global optimization per iteration. In a final step, the resulting binary indicator function is smoothed inspired by the Smooth Signed Distance Reconstruction technique presented in [CT11].

A slightly different variational formulation has been used for the reconstruction of mirroring objects in [WORK13]. Instead of considering directly measured volumetric surface consistency information obtained from triangulation techniques such as structured light, only normal information is used as input in this approach. Individual normal fields obtained for the individual views are combined to one common normal field which contains information about the best local normal and the surface consistency. The underlying assumption is that the normal hypotheses from the different views agree with each other and with the true surface normal at points close to the true surface and deviate further away from the surface. Instead of modeling the probability density of normals under occlusions via a parametric model, which is challenging due to the dependency on the geometry of the considered object as well as on the placement of the involved cameras and screens, only the simplifying assumption is used that the density is highest for the actual surface normal. This assumption is valid as the actual surface normal is consistent over all views where the respective surface point has been observed, whereas outliers are not consistent over several views. Then, the estimation of surface normals can be performed by searching for the largest mode of the underlying probability density function for which mean-shift clustering [Che95] is used as a non-parametric technique because it neither requires assuming a model nor creates discretizetion artifacts. The resulting local normal estimate $\mathbf{N}(\mathbf{x}) = \arg\min_{\mathbf{n}} p_{\mathbf{x}}(\mathbf{n})$ corresponds to the centroid of the highest mode of the probability density function and the density at the centroid is used as a surface consistency measure $c(\mathbf{x}) = p_{\mathbf{x}}(\mathbf{N}(\mathbf{x}))$. As only normal information is used, the surface reconstruction is formulated as a variational energy minimization problem similar to [CLL07] according to

$$\min_{V} \left\{ -\lambda_1 \int_{\delta V} \langle c\mathbf{N}, \mathbf{n} \rangle \, \mathrm{d}A + \lambda_2 \int_{\delta V} \alpha \, \mathrm{d}A \right\}. \tag{38}$$

The parameters λ_1 and λ_2 represent weighting coefficients, c denotes a scalar field of surface consistency and α denotes a regularization parameter. Furthermore, the consistency-scaled vector field $c\mathbf{N}$ represents information about both the local probability of surface presence and the local normal information for the points in the volume. While the first term in the functional (38) is minimized for high consistency values and a surface which is perpendicular to the observed normals $\mathbf{n}$, the second term represents a regularization term which enforces a minimal surface area. The latter term avoids overfitting by increasing the cost for oscillating surfaces. Similar to [WRO*12], an iterative optimization procedure with an initialization of the utilized octree at a coarse level is used and the memory efficient continuous min-cut [Str83, YBT10] is applied for a global optimization per iteration. The resulting binary indicator function is smoothed in a similar way as in the approach in [WRO*12].

Multi-resolution Reconstruction Muecke et al. [MKG11] present another graph-cut based approach. To deal with multi-resolution input data their method uses additional information about the real world size of the sample point similar to the technique described in Section 6.2.3. Depending on this size they build a *global confidence map* inside the volumetric data structure and extract the surface using s-t cuts in a coarse to fine approach on different levels of the octree. See Figure 25 for an overview of the algorithm. This multi-resolution approach can create very high quality surfaces at varying scales. However, in the presence of outliers the top-down approach of first reconstructing on coarse octree levels and only later refining details if possible is not guaranteed to extract the best possible surface in all scenarios. This can be seen in Figure 24 where the technique is compared to the approach by Fuhrmann et al. [FG14].

6.2. Implicit Surface Reconstruction

The idea of an implicit surface representation relies on the fitting of an implicit function to the data which is typically given by values less than zero outside of the volume of the considered object and values greater than zero in its inside. The surface itself is given by the zero-set. Typically, the estimated implicit function is evaluated on regular voxel grids or adaptive data structures.

Among the most popular approaches of this category of implicit function fitting techniques are the Poisson Surface Reconstruction [KBH06], the Smooth Signed Distance Surface Reconstruction [CT11] and the Floating-Scale Surface Reconstruction [FG14].

6.2.1. Poisson Surface Reconstruction

The Poisson Surface Reconstruction [KBH06, KH13] is a global approach for the fitting of a 3D implicit function χ to an oriented point cloud by solving a Poisson equation, which is typically followed by the extraction of an appropriate isosurface based on the Marching Cubes algorithm [LC87].

In more detail, an oriented point set $\mathbf{p}_i \in \mathbf{V}$ of K points with attached normals $\mathbf{p}_i$ can be considered as a sparse sampling of a continuous vector field where the gradient of the indicator function

$$\chi_M(\mathbf{p}_i) = \begin{cases} 1, & \text{if } \mathbf{p}_i \in M \\ 0, & \text{otherwise} \end{cases} \tag{39}$$

should agree with the normals $\mathbf{n}_i$ of the input sample points $\mathbf{p}_i$, i.e. $\nabla\chi_M(\mathbf{p}_i) = \mathbf{n}_i$. Here, it has to be considered that a vector field can be gradient of a function if the necessary condition, that the vector field is conservative and, hence, its curl $\nabla \times \mathbf{V}$ is equal to zero, is fulfilled.

If the system is overdetermined, the normal equations are considered, i.e. the Poisson equation $\Delta\chi_M = \nabla \cdot \mathbf{V}$ is solved by applying the divergence operator to $\nabla\chi_M = \mathbf{V}$, which results in $\nabla \cdot \nabla\chi_M = \nabla \cdot \mathbf{V}$, since solving the Poisson equation amounts in finding an implicit function whose gradient approximates the given vector field best in the L_2 norm sense.

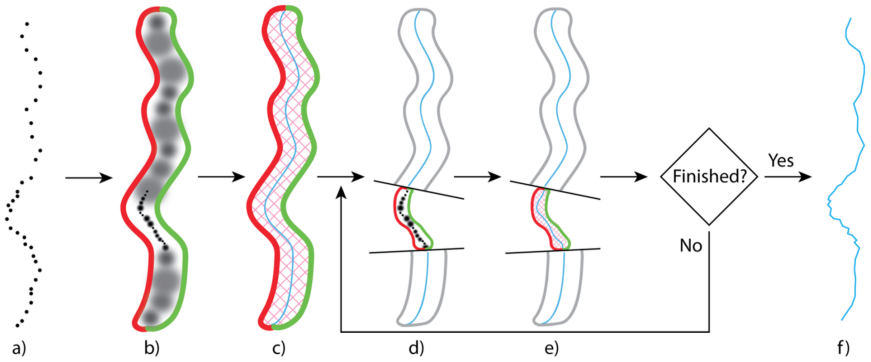

Figure 25: Illustration of the combinatorial algorithm by Muecke et al. [MKG11]. *From left to right*: The input samples (a) are inserted into an octree and a global confidence map is generated inside the occupied region, also called crust, and the boundaries are labeled inside and outside according to the normals of the samples (b). Next, a global surface is extracted with an s-t cut by cutting along high-confidence edges between voxels (c). In regions with high-resolution samples this process is repeated on a finer octree level to create a more detailed surface (d, e). After all fine regions have been reconstructed the final surface is created by stitching the meshes from different octree levels (f).

As the indicator function χ_M might be discontinuous and its gradient might not exist at some sample points, a smoothed indicator function $\widetilde{\chi}_M$ of the smoothed surface normal field is typically used [KBH06], i.e. a discrete point set is used to compute an approximation of the continuous vector field. Then, the relationship between the gradient of a smoothed indicator function and the surface normal field is given according to

$$\nabla(\chi_M * \widetilde{F})(\mathbf{q}) = \int_{\partial M} \widetilde{F}_{\mathbf{p}}(\mathbf{q}) \mathbf{n}_{\partial M}(\mathbf{p}) d\mathbf{p}, \tag{40}$$

where $\widetilde{F}(\mathbf{q})$ denotes a smoothing filter, $\widetilde{F}_{\mathbf{p}}(\mathbf{q})$ represents $\widetilde{F}_{\mathbf{p}}(\mathbf{q}) = F(\mathbf{q}-\mathbf{p})$, ∂M is the boundary of a solid M and $\mathbf{n}_{\partial M}(\mathbf{p})$ describes the inward surface normal at point $\mathbf{p} \in \partial M$.

The integral over the surface ∂M is approximated by a discrete summation of patches centered at the sample points $\mathbf{p}_i$, i.e. the integral over a patch $P_{\mathbf{p}_i}$ is approximated using the value at the coordinates $\mathbf{p}_i$ which is scaled by the area $|P_{\mathbf{p}_i}|$ of the patch. As a consequence, the computation of the vector field reduces to

$$\nabla(\chi_M * \widetilde{F})(\mathbf{q}) \approx \sum_{i=1}^{K} |P_{\mathbf{p}_i}| \widetilde{F}_{\mathbf{p}_i}(\mathbf{q})\, \mathbf{n}_i. \tag{41}$$

In order to allow a fast, memory efficient computation, multi-resolution approaches for representing the solution are used. As an accurate solution to the Poisson equation is only necessary in the vicinity of the surface ∂M, a discrete representation on adaptive grid structures such as octrees built based on the input point set S are typically used. Consequently, the function space consists of base functions F_o, chosen to be the n-th convolution of the box filter (which approximates the Gaussian as n is increased). Basis functions, associated with each leaf node $o \in O$ of the octree, are translated to the center of the node, afterwards scaled with respect to the size of the node and have a large support in the local neighborhood. Therefore, the vector field can be computed as the linear combination of these (smooth) basis functions which are associated with the octree cells

$$\mathbf{V}(\mathbf{q}) \equiv \sum_{i=1}^{K} \sum_{o \in \mathcal{N}_D(\mathbf{p}_i)} \alpha_{o,i} F_o(\mathbf{q}) \mathbf{n}_i \tag{42}$$

where $\mathcal{N}_D(\mathbf{p}_i)$ denotes the neighborhood with depth D at the sample at $\mathbf{p}_i$. Trilinear interpolation weights are represented by $\alpha_{o,i}$ and the interpolation across eight nearest neighbors allows for sub-node precision.

Then, the Poisson equation is solved in the function space defined above. Because the functions $\Delta\widetilde{\chi}$ and $\nabla \cdot \mathbf{V}$ are not necessarily in that function space, the problem is solved for $\widetilde{\chi}$ such that the projection of $\Delta\widetilde{\chi}$ to the function space is closest to the projection of $\nabla \cdot \mathbf{V}$

$$\langle \Delta\chi, F_o \rangle = \langle \nabla \cdot \mathbf{V}, F_o \rangle, \quad \forall o \in O \tag{43}$$

The resulting linear system is sparse and can be solved using adaptive multi-grid solvers. In case of non-uniform samples, the width of the smoothing kernel has to be adapted.

The Poisson Surface Reconstruction provides a global solution that is robust to noise, can handle a non-uniform sampling of the oriented input point cloud and is efficient regarding memory consumption and speed. Besides the original implementation [KBH06], several of its variants such as a streaming technique presented in [BKBH07] as well as the optimized, parallel or GPU-based implementations [ZGHG11, BKBH09] have led to a widespread use of this technique.

6.2.2. Smooth Signed Distance Surface Reconstruction

Similar to the approach of the Poisson Surface Reconstruction [KBH06], the approach in [CT11] is also based on an implicit surface representation. Again, an oriented point cloud consisting of K points $\mathbf{p}_i$ with attached surface normals $\mathbf{n}_i$ serves as input to the technique which is referred to as Smooth Signed Distance Surface Reconstruction. As solving for an implicit function that satisfies $f(\mathbf{p}_i) = 0$ and $\nabla f(\mathbf{p}_i) = \mathbf{n}_i$ for all the points $\mathbf{p}_i$ using an interpolatory scheme involves parameterized families of functions with

many degrees of freedom in the presence of measurement noise, an approximating scheme is followed where the aforementioned constraints are fulfilled in a least-squares sense. The considered energy minimization framework is based on the functional

$$\arg\min_f \lambda_1 E_1(f) + \lambda_2 E_2(f) + \lambda_3 E_3(f), \tag{44}$$

where $E_1(f) = \frac{1}{K}\sum_{i=1}^{K} f(\mathbf{p}_i)^2$ penalizes the distance of the reconstruction to coordinates of the input points, $E_2(f) = \frac{1}{K}\sum_{i=1}^{K} \|\nabla f(\mathbf{p}_i) - \mathbf{n}_i\|^2$ penalizes the deviations regarding the surface normals and $E_3(f) = \frac{1}{|V|}\int_V \|H(f(\mathbf{x}))\|^2 d\mathbf{x}$ represents a regularization term that enforces local smoothness of the gradients, i.e. a low curvature. The weights λ_1, λ_2 and λ_3 are non-negative weights for the individual terms, V denotes the volume and $H(f(\mathbf{x})) = \left[\frac{\partial \nabla f(\mathbf{x})}{\partial x_2} \frac{\partial \nabla f(\mathbf{x})}{\partial x_2} \frac{\partial \nabla f(\mathbf{x})}{\partial x_3}\right]$ denotes the Hessian matrix of the function $f(\mathbf{x})$. The normalization factor $\frac{1}{N}$ removes the dependency of the term on the number of data points. While the data terms $E_1(f)$ and $E_2(f)$ dominate the energy in the vicinity of the data samples and enforce the function to approximate the signed distance function, the regularization term $E_3(f)$ dominates away from the data samples and makes the gradients constant. In their analysis [CT11], the authors consider families of functions

$$f(\mathbf{x}) = \sum_{\alpha \in \Gamma} f_\alpha \Phi_\alpha(\mathbf{x}) \tag{45}$$

that are linearly parameterized by a finite set of parameters to facilitate the computation of a unique solution of the problem in terms of solving a linear system of equations. The parameter $\alpha = 1, \ldots, K$ represents an index belonging to a finite set Γ, $\Phi_\alpha(\mathbf{x})$ represents the basis functions and f_α denotes the corresponding coefficient. In their implementation, an octree is used to partition the oriented input point cloud recursively into subsets and the solution is first calculated for a coarser level and successively refined to the desired level using conjugate gradient solvers. Finally, a dual marching cubes algorithm is used to obtain a polygonal approximation of the surface.

6.2.3. Floating Scale Surface Reconstruction

The idea of Fuhrmann et al. [FG14] is to utilize scale information during surface reconstruction. As the input point cloud is usually created by a 3D scanning device or an image-based reconstruction algorithm such a scale value can be attached to each point. Scanning devices have a spatial resolution that can be used and for image-based algorithms one can easily define the size of the footprint of every pixel that has been reconstructed. This physically motivated scale changes depending on the distance of the surface to the scanning device. Distant surfaces are sampled at a lower resolution (higher scale value) and surfaces close to the acquisition device are sampled at a higher resolution (lower scale value). Depending on this scale value weights are assigned to the input points with the goal to always represent the highest possible resolution without blurring details in areas where parts of the input points represent surface samples at a high scale and other parts represent the same surface at a lower scale. The approach constructs an implicit function $F(\mathbf{x}) : \mathbb{R}^3 \rightarrow \mathbb{R}$ using a sum of basis functions f_i. Each of these basis functions is weighted by a function w_i which has a compact

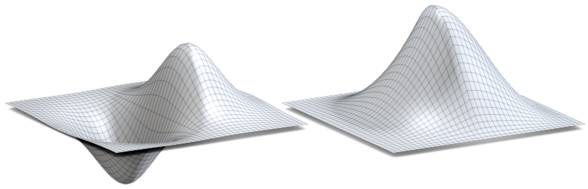

Figure 26: Illustration of FSSR basis functions in the 2D interval $[-3\sigma, 3\sigma]^2$. Left: $f(x,y) = f_x(x) f_y(y)$. Right: $w(x,y) = w_x(x) * w_y(y)$

support so only a small number of samples need to be evaluated in order to construct the final implicit function. F is defined as

$$F(x) = \frac{\sum_i c_i w_i(\mathbf{x}) f_i(\mathbf{x})}{\sum_i c_i w_i(\mathbf{x})} \tag{46}$$

where c_i is an optional confidence that can be attached to each input sample but can also be set to a uniform constant $c_i = 1$. The basis functions f_i are constructed so that every sample contributes the same volume to the implicit function and depending on the scale of the sample this volume is distributed in a certain area. In a local coordinate system of a sample, where the position of the sample is at the origin and the normal of the sample points in the direction of the positive x-axis, the basis function is defined as

$$f_x(x) = \frac{x}{\sigma^2} e^{\frac{-x^2}{2\sigma^2}} \; f_y(x) = f_z(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{\frac{-x^2}{2\sigma^2}} \tag{47}$$

$$f(\mathbf{x}) = f_x(x) f_y(y) f z(z) = \frac{x}{\sigma^4 2\pi} e^{\frac{1}{2\sigma^2}(x^2+y^2+z^2)} \tag{48}$$

where σ equals the scale value of the input sample. The influence of these basis function is practically zero beyond 3σ and the weighting functions are therefore defined as

$$w(\mathbf{x}_i) = w_x(x) * w_{yz}(\sqrt{y^2+z^2}) \tag{49}$$

$$w_x(x) = \begin{cases} \frac{1}{9}\frac{x^2}{\sigma^2} + \frac{2}{3}\frac{x}{\sigma} + 1 & x \in [-3\sigma, 0) \\ \frac{2}{27}\frac{x^3}{\sigma^3} - \frac{1}{3}\frac{x^2}{\sigma^2} + 1 & x \in [0, 3\sigma) \\ 0 & \text{otherwise} \end{cases} \tag{50}$$

$$w_{yz}(r) = \begin{cases} \frac{2}{27}\frac{x^3}{\sigma^3} - \frac{1}{3}\frac{x^2}{\sigma^2} + 1 & r < 3\sigma \\ 0 & \text{otherwise} \end{cases} \tag{51}$$

$$r = \sqrt{y^2+z^2} \tag{52}$$

Figure 26 shows 2D illustrations of the basis and weighting functions. The final function F is now sampled using an octree data structure. The octree is build by iteratively inserting all input samples starting from a root node around the first sample, expanding the octree if necessary, and subdividing an octree node if its side length is bigger than the scale value of the inserted sample. The implicit function is then evaluated at all voxels using *primal sampling*, i.e., the voxels are located at the corners of the octree nodes. After the implicit function is sampled inside the octree the final surface can be extracted as the zero-set of F using an appropriate marching cubes algorithm [KKDH07].

6.2.4. Global Multiscale Surface Reconstruction

The recent method by Ummenhofer et al. [UB15] also uses scale values in addition to the positions and normals of the input points. Similar to Fuhrmann et al. [FG14] they insert all points into an octree and treat each point as a singed distance function with a compact support window that is defined via weight functions w_i. Their technique, however, uses a dual sampling approach and they aggregate the values of the distance functions f_i and the orientations of the points $\mathbf{g}_i$ at the voxel centers. Depending on the voxel center $\mathbf{c}$ and the voxel size $s(\mathbf{c})$ they define for each input point (with position $\mathbf{p}$, normal $\mathbf{n}$, and scale σ)

$$w_i(\mathbf{c}) = \frac{4}{3}\pi\sigma_i^3 R_h(\mathbf{p}_i - \mathbf{c}) \tag{53}$$

$$f_i(\mathbf{c}) = \langle \mathbf{n}_i, \mathbf{c} - \mathbf{p}_i \rangle \tag{54}$$

$$\mathbf{g}_i(\mathbf{c}) = \mathbf{n}_i \tag{55}$$

R_h is the aggregation window for each voxel and defined using $h = 3s(\mathbf{c})$ as

$$R_h(\mathbf{r}) = \begin{cases} \frac{315}{64\pi h^9}(h^2 - \|r\|^2)^2 & \text{if} \|r\| \leq h \\ 0 & \text{otherwise} \end{cases} \tag{56}$$

The approach now solves for a global solution, defined as a signed distance function $u(\mathbf{x})$ and a normal vector field $\mathbf{v}(\mathbf{x})$, by minimizing an energy function

$$E(u,\mathbf{v}) = \lambda_1 E_{\text{data}_u} + \lambda_2 E_{\text{data}_v} + \alpha_1 E_{\text{coupling}} + \alpha_2 E_{\text{smooth}} \tag{57}$$

with

$$E_{\text{data}_u}(u) = \int \frac{1}{s}\sum_n w_n |u - f_n| d\mathbf{x} \tag{58}$$

$$E_{\text{data}_v}(\mathbf{v}) = \int \sum_m w_m \|\mathbf{v} - \mathbf{g}_m\| d\mathbf{x} \tag{59}$$

$$E_{\text{coupling}}(u,\mathbf{v}) = \int \|\nabla u - \mathbf{v}\|^2 d\mathbf{x} \tag{60}$$

$$E_{\text{smooth}}(\mathbf{v}) = \int s\|\mathbf{J}_\mathbf{v}\| d\mathbf{x} \tag{61}$$

where $\mathbf{J}_\mathbf{v}$ is the Jacobian of the vector field. To solve this function they use a finite volume discretization based on the dual sampling of the octree and iterative reweighted least squares. Once the final implicit function has been computed for the octree the surface is extracted using the dual contouring algorithm [JLSW02] which allows for further improvement of the vertex positions of the final mesh. Compared to other strategies this can produce much sharper corners. Compared to Fuhrmann et al. [FG14] the global optimization of the implicit function can deal with an increased amount of outliers in the input point set and still produce a smooth, consistent result.

7. Advances in Reflectance Acquisition

For an adequate acquisition of surface reflectance, the complexity of visual surface reflectance has to be considered in a similar way as in the context of geometry acquisition. The categorization of materials as discussed in Section 3.2 indicates that visual material appearance is characterized by different phenomena of light exchange with a particular object surface of interest which may also be analyzed when focusing on reflectance acquisition. In particular, diffuse and specular components as well as potentially occurring subsurface scattering or refraction characteristics have to be considered in the reflectance models, and the respective reflectance acquisition is typically designed according to the assumed underlying model. Many different models have been proposed in literature to model surface reflectance behavior, each focusing on accurately representing a certain subset of the possible materials.

In the following sections, we first discuss the different reflectance representations which is followed by a review regarding the individual appearance acquisition techniques. In particular, we group the techniques according to the reflectance representation that has been considered in the respective investigations and provide a particular detailed discussion regarding BTF acquisition techniques.

7.1. Reflectance Representations

Efficiently modeling surface reflectance behavior is also coupled with the use of an adequate model, which should have as few parameters as possible to still enable a faithful depiction of the material in a synthetic image with an acceptable acquisition time. Modeling e.g. the surface reflectance behavior of a diffuse object, where the incoming light is reflected uniformly into a hemisphere on the local surface patch, requires considering different material characteristics than modeling surface reflectance of mirrors which, in turn, is determined by an almost ideal direct reflection of the incoming light. Similarly, modeling reflectance behavior for materials with both diffuse and specular components or translucent and transparent materials requires considering the respectively relevant characteristics of the individual materials. In this regard, reflectance acquisition strongly depends on the representation used to model the reflectance of a particular material, as some parameters might not have to be measured. For e.g. a diffuse material, there is no need to capture the full 12-dimensional function $\rho(\mathbf{x}_i,\theta_i,\varphi_i,t_i,\lambda_i,\mathbf{x}_r,\theta_r,\varphi_r,t_r,\lambda_r)$ as only capturing two-dimensional textured bump maps $\rho_{\text{Bump Map}}(\mathbf{x})$ may be sufficient. Furthermore, measuring the four-dimensional bidirectional reflectance distribution function (BRDF) $\rho_{\text{BRDF}}(\theta_i,\varphi_i,\theta_r,\varphi_r)$ is typically sufficient for homogeneous materials without subsurface scattering. For instance, when considering a flat piece of specular metal, parametric BRDF models might be a good choice which means that the dependency of material appearance with respect to the incoming light direction and the outgoing light direction has to be considered during acquisition. In fact, a multitude of BRDF models has been presented in the literature for surfaces with different complexity with respect to the optical reflectance behavior. In contrast, the presence of spatially varying material characteristics might require the measurement of spatially varying BRDFs (SVBRDFs) $\rho_{\text{SVBRDF}}(\mathbf{x},\theta_i,\varphi_i,\theta_r,\varphi_r)$ or bidirectional texture functions (BTFs) $\rho_{\text{BTF}}(\mathbf{x},\theta_i,\varphi_i,\theta_r,\varphi_r)$. For brushed metal, which fulfills the requirements of the representation in terms of BRDFs because of the direct reflection of the incoming light, the typically strong specular reflectance component would require extremely dense measurements to capture the characteristics of the specular highlights. This makes data-driven reflectance acquisition rather impractical, and, hence, SVBRDFs based on analytical mod-

els are typically fitted. In contrast, when considering materials such as cloth samples that exhibit spatially varying mesoscopic effects such as subsurface scattering, interreflections and self-shadowing, parametric models are not sufficient to capture these effects. Instead, it is a better choice to use BTF models which are well-suited to represent such effects.

Depending on the parameters of the respective material model, the acquisition device has to be designed in a way that material appearance can be captured under the involved parameter configurations. The more parameters have to be considered during acquisition, the more data has to be acquired which increases the time required for the post-processing of the data as well. For materials following simple homogeneous BRDF models, only four parameters have to be measured, which might be done in only a few minutes. In contrast, materials with e.g. a spatially varying reflectance behavior require the measurement of six parameters in order to obtain an SVBRDF or BTF representation. In addition, the sampling density of the involved view-light configurations that need to be measured also needs to be taken into account. For dense measurements of several thousands of view-light configurations as used in e.g. [SWRK11, SSWK13], even fine details of mesoscopic reflectance can be adequately captured. However, even when acquiring a densely sampled BTF, the acquisition and processing times might vary significantly depending on the complexity of the involved materials, which influences the acquisition parameters such as the number of exposures used during the measurement. Example reconstructions of the optical reflectance properties obtained using different methods are shown in Figure 27. Obviously, the use of simple texture is not adequate to faithfully reproduce the appearance characteristics of the shown cultural heritage object. In contrast, the acquisition of the six-dimensional BTF enables a subsequent photo-realistic depiction of the object.

7.2. Texture Acquisition

After discussing the representation of textures on the surface geometry, we will briefly review strategies for the registration of the acquired data with respect to the surface geometry as well as strategies for color adjustment.

7.2.1. Representation

As discussed in Section 3.3, simple static textures or texture maps $\rho_{\text{Texture Map}}(\mathbf{x})$ describe the two-dimensional distribution of a certain type of information that might be given by anything we can imagine to be represented in terms of images. When considering reflectance acquisition and assuming that the multiple images have already been registered and the surface geometry is known as a result from techniques mentioned in Section 5, the reflectance behavior of an object surface is stored in a two-dimensional domain parameterized over the object surface. Therefore, only one single color value has to be acquired for a certain spatial position $\mathbf{x}$ on the object surface regardless of variations in viewpoint or illumination. Typically, the viewpoint and the illumination conditions remain almost constant during a controlled acquisition and often a rather orthogonal view on the object surface and diffuse illumination are used. However, when e.g. internet photo collections are used for the reconstruction of a particular scene or object, no such controlled setting is given and, consequently, varying viewing conditions, illumination conditions and camera conditions may occur.

In the simplest form, the measured data is stored in a 2D texture atlas as given e.g. by an image. Depending on the imaging sensor and acquisition parameters such as the number of measured spectral bands, such an image might contain radiance values in terms of a gray-scale representation or (possibly multi-spectral) color representation. Furthermore, the surface reflectance of the object of interest might be acquired under different exposure times to measure and store the high dynamic range of the surface reflectance.

7.2.2. Registration

In a typical acquisition process, images of a scene or object of interest are acquired under different view configurations and the acquired images have to be registered with respect to the surface geometry. In order to obtain a plausible impression, dense measurements have to be performed. Otherwise the low resolution does not allow a sufficiently accurate texture acquisition. As e.g. shown in [FFGG*10, SSS*08], where only the colors of the vertices in a mesh-based surface representation are used, low resolutions of the underlying mesh induce severe problems for texture acquisition. Furthermore, taking the mean of the color values observed for a particular surface position in the scene acquired from multiple views has been applied in e.g. [ASS*09] to obtain colorized point clouds and in e.g. [FFGG*10] to obtain colorized meshes, but this does not result in an adequate visual fidelity. Most of the approaches on surface texturing therefore focus on identifying certain views that should be selected for generating the texture map, which is typically followed by a color adjustment step to reduce possibly occurring seams between neighboring texture patches.

Several approaches focus on the selection of a single view to texture a particular face in the mesh-based geometry representation of the scene or object [LI07, VSJ07, GWO*10, GDDA13]. In [LI07], the problem of texturing the geometry of an object or scene based on registered images is approached by first projecting the individual images onto the surface geometry and a single view per face is selected based on maximizing the quality of the resulting mosaiced texture in a Markov random field energy optimization. While the used data term models the quality of the individual images acquired from different views, an additional smoothness term is used to penalize seams between the patches. The terms used in this energy functional have been extended in [APK08, GWO*10, WMG14]. In [APK08], the data term is specified by the size of the projection of a particular face into an individual registered image, which allows to incorporate image resolution, view angle and view proximity into the optimization. As image blur is not considered, the extension presented in [GWO*10] focuses on the use of gradient magnitudes integrated over the footprint on which a particular face has been projected. Larger sizes of the projected footprint of the face in the respective image as obtained for orthogonal views onto the surface, close-range views and high image resolutions as well as strong gradients due to being in-focus will increase the value of the data term. To reduce the influence of occluders that might be present in individual images, an additional photo-consistency check has been used in [WMG14]. This technique is based on a modified version of the mean-shift algorithm [CM02] and is less sensitive to outliers

Figure 27: Comparison of different reflectance reconstruction approaches: For the object depicted in the photograph on the left, the remaining three images show reconstructions obtained by using HDR-textures, PTMs and BTFs (from left to right). HDR-textures and PTMs are not capable of accurately reproducing the respective reflectance behavior of the considered objects. In contrast, the reconstruction based on BTFs preserves the reflectance properties of the object rather well. However, there are still some slight differences between BTF and the ground truth in regions with large contributions of indirect illumination [SWRK11].

than mean or medoid colors as applied in [SSS*08, GKKP07]. This allows the robust detection of views that are photo-consistent for a particular face.

In contrast, an alternative is given by the selection of several views and the adequate per-face blending of the respective information. In [Deb96], view-dependent texture mapping has been used to blend textures based on the closest neighbors in the set of acquired views to the view specified for the rendering. Additionally, a weighted averaging is used to enforce smoothness regarding the colors of neighboring patches. An alternative has been proposed in [Bau02], where the acquired images are decomposed into a high-frequency image and a low-frequency image. Subsequently, these are blended independently per frequency domain and finally combined to obtain the final texture. The blending in frequency space has also been used in [APK08, CZCW12]. Adequate weights may also be determined based on the view angle and the proximity to the model [GKKP07, CCCS08] and proximity to depth discontinuities has additionally been used in [CCCS08]. As stated in e.g. [WMG14], problems of such techniques are given by the fact that inaccuracies in the camera calibration or the reconstructed geometry might result in misaligned borders of the texture patches or ghosting which both cause visible seams. In addition, scale dependency has to be taken into account as an individual face might be projected onto a rather small footprint or an extremely large footprint depending on the camera pose with respect to the object. Otherwise, details in images acquired under close-range viewpoints are likely to get blurred due to the influence of information of images captured under distant viewpoints. Furthermore, a perfect alignment of the images with respect to the 3D geometry is still challenging. For this purpose, a correction of this local misalignment has been proposed in [DMC*12], where optical flow is used to compute the local displacements instead of a global mapping.

7.2.3. Color Adjustment

Many of the currently available laser scanning devices are equipped with an additional camera so that colored point clouds can be acquired directly, i.e. for each of the reconstructed 3D points there is additional RGB information available as observed under a single viewpoint. In a multi-view scenario such as given for multi-view stereo techniques, the determination of a single still plausible color value per surface point becomes significantly more challenging. The reason for this is given in the nature of such methods that typically try to get an as-complete-as-possible reconstruction of the observed scene content based on multiple images, where parts of the scene are not only observed in a single image but in several images acquired under different viewing conditions where the illumination conditions might be different as well. The color values measured for a particular surface point under different view-light configurations might be rather different due to a more complex reflectance behavior and also due to possibly varying camera characteristics or illumination conditions. Hence, there is a need to determine one single color value for each surface position from all the available observations in order to derive a texture map for the surfaces of objects within the scene. In this process, care has to be taken so that an as-plausible-as-possible impression of the scene content can be preserved based on this rather simple representation.

In order to compensate for color discontinuities between neighboring texture patches arising from different characteristics of the involved cameras (e.g. response curves, exposure times, etc.) or different illumination conditions in the individual images, an additional color adjustment step is typically applied. Local photometric adjustments have been proposed in [VSJ07], where the color at the discontinuity is determined by the mean of the patches and a smooth transition to this value is reached by heat diffusion. Global photometric adjustments have been proposed in e.g. [LI07] based on luminance correction terms that are added to the vertex luminances so that there are only small luminance differences at the

seams and the derivatives of adjustments are small within a patch. Instead of only using color values at the projections of vertices of the object geometry into the individual images as used in [LI07], a more sophisticated approach has been presented in [WMG14] by also using color values along all adjacent seam edges.

7.3. BRDF Models and Their Acquisition

The texture-based representation of surface reflectance described in the previous section is typically not sufficient to preserve a realistic depiction of surface reflectance behavior of materials. The reason for this is the strong dependency of material appearance on the viewing conditions and illumination conditions, which has already been subject of several studies [Pho75, TS67, NRH*77]. In particular, the probably first formal description of the dependency of material appearance on view direction and light direction has been introduced in [NRH*77] with the bidirectional reflectance distribution function (BRDF) $\rho_{\text{BRDF}}(\theta_i, \varphi_i, \theta_r, \varphi_r)$. As already discussed in [NRH*77] and in Section 3.3, the BRDF is only a rather special case of the more general eight-dimensional BSSRDF $\rho_{\text{BSSRDF}}(\mathbf{x}_i, \theta_i, \varphi_i, \mathbf{x}_r, \theta_r, \varphi_r)$ and is tailored to flat and opaque materials/homogeneous opaque materials [DF97] as a result of the reduced parameter set. While the full BRDF has to be measured for materials with anisotropic reflectance behavior, only a subset of these configurations, namely only the azimuthal angle between the view direction and the light direction, has to be measured for isotropic materials. Additionally, if multi-spectral BRDFs $\rho_{\text{BRDF, multi-spectral}}(\theta_i, \varphi_i, \theta_r, \varphi_r, \lambda)$ have to be measured, there is an additional dependency on the wavelength λ. BRDF setups therefore need to be designed in a way that allows a systematic measurement of the dependency of surface reflectance on all its relevant parameters and several concepts have been proposed in literature. In the following sections, we first discuss the representation of BRDFs and, then, group the BRDF acquisition techniques according to the underlying measurement principle into gonioreflectometer-like setups, mirror-based setups and acquisition approaches that exploit geometric configurations by the use of curved objects.

7.3.1. Representation

Once the data has been acquired, either parametric BRDF models or non-parametric BRDF models are fitted to the measured data. The fitting of parametric BRDF models that are given by analytical functions can be achieved based on models that are efficient to compute such as the Blinn-Phong model [Bli77], the Lafortune model [LFTG97], the Ashikhmin model [AS00] or the directional statistics BRDF model [Nis09], based on micro-facet models such as the Ward model [War92], the Cook-Torrance model [CT82], the Schlick model [Sch94] or based on wave-optics-based BRDF models [HTSG91]. The result of the fitting procedure is given by estimates for the parameters of the chosen BRDF model. In contrast, non-parametric BRDF models typically rely on representing the surface reflectance behavior by linear combinations of certain basis functions such as spherical harmonics [WAT92, RH01], Zernike polynomials [KvD98], spherical wavelets [SS95, LF97, MPBM03b], radial basis functions [ZREB06, WWHL07], rational functions [PSCS*12] or measured BRDFs [MPBM03b, WWHL07, RWS*11, ASOS13]. However, a tabulated representation might be used as well [MPBM03b, LBAD*06]. This allows such non-parametric approaches to represent a more general surface reflectance behavior and also offers a higher accuracy in comparison to parametric models. In the scope of these course notes, we will focus on the review of the different approaches that have been proposed to acquire the respective data. Details regarding the individual BRDF models can be found in literature such as the survey in [HF13].

7.3.2. Gonioreflectometer-like Setups

Most of the BRDF acquisition approaches rely on gonioreflectometer-like setups. The underlying measurement principle is based on the possibility to independently vary the viewing conditions and the illumination conditions with respect to the surface of the material sample to be acquired as illustrated in Figure 28. This can be achieved by changing the pose of the material sample relatively to the involved detector and light source or by changing the position of the detector, that measures the radiance coming from the material sample, and the position of the light source, that illuminates the material sample, relatively to the pose of the material sample. Typically, such approaches offer the possibility to sample almost arbitrary view-light configurations. In addition, point light sources are applied to allow for a local illumination on the surface so that no superposition of the reflected radiance and scattering effects is acquired but only the reflected radiance.

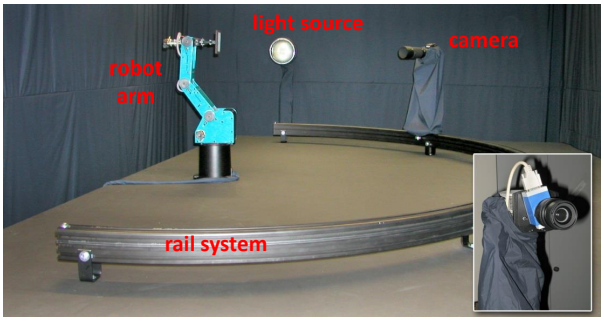

Figure 28: Gonioreflectometer setup used in [SSK03]: A static light source is used to illuminate the sample that can be rotated based on a robot arm. Furthermore, the camera that observes the material sample can be moved based on a rail system.

In [MCS90], a gonioreflectometer setup has been used which allows a mechanical manipulation of the object pose, the detector position and the position of the used fiber-optics light source. The object is placed on a turntable and can be rotated around one axis, there is one degree of freedom for setting the light source position, and the detector can be moved. This allows the setup to measure surface reflectance and transmittance based on four degrees of freedom. Unfortunately, the measurement is time-consuming as a result of the sequential acquisition of all the configurations that are needed. In [BM97], the use of a gonioreflectometer setup has been proposed to capture multispectral BRDFs. The light coming from a quartz halogen lamp with a spectral filter under a maximum elevation angle of 60° is reflected at the surface of the material sample

and observed by a photodiode detector with a spectral resolution of 10nm under elevation angles up to 85°.

The setup in [MWBS96] relies on varying only the azimuthal angle between the light direction and the view direction and therefore allows to measure only isotropic materials. Similarly, the setup in [Foo97] also only allows the acquisition of isotropic BRDFs. While the pose of the material sample can be manipulated via a rotation and lifting stage that allows rotating the object and changing its vertical position, the light source position can be rotated around the material sample with one degree of freedom. Despite the four realized degrees of freedom, only the reflectance of isotropic materials can be acquired as the mounting of components does not allow to measure independent measurements of view-light configurations.

Furthermore, the setup in [WSB*98] involves two rotation arms that are equipped with a photodiode serving as detector and a laser diode serving as light source respectively. These arms can be remotely controlled and the material sample is additionally placed on a turntable. While this setup allowed an improvement of the measurement accuracy at the time of its publication, occlusions of the arms occurring near the retro-reflection represent limitations of this setup. The trend of using controllable arms has been further investigated in [Ke99], where a robotized arm is involved that allows a manipulation of the camera with five degrees of freedom. Additionally, the material can be rotated around one axis and translated along two directions in this setup. Even if the number of degrees of freedom in this setup seems rather large, the use of a single fixed light source position limits the number of measurements that can be performed to only a partial BRDF measurement.

Unfortunately, performing successive measurements of all the configurations of view direction and light direction relative to the surface of the material sample with a gonioreflectometer setup is a time-consuming process although the aforementioned automatic positioning systems already greatly improve the efficiency in comparison to manual adjustments of these parameters. As a consequence, several investigations have focused on speeding-up the measurement process. One possibility to achieve this is given by an as-sparse-as-possible acquisition. In [LSL03], a view planning has been proposed based on an iterative reduction of the parameters in the fitted BRDF model to reach a significantly faster acquisition. In contrast, a different approach has been followed in [SSR*06], where parallel measurements have been performed by using four cameras simultaneously. These are mounted on fixed arms and the material sample can only be tilted in the setup while being illuminated by a halogen lamp that is mounted on an arm that only allows a planar variation of the light source position. While the limitations of this acquisition process only allow a sparse measurement of the full four-dimensional BRDF, the missing data is estimated by interpolation using a fitted parametric BRDF model. This allows to perform a fast BRDF measurement. More recently, a denser and also multi-spectral acquisition of isotropic BRDFs has been proposed in [ZWWW10], where the corresponding setup is based on a rotation stage that allows to rotate the sample around three axes. The sample is illuminated by a halogen lamp and the radiance coming from the material sample is measured with a fiber-optic spectrometer.

Another gonioreflectometer-like setup has been proposed in [LFD*08]. The material sample can be panned and tilted with the sampleholder and is illuminated by a static xenon lamp with special optics and filters. Furthermore, an involved bench with an attached photodetector can be tilted relatively to the sampleholder on the turntable. This setup not only allows the measurement of anisotropic BRDFs or BTDFs, but also the measurement of the more general BSDF. Furthermore, HDR measurements can be performed to better capture the dynamic range of the incoming radiance.

Furthermore, multi-spectral HDR BRDF/BTF acquisition has been approached in [HGH06] with a highly accurate setup, where the material sample is held by a robot arm that allows pose variations along five degrees of freedom. The sample is illuminated by a spatially uniform light source that can be rotated around the arm based on a rotation stage. Several detectors – a photodiode with the spectral range of 250nm to 1700nm, a line-scan imaging sensor with a spectrometer (360nm to 830nm) and a HDR luminance camera with a spectral resolution of 28μm [HAH*12] – are used in the scope of this setup to allow a robust acquisition of the appearance characteristics of the material sample. Highly accurate angular measurements with an accuracy of 0.002° have been reported.

A setup design for the fast BRDF acquisition based on rather specialized hardware has been proposed in [NZV*11]. The material sample is illuminated by a laser-projector and indirect reflections on a Lambertian material can be observed by the detector. As a result, a multitude of view-light configurations of the BRDF can be measured simultaneously.

7.3.3. Mirror-Based BRDF Acquisition

A significantly different approach in comparison to gonioreflectometer-like setups has been introduced with mirror-based acquisition setups. The key idea of such setups is the reduction of degrees of freedom that need to be accounted for by the mechanical adjustments involved in gonioreflectometer-like setups by performing several measurements in parallel. This can be achieved by an adequate arrangement of mirrors in the acquisition setup so that the sample surface is not only directly observed from a single view-light configuration but from a multitude of view-light configurations at the same time. Similar to gonioreflectometer setups, mirror-based setups allow to acquire the BRDF directly without the superposition with additional scattering characteristics. One of the first of such mirror-based setups has been proposed in [War92]. The setup consists of a half-silvered hemispherical mirror under which a material sample is placed and captured by a fish-eye camera. This procedure requires two degrees of freedom less to be measured based on mechanical manipulations. The material sample can be rotated with one degree of freedom and a light source is mounted on a pivoting arm that allows a movement along a one-dimensional path. The light beam emitted by the light source is collimated. Furthermore, care has been taken that no direct illumination is received by the involved camera. This even allows the measurement under retro-reflection directions, which e.g. cannot easily be measured by gonioreflectometers due to occlusions of camera and light source. Moreover, the radiance emitted by the material sample can be simultaneously measured

from a multitude of viewing directions. This results in rather short acquisition times of only a few minutes. However, the acquisition relies on a manual positioning of both the material sample and the light source. Further drawbacks can be recognized in the inaccurate measurements near grazing angles and the distortion of the fish-eye lens.

In the setup used in [Dan01], the material sample to be acquired is placed in the focal point of a parabolic mirror. Directional illumination is generated using fiber-optics, and the generated light ray is controlled with a specific aperture and directed to illuminate the mirror at a certain position, where the light is reflected onto the material sample. The incoming light is reflected at the surface of the material sample towards the parabolic mirror where it is re-directed and finally observed by a camera. The use of a beam splitter allows the realization of coincident view directions and light directions and, therefore, avoids occlusion problems regarding camera or light source that might e.g. occur in gonioreflectometer setups for some configurations. This principle again allows the simultaneous measurement of a multitude of configurations with different view directions under a single illumination condition and, hence, a speed-up of the BRDF acquisition process. Unfortunately, only a rather limited range of inclination angles between 23° and 37° for both view and light direction can be acquired due to the constraints given by the setup.

An approach similar to the setup in [Dan01] has been proposed in [MSY07], where the material sample is also placed in the focal point of an, in this case, ellipsoidal mirror. The main difference of this approach is given by the use of multiplexed illumination. For this purpose, a projector is used to project Hadamard patterns via the mirror onto the surface of the object where they are observed with a camera based on the reflections in the mirror. As in [Dan01], a beam splitter allows the coincidence of the optical axis of the camera and the light direction and, as a result, the projected patterns on the surface of the material sample can be observed simultaneously from a multitude of viewing directions. However, while the illumination multiplexing also allows to measure several illumination configurations at the same time by exploiting the linearity of light superposition, typical problems of such multiplexing techniques are given by a lower signal-to-noise ratio, i.e. the noise component is usually clearly visible and the dynamic range of the acquired data is rather low.

The use of such more advanced illumination techniques has also been investigated in [GAHO07], where a sequence of basis functions is used to illuminate the respective material sample from several directions. As in [Dan01] and in [MSY07], a beam splitter is used in the setup to realize a coincidence of the optical axis of the involved camera and the light direction before the light enters the more sophisticated mirror construction, where it is first reflected at a parabolic mirror and subsequently re-directed towards the material sample based on a dome-like mirror. As a result of the constructed setup, the parabolic mirror occludes the view and illumination from directions close to the surface normal and, hence, an extrapolation based on spherical harmonics, that are fitted to the measured values, has to be used instead.

7.3.4. Exploiting Geometric Configurations for BRDF Acquisition

Another strategy to perform a simultaneous measurement of multiple view-light configurations is based on using curved material samples that are observed in the image taken by a camera, where the geometry of the material sample is typically known a-priori or separately measured. As typically several neighboring surface points are measured simultaneously, image-based techniques do not allow an actual BRDF acquisition but only the acquisition of the superposition of both the BRDF and a scattering induced by the effects at points in the local neighborhood.

In the setup used in [KMG96], the light source is placed with a close range with respect to the material sample so that many configurations with varying light directions can be observed in images taken by a static camera that is also placed in the near-field. The light source can additionally be moved which allows a dense sampling of the view-light configurations in the upper hemisphere with respect to the surface normal of the material sample.

In [LL95], an image sequence is acquired in a setup where the optical axis of the involved camera coincides with the ray of the incoming light emitted by a light source. Based on the acquired image sequence, an approximation of the surface of the curved material sample is estimated together with an isotropic BRDF. Anisotropic surface reflectance cannot be acquired because of the collinearity of view direction and light direction.

In order to acquire the full anisotropic BRDF of velvet, a cylindrical geometry with four attached velvet samples has been used in [LKK98] which has been observed by a fixed CCD camera under illumination by a movable light source. In particular, the attached velvet samples have been cut from the same material sample and attached to the cylinder so that each of the stripes has a different orientation of the fibers. The cylinder geometry allows the simultaneous measurement of a multitude of view-light configurations. Additional scattering characteristics are observed in a second setup that is based on illuminating the velvet sample from the back using a halogen light source emitting white light and illuminating the fiber tips with green laser light. A microscopic objective and a focusing attachment are used in front of the camera to allow the measurement of the light scattered from the velvet fibers. This approach requires the material samples to be bendable so that they can be attached to the cylindrical shape.

Cylindrical material samples have also been used in [MWLT00] as well as spherical material samples to exploit the possibility of a parallel acquisition of several view-light configurations. The setup is based on one camera that observes the material sample placed on a rotation stage while a second camera is used to track the position of a moving light source. Based on this setup, isotropic BRDFs have been acquired. In [MWL*99], this setup has been extended by a laser scanner so that differently shaped material samples can be used instead of relying on an a-priori known spherical or cylindrical geometry. Measuring the object geometry has also been followed in [SWI97] and, based on the reconstructed geometry, the diffuse and specular components of the surface reflectance behavior are estimated which, in turn, serve for the fitting of a BRDF model.

In [MPBM03a], isotropic BRDFs of spherical material samples

are measured in a gonioreflectometer-like setup. The known spherical geometry of the material sample allows to reduce the number of degrees of freedom that have to be realized by mechanical movements of the involved components from four to two, as the radiance coming from the spherical sample towards the camera can be measured in different view directions within a single image. Furthermore, HDR-based BRDF acquisition approaches that also rely on spherical material samples have been presented in [KKP*08] and [GCG*05].

The approach presented in [ND06] is based on cutting the material sample into a set of 20 smaller pieces that are placed on a known geometric shape, i.e. a cylinder, in different orientations, where they are observed by a static camera while the light source can be moved around the material sample. This allows to capture anisotropic BRDFs and, hence, only bendable material samples can be measured. In order to obtain an acceptable azimuthal resolution in the measurements, many pieces of the sample need to be distributed on the known geometry in different orientations.

7.3.5. Arrays of Detectors or Light Sources

A dome-like arrangement of uniformly distributed LEDs has been introduced in [BEWW*08] for the acquisition of multi-spectral BRDFs. The LEDs are equipped with telecentric lenses and are used both as light source and detector which circumvents problems regarding occlusions of detector and light source as occurring in gonioreflectometer-like setups. However, the use of 65 LEDs at fixed positions only allows a rather coarse sampling of view-light configurations to be measured so that details in the surface reflectance behavior, such as highlights, might not be adequately captured.

Sparse BRDF measurements have been performed in [DWT*10] in a setup with six LED light sources that are used to illuminate the material sample. The material sample is simultaneously observed from a multitude of different view directions with a maximum elevation angle of 48° which is achieved by using two specific condenser lenses. This setup has been realized within a rather small and easily portable device. Based on sparse measurements, a micro-facet BRDF model is fitted. However, one of the LEDs is placed between the condenser lenses and therefore occludes some of the view directions. Such occlusions are avoided in the extension of this setup presented in [LDW*10], where the optical axis is not perpendicular to the surface of the material sample as in [DWT*10] but rotated by 20°. The LEDs are placed to sample the hemisphere where there are no sensing lenses. The missing data is obtained by exploiting BRDF reciprocity. This means that only a BRDF slice is measured, where again a micro-facet BRDF model is fitted similar to the approach in [DWT*10].

7.4. SVBRDFs and Their Acquisition

After a discussion of the SVBRDF representation, we discuss the individual methods that can be categorized into the strategies for SVBRDF acquisition that are either based on a sparse sampling of the measured view-light configurations or based on a dense sampling.

7.4.1. Representation

As already discussed in Section 3.3, spatially varying bidirectional reflectance distribution functions (SVBRDFs) $\rho_{\text{SVBRDF}}(\mathbf{x}, \theta_i, \varphi_i, \theta_r, \varphi_r)$ model the dependency of surface reflectance behavior on the light direction (θ_i, φ_i), the view direction (θ_r, φ_r) and spatial positions $\mathbf{x}$ of the parameterized surface. The key difference when acquiring SVBRDFs instead of BRDFs is given by the spatial variation in the surface reflectance behavior that is modeled as a spatial distribution of mutually independent BRDFs over the surface of the respective material sample. As these BRDFs can be represented based on analytical functions, it may be sufficient to fit the parameters of the SVBRDF model based on only a few measurements. In contrast to such sparse sampling strategies, several acquisition techniques still use a dense sampling of view-light configurations to allow a more robust fitting of the model parameters.

7.4.2. Sparse Sampling Strategies

An image-based approach has been presented in [LKG*01], where a camera observes the spatially varying appearance of an object from different viewpoints under illumination by a movable light source. The position of the light source is estimated based on highlight observations on accurately aligned mirroring calibration spheres and the surface geometry of the object is acquired using a laser scanner. First, reflectance values are gathered for the particular surface points based on their back-projection into the individual images acquired under different light source positions and taking the color values at the intersection. Then, analytical Lafortune BRDF models [LFTG97] are fitted to these radiance values acquired for different surface points using Levenberg-Marquardt optimization [Lev44] due to the nonlinearity of the Lafortune BRDF model. Additionally, the covariance matrix obtained in the optimization reveals information regarding the direction with the maximum variation in the observations in the respective parameter space. In a subsequent clustering step, clusters are iteratively split along the principal axis of the variance if their intra-cluster variance is too large. Finally, the spatial variations in appearance are modeled per texel as a weighted sum of these clustered Lafortune models. This approach has been extended in [LKG*03] by a per-texel normal refinement using photometric stereo. As the view configurations and light source configurations are set in a manual procedure, a further extension of this approach has been introduced in [LLSS03] with a planning strategy for iteratively specifying the view-light configurations to be measured by minimizing the uncertainty of the involved parameters based on the already acquired data. Such adaptive measurement approaches allow a systematic acquisition of the data that is needed most urgently to perform an as-sparse-as-possible and still accurate acquisition.

A quite similar strategy has been followed in [McA02] with a gonioreflectometer-based setup where the sampleholder allows a 3D rotation of flat material samples. The latter are observed by a static camera under illumination by a light source that can be rotated around the sample. This allows to acquire radiance values per surface point under different illuminations based on a homography-based registration of the acquired images to the texture map space on the flat material sample. Based on these acquired values, a

BRDF is obtained per surface point based on fitting Lafortune models similar to the approach in [LKG*01, LKG*03].

The principle of the approach presented in [GCHS05] is similar to the techniques presented in [McA02, LKG*01, LKG*03] and differs mainly in the BRDF model used in the fitting step, where the Ward model [War92] is used. The surface reflectance has been represented by a linear combination of two to three basis BRDFs with unknown coefficients. Furthermore, the local surface normals are also estimated during the optimization in addition to the local BRDF parameters and the local blending weights. Similarly, in [RK09] both a heightfield and a SVBRDF have been estimated using a combination of multi-view stereo and photometric stereo has been used to directly reconstruct surface heightfields and a SVBRDF is recovered as well based on a non-linear optimization. While this approach also takes interreflections into account, it is limited to planar samples and not applicable to objects with a complex surface geometry.

Based on only ten images depicting a flat material sample acquired with a static camera under different light directions, an approximate SVBRDF has been acquired in [KSS*04]. While the reported manual acquisition procedure can be performed within a few minutes, this rather coarse sampling does not allow to adequately capture anisotropy in surface reflectance behavior or view-dependent effects in material appearance.

A further strategy for the fast acquisition of isotropic SVBRDFs based on an array of cameras and light sources has been followed in [WMP*06]. In particular, the used light stage setup consists of 150 sequentially triggered light sources and 16 cameras that simultaneously capture the received radiance coming from the object. Additionally, the surface geometry is acquired with a projector and four additional cameras. An acquisition time of 25s has been reported for the sparse acquisition of an isotropic SVBRDF similar to the approach in [LKG*01] but without the use of the respective clustering approach applied in [LKG*01] before the fitting step.

The setup presented in [DWT*10] offers not only the possibility to acquire BRDFs as described in Section 7.3.4 but also allows the acquisition of SVBRDFs. This setup is based on the use of an ocular condenser lens and a field condenser lens through which the camera observes the material sample simultaneously under several view-light configurations and under illumination by six light sources at different positions. The extension towards sparse SVBRDF acquisition relies on measuring surface texture reflectance maps, into which the spatially measured BRDFs are embedded. While the resolution of these acquired SVBRDFs is rather high, both isotropic specularity and anisotropic specularity are captured. However, surface reflectance can only be measured for a rather small range of viewing angles up to an elevation angle of 48°. Hence, the accuracy of the fitted BRDF might be rather limited for larger viewing angles.

In [GCP*09, GCP*10], a polarization-based approach has been proposed to acquire sparse SVBRDFs from only a few images. The setups that are discussed include a setup with a single camera with an attached filter wheel observing a material sample illuminated by second-order gradient illumination generated via a spherical distribution of light sources with attached linear polarizers. In particular, both cross and parallel polarization are considered to allow the separation of diffuse and specular reflections. This allows the illumination of the material sample from the full sphere of incident light directions and is not constrained regarding the surface geometries that can be handled. However, only up to rough specular materials can be handled due to the limited sampling density. Moreover, a setup where flat material samples are illuminated by an LCD monitor with gradient illumination from a close distance and observed by a camera is presented. The LCD monitor is precisely calibrated so that the particular directions of the incoming light, determined by considering the individual pixels of the monitor as light sources, relatively to the material surface are known. Again, a separation of diffuse and specular reflectance can be obtained based on the linear polarized light emitted from the LCD monitor and the radiance observed from the material sample with the camera under different orientations of the polarization. Furthermore, a setup has been used, where the object is placed closed to the center of a rough specular hemisphere. The hemisphere is coated with a specific rough specular paint and the radiance coming from the object to the camera is observed through a hole in the apex of the hemisphere. The illumination is generated based on a projector with a fish-eye lens that illuminates the hemisphere from where the light is reflected onto the material sample. A precise calibration of the setup allows to assign the light directions with the respective projector pixels they come from.

The reflectance sharing approach presented in [ZREB06] focuses on the acquisition of SVBRDFs based on a rather sparse set of images. The object geometry is assumed to be known from a prior processing step and the received radiance values at pixels, where the material is observed in the image, are carefully assigned to individual spatially distributed BRDFs based on a precise calibration of the setup. Then a decomposition of diffuse and specular components is performed based on polarization. Polarization filters are placed in front of the camera and the light source, and the specular component can be obtained according to $I_{\text{specular}} = I_{\text{aligned}} - I_{90^\circ}$. Here, I_{aligned} is captured with aligned filters and represents the sum of diffuse and specular components. In contrast, I_{90° is captured under a rotation of 90° of the filter in front of the camera so that only the diffuse component can be observed. Per image, a single two-dimensional slice of the isotropic SVBRDF is densely sampled. After a scattered-data-interpolation of these samples in both the spatial and the angular domain, an SVBRDF is obtained. As discussed in [WLL*07], an accurate SVBRDF can typically be obtained despite this sampling as densely sampled SVBRDF regions capture rapid angular variations rather well.

The investigations in [HFB*09] focus on the acquisition of SVBRDFs from collections of images taken under uncontrolled conditions. A multi-view stereo reconstruction has been performed to estimate the scene geometry in an initial step. In a further step, both the light conditions and the spatially varying reflectance behavior have been alternatively updated.

Less controlled capture conditions have also been considered in [PCDS12], where a hand-held video camera has been moved around the object under fixed general illumination. The object geometry is assumed to be known from a previous geometry acquisition procedure. The video data has first been aligned with the object

geometry and, then, both illumination conditions and the spatially varying reflectance have been computed. In this method, the reflectance characteristics are represented based on a spatially varying Phong model with a diffuse component and an additional specular lobe.

Aiming for a rather light-weight acquisition setup that allows a rather fast, sparse acquisition of near-planar material samples using inexpensive hardware, the approach presented in [AWL13] makes use of a rigid setup of a display used for illumination and a camera. SVBRDFs have been acquired based on the acquisition of only a few photographs under certain Fourier basis patterns emitted by the display. In addition, the reflectance function has been modeled as a Mixture of Gaussians.

The approach presented in [RPG14] focused on SVBRDF acquisition of planar material samples based on mobile hardware. Materials with rough specular BRDFs can be acquired using the camera and flash light available in standard mobile phones. For highly specular materials, it is required to use an extended light source and, hence, the combination of a camera and a display of a tablet device has been used, where the polarization of the display has also been used to separate the diffuse component and the specular component.

A further light-weight acquisition approach has been introduced in [AWL15]. This technique exploits self-similarities, i.e. that several points on the object surface have similar reflectance properties, given in the reflectance behavior of many materials. Flat material samples have been captured from only two images taken with a mobile phone with and without a headlight flash. Based on an initial gathering of reflectance observations, a regularized texture statistics transfer has been performed and SVBRDFs have been fitted to the respective data.

7.4.3. Dense Sampling Strategies

The setup used in [Dan01] for BRDF acquisition has been extended in the scope of the investigations in [DW04] to allow the acquisition of anisotropic SVBRDFs. The material sample is translated so that different surface positions can be moved to the focal point of the involved parabolic mirror. Besides this translation stage, only the aperture used to direct the light ray onto the material surface via the mirror can be moved in this setup. The local BRDFs are measured for the individual surface points rather fast as a multitude of view-light configurations can be measured simultaneously in this setup as already described in Section 7.3.3. However, there are the same limitations of a rather restricted inclination angle for both view and light direction as in [Dan01]. In addition, the limited height of the material sample and the rather long acquisition times for dense spatial measurements are further problems of this approach.

Based on an accurate gonioreflectometer setup with four degrees of freedom, the investigations in [MWAM05] focus on the acquisition of SVBRDFs for lacquered wood. The pivoting arms are equipped with a camera and a light source respectively, and the material samples are placed on a sampleholder on a rotation stage. The images are rectified to obtain a stack of aligned images and a reflectance model tailored to wood is fitted per pixel. This setup has also been used to acquire SVBRDFs for flat material samples in [LBAD*06] and might furthermore be used to acquire BTFs.

In [HLZ10], the acquisition of SVBRDFs has been approached based on a gonioreflectometer-like setup. A light source and a camera with a beam-splitter have been mounted to two moveable robot arms. The object can thus be acquired under a coaxial arrangement of view direction and light direction. This allows the acquisition of retro-reflections and reciprocal image pairs. An additional geometry reconstruction has been performed based on the use of a fringe projector. However, rather long acquisition times have been reported. A representation in terms of a set of basis BRDFs has been used to locally fit the respective BRDF parameters.

In order to overcome the problems of data-driven techniques regarding irregular angular samplings of the reflectance function with a limited angular resolution, a compact data-driven reflectance representation has been introduced in [RSK12]. The technique is based on directly fitting a sum of separable functions to the irregularly sampled data and allows an accurate reproduction of the surface reflectance behavior in real-time. Its data-driven nature of this approach allows to replace the constraints given by analytical models with a more flexible regularization and to avoid the selection of a suitable model.

The investigations in [KNRS13,NJRS13,NKRS15] are based on the "OrCam" acquisition device that allows both a reconstruction of the full 3D geometry of the object and a respective SVBRDF model. In this device, 633 static light sources are mounted equidistantly distributed on a spherical gantry. The object is placed on a large glass plate that can be rotated and can be observed by an array of seven cameras that allows an acquisition under different inclination angles. This allows an acquisition of the object from view-light configurations in both the upper and the lower hemisphere. In total, 133 different view directions under illumination by 19 different illumination patterns are considered during the acquisition of 3D objects. The respective illumination patters are generated by activating certain light sources. The acquisition of these view-light configurations takes about 1h. The rather large diameter of this setup and the use of wide-angle lenses allows the acquisition of objects within an effective volume with a diameter of about 80cm. Local BRDFs have been fitted based on the isotropic Ward model [War92]and a non-linear optimization [Lev44].

Both the geometry and spatially varying reflectance characteristics of 3D objects have been acquired in [TFG*13] from observation under continuous spherical harmonic illumination conditions. The latter has been generated with a semi-circular rotating arc of 105 controllable LEDs. The respective observations can be used to perform a separation of diffuse and specular reflections and to compute surface normals. Furthermore, anisotropic roughness parameters can be acquired as well.

A different approach to acquire SVBRDFs has been presented in [CDP*14] with a generalized linear light source reflectometry framework. This method allows the acquisition of anisotropic surface reflectance behavior and the local shading frame based on observations of the material sample under illumination by a constant linear light source and two phase-shifted sinusoidally modulated linear light sources. In this setup, the camera is either fixed rela-

tively to the linear light source or fixed with respect to the material sample.

7.5. BTFs and Their Acquisition

Similar to SVBRDFs, Bidirectional Texture Functions (BTFs) $\rho_{\mathrm{BTF}}(\mathbf{x},\theta_i,\varphi_i,\theta_r,\varphi_r)$ that have been introduced in [DvGNK97] also model material appearance depending on particular spatial positions $\mathbf{x}$ on the object surface, the directions of the incoming light (θ_i,φ_i) and the direction of the outgoing light (θ_r,φ_r). However, SVBRDFs cannot represent non-local light transport occurring due to e.g. interreflections or subsurface scattering. The reason for this is given by the nature of BRDFs that only allow to model local phenomena. BTFs can be interpreted as a special case of eight-dimensional reflectance fields [DHT*00] given by $\rho_{\mathrm{RF}}(\mathbf{x}_i,\theta_i,\varphi_i,\mathbf{x}_r,\theta_r,\varphi_r)$. However, for BTFs, the light sources that determine the incoming light field are assumed to be infinitely far away so that the incident radiance $L_i(\mathbf{x}_i,\theta_i,\varphi_i)$ for a given light direction (θ_i,φ_i) is equal for each 3D point $\mathbf{x}_i$ in the scene, i.e. $L_i(\mathbf{x}_i,\theta_i,\varphi_i) = L_i(\theta_i,\varphi_i)$, and, hence, its dimensionality can be reduced by two dimensions. Although this assumption is not completely fulfilled in real measurement setups, the light sources are typically significantly farer away from the object surface than the size of its geometric details which makes this assumption nevertheless reasonable. The dependency of BTFs on spatial surface positions requires them to be parameterized over a certain reference geometry. For this purpose, a planar proxy geometry can be used or the geometry of the physical object surface can be reconstructed using e.g. structured light techniques or laser scanners. If an approximate geometry is used that deviates from the true geometry, the surface details not included in the proxy geometry might cause non-local effects such as interreflections, self-shadowing, self-occlusions or subsurface scattering characteristics due to the influence of other surface details nearby. Such structures may be given by scratches, engravings, grain of wood or the embossing of leather and might not be fully resolved in the digitized material representation so that the surface orientation and with it also the direction-dependent reflectivity might vary within one spatial sampled unit [Sch14]. Detailed surveys regarding BTFs are given in [MMS*04, FH09, HF13, SSW*14]. In addition, a detailed evaluation of different acquisition principles is given in [SSW*14]. In this part, we follow the main trends in BTF acquisition as also discussed in these surveys.

Similar to the acquisition of BRDFs, BTF acquisition setups also rely on the arrangements of light sources and imaging sensors relatively to the material sample that has to be acquired. However, more degrees of freedom have to be realized in the setup as six parameters have to be taken into account during the acquisition of BTFs. Most of the BTF acquisition setups that have been introduced so far can be categorized into the categories of gonioreflectometer-based setups (see Section 7.5.1), mirror-based setups (see Section 7.5.2) and arrays of cameras and/or light sources (see Section 7.5.3).

7.5.1. Gonioreflectometer-Based BTF Acquisition

Gonioreflectometer-like setups as often used for BRDF acquisition have also been used in seminal work on BTF acquisition [DvGNK97], where the CUReT BTF database with 61 material measurements has been introduced. Each of the individual material samples has been measured in a gonioreflectometer setup where it is positioned by a robotic arm while being illuminated by a static light source and observed by a camera from different, manually configured viewpoints. The involved manual effort severely limits the number of measured view-light configurations. Only seven different view directions are taken into account during the acquisition process and the materials are assumed to exhibit isotropic surface reflectance. In order to also take anisotropic reflectance behavior into account, material samples have to be measured multiple times with different azimuthal orientation. In total, each material is measured under 205 different view-light configurations with a video camera in about 1h or 410 different view-light configurations in about 2h if the material has been acquired under two azimuthal orientations. However, this represents a rather coarse sampling of the possible configurations so that high-frequency details in surface reflectance behavior might not be captured adequately.

The setups in [KTT06, SSK03] are rather similar as also a static illuminant is used to illuminate a material sample that is positioned by a robot arm. However, in [KTT06], the camera is also re-positioned automatically using another robot arm. In addition, the use of a spectral filter wheel in front of the light source allows to consider eight different spectral bands and, hence, a multi-spectral acquisition. In total, the eight spectral responses under 71 view-light configurations have been measured for fixed azimuthal angles of $\varphi_i = 180^\circ$ and $\varphi_i = 0^\circ$ and varying inclination angles θ_i. In contrast, the camera is automatically positioned based on a half-circular rail system in the setup presented in [SSK03]. Using this setup, BTFs have been measured with a reasonable uniform angular sampling of the hemisphere above the material sample with 81 view directions and 81 light directions as well as with a reasonable spatial resolution in about 14h. This setup has been extended in [RSK10] by using 32 spectral filters in front of the light source in order to perform a multi-spectral BTF acquisition in about 60h.

A rather similar setup design is given by mounting the camera at a fixed position and moving the material sample and the light source. The setup in [McA02] is based on two rotation stages that allow to specify the pose of the material sample that is observed by a static camera in HDR and illuminated by a light source mounted at a moveable arm. The acquisition of $311 - 7{,}650$ view-light configurations takes about 45min to 36h. In order to allow hyperspectral material measurements with 16 spectral bands, the rather lightweight setup in [TAN*05, TSA*05] extends this concept by the use of a band-pass filter wheel in front of the light source. Measurements of 6,500 view-light configurations can be performed in about 13h. However, only BRDFs have been measured with this setup although it is also applicable to measure BTFs. The compact setup size of $80\text{cm} \times 80\text{cm} \times 80\text{cm}$ restricts the sample size to $4\text{cm} \times 4\text{cm}$. Based on a pan-tilt head, the material sample can be moved so that a static camera observes it under different viewing directions while being illuminated by a moveable light source from positions in the upper hemisphere above the material sample [KMBK03]. In these investigations, 120 light directions and 90 viewing directions have been sampled which results in a total of 10,800 view-light configurations that can be measured in about 10h as reported in [HF11]. The used LDR video camera has a rather low

spatial resolution of 640×480 pixels and results in a low dynamic range of the measured data.

Instead of moving either the light source or the camera, several setups also exploit moving both types of setup components using robot arms (see e.g. [FVH*13, HLZ10]). In turn, this allows to avoid the need for manipulating the pose of the material sample. In [FVH*13], an angular sampling of 81×81 uniformly sampled view-light configurations similar to the sampling in [SSK03] has been used to capture HDR data and an extremely high maximum resolution of 1,071DPI as well as a high angular accuracy of 0.03° have been achieved. While the camera is mounted on a robot arm that can only be tilted, the arm with the light source can also be panned. The acquisition of a single BTF takes about 18h. In contrast, the setup concept in [HLZ10] is based on using a gonioreflectometer-like setup with a light source and a camera with a beam-splitter on two robot arms. This allows the acquisition of the material sample under a coaxial arrangement of view direction and light direction and, hence, also includes the possibility to acquire retro-reflections and to acquire reciprocal image pairs. In addition, geometry reconstruction as well as the separation of direct and global components of light transport can be achieved based on the use of fringe projectors as light sources. However, the sequential acquisition is rather slow as the material sample can be acquired under only 84 view-light configurations in about 5h.

The main advantage of gonioreflectometer-like setups is their flexibility regarding the acquisition of almost arbitrary view-light configurations based on the involved robot arms, tilt-pan heads or rotation stages as well as the low number of components which allows the use of more costly high-quality components that offer better optics, higher resolutions, etc.. For instance, the use of industrial cameras is more adequate regarding the huge number of images that have to be acquired. Furthermore, such setups can be rather easily extended to allow a multi-spectral acquisition by placing remotely-controllable spectral band-pass filters in front of the light sources and possibly also in front of the cameras. However, relying on moving parts comes at the cost of being susceptible to inaccuracies in the positioning of the components which results in inaccurately registered measurement data. In addition, the sequential measurements under different view-light configurations result in rather long acquisition times.

7.5.2. BTF Acquisition Using Mirror-Based Setups

Because of the dependency of BTFs on six parameters, the space of view-light configurations that have to be sampled is rather large and a subsequent acquisition as typically performed with gonioreflectometer-like setups results in long, rather impractical acquisition times. For this reason, mirror-based concepts have also been used for BTF acquisition similar to the developments regarding setup designs for BRDF acquisition. The advantage of such setups is given by their ability to capture several view-configurations simultaneously based on a more sophisticated arrangement of mirroring surfaces which allows to speed-up the acquisition process significantly.

Several acquisition approaches rely on the use of curved mirrors. The setups in [Dan01, DW04, WD06] involve a parabolic mirror similar to the BRDF acquisition device in [Dan01]. Directional illumination in the form of a light beam is generated using fiber-optics and the direction of the light beam is controlled by a specific aperture. A light beam is first sent towards the mirror, where it is reflected to illuminate the material sample from a small solid angle. On the material surface, the incoming light is again reflected towards the mirror where it is re-directed towards the camera. As the observed surface point is placed in the focal point of the parabolic mirror, the camera captures multiple view directions under the same illumination simultaneously. However, light directions and spatial dimension still have to be sampled sequentially. The acquisition of the spatial variations in the surface reflectance at different surface points is realized with a translation stage that allows to move the mirror above the material sample. Similar to the BRDF acquisition setup in [Dan01], the main drawback is given by the rather limited range of directions between $\theta \leq 23^\circ$ and $\theta \leq 37^\circ$ elevation (depending on the azimuth angle) which does not allow to capture reflectance under grazing angles. As only a horizontal translation of the mirror is taken into account, the setup only allows the acquisition of flat material samples and the position of the focal point would have to be adjusted regarding the surface of non-planar material samples. Unfortunately, the sampling of 200×200 surface points under a single light direction already requires an acquisition time of about 1h which makes a dense acquisition rather impractical.

Other techniques exploit arrangements of several piecewise planar mirrors as e.g. presented in [HP03]. This allows to observe the material sample from an individual viewing direction for each of the piecewise planar mirrors and, hence, from several viewing directions simultaneously in an image acquired by a camera. This means that a complete outgoing light field is captured within a single image. One advantage offered by this strategy is given by the parallel acquisition of the spatial domain which offers the possibility to achieve a fast acquisition of the complete BTF in about 1h (see [HP03]) which is significantly faster than with a successive pointwise acquisition as in [Dan01, DW04, WD06].

The acquisition setups in [LCV*04, GTLL06, MTK*11, TMY12] follow this principle and are based on parabolic arrangements and ellipsoidal arrangements of the piecewise planar mirror facets. Similar to the setup in [HLZ10], the combination of a beam-splitter and a projector allow a coaxial camera-projector arrangement. To illuminate the material sample, only the projector pixels that are projected to a specified mirror are activated. The accurate calibration of the camera, projector and mirrors are crucial for a reliable acquisition and no further mechanical manipulations need to be performed during the acquisition. However, only sparse samplings have been acquired. For instance, in [MTK*11], direct reflections from 50 mirror facets are used which results in the acquisition of 2,500 view-light configurations.

Furthermore, kaleidoscope-based setups have been presented in [HP03, IRM*12]. These investigations focus on reducing the number of the mirror facets by exploiting recursive interreflections in a setup with three planar mirrors to obtain observations from a multitude of virtual viewpoints simultaneously with a single camera. The individual light directions are controlled using digital projectors. In [HP03], different samplings of 22 light directions and 22 view directions as well as 79 light directions and 79 view directions

have been acquired based on different angles of taper. While only flat material samples have been acquired with the setup in [HP03], a kaleidoscope-based setup has been used in [IRM*12] to acquire 3D objects, where the object is placed on an additional mirror so that the lower hemisphere can be captured as well. In total, a sampling of 35,424 view-light configurations can be measured based on 246 virtual views and 144 virtual light sources. However, only 8,856 view-light configurations can be used for flat material samples, for which the lower hemisphere cannot be measured.

As typically only one camera and one light source are used in mirror-based setups, such setups allow the use of high-quality hardware at still acceptable costs. Furthermore, moving parts are typically avoided in mirror-based setups such as [HP03, LCV*04, GTLL06, MTK*11, IRM*12] which means that the setup can be calibrated beforehand. However, the calibration itself is more complex for these setups as several levels of indirection might have to be taken into account as argued in [SSW*14]. Furthermore, the trade-off between angular resolution and spatial resolution of such setups is rather problematic regarding appearance acquisition where both fine surface details and the variations of material appearance under a dense set of view-light configurations have to be measured.

7.5.3. BTF Acquisition based on Camera Arrays and Light Source Arrays

A further possibility to reduce the acquisition time in terms of performing parallel measurements is given by the use of camera arrays and light source arrays. This reduces the mechanical manipulation of components that need to be performed during the acquisition process and therefore introduces less inaccuracy in the registration of the measured data.

In [CDMR04], several light sources have been mounted on an arc and the material sample is observed by a camera mounted on a moveable arm. However, only three viewing directions and ten illumination directions or four viewing directions and eight illumination directions have been used when acquiring the Rutgers Skin Texture Database that depicts skin diseases under a sparse sampling of view-light configurations.

A denser sampling of the light directions has been used in the "Light Stage" introduced in [DHT*00]. While this setup allows the acquisition of reflectance fields based on only two cameras observing the object of interest that is illuminated by a moveable light source on a two-axes rotation system, it has been extended in [HCD01] by additional light sources. The 27 used light sources are mounted on a rotating arc. A further extension of the "Light Stage" has been presented in [DWT*02, WGT*05], where 156 fixed light sources are arranged in a dome-like manner. By using high-speed cameras, the acquisition can be performed in only a few seconds or even less. However, the use of expensive cameras at static positions allows only an insufficient sampling of the view directions. Similar to the setup in [DWT*02, WGT*05], a dome-like arrangement of 150 uniformly distributed light sources has been used in the setup presented in [WMP*05, WMP*06], where 16 cameras have been used to improve the sampling of the view directions. However, considering a sampling of only 16 view directions is still not sufficient for the bi-directional representations of materials with BTFs. For this reason, the setup design in [HQS10, HWQ13] also relies on a light dome with 238 light sources, however, the twelve involved cameras have been mounted on a vertical arc, i.e. all of the view directions have the same azimuthal and a different inclination angle, which allows to acquire a five-dimensional BTF slice for isotropic materials. In the setup presented in [WLDW11], a light dome of 290 light sources is used where 20 cameras have been mounted on a horizontal ring, i.e. the viewing directions have different azimuthal angles and the same inclination angles. However, using only view directions with similar inclination angles does not result in an acquired BTF slice that is adequate to model surface reflectance behavior.

A different setup design has been proposed in [FKIS02], where two arcs have been equipped with six light sources and five cameras respectively. As the object is placed on a turntable and the light arc can be rotated as well, both view directions and light directions can be sampled on the upper hemisphere. As intervals of 30° between azimuthal angles and between inclination angles have been chosen, a total of $72 \times 60 = 4{,}320$ view-light configurations has been sampled during the acquisition of 3D objects. However, only the view-light configurations in the upper hemisphere are meaningful when acquiring flat material samples. The setup design in [TWL*05] is rather similar, however, eight light sources and eight industry cameras are used and 7,056 view-light configurations are captured within 2h. A similar setup design with four light sources and six cameras has been used in [MPZ*02, MPN*02], where two additional screens have been placed below and behind the object so that the transmitted light can be densely measured. In addition, matting images have been captured. While the focus in the corresponding investigations was not placed on BTF acquisition, a part of the measurement can be seen as a BTF with a sampling of 60×216 $(= 12{,}960)$ view-light configurations. However, this sampling is rather non-uniform and the involved mechanical movement of components makes the acquisition process rather slow so that about 14h are needed to capture all the view-light configurations. To take different object sizes better into account, different sets of lenses have been used.

More recently, the "OrCam" acquisition device has been presented in [KNRS13, NJRS13]. This setup is based on a spherical gantry with 633 equidistantly distributed static light sources. Similar to the setups in [FKIS02, MPZ*02, MPN*02, TWL*05, SSWK13], the object is placed on a turntable so that it can be captured from different azimuthal angles with an array of seven cameras. This camera array can be rotated so that an acquisition under different inclination angles can be performed. As the turntable consists of a large glass plate, the object can also be acquired from view-light configurations in the lower hemisphere. 3D objects have been acquired from 133 different view directions under illumination by 19 different illumination patterns, where the illumination patters are generated by activating certain light sources. In order to capture the resulting 2,527 view-light configurations, an acquisition time of about 1h is needed. As a result of the rather large diameter of the spherical device and the use of wide-angle lenses, this setup is designed particularly regarding larger objects, i.e. the effective volume that can be considered during the acquisition has a diameter of 80cm and the reported resolution is about 127 DPI.

Closely related is the "KULETH Dome" presented in [NZVG05, NZVG06]. A dome-like arrangement of 169 light sources that have been mounted on a quarter-spherical gantry is used to illuminate the object that is placed on a rotation stage in the center of the setup from different light directions. The involved camera has been mounted on a robot tilt arm that allows to specify different viewing directions. As only one camera is used in this setup, view directions cannot be captured in parallel. However, the use of an array of fixed light sources allows a reduction of mechanical manipulations of the poses of the components. A highly parallelized BTF acquisition device has been presented with the "Multiview Dome device" (see Figure 29) in [MMS*04, SWRK11]. This setup is designed to observe the object from a multitude of viewpoints and under many illumination directions without the need to re-position parts of the setup. Accordingly, the Multiview Dome is equipped with 151 consumer cameras (resolution of 12MP, 8-bit, RGB) which are uniformly distributed on a hemispherical gantry with a diameter of approximately 150cm. The flashes of the cameras serve as light sources. This allows to capture 151 images per selected light source simultaneously. In order to acquire the 3D geometry of objects, nine projectors with a resolution of 800×600 pixels are additionally mounted at different locations on the gantry and used for projecting a series of structured light patterns. By exploiting the capabilities of high dynamic range imaging, even the acquisition of objects with complex surface materials and mesostructures can be achieved which is a very challenging task for conventional 3D measurement systems. As a result, the Multiview Dome is capable of acquiring both a precise geometry of the object surface together with a faithful representation of the surface reflectance behavior of the object in the form of a bidirectional texture function. As shown in Figure 30, the BTF is parameterized on the surface geometry that has been reconstructed in a prior step. When measuring flat samples, the spatial resolution is approx. 53.0μm per pixel and the angular resolution is 7.5° in zenithal and at most 15° degrees in azimuthal direction. Typical acquisition times for the total set of 22,801 view-light configurations are in the range of 1.8h to 4.0h.

Following the demand for a moveable acquisition device which can be brought to different cultural heritage institutions, a moveable version of the Multiview Dome (see Figure 31) has been presented in [SSWK13, SSW*14] which can be readily disassembled into four transportable parts and is capable of acquiring artifacts with a maximum size of 100mm $\times$ 100mm $\times$ 100mm. In the moveable Multiview Dome device, eleven industry cameras with a resolution of $2{,}048 \times 2{,}048$ pixels (12-bit, RGB) are mounted along an arc on a hemispherical gantry (diameter approx. 2m), covering zenith angles between 0° and 75° with an angular spacing of 7.5°. In order to be able to deal with objects of different sizes, two full sets of high quality Zeiss prime lenses with 50mm and 100mm are available for the cameras. Using the 100mm lenses leads to a per-pixel surface resolution of about 67.5μm. In case of the 50mm lenses, the per pixel resolution is about 125.0μm per pixel. In addition, 198 LED light sources are uniformly distributed on the hemispherical gantry in whose center the artifact to be acquired is placed on a highly accurate turntable which allows a dense azimuthal sampling. The use of a camera arc in combination with the turntable makes the acquisition significantly slower than the one with the "Multiview Dome device" in [MMS*04, SWRK11, SSW*14] and an acquisition takes 4h to 12h. Additionally, a precise reconstruction of the 3D geometry is achieved via the placement of four projectors with a resolution of 800×600 pixels, which serve for projecting structured light patterns onto the object surface. The moveable Multiview device allows to measure both full 3D shape and bidirectional reflectance. For a rather dense angular sampling of view and light directions, the acquisition process takes a reasonably short amount of time of six to 12h per pose of the object.

Improving the acquisition speed using the setup from [SSWK13] has been investigated in [dBSHK14, dBSHK15]. In [dBSHK14], the similarity of surface structure and reflectance properties exhibited by similar material types is exploited based on clustering a multitude of BTFs and fitting linear models per cluster. This allows a sparse reconstruction based on a linear least-squares fitting. Furthermore, multiplexed illumination by using illumination patterns is used to reduce the required exposure times during the acquisition of flat material samples in [dBSHK15]. The linearity of light superposition allows to obtain a BTF based on the multiplexed data. However, the illumination multiplexing leads to a reduced signal-to-noise ratio. In order to remove noise from the reconstruction, a linear model has been computed based on an existing database of BTFs which can be used for denoising purposes.

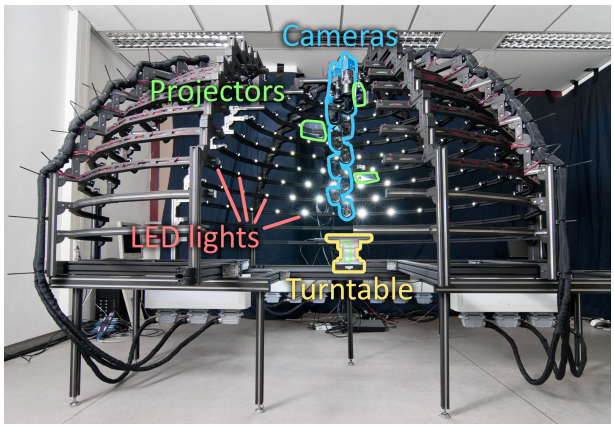

Figure 31: Multiview BTF acquisition device introduced in [SSWK13]: 11 cameras are mounted in an arc-like arrangement. The object is placed on a turntable and can be illuminated by 198 light source positions placed on a hemispherical gantry. In addition, 4 projectors are included in the setup to allow a structured light based geometry reconstruction.

The different camera array and light array setups clearly show the benefit of parallel acquisition regarding a reduction of the acquisition time in comparison to sequential gonioreflectometer-based acquisition, and there is also no need for a trade-off between spatial and angular resolution as for mirror-based setups. The reduction of moving parts allows a calibration of the setup components in a calibration step performed before the actual measurement and is therefore less susceptible to positioning errors. So far, three devices (the ones in [MMS*04, SWRK11], [WMP*05, WMP*06] and [HQS10, HWQ13]) perform an acquisition of all view directions on the full hemisphere simultaneously without moving parts which makes them being the fastest BTF acquisition devices as re-

Figure 29: Multiview BTF acquisition device used in [MMS*04, SWRK11]: An array of 151 cameras on the upper hemisphere above the object allows a highly-parallelized acquisition. The camera flashes have been used as light sources so that a dense sampling of 22,801 view-light configurations can be achieved. The use of projectors in the setup as used in [WSRK11, SWRK11] allows a structured light based geometry reconstruction.

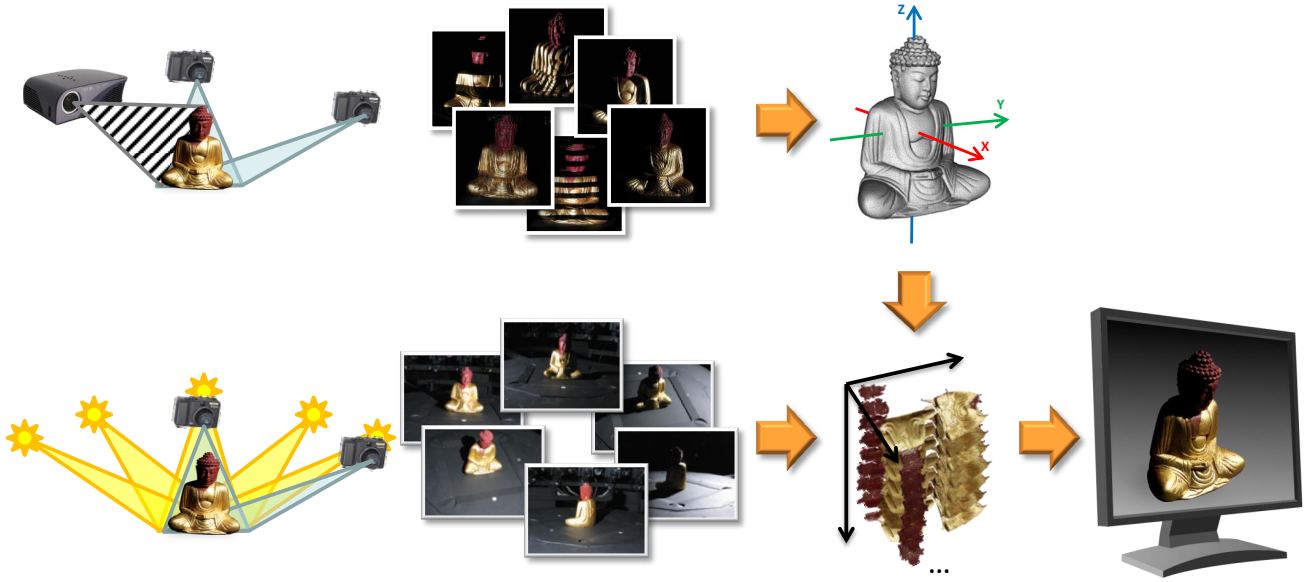

Figure 30: Integrated BTF acquisition pipeline as presented in [SWRK11, SSWK13]: The object geometry is reconstructed using structured light similar to [WSRK11] and the BTF is stored as rectified textures that are parameterized over the reconstructed object surface.

ported in the survey in [SSW*14]. Nevertheless, it has to be mentioned that only the setup in [MMS*04, SWRK11] acquires a dense sampling of the view-light configurations in the upper hemisphere while the remaining setups only consider a rather insufficient sampling of view-light configurations. However, as significantly more components are involved than for gonioreflectometer-like setups or mirror-based setups, it is extremely expensive to equip such setups with high-quality components, and control issues as well as synchronization have to be taken into account as well [SSW*14]. The obvious compromise seems to be seen in the combination of smaller arrays of cameras and light sources with a turntable as followed in [FKIS02, MPZ*02, MPN*02, TWL*05, SSWK13] and possibly an additional movable tilt arm [NZVG05, KNRS13, NJRS13] at the cost of an increasing acquisition time.

As also discussed in [SSW*14], the reduction of the number of moving components to at most a turntable rotation and the typically rather large setups required to mount all the components have led to the use of this type of setup for the reflectance acquisition of 3D objects. Furthermore, camera and light arrays also allow the integration of geometry acquisition into the setups. For instance, the shape-from-silhouette technique has been used in several setups [FKIS02, MBK05, MPZ*02, MPN*02] to acquire a visual hull of the object geometry. However, this technique does not allow a robust geometry reconstruction as object silhouettes cannot easily be extracted in the individual images and concavities in the surface geometry cannot be handled by this approach. In turn, an inaccurate geometry reconstruction also affects the reconstructed reflectance on the object surface. For this reason, laser scanners [FKIS02] or structured light systems [WMP*05, WMP*06, SWRK11, SSWK13, KNRS13, NJRS13]

have been integrated into the individual setups which results in accurately captured geometry information.

7.6. BSSRDF Models and Their Acquisition

While local subsurface scattering can already be modeled by data-driven BTFs to some degree, the more general representation of materials exhibiting subsurface scattering characteristics requires a model that is based on even more parameters. The bidirectional scattering surface reflectance distribution function (BSSRDF) $\rho_{\text{BSSRDF}}(\mathbf{x}_i, \theta_i, \varphi_i, \mathbf{x}_r, \theta_i, \varphi_i)$ has been formally introduced in [NRH*77] as such a more general model. While this model describes material appearance depending on the light direction (θ_i, φ_i) and the viewing direction (θ_r, φ_r), it not only models the direct reflection of the light at a local point $\mathbf{x}_i$ on the object surface where the light hits the surface, but instead takes into account that light enters the material at the point $\mathbf{x}_i$, travels inside the material where it is scattered and exits the material at the surface point $\mathbf{x}_r$. As e.g. described in [WLL*07], the definition of the BSSRDF is given by the quotient of the differential reflected radiance $\mathrm{d}L(\mathbf{x}_r, \theta_r, \varphi_r)$ over the incident flux $\mathrm{d}\Phi(\mathbf{x}_i, \theta_i, \varphi_i)$, i.e.

$$\rho_{\text{BSSRDF}}(\mathbf{x}_i, \theta_i, \varphi_i, \mathbf{x}_r, \theta_r, \varphi_r) = \frac{\mathrm{d}L(\mathbf{x}_r, \theta_r, \varphi_r)}{\mathrm{d}\Phi(\mathbf{x}_i, \theta_i, \varphi_i)}. \tag{62}$$

However, not only subsurface scattering can be modeled using BSSRDFs. It is also well-known that including $\mathbf{x}_i$ and $\mathbf{x}_r$ as additional parameters to $\theta_i, \varphi_i, \theta_r$ and φ_r allows to represent the near-field reflectance field, where incident and outgoing light rays are coupled [WLL*07]. In comparison to e.g. BTFs, where far-field illumination and thus a directional illumination is assumed, it is often desirable to consider illumination from a rather close range instead. In this scenario, directional illumination does not sufficiently model the light exchange between several objects within a scene. Instead, the light paths within the scene need to be taken into account which requires to consider the parameters $\mathbf{x}_i$ and $\mathbf{x}_r$ as well. Obviously, the dependency of material appearance on a rather high number of eight parameters that need to be taken into account during the BSSRDF acquisition results in typically long acquisition times. This, in turn, makes the acquisition of BSSRDFs rather impractical. In the following sections, we provide a survey on individual techniques depending on the respective principles.

7.6.1. Techniques Based on the Separation of Diffuse and Specular Components

As also discussed in [HF13], one option towards a more practical acquisition can be seen in the possibility to model material appearance based on two separated components. The diffuse component describes non-local reflection and is determined by the characteristics of the scattering of light within the material. Furthermore, the specular component describes local reflection effects, i.e. the direct reflection of light at the material surface, which can be explained by approximate BRDF models. As a result, the separation of these components seems to offer the potential for an easier acquisition and modeling of the reflectance behavior of such materials, and several studies have investigated respective separation approaches for the acquisition of BSSRDFs.

Based on modeling the reflectance of dielectrics with a dichromatic reflectance model, a separation of diffuse and specular components has been approached in [Sha85] by exploiting that the spectral characteristics of the diffuse component are determined by material properties and the spectral characteristics of the specular component are determined by the illumination characteristics. Alternatively, polarization-based approaches also allow the separation of both components for dielectrics. Such techniques typically rely on manipulating the characteristics of the propagated light based on different types of polarization filters with different refractive indices that can be placed in front of the cameras and light sources. This allows to influence the polarization of the emitted and observed light e.g. being linearly-polarized or circularly-polarized. Polarization filters can be used in front of the cameras, as e.g. in the setup in [WB91], so that the non-polarized or polarized parts of the received light can be measured separately depending on the filter orientation. Furthermore, the investigations in [WB91] show that the specular component typically exhibits a polarization depending on the filter type, while the diffuse component is non-polarized. When such filters are not only placed in front of the cameras but also in front of the light sources, the polarization characteristics can be specified for the incoming light. This helps to analyze the light transport in subsurface scattering materials, where the light leaving the materials is non-polarized due to the scattering, while the directly reflected light keeps its polarization. Depending on the filter orientation, the camera receives the non-polarized, partially polarized or polarized light, and, hence, the separation of the diffuse and specular component can be performed based on only a few images.

In [NFB97], both color and polarization information have been used to separate the diffuse and specular components per image pixel. The polarization of the reflected light is used to locally compute the color of specular reflections. This results in further constraints regarding a linear color subspace in which the local diffuse components are specified. In contrast, a more simple approach based on a static camera and a projector has been introduced in [NKGR06], where high-frequency illumination patterns are projected onto the object surface with the projector to separate the direct and global components. The key observation is that both components are involved for the regions of lit parts and that only the diffuse component is available for the dark pattern regions. The authors used shifted checkerboard patterns as well as phase shift patterns. For the latter, only three images have to be acquired under illumination by different phase shifts to allow the separation of diffuse and specular components.

7.6.2. Acquisition of Subsurface Scattering Characteristics

Other approaches have focused on the acquisition of subsurface scattering characteristics for materials with a spatially homogeneous reflectance behavior. This allows the use of easier acquisition setups in comparison to the ones that are required to measure spatially varying BSSRDFs. For instance, a reflectance model that models the standard surface reflectance and an additional subsurface reflectance due to backscattering in a layered turbid media has been presented in [HK93]. Anisotropic surface reflectance can be represented as the model incorporates directional scattering within the layer, and this model has been applied to biological and in-

organic materials with low indices of refraction. For such materials, a rather large fraction of the incoming light is scattered within the material. Another simple analytical BSSRDF model based on single scattering characteristics has been presented in [JMLH01] together with a dipole model for multiple scattering. A static camera is used as a detector to measure the scattering in the material induced under illumination by a focused light ray. A generalization of this technique has been proposed in [FGCS05], where the dipole model is replaced by a sum of exponential functions and the object geometry is used. This allows the modeling of inhomogeneous materials. Furthermore, a more complex setup in comparison to the one presented in [JMLH01] has been used in [WMP*06] to measure the subsurface scattering in human skin. A set of 28 fibers is arranged in an array-like structure and each of the fibers is capable of serving both as light source and as detector. Each of the fibers successively illuminates the material sample, while the remaining fibers capture the subsurface scattering characteristics. Based on this data, a rather simple BSSRDF model is fitted similar to [JMLH01].

The more complex scenario of materials with spatially varying subsurface scattering characteristics has been studied in several investigations. In [GLL*04], the acquisition of translucent objects with a high scattering albedo is approached based on the assumption that the light traveling inside the material has enough scattering events so that the influence of the light direction of the incident light on the direction of the outgoing light can be neglected. This results in a four-dimensional diffuse scattering function $\rho(\mathbf{x}_i, \mathbf{x}_r)$ that describes the scattering depending on the parameters $\mathbf{x}_i$ and $\mathbf{x}_r$. The object geometry is assumed to be known from a previous geometry reconstruction step that is based on covering the object with a specific dust, performing a laser scan and removing the dust again. Therefore, the acquisition setup proposed in [GLL*04] is based on a static HDR video camera that observes an object placed on a turntable and illuminated by a laser projector. The projected laser is controlled so that it sweeps over the object surface and three different wavelengths are measured per surface point. The scattering function $\rho(\mathbf{x}_i, \mathbf{x}_r)$ can be modeled as a four-dimensional tensor, and for each individual measurement, one slice in this tensor is filled. However, occlusions or self-shadowing occurring due to the complexity of the surface shape of the object might cause missing values that are filled with an inpainting technique. It is also worth mentioning that a hierarchical representation has been chosen for the four-dimensional tensor so that high sampling rates are only available close to points $\mathbf{x}_i$, while a sparse sampling is sufficient for distant points. Rather long acquisition times of several hours have been reported for measurements based on this approach. Focusing on planar samples, the acquisition of four-dimensional slices of the full spatially varying BSSRDF has been investigated in [PvBM*06] based on a more efficient acquisition. The latter is achieved by the use of a projector to simultaneously sweep a grid of light points over the surface of the material sample, and an acquisition takes about three hours.

Acquiring the appearance of translucent materials based on multi-layer models has been studied in [DJ05, DWd*08]. While a specific light diffusion model including surface roughness and spatially varying refractive indices has been introduced in [DJ05] to model multi-layered materials, a layered model tailored to human skin has been proposed in [DWd*08]. The latter approach is based on the acquisition of multi-spectral data obtained by placing different spectral narrow-band filters in front of the light sources and allows the modeling of spatial variations in absorption and scattering characteristics. Furthermore, the appearance acquisition of human skin in terms of a layered model has also been investigated in [GHP*08] with a single-view setup. Only 20 photos are used and the model includes specular reflectance, single scattering, shallow subsurface scattering and deep subsurface scattering. Based on layers with known scattering parameters, the technique in [DWP*10] focuses on the adjustment of the thickness of the individual layers to produce certain BSSRDF characteristics.

As the illumination of translucent materials by a laser beam may result in large intensity differences between specular radiance and global scattering radiance due to the highly concentrated light, a different approach has been proposed in [TWL*05], where global and local scattering are separately captured based on different types of illumination. The acquisition setup consists of an arc that is equipped with eight cameras that allow to capture an object placed on a turntable in HDR. The object can be illuminated by eight halogen lamps and a laser scan unit mounted on a second moveable light arc. As the rather uniform illumination of the lamps results in a relatively narrow dynamic range and reduces the effect of obscuring texture details due to specular reflections, these illuminants are used for the measurement of a BTF that captures the local scattering. Additionally, laser stripes have been projected onto the object with the laser scan unit to acquire the global subsurface scattering characteristics and to allow a reduction of the required measurements in comparison to using laser points. The global subsurface scattering characteristics have been modeled using both a mesostructure entrance function and a mesostructure exiting function that describe mesostructure effects of light exchange when entering or exiting the material.

A rather simple projector-camera setup has been proposed in [TGL*06] to acquire spatially varying subsurface reflectance properties of a human face. First, per-pixel profiles are obtained based on projecting a sequence of phase-shifted periodic stripe patterns, which reveals how light scatters from adjacent locations. The face geometry has been acquired with a structured light approach. In order to remove the influence of the interreflected light, the minimum of each profile is subtracted before the observed reflectance profiles are matched to scattering properties predicted by a scattering model. This approach allows to generate realistic depictions of a material under complex incident illumination with acquisition times of less than a minute.

7.7. Light Fields and Reflectance Fields

Unfortunately, illumination conditions encountered when acquiring materials in uncontrolled conditions of our daily life environments are much more complex in comparison to the ones given in controlled lab environments. As a consequence, relying on directional illumination conditions during the acquisition process as typically applied for the acquisition of BRDFs, SVBRDFs, BTFs and BSSRDFs is not sufficient anymore. As discussed in Section 3.3, *reflectance fields* [DHT*00] model material appearance depending on the radiant light field $\rho_{\mathrm{LF},r}(\mathbf{x}_r, \theta_r, \varphi_r)$ and on the incident light field

$\rho_{LF,i}(\mathbf{x}_i, \theta_i, \varphi_i)$. As a result, there is a dependency on eight parameters. However, it is rather impractical to densely sample the configurations of such high-dimensional spaces during the acquisition. Instead, simplifying assumptions are typically applied. Regarding the incoming light field, the assumption of far-field illumination allows to represent them via simple two-dimensional images such as environment maps that only depend on the parameters θ_i and φ_i. In practice, this can be achieved by using directional light sources or the projection of light patterns. Furthermore, the dimensionality of reflectance fields can be reduced by fixing the view direction which results in a four-dimensional function. If the reflectance field is further parameterized over the true object surface, i.e. $\rho_{SRF}(\mathbf{x}, \theta_i, \varphi_i)$, the term surface reflectance field is used.

The dependency on a rather large number of parameters makes the acquisition of reflectance fields rather impractical. Therefore, it is desirable to reduce the number of configurations that have to be measured. This can be achieved based on exploiting different principles such as the linearity of light transport or symmetry of light transport [BCNR10, Leh07]. The linearity of light transport allows to model the incoming illumination by a set of a finite number of directional lights or other more complex basis functions based on superposition. In contrast, the symmetry of light transport allows to swap the light sources with the detectors.

In the following sections, we review techniques that focus on the acquisition of surface light fields and surface reflectance fields.

7.7.1. Acquisition of Surface Light Fields

The acquisition of four-dimensional surface light fields is typically approached by taking images with a single camera at different viewpoints [LH96, GGSC96, WAA*00, LLW*08] or by using arrays of cameras [WJV*05] so that the appearance of the material sample of interest can be captured under fixed illumination from multiple viewpoints and the acquired images are registered to the surface geometry which may be obtained from an initial geometry acquisition.

Unlike such conventional surface light field estimation techniques, the technique presented in [PDCS13] does not rely on regularly sampled viewing conditions with respect to the object of interest. Based on an irregular sampling of viewpoints obtained from an acquisition where a video camera is moved around the object, which results in a dense sampling of viewpoints along the trajectory and no view conditions with other directions, both an approximation of a surface light field and a rough approximation of the main light directions in the scene have been estimated based on a known object geometry. After an initial video-to-geometry registration where the measured color values are registered and associated with a quality value, a rather rough estimation of the directions of the main light sources is performed. This can be achieved by assuming distant illumination and estimating an approximate environment map based on its specular reflection in a mirroring sphere. The respective color samples on the sphere that have a luminance value above a fixed threshold are expected to result from highlights and k-means clustering is applied to estimate the main light directions. Then, a separate estimation of the diffuse color and the view-dependent effects per surface point is performed. In particular, the diffuse color is computed using statistical operations. Other residual surface light field effects are fitted in terms of a linear combination of spherical functions. The final rendering results show a high fidelity and similarity with the input video frames, without ringing and banding effects.

Further work has focused on the use of compressive sensing techniques for light field acquisition [KGV12, MWBR13]. The technique presented in [MWBR13] allows the reconstruction of high-resolution light fields from a single coded camera image. A dictionary of light field atoms is used which allows a sparse representation of the light field to be measured. The high-resolution light fields can then be reconstructed using nonlinear sparse coding techniques.

7.7.2. Acquisition of Surface Reflectance Fields

An image-based approach that has been presented in [KBMK01] relies on the acquisition of spatially varying reflectance under varying illumination conditions from a single viewpoint. The acquisition of reflectance fields can be approached with reflectance transformation imaging (RTI) as described in [DHT*00, MGW01]. The object is acquired from a single static viewpoint under varying illumination conditions and RTI allows to store per-pixel reflectance functions for the acquired data in a compact representation. This sampling of the reflectance field can be achieved based on acquisition setups such as the "Light Stage" [DHT*00], where a single static camera takes photos of an object that is illuminated using single light sources that are uniformly distributed around the object. If the object is acquired from a single viewpoint and under varying illumination, the respective acquisition setup can be simple and inexpensive as often only a single camera and a single light source are used [MMSL06]. Using this technique, both large-scale objects [DCCS06] and small objects [MVSL05] can be acquired. So far, RTI has been widely applied for relighting surface details [MGW01] as well as in the domain of cultural heritage [MMC*08, PSM06], and in the scope of applications for the movie industry such as the relighting of human faces [DHT*00, WGT*05]. In addition, the resolution during the image-based relighting is identical to the resolution of the used camera which allows a rather accurate depiction of fine surface details under varying illumination conditions and an easy editing of the data based on image operations. Often used editing operations include specular enhancement and diffuse gain [MGW01], contrast enhancement [MWGA06] and further dynamic shading enhancement techniques as proposed in [PCC*10]. If the acquired reflectance functions per pixel are modeled using polynomials, the technique is also named as polynomial texture maps [MGW01]. However, further basis functions such as spherical harmonics and hemi-spherical harmonics might be used as well [MMC*08].

Several other approaches [HED05, FBLS07, WZT*08, AZK08] also focus on the acquisition of the respective material sample from a single viewpoint under varying illumination. In [HED05], the acquisition of reflectance fields has been approached based on the Helmholtz principle [ZBK02] which states that, for many materials, the received radiance value remains the same when the positions of light sources and cameras are exchanged. Translucent objects can be placed inside a dome-like setup with densely sampled light source positions, while a camera is mounted on top of the

dome and acquires the reflectance and transmission characteristics of the material respectively.

The setup in [FBLS07] uses a more sophisticated strategy to allow (far-field) illumination from the full upper hemisphere above the material sample which is achieved based on two spotlight emitters that emit light towards an enclosing tent from where it is reflected into the scene. The material sample is observed in HDR by a single camera at a fixed position.

Similarly, the material sample can also be observed from a single viewpoint under varying illumination [WZT*08]. However, in these investigations, the illumination is generated based on an array of LEDs that can be translated over the material sample. Using this setup, a two-dimensional slice of the BRDF for a single view can be acquired per surface point. Finally, an isotropic micro-facet BRDF model is fitted to the measured normal distribution function of the micro-facets and missing values for different view conditions are obtained by synthesis. This approach allows rather accurate results for flat materials although only a BRDF slice is acquired.

In [AZK08], a surface reflectance field has been acquired by measuring the material sample under varying, known illumination conditions and a fixed viewpoint. Then, the acquired data has been used for the fitting of an approximate reflectance field based on a linear combination of isotropic reflectance functions. In addition, normals can be estimated as well using this approach. Consequently, this approach considers the simultaneous estimation of geometry and reflectance. This allows a more realistic relighting of the material sample under novel view conditions but is only applicable to materials with isotropic surface reflectance. In addition, only surfaces with a very limited number of base materials can be handled.

Based on compressive sensing, a sparse reflectance field acquisition has been proposed in [PD05,PML*09,SD09], where an illumination in terms of controlled noise patterns is used. To fill the data for missing viewing and illumination conditions that have not been taken into account during the acquisition, a fitting of basis functions is used. Unfortunately, the strategy of projecting noise patterns results in a low dynamic range and therefore the reconstruction is rather noisy.

The investigations in [MPN*02] aim at the acquisition of translucent materials under real-world illumination. Parallel measurements have been achieved based on using six static cameras mounted on an arc to observe the translucent object that is placed on a turntable and in front of a screen and illuminated by four point light sources mounted on a rotating arm. Furthermore, a second screen is attached to the turntable. The screens have been used to generate more realistic environments during the acquisition, and both illumination- and view-dependent reflectance fields have been captured.

8. Applications, Novel Trends and Conclusions

In the scope of this section, we discuss some typical applications of geometry and reflectance acquisition techniques (see Section 8.1). Furthermore, we also discuss novel trends in the domain of acquisition (see Section 8.2) and conclude the course notes in Section 8.3.

8.1. Applications

The aforementioned techniques for geometry and reflectance acquisition are of great importance for numerous applications, where digitized 3D content is used e.g. for virtual prototyping, reverse engineering, advertisement and digital preservation. In this part, we focus on several applications where the digitization of objects has opened up new possibilities.

8.1.1. Industrial Applications

Creating digital content has an enormous importance in industry. Many product prototypes need to be visualized in the scope of company-internal discussions regarding product development or even advertised to the envisioned customers before their real counterparts are actually assembled for the first time. This raises the demand for a high-quality visualization of the virtual surrogates to produce images that convince customers to order the products.

In car industry, digitized geometry and material models have already reached maturity in the scope of the product development chains. The digitized models are typically used for a virtual prototyping of novel products where a variety of construction and design alternatives can easily be tested. Therefore, the product development period can be reduced.

While virtual prototyping has found its way into the car industry already several years ago, a more recent trend can be recognized in the several upcoming virtual try-on solutions that are offered for advertising fashion products and allow the customer to explore how a dress might look on a virtual avatar that has the body size of the customer or how certain eyewear appears on the customer's face. Similarly, based on a scan of the customer's head, different hair cuts might be simulated to support the customer in his choice.

The availability of digitized real-world 3D content also offers new possibilities for the movie and video game industry, where there has always been the need for 3D content. Today, many of the modeling aspects are still performed by designers and artists such as e.g. the creation of realistic scenes and the specification of the lighting conditions. The availability of accurately digitized real-world models or maybe even only digitized materials may facilitate their work as they might perform the necessary editing based on such *templates*.

8.1.2. Digitization of Cultural Heritage Artifacts

The progress towards accurate digitization techniques has gained particular attention in the field of cultural heritage where the virtual models of accurately digitized artifacts might serve as surrogates for their physical counterparts. However, this also means that the focus is placed on an as-accurate-as-possible appearance reproduction of an artifact including subtle details regarding both its shape and its reflectance behavior. The availability of accurate virtual surrogates not only offers the possibility to make cultural heritage artifacts easily, i.e. simultaneously and location-independently, accessible to a broad audience of researchers, curators, teachers or lecturers and even the general public, but also to prevent fragile artifacts from being damaged due to inappropriate use. An easy inspection might be possible based on e.g. disseminating objects via the internet, letting the user manipulate both the pose of the object and the

scene illumination as presented in e.g. [SRWK11]. This interactive artifact exploration obviously transports a better user experience than simple photographs as it probably closely follows our natural way of inspecting objects and allows us to perceive fine details such as scratches or wear that are revealed only under certain view-light conditions and contribute to a deeper understanding regarding the manufacturing process of objects and the used materials. Furthermore, the respective artifacts can easily be showcased in synthetic scenes that model different contexts such as excavation sites or reconstructed historical sites to even increase the user experience. In addition, exchanges of artifacts between museums might be facilitated as the states of the objects of interest might be documented by digitizing them before they are lent and after they are returned to allow a proper documentation. Similarly, other time-varying aspects might be documented as well such as appearance changes due to restoration or decay.

The examples of digitized artifacts that are shown in the following are obtained by applying the techniques for geometry and reflectance acquisition described in [WSRK11, SWRK11, SSW*14]. Such artifacts might include fossils that reveal insights about extinct species that populated the earth millions of years ago. Figure 32 shows exemplary reconstructions of some ammonites and the holotype of a Scaphognathus crassirostris. As shown for the holotype of a Scaphognathus crassirostris, the accurate acquisition of both geometric details and details in surface reflectance allows to reproduce soft tissue accurately within the virtual surrogate. But also local subsurface scattering characteristics can be reliably reproduced as demonstrated for one of the shown ammonites.

The second example shows moulages with their corresponding virtual surrogates. Moulages have been used from the 17th to the early 20th century in order to showcase certain diseases in terms of anatomic wax models that include even fine details such as scurf or hairs and been used to educate medical staff. The advantage of having such real objects for depicting a certain disease in comparison to a photograph is given by the fact that the 3D model can be interactively explored to reveal fine details that are only visible in particular view-light configurations and allow to better understand the disease and its current state for diagnostics. Unfortunately, moulages are rather fragile objects as most of their parts are made of wax. This makes them not only sensitive to variations in temperature and air humidity but also to too much light exposure or too much agitation. This severely limits the interaction possibilities with such objects in the scope of lectures where they are still used nowadays, e.g. at the University of Zurich in Switzerland. Being manufactured manually from the materials that had been available at the respective time, moulages also have a strong connection to art history [Gei09]. While several reflectance models have been studied for the realistic virtual modeling of human skin [PB11], the investigations in [SSWK13] demonstrate that BTFs also allow to faithfully model diseased skin as both the local subsurface scattering in the wax and skin details such as pores, scars or hairs can faithfully be reproduced. Similarly, many other cultural heritage artifacts may be digitized. An example of a shiny figurine and the corresponding digitized model are shown in Figure 34.

8.1.3. Art

While the moulages mentioned in the previous section already have some connection to arts, the access to the many digitized objects offers significantly more opportunities. Regarding souvenirs, characteristics sites such as famous buildings are often represented in different ways ranging from manufactured models, that might be based on a 3D printed base model, to the depiction of these objects in terms of a point cloud that is burnt inside a block of glass. In a similar way, the data obtained by scanning people can be either printed using 3D color printers to obtain 3D miniatures of the scanned persons [SBKC13] or burnt inside a block of glass. The capabilities of 3D printing might also allow a fast access to printing certain *templates* for parts of objects or avatars that might be further edited manually by designers. This might improve the creativity as such base parts can be easily created in an automatic way.

8.1.4. Medical Applications

Even in medical applications, there is a trend of including more and more information based on an adequate acquisition. The accurate acquisition of the body geometry including the size of the individual parts and the positions of joints can be used to make a spine analysis or a diagnosis regarding postural deformity. Furthermore, acquisition might help in several applications such as orthopedics, dental prosthetics plastic surgery or dermatology. For instance, 3D printing might open the possibility towards obtaining tooth surrogates or surrogates for prostheses that exactly fit to the requirements of an individual person. In addition, the scanning techniques allow the diagnosis and analysis of skin diseases.

8.2. Novel Trends in Acquisition

While a huge multitude of techniques for both geometry and reflectance acquisition have been developed so far, current research focuses in particular on the following aspects:

- the development of novel methods for the acquisition of objects with a complicated surface reflectance behavior (see Section 8.2.1)
- the development of techniques to increase the efficiency of such automated acquisition pipelines (see Section 8.2.2)

8.2.1. Novel Methods for Acquiring Objects with Challenging Optical Properties

While numerous robust techniques have been developed for the acquisition of diffuse to glossy objects, the acquisition of objects with a more complex surface reflectance behavior is still challenging. While the geometry acquisition of mirroring surfaces, surfaces that exhibit subsurface scattering or transparent surfaces of glass already shows some rather promising results, one of the key challenges is given by strongly inhomogeneous materials with significantly different surface characteristics. For instance, a translucent or transparent object with spatially varying refractive indices might additionally contain inclusions of other materials with a different type of reflectance. In this scenario, it might be interesting to not only perform a 3D reconstruction of the outer surface but also a reconstruction of the geometry of the inclusions. Furthermore, a

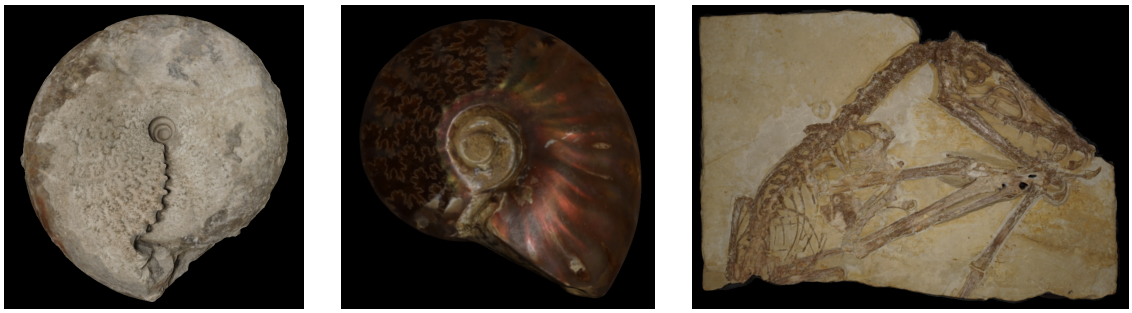

Figure 32: Several virtual surrogates of digitized fossils.

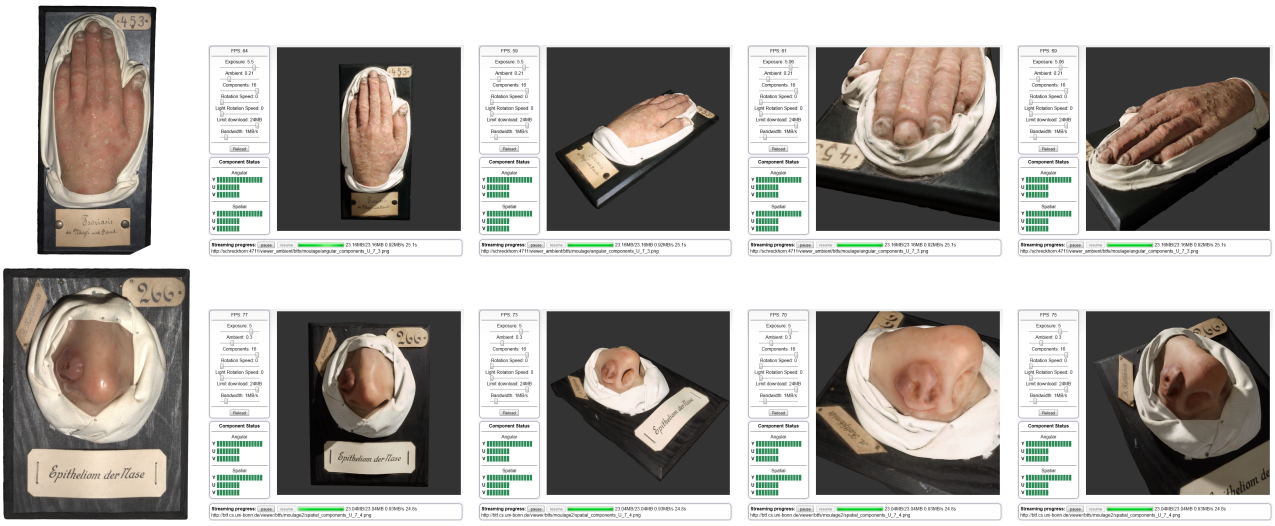

Figure 33: Digitized moulages visualized in a web-based viewer [SRWK13]. Image taken from [SRWK13].

Figure 34: A Buddha figurine (left) and its virtual surrogate (right).

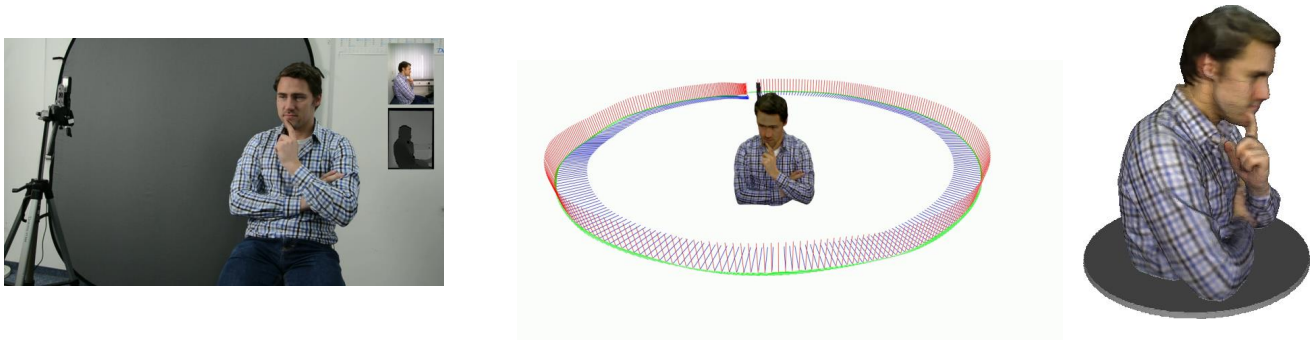

Figure 35: Generation of 3D miniatures of scanned persons [SBKC13]. The reconstructed model can be printed using 3D color printers as demonstrated in [SBKC13]. Image courtesy of Jürgen Sturm (FabliTec UG).

current trend can be seen in the development of "in the wild" solutions that allow to acquire the geometry with cheap, possibly mobile hardware under uncontrolled conditions regarding viewpoint and illumination conditions. However, this means that fast algorithms are needed that allow the direct processing of the data on the possibly mobile hardware or to rely on cloud processing. In the latter case, the data has to be transmitted onto a cluster where it is processed and the result is streamed to the mobile device. In the context of diffuse surfaces, respective techniques have long reached maturity. However, more research has to be performed in order to also handle more complex materials.

In the context of appearance acquisition, there is also a trend towards acquisition based on light-weight setups with either customer hardware such as flatbed-scanners or mobile hardware. In addition, modeling of inhomogeneous subsurface scattering is still one of the main problems in reflectance acquisition. The acquisition of such objects with a more complex reflectance behavior, where a dense sampling in a high-dimensional parameter space has to be taken into account, is still rather impractical but might benefit from the developments regarding the hardware components that can be used in the acquisition setup and for the processing of the data. In addition, reflectance acquisition for translucent or transparent objects with spatially varying characteristics and inclusions of other materials within the outer object geometry is also challenging.

Furthermore, mass digitization of whole object collections might necessitate a significant reduction of the manual interaction. Most of the techniques rely on a positioning of the object in the acquisition setup by the acquisition expert. In order to increase the efficiency of the acquisition, an automatic supply might allow to include an automatic object placement in addition to an automatic acquisition as e.g. presented in [SRT*14].

8.2.2. Increasing the Efficiency of Automated Acquisition Pipelines

As discussed in the scope of Section 5 and Section 7, there is a huge multitude of acquisition techniques where each of the individual methods is typically tailored to only a certain type of surface reflectance behavior. Consequently, today's acquisition methods still meet their limitations when prior knowledge about the surface reflectance behavior of the considered object is not available.

In Section 5, we have seen that e.g. in the context of 3D geometry reconstruction the respective techniques typically depend on some basic assumptions about material reflectance. As a consequence, the developed techniques typically consider only a fraction of the possible surface materials and are – by design – not capable of handling arbitrary surface reflectance. Multi-view stereo techniques and photometric stereo techniques are typically based on assuming Lambertian reflectance behavior and only a few extensions allow the acquisition of objects with a more complex reflectance behavior. Furthermore, structured light systems are well-suited for the geometry acquisition of objects with a reflectance behavior ranging from diffuse to even specular as long as a sufficient diffuse reflectance component is present. Further reconstruction techniques are tailored to mirroring surfaces, translucent surfaces or refractive surfaces.

In a similar way, the different acquisition strategies based on the different reflectance models are also tailored to different fractions of the individual types of surface reflectance. While smooth, homogeneous materials can be represented well with analytical BRDF models as their appearance behavior mainly depends on the direction of the incoming light and the view direction, materials exhibiting mesoscopic effects of light exchange on surface structures imaged to a size of approximately one pixel cannot be modeled by using simple BRDF models. Instead data-driven BTFs might be used in such cases as they consider the spatial material variations in addition to the view direction and the direction of the incoming light.

In many cases, the materials of the objects to be acquired are not known a priori or objects might be made of several materials with different types of surface reflectance, and still an efficient-as-possible technique might be desirable. Consequently, the automatic selection of appropriate acquisition methods that provide reliable information represents an important prerequisite. One simple strategy might be the naïve application of several different techniques and a subsequent merging of their individual results. This requires automatically derived information regarding which of the individual techniques is suited best for a certain part of the object surface. However, in many cases, simply applying all the available acquisition techniques is highly inefficient regarding acquisition time, and hardware components are stressed unnecessarily as many of

the taken images do not have an influence on the final reconstruction and, hence, have to be neglected. For more efficient geometry and reflectance acquisition procedures in case of missing information about the material properties, it might therefore be desirable to automatically select only the appropriate techniques instead of applying several different methods.

In the scope of this section, we will more closely discuss the following approaches towards more efficient acquisition pipelines:

- A theoretical concept for combining different 3D shape acquisition techniques to an automatic, efficient acquisition system where the local consistencies obtained with the individually applied methods are explored in order to select the locally most suitable acquisition method. The resulting conceptual framework is based on establishing consistencies using cues that are individual to different acquisition techniques. Once such different consistencies have been extracted, they can be integrated into the same efficient optimization framework to get a 3D reconstruction for the considered object.
- A novel automatic, efficient geometry and reflectance acquisition framework, in which a highly accurate material recognition step is used to select the required acquisition techniques for both geometry and reflectance.

8.2.2.1. Combinations of Multiple Techniques for 3D Shape Acquisition of Objects With Unknown Surface Reflectance

Knowing about the variety of developed 3D geometry acquisition techniques, where each of them is capable of handling different types of surface reflectance, one obvious question seems to be whether these different methods can be combined in a single automatic acquisition pipeline. Such a pipeline would not only allow the acquisition of objects consisting of a single a-priori unknown material but also allow the acquisition of objects with heterogeneous surface reflectance behavior where different acquisition techniques have to be used for the individual parts made of different materials. Achieving such a combination of methods seems to be much easier if the respective techniques are based on rather similar approaches or similar representations. In the scope of this section, we provide a discussion of concepts for combining different methods that are tailored to different types of surface reflectance.

As described in Section 5, the geometry information reconstructed via structured light systems can be transformed to a corresponding volumetric consistency and, hence, to a volumetric representation. Furthermore, representations in terms of volumetric consistencies can be derived by different methods via photometric stereo, photo-consistency or shape-from-coded illumination. This means that possibly rather similar or even the same optimization frameworks might be used for such reconstruction methods. Therefore, the task of combining these methods might be approached by combining the individual consistency measures derived per point in the volume using each of the acquisition techniques.

For instance, the structured light information and the Helmholtz normal consistency information measured as described in [WRO*12] might be complemented with the normal consistency information for mirroring objects as discussed in [WORK13] by adapting the energy functional. Another possibility to combine methods for Lambertian and mirroring objects could be the fusion of the geometry acquisition techniques based on multi-view normal field integration techniques. In fact, combining the normal information obtained from photometric stereo measurements and the normal information derived for mirroring objects following the technique presented in [WORK13] leads to a purely normal-based geometry reconstruction process without the need for relying on other constraints given e.g. by silhouettes. For instance, the object depicted in Figure 36 exhibits heterogeneous surface reflectance characteristics as it has both rather matte and almost ideally mirroring parts. Obviously, normal hypotheses obtained via photometric stereo or Helmholtz stereopsis exhibit a rather small variance at points close to the surface whereas normal hypotheses derived via the technique for mirroring surfaces will typically not be consistent at the true surface for the rather diffuse parts. In contrast, on mirroring parts, the observation will be different as the normal hypotheses derived via the technique for mirroring surfaces will typically show a highly consistent, dominant mode whereas the normal hypotheses obtained via photometric stereo or Helmholtz stereopsis will not form a dominant cluster but show a large variance at the true surface. While both of these types of normal information are separately used to acquire the geometry of Lambertian and mirroring surfaces, their combination to handle inhomogeneous objects can be approached by simply combining the local normal consistencies obtained for the individual methods in the volume. In order to make a robust reconstruction possible, an adequate combination of the individual consistencies per volumetric point could be performed by e.g. taking only the normal information with the higher consistency. Consequently, it is not only possible to identify the local material type and the adequate acquisition technique based on the consistencies obtained for the individual methods, but also to select the local surface normal computed via the respective technique. Subsequently, based on the local normal estimates and their respective consistency estimate resulting after the combination of the individual methods, the same min-cut-based optimization framework as discussed in [WORK13] can directly be applied for the final reconstruction. Therefore, in addition to providing information regarding material properties inferred from the consistencies, combining these methods would allow the reconstruction of the full 3D shape for the range from diffuse to mirroring objects with complex surface geometry.

Other investigations also derive normal information for specular objects (e.g. [BS03, CGS06, YIX07, NWR08, FCM*08, BHB11]), transparent objects (e.g. [MK05, KS05, MK07, YIX07, KS08, YWT*11]) or translucent objects (e.g. [CGS06, DMZP14, IMMY14]), and this normal information could probably also be integrated into a geometry acquisition framework. Particularly the integration of volumetric normal information for translucent objects or glass objects and the fusion of the respective consistencies might even improve the range of material types for which the shape of the respective objects can reliably be acquired.

Unfortunately, the aforementioned concepts for combining the individual methods require full measurements of the involved objects with all the different acquisition techniques that are available for the individual material types as illustrated in the naïve, automatic acquisition pipeline in Figure 5. As a result, the total acqui-

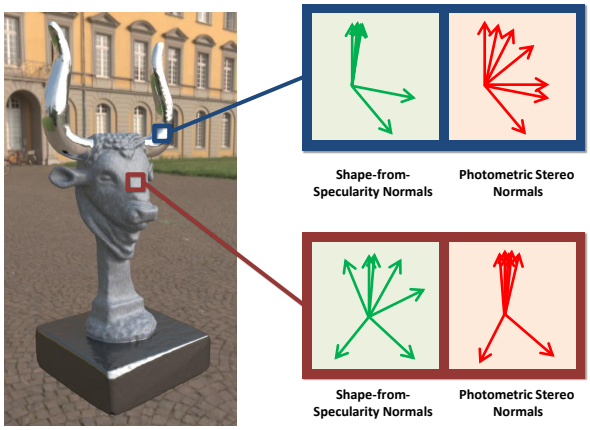

Figure 36: Normal consistency for different acquisition methods: The shown object consists of rather matte and mirroring parts. Normal hypotheses obtained via photometric stereo or Helmholtz stereopsis (red) show a small variance at surface points on the rather matte parts, whereas normal hypotheses derived via techniques for mirroring surfaces (green) will not be consistent. For points on the mirroring parts, the normal hypotheses derived via techniques for mirroring surfaces (green) form a dominant mode with a small variance, whereas the normal hypotheses obtained via photometric stereo or Helmholtz stereopsis (red) are not consistent on the surface.

sition time t_Σ can be computed as the sum of the acquisition times t_i with $i = 1,\ldots,N$ needed for the N individual measurements according to

$$t_\Sigma = \sum_{i=1}^{N} t_i. \tag{63}$$

Furthermore, hardware components might also be stressed unnecessarily as some of the measurements are not required for the final reconstruction and some of the computational effort is not required as well. This renders such an acquisition strategy rather impractical. As we would favor a smart acquisition process only involving the required effort regarding measurement time and hardware usage, our investigations in the scope of the next section will focus on the information that can be inferred visually from images of a material surface and show whether this information can be used for a more efficient acquisition.

8.2.2.2. An Efficient Acquisition Pipeline Based on an Initial Material Recognition

Having adequate acquisition techniques for different surface types, the question arises whether a more efficient acquisition can be realized if there is no prior knowledge available regarding the surface materials of the object to be measured. As pointed out in [Rui13], cost-efficient measurements of complex objects or materials will probably necessitate the utilization of some kind of prior information. However, an open question is whether typical assumptions such as analytical BRDF models, smoothness, low-rank, etc. provide the best priors or whether data-driven priors represent a superior approach. In the scope of this section, we discuss a concept for a more efficient acquisition of geometry and reflectance. In particular, we demonstrate how different of the aforementioned techniques can be composed to an efficient, fully automatic acquisition framework where the required acquisition techniques are selected based on a prior material recognition (see Figure 37). In order to let the involved material recognition system act as an automatic assistance system for guiding the subsequent acquisition process in such a way that only the required acquisition methods are involved, the approach presented in [WK15a] relies on using a-priori information in the form of a database of material measurements with additional annotations regarding the methods that have to be chosen for the acquisition of geometry and reflectance. The key idea is the classification of a measured material based on only a rather small set of photos and, depending on the annotations of the closest match in the database, corresponding methods can easily be determined. For an almost mirroring metal, for instance, a shape-from-specularity approach could be used for geometry reconstruction and a BRDF measurement could be started to capture surface reflectance properties. If the considered material sample is classified as a material with strong mesoscopic effects as e.g. given in leather, a method adequate for glossy materials such as the technique described in [WRO*12] could be used in combination with a BTF acquisition. As a result, a significant amount of acquisition and processing time can be saved and the hardware components are less stressed as e.g. less images need to be taken. The crucial prerequisite of such a strategy is the availability of a reliable material recognition framework and the availability of a database of materials with annotations regarding their adequate acquisition processes as well as the availability of acquisition methods that are suitable for the different surface materials. Hence, one core component of the framework can be identified in a material database which contains images of a multitude of material samples taken under different viewing and lighting conditions, which are expected to be met during the acquisition with standard devices. For material recognition, a huge variety of methods has been proposed in the past. For a respective survey on material recognition, we refer to [WK15b]. As typically information from several viewpoints might be acquired easily by typical acquisition setups, the multi-view material recognition approach presented in [WK15a] might be a good choice for a rather accurate material recognition.

In order to get an impression about the amount of time that can be saved with an efficient automatic acquisition system based on a smart selection of acquisition techniques for the involved materials, we provide a short discussion with respect to the acquisition times. The structured light based geometry acquisition technique described in [WSRK11] requires images that depict observed stripe patterns as projected onto the object surface by the involved projectors. The respective acquisition typically takes approximately 1.5 hours when using the highly parallel setup presented in [SWRK11] and in the range from approximately 1.5 hours to 3.0 hours in the turntable-based acquisition device introduced in [SSWK13]. In contrast, the geometry acquisition method introduced in [WRO*12] requires a full structured light measurement as well as a full reflectance measurement as input. Therefore, the acquisition time is

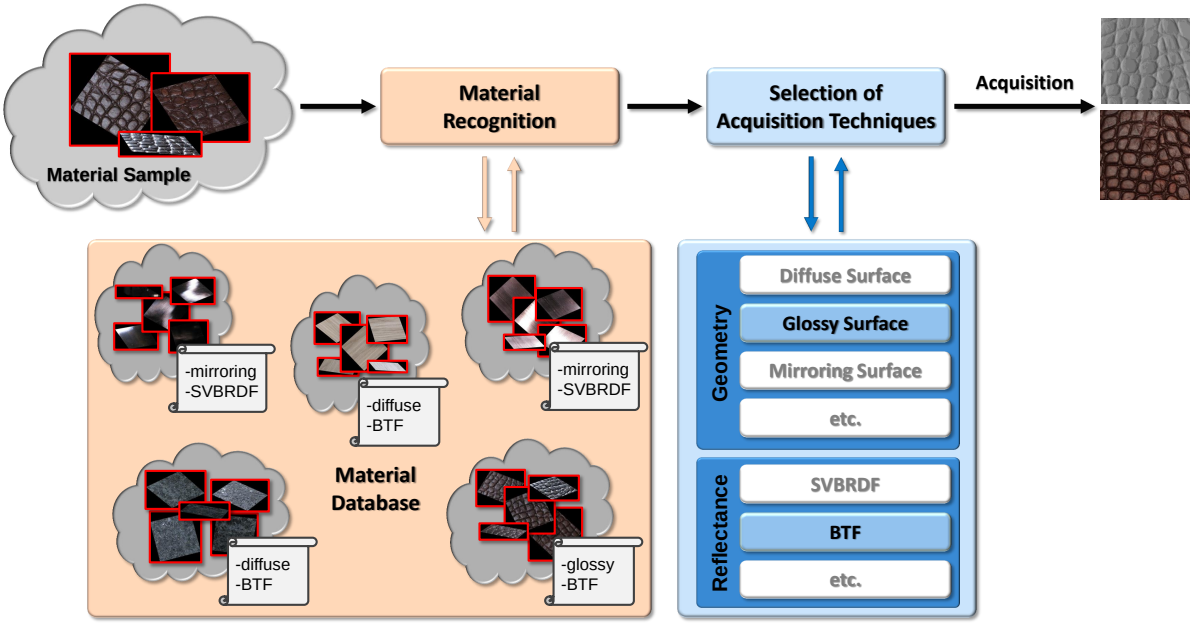

Figure 37: Efficient acquisition pipeline: Based on a few input images of the material that should be acquired, an initial material recognition is used to find adequate acquisition parameters in a database of reference materials with respective annotations. Subsequently, these attached annotations for the closest material in the database can be used to guide the acquisition process by selecting the adequate acquisition techniques.

significantly longer as the reflectance acquisition usually requires approximately 4 hours to 10 hours depending on the complexity of the surface materials. Furthermore, the geometry acquisition using the method presented in [WORK13] takes approximately 1.5 hours to 2.0 hours when using illuminations from two to three different positions of the involved screens. When analyzing the required processing times, the higher complexity of the consistency-based techniques in [WRO*12, WORK13] becomes apparent in processing times of approximately 8 hours to 12 hours on a workstation with two Intel Xeon 5645 CPUs with 2.4 GHz or approximately 4 hours to 5 hours on a workstation with two Intel Xeon E5 – 2650 CPUs with 2.0 GHz. In contrast, the triangulations required for the structured light approach in [WSRK11] and the subsequent Poisson reconstruction are carried out in approximately 1.5 hours.

Furthermore, the acquisition time as well as the processing time required for the reflectance reconstruction also strongly depend on the respective material. For materials following simple homogeneous BRDF models, only four parameters have to be measured, which might be performed in only a few minutes. In contrast, materials with e.g. a spatially varying reflectance behavior require the measurement of six parameters in order to get a SVBRDF or BTF representation. In addition, the sampling density of the involved view-light configurations that need to be measured also needs to be taken into account. For dense measurements of several thousands of view-light configurations as used in e.g. [SWRK11, SSWK13], even fine details of mesoscopic reflectance can be adequately captured. However, even when acquiring the computationally demanding BTF, the acquisition and processing times might vary significantly depending on the complexity of the involved materials, which influences the acquisition parameters such as the number of exposures used during the measurement. This can be seen e.g. in the evaluation given in [Sch14], where the acquisition parameters and the processing parameters for different objects are given as listed in Figure 38, Figure 39 and Figure 40.

As a result, the automatic selection of appropriate geometry and reflectance acquisition techniques with the respective postprocessing methods might save several hours.

8.2.3. Acquisition *in the Wild*

A further trend can be seen in the development of *in the wild* solutions that allow to acquire the geometry with cheap consumer hardware under uncontrolled conditions regarding viewpoint and illumination conditions. In the context of diffuse surfaces, respective techniques have reached a very mature state. Multi-view stereo algorithms like Goesele et al. [GSC*07] and Furukawa et al. [FCSS10] can handle images that have been downloaded from large internet databases such as Flickr. These images are usually captured from vastly different viewpoints and under various lighting conditions. Fuhrmann et al. [FLG14] used the approach from Goesele et al. together with their surface reconstruction (see Section 6.2.3) to build a complete reconstruction pipeline which is available as an open source software called *MVE*. This software can create high quality geometry reconstructions from a simple unordered set of input images. Such a front-to-end pipeline was not available before in a single open source software solution. As the system is built with very robust components, it is able handle a variety of input

Donkey[1,2]
specular brass
3/4[1] or 2/3[2]
2:49[1] / 1:55[1] or 3:18[2] / 4:50[2]
12:35[1] or 14:38[2]

Minotaur[1]
bronze, green paint, marble
4/4
2:45 / 2:56
10:23

Terracotta Soldier[1]
black terracotta
2/2
1:45 / 1:03
15:17

Buddha[1]
red paint, gold leaf, wood
3/3
2:41 / 1:30
22:55

Strawberry[1]
strawberry skin and leafs
2/2
0:46 / 0:57
8:26

Pudding Pastry[1]
pastry, sugar-coating, vanilla pudding
3/3
1:12 / 1:21
31:58

Apple[1]
apple skin and flesh
3/3
1:14 / 1:21
9:33

Almond Horn[1]
almonds, pastry, chocolate
3/4
1:12 / 2:14
15:57

Crispy Roast Pork[1]
pork, bacon, crust
3/3
1:00 / 1:22
26:38

Billiard Ball[2]
red, black and white phenolic
?[3]
?[3]
?[4]

Santa[2]
mixed glossy paints
4/4
3:59 / 6:10
?[4]

Psoriasis Moulage[2]
wax, paint, fabric, lacquered wood, paper
3/3
3:07 / 7:02
34:51

Figure 38: Objects digitized in the scope of [Sch14]: Listed are the respective apparent materials of the individual objects, the number of used exposures (geometry acquisition / reflectance acquisition), the acquisition times (geometry acquisition / reflectance acquisition) as well as the total processing times ([1]Measured in the Dome 1 device, [2]Measured in the Dome 2 device, [3]Original measurement data damaged, [4]No information available due to data loss).

Chess Piece[1,2]
resin, matte white finish
1/3[1] or ?[2,3]
0:26[1] / 1:29[1] or ?[2,3]
12:52[1] or ?[2,4]

Tennis Ball[2]
synthetic fabric (fluorescent)
4/3
1:10 / 7:08
8:56

Shoe[2]
synthetic fabric, rubber, plastic
?[3]
?[3]
?[2,4]

Mug[2]
ceramics
2/3
0:50 / 2:52
28:57

Ganesha[2]
labradorite
3/4
1:47 / 9:12
22:39

Paintbrush[2]
lacquered wood, metal, bristles
3/5
1:01 / 8:27
10:18

Micrometer[2]
polished and rough metal, plastic
3/5
2:04 / 16:21
10:31

Fish[2]
gold and red paint
3/4
1:44 / 7:34
36:12

Inkwell[2]
silver
2/4
0:51 / 5:25
10:41

Teal[2]
feathers, beak, green paint
2/2
1:13 / 4:22
74:48

Epithelioma Moulage[2]
wax, paint, fabric, lacquered wood, paper
3/4
5:21 / 13:20
22:12

Pyramid[2]
clay, patina
1/2
1:39 / 3:01
32:43

Figure 39: Objects digitized in the scope of [Sch14]: Listed are the respective apparent materials of the individual objects, the number of used exposures (geometry acquisition / reflectance acquisition), the acquisition times (geometry acquisition / reflectance acquisition) as well as the total processing times ([1]Measured in the Dome 1 device, [2]Measured in the Dome 2 device, [3]Original measurement data damaged, [4]No information available due to data loss).

Ammonite 1[2]
fossil
2/3
1:04 / 6:24
19:36

Ammonite 2[2]
ammolite
3/5
3:20 / 20:13
9:28

Rhinoceros Teeth[2]
fossil
2/3
2:19 / 4:57
18:10

Figure 40: Objects digitized in the scope of [Sch14]: Listed are the respective apparent materials of the individual objects, the number of used exposures (geometry acquisition / reflectance acquisition), the acquisition times (geometry acquisition / reflectance acquisition) as well as the total processing times ([1]Measured in the Dome 1 device, [2]Measured in the Dome 2 device, [3]Original measurement data damaged, [4]No information available due to data loss).

Figure 41: Reconstruction by MVE [FLG14] of an Internet dataset downloaded from Flickr. The bottom row shows 3 images of a total of 871 input images. The top row shows the reconstruction rendered with shading (left) and with texture (right). Photo credits: Flickr users *Vince O'Sullivan, Andy Hay, Ecyrd,* Creative Commons License

data. Optionally, the recovered geometry can also be textured by an appropriate method such as Waechter et al. [WMG14]. Figure 41 shows an example reconstruction from the *Trevi Fountain* in Rome. For this dataset more than 800 images have been downloaded and processed in about 12 hours. Another approach that utilizes internet data was presented by [ALFG12]. Building on a photometric stereo technique their algorithm can recover geometry and a sparse set of basis materials in an outdoor scene that is captured by a static camera, e.g., a webcam, over a long period of time. Using the sun as a varying light source they calculate the direction of the incident illumination from the time and date of each image. To get enough variation in these directions the images have to be captured over the course of at least 6 months. During this time frame, tens of thousands of images are captured. As a reconstruction from such a large number of images is rather impractical because of computation time and memory requirements, an appropriate image subset is typically selected first. With a special image selection strategy the algorithm automatically selects images with a clear sky that are distributed over the whole time period. These images are then aligned with subpixel accuracy to account for small camera movements and are then used for a photometric reconstruction. The normals are first initialized with a basic photometric stereo algorithm and all pixels are clustered into a small amount of material subsets according to color. Next, an iterative optimization recovers surface normals as well as a mixture of a set of basis materials for each pixel in the image. This is done by alternating two optimization schemes. One uses an isotropic Ward reflection model [War92] and recovers materials for the current fixed distribution of material weights and normals. The other uses the fixed set of optimized materials and updates the normals and material distribution for each pixel. The final result is a normal map of the scene and mixtures of materials that can be used for relighting or further processing.

8.3. Conclusions

In the scope of this course, we have given a thorough overview on the acquisition ecosystem with its components. In particular, we have reviewed the advances in both geometry acquisition and reflectance acquisition for objects with different types of reflectance behavior ranging from diffuse over opaque to specular surfaces or even translucent and transparent surfaces.

Acknowledgements

This work was partially funded by the European Commission's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 323567 (Harvest4D).

References

[AB91] ADELSON E. H., BERGEN J. R.: The plenoptic function and the elements of early vision. In *Computational Models of Visual Processing*. 1991, pp. 3–20. 10

[AEB*12] AUBRETON O., EREN G., BOKHABRINE Y., BAJARD A., TRUCHETET F.: Estimation of surface normal vectors based on 3d scanning from heating approach. In *Proceedings of SPIE* (2012), vol. 8290, pp. 82900X–1 – 82900X–6. 29

[AFG13] ACKERMANN J., FUHRMANN S., GOESELE M.: Geometric point light source calibration. In *Proceedings of Vision, Modeling and Visualization (VMV)* (2013), pp. 161–168. 14

[ALFG12] ACKERMANN J., LANGGUTH F., FUHRMANN S., GOESELE M.: Photometric stereo for outdoor webcams. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2012), pp. 262–269. 60

[Alo88] ALOIMONOS J.: Shape from texture. *Biological Cybernetics 58* (1988), 345–360. 16

[AMKB04] AGARWAL S., MALLICK S. P., KRIEGMAN D., BELONGIE S.: On refractive optical flow. In *Proceedings of the European Conference on Computer Vision (ECCV)* (2004), pp. 483–494. 28

[APK08] ALLÈNE C., PONS J.-P., KERIVEN R.: Seamless image-based texture atlases using multi-band blending. In *Proceedings of the International Conference on Pattern Recognition (ICPR)* (2008), pp. 1–4. 35, 36

[AS00] ASHIKHMIN M., SHIRLEY P.: An anisotropic phong brdf model. *Journal of Graphics Tools 5*, 2 (2000), 25–32. 37

[ASOS13] ALI M. A., SATO I., OKABE T., SATO Y.: Toward efficient acquisition of brdfs with fewer samples. In *Proceedings of the Asian Conference on Computer Vision (ACCV)* (2013), vol. 4, pp. 54–67. 37

[ASS*09] AGARWAL S., SNAVELY N., SIMON I., SEITZ S. M., SZELISKI R.: Building Rome in a day. In *Proceedings of the International Conference on Computer Vision (ICCV)* (2009), pp. 72–79. 35

[AVBSZ07] ADATO Y., VASILYEV Y., BEN-SHAHAR O., ZICKLER T.: Towards a theory of shape from specular flow. In *Proceedings of the International Conference on Computer Vision (ICCV)* (2007), pp. 1–8. 25

[AWL13] AITTALA M., WEYRICH T., LEHTINEN J.: Practical svbrdf capture in the frequency domain. *ACM Transactions on Graphics 32*, 4 (2013), 110:1–110:12. 42

[AWL15] AITTALA M., WEYRICH T., LEHTINEN J.: Two-shot svbrdf capture for stationary materials. *ACM Transactions on Graphics 34*, 4 (2015), 110:1–110:13. 42

[AX08] ALIAGA D. G., XU Y.: Photogeometric structured light: A self-calibrating and multi-viewpoint framework for accurate 3D modeling. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2008), pp. 1–8. 21

[AX10] ALIAGA D. G., XU Y.: A self-calibrating method for photogeometric acquisition of 3D objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 32*, 4 (2010), 747–754. 21, 23

[AZK08] ALLDRIN N., ZICKLER T., KRIEGMAN D.: Photometric stereo with non-parametric and spatially-varying reflectance. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2008), pp. 1–8. 50, 51

[BAG12] BELJAN M., ACKERMANN J., GOESELE M.: Consensus multi-view photometric stereo. In *Pattern Recognition* (2012), vol. 7476 of *Lecture Notes in Computer Science*, pp. 287–296. 19, 27

[Bau02] BAUMBERG A.: Blending images for texturing 3d models. In *Proceedings of the British Machine Vision Conference (BMVC)* (2002), pp. 38.1–38.10. 36

[BCJS06] BIRKBECK N., COBZAS D., JAGERSAND M., STURM P.: Variational shape and reflectance estimation under changing light and viewpoints. In *Proceedings of the European Conference on Computer Vision (ECCV)* (2006), pp. 536–549. 23

[BCNR10] BAI J., CHANDRAKER M., NG T.-T., RAMAMOORTHI R.: A dual theory of inverse and forward light transport. In *Proceedings of the European Conference on Computer Vision (ECCV)* (2010), pp. 1–8. 50

[BEN03] BEN-EZRA M., NAYAR S. K.: What does motion reveal about transparency? In *Proceedings of the International Conference on Computer Vision (ICCV)* (2003), vol. 2, pp. 1025–1032. doi:10.1109/ICCV.2003.1238462. 28

[BETvG08] BAY H., ESS A., TUYTELAARS T., VAN GOOL L.: Speeded-up robust features (SURF). *Computer Vision and Image Understanding 110*, 3 (2008), 346–359. 13, 18

[BEWW*08] BEN-EZRA M., WANG J., WILBURN B., LI X., MA L.: An led-only brdf measurement device. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2008), pp. 1–8. 40

[BFI*14] BHANDARI A., FEIGIN M., IZADI S., RHEMANN C., SCHMIDT M., RASKAR R.: Resolving multipath interference in Kinect: An inverse problem approach. In *Sensors* (2014), pp. 614–617. 20

[BHB11] BALZER J., HOLER S., BEYERER J.: Multiview specular stereo reconstruction of large mirror surfaces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2011), pp. 2537–2544. 26, 27, 55

[BJK07] BASRI R., JACOBS D., KEMELMACHER I.: Photometric stereo with general, unknown lighting. *International Journal of Computer Vision (IJCV) 72*, 3 (2007), 239–257. 19

[BKBH07] BOLITHO M., KAZHDAN M., BURNS R., HOPPE H.: Multilevel streaming for out-of-core surface reconstruction. In *Proceedings of the Fifth Eurographics Symposium on Geometry Processing (SGP)* (2007), pp. 69–78. 32

[BKBH09] BOLITHO M., KAZHDAN M., BURNS R., HOPPE H.: *Advances in Visual Computing: 5th International Symposium, ISVC 2009, Proceedings, Part I*. Springer Berlin Heidelberg, 2009, ch. Parallel Poisson Surface Reconstruction, pp. 678–689. 32

[Bli77] BLINN J. F.: Models of light reflection for computer synthesized pictures. In *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)* (1977), pp. 192–198. 37

[BM97] BEECROFT M. T., MATTISON P. R.: Design review of an in-situ bidirectional reflectometer. In *Proceedings of SPIE* (1997), vol. 3141, pp. 196–208. 37

[BS87] BECKMANN P., SPIZZICHINO A.: *The Scattering of Electromagnetic Waves from Rough Surfaces*. Artech House Publishers, 1987. 25

[BS03] BONFORT T., STURM P.: Voxel carving for specular surfaces. In *Proceedings of the International Conference on Computer Vision (ICCV)* (2003), pp. 591–596. 26, 27, 55

[BW10] BALZER J., WERLING S.: Principles of shape from specular reflection. *Measurement 43* (2010), 1305–1317. 25

[BY84] BRADY M., YUILLE A.: An extremum principle for shape from contour. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 6* (1984), 288–301. 16

[CCCS08] CALLIERI M., CIGNONI P., CORSINI M., SCOPIGNO R.: Masked photo blending: Mapping dense photographic data set on high-resolution sampled 3d models. *Computers & Graphics 32*, 4 (2008), 464–473. 36

[CDMR04] CULA O., DANA K., MURPHY F., RAO B.: Bidirectional imaging and modeling of skin texture. *IEEE Transactions on Biomedical Engineering 51*, 12 (2004), 2148–2159. 45

[CDP*14] CHEN G., DONG Y., PEERS P., ZHANG J., TONG X.: Reflectance scanning: Estimating shading frame and brdf with generalized linear light sources. *ACM Transactions on Graphics 33*, 4 (2014), 117:1–117:11. 42

[CGS06] CHEN T., GOESELE M., SEIDEL H.-P.: Mesostructure from specularity. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2006), vol. 2, pp. 1825–1832. 13, 21, 26, 29, 55

[Che95] CHENG Y.: Mean shift, mode seeking, and clustering. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 17*, 8 (1995), 790–799. 31

[CJ82] COLEMAN E. N., JAIN R.: Obtaining 3-dimensional shape of textured and specular surfaces using four-source photometry. *Computer Graphics and Image Processing 18*, 4 (1982), 309–328. 19, 27

[CJ08] CHUNG H.-S., JIA J.: Efficient photometric stereo on glossy surfaces with wide specular lobes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2008), pp. 1–8. 24

[CK11] CREMERS D., KOLEV K.: Multiview stereo and silhouette consistency via convex functionals over convex domains. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 33*, 6 (2011), 1161–1174. 18, 23

[CLFS07] CHEN T., LENSCH H., FUCHS C., SEIDEL H.-P.: Polarization and phase-shifting for 3d scanning of translucent objects. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2007), pp. 1–8. 29

[CLL07] CHANG J. Y., LEE K. M., LEE S. U.: Multiview normal field integration using level set methods. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2007), pp. 1–8. 18, 19, 31

[CM02] COMANICIU D., MEER P.: Mean shift: A robust approach toward feature space analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 24*, 5 (2002), 603–619. 35

[CSL08] CHEN T., SEIDEL H.-P., LENSCH H.: Modulated phase-shifting for 3d scanning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2008), pp. 1–8. 29

[CT82] COOK R. L., TORRANCE K. E.: A reflectance model for computer graphics. *ACM Transactions on Graphics 1*, 1 (1982), 7–24. 37

[CT11] CALAKLI F., TAUBIN G.: SSD: Smooth signed distance surface reconstruction. *Computer Graphics Forum 30*, 7 (2011), 1993–2002. 30, 31, 32, 33

[Cur97] CURLESS B. L.: *New methods for surface reconstruction from range images*. PhD thesis, Stanford University, Stanford, USA, 1997. 14

[CZCW12] CHEN Z., ZHOU J., CHEN Y., WANG G.: 3D texture mapping in multi-view reconstruction. In *Advances in Visual Computing* (2012), vol. 7431 of *Lecture Notes in Computer Science*, pp. 359–371. 36

[Dai09] DAI Z.: A Markov random field approach for multi-view normal integration. Master-Thesis, University of Hong Kong, 2009. URL: http://books.google.de/books?id=diqwXwAACAAJ. 18, 19

[Dan01] DANA K. J.: BRDF/BTF measurement device. In *Proceedings of the International Conference on Computer Vision (ICCV)* (2001), vol. 2, pp. 460–466. 39, 42, 44

[dBSHK14] DEN BROK D., STEINHAUSEN H. C., HULLIN M. B., KLEIN R.: Patch-based sparse reconstruction of material btfs. *Journal of WSCG 22*, 2 (2014), 83–90. 46

[dBSHK15] DEN BROK D., STEINHAUSEN H. C., HULLIN M. B., KLEIN R.: Multiplexed acquisition of bidirectional texture functions for materials. In *Proceedings of SPIE* (2015), vol. 9398, pp. 93980F–1 – 93980F–9. 46

[DCCS06] DELLEPIANE M., CORSINI M., CALLIERI M., SCOPIGNO R.: High quality ptm acquisition: Reflection transformation imaging for large objects. In *Proceedings of the International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage (VAST)* (2006), pp. 179–186. 50

[Deb96] DEBEVEC P. E.: *Modeling and rendering architecture from photographs*. PhD thesis, University of California, Berkeley, USA, 1996. 36

[DF97] DEYOUNG J., FOURNIER A.: Properties of tabulated bidirectional reflectance distribution functions. In *Proceedings of Graphics Interface* (1997), pp. 47–55. 37

[DFS08] DUROU J.-D., FALCONE M., SAGONA M.: Numerical methods for shape-from-shading: A new survey with benchmarks. *Computer Vision and Image Understanding 109*, 1 (2008), 22–43. 16

[DGC*11] DORRINGTON A. A., GODBAZ J. P., CREE M. J., PAYNE A. D., STREETER L. V.: Separating true range measurements from multi-path and scattering interference in commercial range cameras. *Proceedings of SPIE 7864* (2011), 786404–1 – 786404–10. 20

[DHB10] DROESCHEL D., HOLZ D., BEHNKE S.: Multi-frequency phase unwrapping for time-of-flight cameras. *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems* (2010), 1463Ű–1469. 20

[DHT*00] DEBEVEC P., HAWKINS T., TCHOU C., DUIKER H.-P., SAROKIN W., SAGAR M.: Acquiring the reflectance field of a human face. In *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)* (2000), pp. 145–156. 11, 43, 45, 49, 50

[DJ05] DONNER C., JENSEN H. W.: Light diffusion in multi-layered translucent materials. *ACM Transactions on Graphics 24*, 3 (2005), 1032–1039. 49

[DLG13] DONG Y., LIN S., GUO B.: *Material appearance modeling: A data-coherent approach*. Springer Berlin Heidelberg, Germany, 2013. 11, 12

[DMC*12] DELLEPIANE M., MARROQUIM R., CALLIERI M., CIGNONI P., SCOPIGNO R.: Flow-based local optimization for image-to-geometry projection. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 18*, 3 (2012), 463–474. 36

[DMZP14] DONG B., MOORE K. D., ZHANG W., PEERS P.: Scattering parameters and surface normals from homogeneous translucent materials using photometric stereo. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2014), pp. 2299–2306. 29, 55

[DP11] DELAUNOY A., PRADOS E.: Gradient flows for optimizing triangular mesh-based surfaces: Applications to 3d reconstruction problems dealing with visibility. *International Journal of Computer Vision (IJCV)* (2011), 100–123. 23

[DPB10] DELAUNOY A., PRADOS E., BELHUMEUR P. N.: Towards full 3D Helmholtz stereovision algorithms. In *Proceedings of the Asian Conference on Computer Vision (ACCV)* (2010), vol. 1, pp. 39–52. 23, 25

[DvGNK97] DANA K. J., VAN GINNEKEN B., NAYAR S. K., KOENDERINK J. J.: Reflectance and texture of real world surfaces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (1997), pp. 151–157. 43

[DW88] DARRELL T., WOHN K.: Pyramid based depth from focus. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (1988), 504–509. 16

[DW04] DANA K. J., WANG J.: Device for convenient measurement of spatially varying bidirectional reflectance. *Journal of the Optical Society of America A 21*, 1 (2004), 1–12. 42, 44

[DWd*08] DONNER C., WEYRICH T., D'EON E., RAMAMOORTHI R., RUSINKIEWICZ S.: A layered, heterogeneous reflectance model for acquiring and rendering human skin. In *ACM SIGGRAPH Asia 2008 Papers* (2008), pp. 140:1–140:12. 49

[DWP*10] DONG Y., WANG J., PELLACINI F., TONG X., GUO B.: Fabricating spatially-varying subsurface scattering. *ACM Transactions on Graphics 29*, 4 (2010), 62:1–62:10. 49

[DWT*02] DEBEVEC P., WENGER A., TCHOU C., GARDNER A., WAESE J., HAWKINS T.: A lighting reproduction approach to live-action compositing. *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH) 21*, 3 (2002), 547–556. 45

[DWT*10] DONG Y., WANG J., TONG X., SNYDER J., LAN Y., BEN-EZRA M., GUO B.: Manifold bootstrapping for svbrdf capture. *ACM Transactions on Graphics 29*, 4 (2010), 98:1–98:10. 40, 41

[DZC12] DAL MUTTO C., ZANUTTIGH P., CORTELAZZO G. M.: *Time-of-Flight Cameras and Microsoft Kinect(TM)*. Springer New York, USA, 2012. 20

[EAM*09] EREN G., AUBRETON O., MERIAUDEAU F., SECADES L. A. S., FOFI D., NASKALI A. T., TRUCHETET F., ERCIL A.: Scanning from heating: 3d shape estimation of transparent objects from local surface heating. *Optics Express 17*, 14 (2009), 11457–11468. 29

[EF15] EIGEN D., FERGUS R.: Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *Proceedings of the International Conference on Computer Vision (ICCV)* (2015), pp. 2650–2658. 23

[EPF14] EIGEN D., PUHRSCH C., FERGUS R.: Depth map prediction from a single image using a multi-scale deep network. In *Proceedings of the Conference on Neural Information Processing Systems (NIPS)* (2014), pp. 2366–2374. 23

[ES04] ESTEBAN C. H., SCHMITT F.: Silhouette and stereo fusion for 3d object modeling. *Computer Vision and Image Understanding 96* (2004), 367–392. 18, 23

[EVC08] ESTEBAN C. H., VOGIATZIS G., CIPOLLA R.: Multiview photometric stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 30*, 3 (2008), 548–554. 19, 23

[FBLS07] FUCHS M., BLANZ V., LENSCH H. P. A., SEIDEL H.-P.: Adaptive sampling of reflectance fields. *ACM Transactions on Graphics 26*, 2 (2007), 10:1–10:18. 50, 51

[FCM*08] FRANCKEN Y., CUYPERS T., MERTENS T., GIELIS J., BEKAERT P.: High quality mesostructure acquisition using specularities. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2008), pp. 1–7. 26, 55

[FCSS10] FURUKAWA Y., CURLESS B., SEITZ S., SZELISKI R.: Towards internet-scale multi-view stereo. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2010), pp. 1434–1441. 57

[FFGG*10] FRAHM J.-M., FITE-GEORGEL P., GALLUP D., JOHNSON T., RAGURAM R., WU C., JEN Y.-H., DUNN E., CLIPP B., LAZEBNIK S., POLLEFEYS M.: Building Rome on a cloudless day. In *Proceedings of the European Conference on Computer Vision (ECCV)* (2010), pp. 368–381. 35

[FG14] FUHRMANN S., GOESELE M.: Floating scale surface reconstruction. *ACM Transactions on Graphics 33*, 4 (2014), 46:1–46:11. 30, 31, 33, 34

[FGCS05] FUCHS C., GOESELE M., CHEN T., SEIDEL H.-P.: An empirical model for translucent objects. *SIGGRAPH Sketches 2005* (2005). 49

[FH09] FILIP J., HAINDL M.: Bidirectional texture function modeling: A state of the art survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 31*, 11 (2009), 1921–1940. 43

[FKG15] FUHRMANN S., KAZHDAN M., GOESELE M.: Accurate Isosurface Interpolation with Hermite Data. In *Proceedings of the International Conference on 3D Vision (3DV)* (2015). 30

[FKIS02] FURUKAWA R., KAWASAKI H., IKEUCHI K., SAKAUCHI M.: Appearance based object modeling using texture database: Acquisition, compression and rendering. In *Proceedings of the Eurographics Workshop on Rendering* (2002), pp. 257–266. 18, 45, 47

[FLG14] FUHRMANN S., LANGGUTH F., GOESELE M.: MVE - A Multi-View Reconstruction Environment. In *Proceedings of the Eurographics Workshop on Graphics and Cultural Heritage (GCH)* (2014). 57, 60

[Foo97] FOO S. C.: *A gonioreflectometer for measuring the bidirectional reflectance of material for use in illumination computation*. PhD thesis, Cornell University, Ithaca, USA, 1997. 38

[FSK*14] FREEDMAN D., SMOLIN Y., KRUPKA E., LEICHTER I., SCHMIDT M.: SRA: Fast removal of general multipath for ToF sensors. In *Computer Vision - ECCV 2014*, Fleet D., Pajdla T., Schiele B., Tuytelaars T., (Eds.), vol. 8689 of *Lecture Notes in Computer Science*. Springer International Publishing, 2014, pp. 234–249. URL: http://dx.doi.org/10.1007/978-3-319-10590-1_16, doi:10.1007/978-3-319-10590-1_16. 20

[FVH*13] FILIP J., VÁVRA R., HAINDL M., ŽID P., KRUPIČKA M., HAVRAN V.: BRDF slices: Accurate adaptive anisotropic appearance acquisition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2013, pp. 1468–1473. 44

[GAHO07] GHOSH A., ACHUTHA S., HEIDRICH W., O'TOOLE M.: Brdf acquisition with basis illumination. In *Proceedings of the International Conference on Computer Vision (ICCV)* (2007), pp. 1–8. 39

[GAVN11] GUPTA M., AGRAWAL A., VEERARAGHAVAN A., NARASIMHAN S. G.: Structured light 3d scanning in the presence of global illumination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2011), pp. 713–720. 29

[GCG*05] GUENTHER J., CHEN T., GOESELE M., WALD I., SEIDEL H.-P.: Efficient acquisition and realistic rendering of car paint. *Proceedings of Vision, Modeling, and Visualization (VMV)* (2005), 487–494. 40

[GCHS05] GOLDMAN D. B., CURLESS B., HERTZMANN A., SEITZ S. M.: Shape and spatially-varying brdfs from photometric stereo. In *Proceedings of the International Conference on Computer Vision (ICCV)* (2005), vol. 1, pp. 341–348. 24, 41

[GCP*09] GHOSH A., CHEN T., PEERS P., WILSON C. A., DEBEVEC P.: Estimating specular roughness and anisotropy from second order spherical gradient illumination. In *Proceedings of the Eurographics Symposium on Rendering (EGSR)* (2009), pp. 1161–1170. 41

[GCP*10] GHOSH A., CHEN T., PEERS P., WILSON C. A., DEBEVEC P.: Circularly polarized spherical illumination reflectometry. *ACM Transactions on Graphics 29*, 6 (2010), 162:1–162:12. 29, 41

[GDDA13] GARCIA-DORADO I., DEMIR I., ALIAGA D. G.: Automatic urban modeling using volumetric reconstruction with surface graph cuts. *Computers & Graphics 37*, 7 (2013), 896–910. 35

[Gei09] GEIGES M. L.: Traces of Marion B. Sulzberger in the museum of wax moulages in Zurich and their importance for the history of dermatology. *Journal of the American Academy of Dermatology 60*, 6 (2009), 980–984. 52

[Geo03] GEORGHIADES A. S.: Incorporating the torrance and sparrow model of reflectance in uncalibrated photometric stereo. In *Proceedings of the International Conference on Computer Vision (ICCV)* (2003), vol. 2, pp. 816–823. 24

[GG03] GODDYN L., GVOZDJAK P.: Binary gray codes with long bit runs. *The Electronic Journal of Combinatorics 10* (2003), 1–10. R27. 29

[GGSC96] GORTLER S. J., GRZESZCZUK R., SZELISKI R., COHEN M. F.: The lumigraph. In *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)* (1996), pp. 43–54. 10, 18, 50

[GHP*08] GHOSH A., HAWKINS T., PEERS P., FREDERIKSEN S., DEBEVEC P.: Practical modeling and acquisition of layered facial reflectance. *ACM Transactions on Graphics 27*, 5 (2008), 139:1–139:10. 49

[Gib50] GIBSON J. J.: *The perception of the visual world*. Houghton Mifflin, Boston, USA, 1950. 16

[GILM07] GOLDLUCKE B., IHRKE I., LINZ C., MAGNOR M.: Weighted minimal hypersurface reconstruction. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 29*, 7 (2007), 1194–1208. 28

[GKKP07] GRAMMATIKOPOULOS L., KALISPERAKIS I., KARRAS G., PETSA E.: Automatic multi-view texture mapping of 3d surface projections. In *3D Virtual Reconstruction & Visualization of Complex Architectures (3D Arch)* (2007). 36

[Gla95] GLASSNER A. S.: *Principles of digital image synthesis*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1995. 7

[GLL*04] GOESELE M., LENSCH H. P. A., LANG J., FUCHS C., SEIDEL H.-P.: Disco: Acquisition of translucent objects. *ACM Transactions on Graphics 23*, 3 (2004), 835–844. 29, 49

[GSC*07] GOESELE M., SNAVELY N., CURLESS B., HOPPE H., SEITZ S.: Multi-view stereo for community photo collections. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2007), pp. 1–8. 57

[GTLL06] GARG G., TALVALA E.-V., LEVOY M., LENSCH H. P. A.: Symmetric photography: Exploiting data-sparseness in reflectance fields. In *Proceedings of the Eurographics Symposium on Rendering (EGSR)*. 2006, pp. 251–262. 44, 45

[GWO*10] GAL R., WEXLER Y., OFEK E., HOPPE H., COHEN-OR D.: Seamless montage for texturing models. *Computer Graphics Forum 29*, 2 (2010), 479–486. 35

[HAH*12] HÖPE A., ATAMAS T., HÜNERHOFF D., TEICHERT S., HAUER K.-O.: ARGon3: 3d appearance robot-based gonioreflectometer at ptb. *The Review of Scientific Instruments 83*, 4 (2012), 11–16. 38

[HCD01] HAWKINS T., COHEN J., DEBEVEC P.: A photometric approach to digitizing cultural artifacts. In *Proceedings of the International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage (VAST)* (2001), pp. 333–342. 45

[HED05] HAWKINS T., EINARSSON P., DEBEVEC P.: A dual light stage. In *Proceedings of the Eurographics Symposium on Rendering (EGSR)* (2005), pp. 91–98. 50

[HF11] HAINDL M., FILIP J.: Advanced textural representation of materials appearance. In *SIGGRAPH Asia 2011 Courses*. 2011, pp. 1:1–1:84. 43

[HF13] HAINDL M., FILIP J.: *Visual texture: Accurate material appearance measurement, representation and modeling*. Advances in Computer Vision and Pattern Recognition. Springer, New York, USA, 2013. 37, 43, 48

[HFB*09] HABER T., FUCHS C., BEKAERT P., SEIDEL H.-P., GOESELE M., LENSCH H. P. A.: Relighting objects from image collections. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2009), pp. 627–634. 41

[HFI*08] HULLIN M. B., FUCHS M., IHRKE I., SEIDEL H.-P., LENSCH H. P. A.: Fluorescent immersion range scanning. *ACM Transactions on Graphics 27*, 3 (2008), 87:1–87:10. 29

[HGH06] HÜNERHOFF D., GRUSEMANN U., HÖPE A.: New robot-based gonioreflectometer for measuring spectral diffuse reflection. *Metrologia 43*, 2 (2006), 11–16. 38

[HK93] HANRAHAN P., KRUEGER W.: Reflection from layered surfaces due to subsurface scattering. In *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)* (1993), pp. 165–174. 48

[HK06] HORNUNG A., KOBBELT L.: Hierarchical volumetric multi-view stereo reconstruction of manifold surfaces based on dual graph embedding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2006), pp. 503–510. 23

[HLZ10] HOLROYD M., LAWRENCE J., ZICKLER T.: A coaxial optical scanner for synchronous acquisition of 3D geometry and surface reflectance. *ACM Transactions on Graphics 29*, 4 (2010), 99:1–99:12. 42, 44

[HMI10] HIGO T., MATSUSHITA Y., IKEUCHI K.: Consensus photometric stereo. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2010), pp. 1157–1164. 19, 27

[HMJI09] HIGO T., MATSUSHITA Y., JOSHI N., IKEUCHI K.: A hand-held photometric stereo camera for 3-d modeling. In *Proceedings of the International Conference on Computer Vision (ICCV)* (2009), pp. 1234–1241. 23

[Hor70] HORN B. K. P.: *Shape from Shading: A Method for Obtaining the Shape of a Smooth Opaque Object from One View*. Tech. rep., Massachusetts Institute of Technology, Cambridge, MA, USA, 1970. 15

[HP03] HAN J. Y., PERLIN K.: Measuring bidirectional texture reflectance with a kaleidoscope. *ACM Transactions on Graphics 22*, 3 (2003), 741–748. 44, 45

[HQS10] HU Y., QI Y., SHEN F.: Modeling spatially-varying reflectance based on kernel nyström. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology (VRST)*. 2010, pp. 91–92. 45, 46

[HS81] HORN B. K. P., SCHUNCK B. G.: Determining optical flow. *Artificial Intelligence 17* (1981), 185–203. 27

[HS05] HERTZMANN A., SEITZ S. M.: Example-based photometric stereo: Shape reconstruction with general, varying brdfs. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 27* (2005), 1254–1264. 24

[HSKK96] HATA S., SAITOH Y., KUMAMURA S., KAIDA K.: Shape extraction of transparent object using genetic algorithm. In *Proceedings of the International Conference on Pattern Recognition (ICPR)* (1996), vol. 4, pp. 684–688. 28

[HTSG91] HE X. D., TORRANCE K. E., SILLION F. X., GREENBERG D. P.: A comprehensive physical model for light reflection. In *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)* (1991), pp. 175–186. 37

[HW11] HERBORT S., WÖHLER C.: An introduction to image-based 3d surface reconstruction and a survey of photometric stereo methods. *3D Research 2*, 3 (2011), 40:1–40:17. 20, 24, 25

[HWQ13] HU Y., WANG S., QI Y.: Efficient reconstruction, decomposition and editing for spatially-varying reflectance data. In *Proceedings of the Graphics Interface Conference (GI)*. 2013, pp. 55–62. 45, 46

[HZ04] HARTLEY R. I., ZISSERMAN A.: *Multiple view geometry in computer vision (2nd edition)*. Cambridge University Press, Cambridge, UK, 2004. 12, 13

[IH81] IKEUCHI K., HORN B. K. P.: Numerical shape from shading and occluding boundaries. *Artificial Intelligence 17*, 1 (1981), 141–184. 16

[Ike81] IKEUCHI K.: Determining surface orientations of specular surfaces by using the photometric stereo method. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 3*, 6 (1981), 661–669. 21, 24, 26

[IKL*10] IHRKE I., KUTULAKOS K. N., LENSCH H. P. A., MAGNOR M., HEIDRICH W.: Transparent and specular object reconstruction. *Computer Graphics Forum 29*, 8 (2010), 2400–2426. 8, 14, 25, 27, 28, 29

[IM06] IHRKE I., MAGNOR M.: Adaptive grid optical tomography. In *Graphical Models* (2006), vol. 68, pp. 484–495. 28

[IMMY14] INOSHITA C., MUKAIGAWA Y., MATSUSHITA Y., YAGI Y.: Surface normal deconvolution: Photometric stereo for optically thick translucent objects. In *Proceedings of the European Conference on Computer Vision (ECCV)* (2014), vol. 2, pp. 346–359. 55

[IRM*12] IHRKE I., RESHETOUSKI I., MANAKOV A., TEVS A., WAND M., SEIDEL H.-P.: A kaleidoscopic approach to surround geometry and reflectance acquisition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. 2012, pp. 29–36. 44, 45

[JCYS04] JIN H., CREMERS D., YEZZI A. J., SOATTO S.: Shedding light on stereoscopic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2004), vol. 1, pp. I–36 – I–42. 23

[JLSW02] JU T., LOSASSO F., SCHAEFER S., WARREN J.: Dual contouring of hermite data. In *Proceedings of the 29th Annual Conference on Computer Graphics and Interactive Techniques* (2002), SIGGRAPH '02. 34

[JMLH01] JENSEN H. W., MARSCHNER S. R., LEVOY M., HANRAHAN P.: A practical model for subsurface light transport. In *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)* (2001), pp. 511–518. 49

[Jut09] JUTZI B.: Investigations on ambiguity unwrapping of range images. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences XXXVIII-3/W8* (2009), 265Ű–270. 20

[Jut12] JUTZI B.: Extending the range measurement capabilities of modulated range imaging devices by time-frequency-multiplexing. *Allgemeine Vermessungs-Nachrichten 2/2012* (2012), 54–62. 20

[Kaj85] KAJIYA J. T.: Anisotropic reflection models. In *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)* (1985), pp. 15–21. 8

[KBH06] KAZHDAN M., BOLITHO M., HOPPE H.: Poisson surface reconstruction. In *Proceedings of the Eurographics Symposium on Geometry Processing (SGP)* (2006), pp. 61–70. 30, 31, 32

[KBKL10] KOLB A., BARTH E., KOCH R., LARSEN R.: Time-of-flight cameras in computer graphics. *Computer Graphics Forum 29*, 1 (2010), 141–159. 20

[KBMK01] KOUDELKA M., BELHUMEUR P., MAGDA S., KRIEGMAN D.: Image-based modeling and rendering of surfaces with arbitrary brdfs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2001), vol. 1, pp. 568–575. 50

[Ke99] KE L.: *A method of light reflectance measurement*. PhD thesis, University of British Columbia, Vancouver, Canada, 1999. 38

[KGV12] KAMAL M. H., GOLBABAEE M., VANDERGHEYNST P.: Light field compressive sensing in camera arrays. In *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing* (2012). 50

[KH13] KAZHDAN M., HOPPE H.: Screened poisson surface reconstruction. In *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)* (2013), vol. 32, pp. 29:1–29:13. 30, 31

[KKDH07] KAZHDAN M., KLEIN A., DALAL K., HOPPE H.: Unconstrained isosurface extraction on arbitrary octrees. In *Proceedings of the Eurographics Symposium on Geometry Processing (SGP)* (2007), pp. 125–133. 33

[KKH*11] KOLEV K., KIRCHGEÍNER N., HOUBEN S., CSISZÁR A., RUBNER W., PALM C., EIBEN B., MERKEL R., CREMERS D.: A variational approach to vesicle membrane reconstruction from fluorescence imaging. *Pattern Recognition 44*, 12 (2011), 2944–2958. 23

[KKP*08] KIM D. B., KIM K. Y., PARK K. S., SEO M. K., LEE K. H.: A fast and accurate image-based measuring system for isotropic reflection materials. In *Proceedings of SPIE* (2008), vol. 7065, pp. 70650I1 – 70650I12. 40

[KMBK03] KOUDELKA M. L., MAGDA S., BELHUMEUR P. N., KRIEGMAN D. J.: Acquisition, compression, and synthesis of bidirectional texture functions. In *Proceedings of the ICCV Workshop on Texture Analysis and Synthesis* (2003), pp. 47–52. 43

[KMG96] KARNER K. F., MAYER H., GERVAUTZ M.: An image based measurement system for anisotropic reflection. *Computer Graphics Forum 15*, 3 (1996), 119Ű–128. 39

[KNRS13] KÖHLER J., NÖLL T., REIS G., STRICKER D.: A full-spherical device for simultaneous geometry and reflectance acquisition. In *Proceedings of the IEEE Workshop on Applications of Computer Vision (WACV)*. 2013, pp. 355–362. 42, 45, 47

[KPC10] KOLEV K., POCK T., CREMERS D.: Anisotropic minimal surfaces integrating photoconsistency and normal information for multiview stereo. In *Proceedings of the European Conference on Computer Vision (ECCV)* (2010), pp. 538–551. 23

[KPDVG05] KONINCKX T. P., PEERS P., DUTRE P., VAN GOOL L.: Scene-adapted structured light. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2005), vol. 2, pp. 611–618. 24

[Kro87] KROTKOV E.: Focusing. *International Journal of Computer Vision (IJCV) 1*, 3 (1987), 223–237. 16

[KS05] KUTULAKOS K. N., STEGER E.: A theory of refractive and specular 3d shape by light-path triangulation. In *Proceedings of the International Conference on Computer Vision (ICCV)* (2005), vol. 2, pp. 1448–1455. 28, 55

[KS08] KUTULAKOS K. N., STEGER E.: A theory of refractive and specular 3d shape by light-path triangulation. *International Journal of Computer Vision (IJCV) 76*, 1 (2008), 13–29. 25, 28, 55

[KSS*04] KAUTZ J., SATTLER M., SARLETTE R., KLEIN R., SEIDEL H.-P.: Decoupling brdfs from surface mesostructures. In *Proceedings of Graphics Interface* (2004), pp. 177–182. 41

[KTM*02] KANITSAR A., THEUSSL T., MROZ L., SRAMEK M., BARTROLI A. V., CSEBFALVI B., HLADUVKA J., FLEISCHMANN D., KNAPP M., WEGENKITTL R., FELKEL P., ROTTGER S., GUTHE S., PURGATHOFER W., GROLLER M. E.: Christmas tree case study: Computed tomography as a tool for mastering complex real world objects with applications in computer graphics. In *Proceedings of the Conference on Visualization* (2002), pp. 489–492. 28

[KTT06] KIMACHI A., TANAKA N., TOMINAGA S.: Development and calibration of a gonio-spectral imaging system for measuring surface reflection. *IEICE - Transactions on Information and Systems E89-D* (2006), 1994–2003. 43

[KvD98] KOENDERINK J. J., VAN DOORN A. J.: Phenomenological description of bidirectional surface reflection. *Journal of the Optical Society of America A 15*, 11 (1998), 2903–2912. 37

[LA09] LOURAKIS M. I. A., ARGYROS A. A.: SBA: A software package for generic sparse bundle adjustment. *ACM Transactions on Mathematical Software 36*, 1 (2009), 1–30. 13

[lab] LABSPHERE: Spectralon© diffuse reflectance targets. http://www.labsphere.com/uploads/pb13025Targets.pdf (accessed on 19 March 2014). URL: http://www.labsphere.com/uploads/pb13025Targets.pdf. 14

[Lau94] LAURENTINI A.: The visual hull concept for silhouette-based image understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 16* (1994), 150–162. 18

[LBAD*06] LAWRENCE J., BEN-ARTZI A., DECORO C., MATUSIK W., PFISTER H., RAMAMOORTHI R., RUSINKIEWICZ S.: Inverse shade trees for non-parametric material representation and editing. In *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)* (2006), pp. 735–745. 37, 42

[LBN08] LADIKOS A., BENHIMANE S., NAVAB N.: Multi-view reconstruction using narrow-band graph-cuts and surface normal optimization. In *Proceedings of the British Machine Vision Conference (BMVC)* (2008), pp. 15.1–15.10. 23

[LC87] LORENSEN W. E., CLINE H. E.: Marching cubes: A high resolution 3d surface construction algorithm. *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH) 21*, 4 (1987), 163–169. 31

[LCV*04] LEVOY M., CHEN B., VAISH V., HOROWITZ M., MCDOWALL I., BOLAS M.: Synthetic aperture confocal imaging. *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH) 23*, 3 (2004), 825–834. 44, 45

[LDW*10] LAN Y., DONG Y., WANG J., TONG X., GUO B.: Condenser-based instant reflectometry. In *Computer Graphics Forum (Proceedings of Pacific Graphics)* (2010), vol. 29, pp. 117–126. 40

[Leh07] LEHTINEN J.: A framework for precomputed and captured light transport. *ACM Transactions on Graphics 26*, 4 (2007), 13:1–13:22. 50

[Lev44] LEVENBERG K.: A method for the solution of certain non-linear problems in least squares. *Quarterly Journal of Applied Mathematics II*, 2 (1944), 164–168. 13, 40, 42

[LF97] LALONDE P., FOURNIER A.: A wavelet representation of reflectance functions. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 3*, 4 (1997), 329–336. 37

[LFD*08] LELOUP F. B., FORMENT S., DUTRÉ P., POINTER M. R., HANSELAER P.: Design of an instrument for measuring the spectral bidirectional scatter distribution function. *Applied Optics 47*, 29 (2008), 5454–5467. 38

[LFTG97] LAFORTUNE E. P. F., FOO S.-C., TORRANCE K. E., GREENBERG D. P.: Non-linear approximation of reflectance functions. In *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)* (1997), pp. 117–126. 37, 40

[LGK10] LIU B., GOULD S., KOLLER D.: Single image depth estimation from predicted semantic labels. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2010), pp. 1253–1260. 22

[LH96] LEVOY M., HANRAHAN P.: Light field rendering. In *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)* (1996), pp. 31–42. 10, 11, 50

[LI07] LEMPITSKY V., IVANOV D.: Seamless mosaicing of image-based texture maps. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2007), pp. 1–6. 35, 36, 37

[LK81] LUCAS B. D., KANADE T.: An iterative image registration technique with an application to stereo vision. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)* (1981), vol. 2, pp. 674–679. 27

[LKG*01] LENSCH H. P. A., KAUTZ J., GOESELE M., HEIDRICH W., SEIDEL H.-P.: Image-based reconstruction of spatially varying materials. *Rendering Techniques* (2001), 103–114. 40, 41

[LKG*03] LENSCH H. P. A., KAUTZ J., GOESELE M., HEIDRICH W., SEIDEL H.-P.: Image-based reconstruction of spatial appearance and geometric detail. *ACM Transactions on Graphics 22* (2003), 234–257. 40, 41

[LKK98] LU R., KOENDERINK J. J., KAPPERS A. M.: Optical properties (bidirectional reflection distribution functions) of velvet. *Applied Optics 37*, 25 (1998), 5974–5984. 39

[LL95] LU J., LITTLE J.: Reflectance function estimation and shape recovery from image sequence of a rotating object. In *Proceedings of the International Conference on Computer Vision (ICCV)* (1995), pp. 80–86. 39

[Llo57] LLOYD S. P.: *Least squares quantization in PCM*. Tech. rep., Bell Laboratories, 1957. 27

[LLSS03] LENSCH H. P. A., LANG J., SÁ A. M., SEIDEL H.-P.: Planned sampling of spatially varying brdfs. *Computer Graphics Forum 22*, 3 (2003), 473–482. 40

[LLW*08] LIANG C.-K., LIN T.-H., WONG B.-Y., LIU C., CHEN H. H.: Programmable aperture photography: Multiplexed light field acquisition. *ACM Transactions on Graphics 27*, 3 (2008), 55:1–55:10. 50

[Low04] LOWE D. G.: Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision (IJCV) 60* (2004), 91–110. 13, 18

[LPK09] LABATUT P., PONS J.-P., KERIVEN R.: Robust and efficient surface reconstruction from range data. *Computer Graphics Forum 28*, 8 (2009), 2275–2290. 30

[LSH14] LIU M., SALZMANN M., HE X.: Discrete-continuous depth estimation from a single image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2014), pp. 716–723. 23

[LSL03] LANG J., SEIDEL H.-P., LENSCH H. P. A.: View planning for brdf acquisition. In *ACM SIGGRAPH 2003 Sketches and Applications* (2003). 38

[LSL15] LIU F., SHEN C., LIN G.: Deep convolutional neural fields for depth estimation from a single image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2015), pp. 5162–5170. 23

[LSLR15] LIU F., SHEN C., LIN G., REID I. D.: Learning depth from single monocular images using deep convolutional neural fields. *CoRR abs/1502.07411* (2015). URL: http://arxiv.org/abs/1502.07411. 23

[LTBEB10] LU Z., TAI Y.-W., BEN-EZRA M., BROWN M. S.: A framework for ultra high resolution 3d imaging. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2010), pp. 1205–1212. 23

[LWDC10] LIU M., WONG K.-Y. K., DAI Z., CHEN Z.: Specular surface recovery from reflections of a planar pattern undergoing an unknown pure translation. In *Proceedings of the Asian Conference on Computer Vision (ACCV)* (2010), vol. 2, pp. 137–147. 26

[Mac67] MACQUEEN J. B.: Some methods for classification and analysis of multivariate observations. In *Proceedings of the Berkeley Symposium on Mathematical Statistics and Probability* (1967), vol. 1, pp. 281–297. 27

[MBK05] MÜLLER G., BENDELS G. H., KLEIN R.: Rapid synchronous acquisition of geometry and appearance of cultural heritage artefacts. In *Proceedings of the International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage (VAST)* (2005), pp. 13–20. 18, 47

[MC09] MAKI A., CIPOLLA R.: Obtaining the shape of a moving object with a specular surface. In *Proceedings of the British Machine Vision Conference (BMVC)* (2009), pp. 39.1–39.10. 19, 27

[McA02] MCALLISTER D. K.: *A Generalized Surface Appearance Representation for Computer Graphics*. PhD thesis, University of North Carolina at Chapel Hill, Chapel Hill, USA, 2002. 40, 41, 43

[MCS90] MURRAY-COLEMAN F., SMITH A. M.: The automated measurement of brdfs and their application to luminaire modeling. 87–99. 37

[MGW01] MALZBENDER T., GELB D., WOLTERS H.: Polynomial texture maps. In *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)* (2001), pp. 519–528. 50

[MHP*07] MA W.-C., HAWKINS T., PEERS P., CHABERT C.-F., WEISS M., DEBEVEC P.: Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. In *Proceedings of the Eurographics Symposium on Rendering (EGSR)* (2007), pp. 183–194. 29

[MK05] MORRIS N. J. W., KUTULAKOS K. N.: Dynamic refraction stereo. In *Proceedings of the International Conference on Computer Vision (ICCV)* (2005), vol. 2, pp. 1573–1580. 28, 55

[MK07] MORRIS N. J. W., KUTULAKOS K. N.: Reconstructing the surface of inhomogeneous transparent scenes by scatter-trace photography. In *Proceedings of the International Conference on Computer Vision (ICCV)* (2007), pp. 1–8. 28, 55

[MKG11] MÜCKE P., KLOWSKY R., GOESELE M.: Surface reconstruction from multi-resolution sample points. In *Proceedings of the Vision, Modeling, and Visualization Workshop 2011* (2011), pp. 105–112. 30, 31, 32

[MMC*08] MUDGE M., MALZBENDER T., CHALMERS A., SCOPIGNO R., DAVIS J., WANG O., GUNAWARDANE P., ASHLEY M., DOERR M., PROENCA A., BARBOSA J.: Image-based empirical information acquisition, scientific reliability, and long-term digital preservation for the natural sciences and cultural heritage. *Eurographics Tutorials* (2008). 50

[MMS*04] MÜLLER G., MESETH J., SATTLER M., SARLETTE R., KLEIN R.: Acquisition, synthesis and rendering of bidirectional texture functions. In *Eurographics 2004 State of the Art Reports* (2004), pp. 69–94. 7, 8, 11, 12, 14, 43, 46, 47

[MMSL06] MUDGE M., MALZBENDER T., SCHROER C., LUM M.: New reflection transformation imaging methods for rock art and multiple-viewpoint display. In *Proceedings of the International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage (VAST)* (2006), pp. 195–202. 50

[MPBM03a] MATUSIK W., PFISTER H., BRAND M., MCMILLAN L.: A data-driven reflectance model. In *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)* (2003), pp. 759–769. 39

[MPBM03b] MATUSIK W., PFISTER H., BRAND M., MCMILLAN L.: Efficient isotropic brdf measurement. In *Proceedings of the Eurographics Workshop on Rendering* (2003), pp. 241–247. 37

[MPN*02] MATUSIK W., PFISTER H., NGAN A., BEARDSLEY P., ZIEGLER R., MCMILLAN L.: Image-based 3D photography using opacity hulls. *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH) 21*, 3 (2002), 427–437. 45, 47, 51

[MPZ*02] MATUSIK W., PFISTER H., ZIEGLER R., NGAN A., MCMILLAN L.: Acquisition and rendering of transparent and refractive objects. In *Proceedings of the Eurographics Workshop on Rendering*. 2002, pp. 267–278. 45, 47

[MRA*12] MERIAUDEAU F., RANTOSON R., ADAL K., FOFI D., STOLZ C.: Non-conventional imaging systems for 3D digitization of transparent objects: Shape from polarization in the IR and shape from visible fluorescence induced UV. In *Proceedings of the International Topical Meeting on Optical Sensing and Artificial Vision (OSAV)* (2012), pp. 34–40. 29

[MSSE*10] MERIAUDEAU F., SANCHEZ SECADES L. A., EREN G., ERCIL A., TRUCHETET F., AUBRETON O., FOFI D.: 3-d scanning of nonopaque objects by means of imaging emitted structured infrared patterns. *IEEE Transactions on Instrumentation and Measurement 59*, 11 (2010), 2898–2906. 29

[MSY07] MUKAIGAWA Y., SUMINO K., YAGI Y.: Multiplexed illumination for measuring brdf using an ellipsoidal mirror and a projector. In *Computer Vision Ű ACCV 2007* (2007), vol. 4844 of *Lecture Notes in Computer Science*, pp. 246–257. 39

[MTK*11] MUKAIGAWA Y., TAGAWA S., KIM J., RASKAR R., MATSUSHITA Y., YAGI Y.: Hemispherical confocal imaging using turtleback reflector. In *Computer Vision Ű ACCV 2010*, vol. 6492 of *Lecture Notes in Computer Science*. 2011, pp. 336–349. 44, 45

[Mur90] MURASE H.: Surface shape reconstruction of an undulating transparent object. In *Proceedings of the International Conference on Computer Vision (ICCV)* (1990), pp. 313–317. 27, 28

[Mur92] MURASE H.: Surface shape reconstruction of a nonrigid transport object using refraction and motion. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 14*, 10 (1992), 1045–1052. 27, 28

[MVGV10] MOONS T., VAN GOOL L., VERGAUWEN M.: 3D reconstruction from multiple images part 1: Principles. *Foundations and Trends in Computer Graphics and Vision 4*, 4 (2010), 287–404. 16, 20

[MVSL05] MUDGE M., VOUTAZ J.-P., SCHROER C., LUM M.: Reflection transformation imaging and virtual representations of coins from the hospice of the grand st. bernard. In *Proceedings of the International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage (VAST)* (2005), pp. 29–39. 50

[MWAM05] MARSCHNER S. R., WESTIN S. H., ARBREE A., MOON J. T.: Measuring and modeling the appearance of finished wood. *ACM Transactions on Graphics 24*, 3 (2005), 727–734. 42

[MWBR13] MARWAH K., WETZSTEIN G., BANDO Y., RASKAR R.: Compressive light field photography using overcomplete dictionaries and optimized projections. *ACM Transactions on Graphics 32*, 4 (2013), 46:1–46:12. 50

[MWBS96] MEISTER G., WIEMKER R., BIENLEIN J., SPITZER H.: In situ BRDF measurements of selected surface materials to improve analysis of remotely sensed multispectral imagery. In *International Archives for Photogrammetry and Remote Sensing* (1996), vol. 31-B7, pp. 493–498. 38

[MWGA06] MALZBENDER T., WILBURN B., GELB D., AMBRISCO B.: Surface enhancement using real-time photometric stereo and reflectance transformation. In *Proceedings of the Eurographics Symposium on Rendering (EGSR)* (2006), pp. 245–250. 50

[MWL*99] MARSCHNER S. R., WESTIN S. H., LAFORTUNE E. P. F., TORRANCE K. E., GREENBERG D. P.: Image-based brdf measurement including human skin. In *Proceedings of the Eurographics Workshop on Rendering* (1999), pp. 131–144. 39

[MWLT00] MARSCHNER S. R., WESTIN S. H., LAFORTUNE E. P. F., TORRANCE K. E.: Image-based bidirectional reflectance distribution function measurement. *Applied Optics 39* (2000), 2592–2600. 39

[MZKB05] MALLICK S. P., ZICKLER T. E., KRIEGMAN D., BELHUMEUR P. N.: Beyond lambert: reconstructing specular surfaces using color. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2005), vol. 2, pp. 619–626. 25

[ND06] NGAN A., DURAND F.: Statistical acquisition of texture appearance. In *Proceedings of the Eurographics Symposium on Rendering (EGSR)* (2006), pp. 31–40. 40

[NFB97] NAYAR S. K., FANG X.-S., BOULT T.: Separation of reflection components using color and polarization. *Proceedings of the International Conference on Computer Vision (ICCV) 21*, 3 (1997), 163–186. 25, 48

[Nis09] NISHINO K.: Directional statistics brdf model. In *Proceedings of the International Conference on Computer Vision (ICCV)* (2009), pp. 476–483. 37

[NJRS13] NÖLL T., J. K., REIS G., STRICKER D.: Faithful, compact and complete digitization of cultural heritage using a full-spherical scanner. In *Proceedings of the International Congress on Digital Heritage* (2013), pp. 15–22. 42, 45, 47

[NKGR06] NAYAR S. K., KRISHNAN G., GROSSBERG M. D., RASKAR R.: Fast separation of direct and global components of a scene using high frequency illumination. *ACM Transactions on Graphics 25* (2006), 935–944. 29, 48

[NKRS15] NÖLL T., KÖHLER J., REIS G., STRICKER D.: Fully automatic, omnidirectional acquisition of geometry and appearance in the context of cultural heritage preservation. *ACM Journal on Computing and Cultural Heritage (JOCCH) 8*, 1 (2015), 2:1–2:28. 42

[NN94] NAYAR S. K., NAKAGAWA Y.: Shape from focus. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 16*, 8 (1994), 824–831. 16

[NRDR05] NEHAB D., RUSINKIEWICZ S., DAVIS J., RAMAMOORTHI R.: Efficiently combining positions and normals for precise 3d geometry. *ACM Transactions on Graphics 24* (2005), 536–543. 23

[NRH*77] NICODEMUS F. E., RICHMOND J. C., HSIA J. J., GINSBERG I. W., LIMPERIS T.: *Geometrical considerations and nomenclature for reflectance*. National Bureau of Standards Monograph #160, U.S. Department of Commerce, 1977. 11, 37, 48

[NS92] NAIR H. N., STEWART C. V.: Robust focus ranging. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (1992), pp. 309–314. 16

[NWN96] NAYAR S. K., WATANABE M., NOGUCHI M.: Real-time focus range sensor. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 18*, 12 (1996), 1186–1198. 17

[NWR08] NEHAB D., WEYRICH T., RUSINKIEWICZ S.: Dense 3D reconstruction from specularity consistency. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2008), pp. 1–8. 26, 27, 28, 55

[NWSS90] NAYAR S. K., WEISS L. E., SIMON D. A., SANDERSON A. C.: Specular surface inspection using structured highlight and Gaussian images. *IEEE Transactions on Robotics and Automation 6*, 2 (1990), 208–218. 24, 26

[NZV*11] NAIK N., ZHAO S., VELTEN A., RASKAR R., BALA K.: Single view reflectance capture using multiplexed scattering and time-of-flight imaging. *ACM Transactions on Graphics 30*, 6 (2011), 171:1–171:10. 38

[NZVG05] NEUBECK A., ZALESNY A., VAN GOOL L.: 3D texture reconstruction from extensive btf data. In *Proceedings of the Texture 2005 Workshop* (2005), pp. 13–19. 46, 47

[NZVG06] NEUBECK A., ZALESNY A., VAN GOOL L.: Light source calibration for IBR and BTF acquisition setups. In *Proceedings of the International Symposium on 3D Data Processing, Visualisation and Transmission (3DPVT)*. 2006, pp. 272–278. 46

[Pak12] PAK A.: Recovering shapes of specular objects in motion via normal vector map consistency. In *Proceedings of SPIE* (2012), vol. 8493, pp. 84930T–1 – 84930T–8. 27

[Pak13] PAK A.: Reconstruction of specular surfaces via probabilistic voxel carving. In *Proceedings of SPIE* (2013), vol. 8791, pp. 87911B–1 – 87911B–8. 27

[PB11] PENNER E., BORSHUKOV G.: Pre-integrated skin shading. In *GPU Pro 2*, vol. 2. 2011, pp. 41–55. 52

[PCC*10] PALMA G., CORSINI M., CIGNONI P., SCOPIGNO R., MUDGE M.: Dynamic shading enhancement for reflectance transformation imaging. *ACM Journal on Computing and Cultural Heritage (JOCCH) 3* (2010), 6:1–6:20. 50

[PCDS12] PALMA G., CALLIERI M., DELLEPIANE M., SCOPIGNO R.: A statistical method for SVBRDF approximation from video sequences in general lighting conditions. *Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering) 31*, 4 (2012), 1491–1500. 41

[PD05] PEERS P., DUTRÉ P.: Inferring reflectance functions from wavelet noise. In *Proceedings of the Eurographics Symposium on Rendering (EGSR)* (2005), pp. 173–182. 51

[PDCS13] PALMA G., DESOGUS N., CIGNONI P., SCOPIGNO R.: Surface light field from video acquired in uncontrolled settings. In *Proceedings of the International Congress on Digital Heritage* (2013), pp. 31–38. 50

[Pho75] PHONG B. T.: Illumination for computer generated pictures. *Communications of the ACM 18*, 6 (1975), 311–317. 37

[PKHK15] PETERS C., KLEIN J., HULLIN M. B., KLEIN R.: Solving trigonometric moment problems for fast transient imaging. In *ACM SIGGRAPH Asia 2015 Papers (conditionally accepted)* (2015). 20

[PML*09] PEERS P., MAHAJAN D. K., LAMOND B., GHOSH A., MATUSIK W., RAMAMOORTHI R., DEBEVEC P.: Compressive light transport sensing. *ACM Transactions on Graphics 28*, 1 (2009), 3:1–3:18. 51

[PSCS*12] PACANOWSKI R., SALAZAR-CELIS O., SCHLICK C., GRANIER X., PIERRE P., ANNIE C.: Rational BRDF. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 18*, 11 (2012), 1824–1835. 37

[PSM06] PADFIELD J., SAUNDERS D., MALZBENDER T.: Polynomial texture mapping: a new tool for examining the surface of paintings. In *Triennial Meeting of the ICOM Committee for Conservation* (2006), vol. 1, pp. 504Ű–510. 50

[PvBM*06] PEERS P., VOM BERGE K., MATUSIK W., RAMAMOORTHI R., LAWRENCE J., RUSINKIEWICZ S., DUTRÉ P.: A compact factored representation of heterogeneous subsurface scattering. *ACM Transactions on Graphics 25*, 3 (2006), 746–753. 49

[RB06] ROTH S., BLACK M. J.: Specular flow and the recovery of surface structure. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2006), vol. 2, pp. 1869–1876. 25

[RH01] RAMAMOORTHI R., HANRAHAN P.: A signal-processing framework for inverse rendering. In *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)* (2001), pp. 117–128. 37

[RK09] RUITERS R., KLEIN R.: Heightfield and spatially varying brdf reconstruction for materials with interreflections. *Computer Graphics Forum (Proceedings of Eurographics) 28*, 2 (2009), 513–522. 23, 24, 41

[RPG14] RIVIERE J., PEERS P., GHOSH A.: Mobile surface reflectometry. In *ACM SIGGRAPH 2014 Posters* (2014), pp. 68:1–68:1. 42

[RS13] REMONDINO F., STOPPA D.: *TOF range-imaging cameras*. Springer, Berlin, Germany, 2013. 20

[RSFM10a] RANTOSON R., STOLZ C., FOFI D., MERIAUDEAU F.: 3d reconstruction of transparent objects exploiting surface fluorescence caused by UV irradiation. In *Proceedings of the IEEE International Conference on Image Processing (ICIP)* (2010), pp. 2965–2968. 29

[RSFM10b] RANTOSON R., STOLZ C., FOFI D., MERIAUDEAU F.: Non contact 3d measurement scheme for transparent objects using UV structured light. In *Proceedings of the International Conference on Pattern Recognition (ICPR)* (2010), pp. 1646–1649. 29

[RSGS10] RITZ M., SCHOLZ M., GOESELE M., STORK A.: High resolution acquisition of detailed surfaces with lens-shifted structured light. In *Proceedings of the International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage (VAST)* (2010), pp. 1–8. 21

[RSK10] RUMP M., SARLETTE R., KLEIN R.: Groundtruth data for multispectral bidirectional texture functions. In *Proceedings of the European Conference on Colour in Graphics, Imaging, and Vision (CGIV)* (2010), pp. 326–330. 43

[RSK12] RUITERS R., SCHWARTZ C., KLEIN R.: Data driven surface reflectance from sparse and irregular samples. *Computer Graphics Forum 31*, 2 (2012), 315–324. 42

[Rui13] RUITERS R. A.: *Data-driven analysis and interpolation of optical material properties*. PhD thesis, Rheinische Friedrich-Wilhelms-Universität Bonn, Bonn, Germany, 2013. 56

[RWS*11] REN P., WANG J., SNYDER J., TONG X., GUO B.: Pocket reflectometry. *ACM Transactions on Graphics 30*, 4 (2011), 45:1–45:10. 37

[SBKC13] STURM J., BYLOW E., KAHL F., CREMERS D.: CopyMe3D: Scanning and printing persons in 3D. In *Proceedings of the German Conference on Pattern Recognition (GCPR)* (2013), pp. 405–414. 52, 54

[SCD*06] SEITZ S. M., CURLESS B., DIEBEL J., SCHARSTEIN D., SZELISKI R.: A comparison and evaluation of multi-view stereo reconstruction algorithms. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2006), vol. 1, pp. 519–528. 17

[Sch94] SCHLICK C.: An inexpensive brdf model for physically-based rendering. *Computer Graphics Forum 13*, 3 (1994), 233–246. 37

[Sch97] SCHILLING A.: Towards real-time photorealistic rendering: Challenges and solutions. In *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS Workshop on Graphics Hardware (HWWS)* (1997), pp. 7–15. 8

[Sch14] SCHWARTZ J. C.: *Acquisition, transmission and rendering of objects with optically complicated material appearance*. PhD thesis, Rheinische Friedrich-Wilhelms-Universität Bonn, Bonn, Germany, 2014. 43, 57, 58, 59, 60

[SCN06] SAXENA A., CHUNG S. H., NG A. Y.: Learning depth from single monocular images. In *Advances in Neural Information Processing Systems*, vol. 18. MIT Press, 2006, pp. 1161–1168. 22

[SCP05] SAVARESE S., CHEN M., PERONA P.: Local shape from mirror reflections. *International Journal of Computer Vision (IJCV) 64*, 1 (2005), 31–67. 26

[SD09] SEN P., DARABI S.: Compressive dual photography. *Computer Graphics Forum 28*, 2 (2009), 609–618. 51

[SF11] SILBERMAN N., FERGUS R.: Indoor scene segmentation using a structured light sensor. In *IEEE International Conference on Computer Vision Workshops (ICCV Workshops)* (2011), pp. 601–608. 22

[SG31] SMITH T., GUILD J.: The C.I.E. colorimetric standards and their use. *Transactions of the Optical Society 33*, 3 (1931), 73–134. 7

[Sha85] SHAFER S. A.: Using color to separate reflection components. *Color Research & Application 10*, 4 (1985), 210–218. 48

[SHKF12] SILBERMAN N., HOIEM D., KOHLI P., FERGUS R.: Indoor segmentation and support inference from rgbd images. *Proceedings of the European Conference on Computer Vision (ECCV)* (2012), 746Ű760. 22

[SJ09] STILLA U., JUTZI B.: Waveform analysis for small-footprint pulsed laser systems. In *Topographic Laser Ranging and Scanning: Principles and Processing* (2009), CRC Press, Boca Raton, USA, pp. 215Ű–234. 19

[SL00] SKOCAJ D., LEONARDIS A.: Range image acquisition of objects with non-uniform albedo using structured light range sensor. In *Proceedings of the International Conference on Pattern Recognition (ICPR)* (2000), vol. 1, pp. 778–781. 24

[SP05] SINHA S. N., POLLEFEYS M.: Multi-view reconstruction using photo-consistency and exact silhouette constraints: A maximum-flow formulation. In *Proceedings of the International Conference on Computer Vision (ICCV)* (2005), vol. 1, pp. 349–356. 23

[SPB04] SALVI J., PAGÈS J., BATLLE J.: Pattern codification strategies in structured light systems. *Pattern Recognition 37* (2004), 827–849. 20

[Sph] SPHEREOPTICS: Reflectance coatings and materials – technical information. http://www.laser2000.fr/fileadmin/Produktgruppenkataloge/SPO_MaterialsCoatingsExtendedGuide2009.pdf (accessed on 19 March 2014). URL: http://www.laser2000.fr/fileadmin/Produktgruppenkataloge/SPO_MaterialsCoatingsExtendedGuide2009.pdf. 14

[SRT*14] SANTOS P., RITZ M., TAUSCH R., SCHMEDT H., MONROY R., DE STEFANO A., POSNIAK O., FUHRMANN C., FELLNER D. W.: CultLab3D - On the Verge of 3D Mass Digitization. In *Eurographics Workshop on Graphics and Cultural Heritage* (2014). 54

[SRWK11] SCHWARTZ C., RUITERS R., WEINMANN M., KLEIN R.: WebGL-based streaming and presentation framework for bidirectional texture functions. In *Proceedings of the International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage (VAST)* (2011), pp. 113–120. 52

[SRWK13] SCHWARTZ C., RUITERS R., WEINMANN M., KLEIN R.: Webgl-based streaming and presentation of objects with bidirectional texture functions. *ACM Journal on Computing and Cultural Heritage (JOCCH) 6*, 3 (2013), 11:1–11:21. 53

[SS95] SCHRÖDER P., SWELDENS W.: Spherical wavelets: Efficiently representing functions on the sphere. In *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)* (1995), pp. 161–172. 37

[SS02] SCHARSTEIN D., SZELISKI R.: A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *International Journal of Computer Vision (IJCV) 47* (2002), 7–42. 17

[SS03] SCHARSTEIN D., SZELISKI R.: High-accuracy stereo depth maps using structured light. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2003), vol. 1, pp. 195–202. 24

[SSK03] SATTLER M., SARLETTE R., KLEIN R.: Efficient and realistic visualization of cloth. In *Proceedings of the Eurographics Symposium on Rendering (EGSR)* (2003), pp. 167–177. 37, 43, 44

[SSN09] SAXENA A., SUN M., NG A. Y.: Make3d: Learning 3d scene structure from a single still image. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 31*, 5 (2009), 824–840. 22

[SSR*06] SUN B., SUNKAVALLI K., RAMAMOORTHI R., BELHUMEUR P., NAYAR S.: Time-varying BRDFs. In *Proceedings of the Eurographics Workshop on Natural Phenomena* (2006), pp. 15–23. 38

[SSS*08] SINHA S. N., STEEDLY D., SZELISKI R., AGRAWALA M., POLLEFEYS M.: Interactive 3d architectural modeling from unordered photo collections. *ACM Transactions on Graphics 27*, 5 (2008), 159:1–159:10. 35, 36

[SSW*14] SCHWARTZ C., SARLETTE R., WEINMANN M., RUMP M., KLEIN R.: Design and implementation of practical bidirectional texture function measurement devices focusing on the developments at the University of Bonn. *Sensors 14*, 5 (2014), 7753–7819. 2, 14, 43, 45, 46, 47, 52

[SSWK13] SCHWARTZ C., SARLETTE R., WEINMANN M., KLEIN R.: DOME II: A parallelized BTF acquisition system. In *Proceedings of the Eurographics Workshop on Material Appearance Modeling: Issues and Acquisition* (2013), pp. 25–31. 13, 35, 45, 46, 47, 52, 56, 57

[STD09] SANSONI G., TREBESCHI M., DOCCHIO F.: State-of-the-art and applications of 3d imaging sensors in industry, cultural heritage, medicine, and criminal investigation. *Sensors 9*, 1 (2009), 568–601. 17, 20

[Ste57] STEINHAUS H.: Sur la division des corps matériels en parties. *Bulletin de l'Académie Polonaise des Sciences, Classe III 4* (1957), 801–804. 27

[Str83] STRANG G.: Maximal flow through a domain. *Mathematical Programming 26*, 2 (1983), 123–143. 23, 27, 31

[Str10] STRANG G.: Maximum flows and minimum cuts in the plane. *Journal of Global Optimization 47*, 3 (2010), 527–535. 23, 27

[SVTA10] SANKARANARAYANAN A. C., VEERARAGHAVAN A., TUZEL O., AGRAWAL A. K.: Specular surface reconstruction from sparse reflection correspondences. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2010), pp. 1245–1252. 26

[SWI97] SATO Y., WHEELER M. D., IKEUCHI K.: Object shape and reflectance modeling from observation. In *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)* (1997), pp. 379–387. 39

[SWN88] SANDERSON A. C., WEISS L. E., NAYAR S. K.: Structured highlight inspection of specular surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 10*, 1 (1988), 44–55. 21, 26

[SWRK11] SCHWARTZ C., WEINMANN M., RUITERS R., KLEIN R.: Integrated high-quality acquisition of geometry and appearance for cultural heritage. In *Proceedings of the International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage (VAST)* (2011), pp. 25–32. 14, 35, 36, 46, 47, 52, 56, 57

[Sze11] SZELISKI R.: *Computer vision: Algorithms and applications*. Springer, London, UK, 2011. 17

[TAN*05] TSUCHIDA M., ARAI H., NISHIKO M., SAKAGUCHI Y., UCHIYAMA T., YAMAGUCHI M., HANEISHI H., OHYAMA N.: Development of BRDF and BTF measurement and computer-aided design systems based on multispectral imaging. In *AIC Colour 05 - 10th Congress of the International Colour Association*. 2005, pp. 129–132. 43

[TBH06] TRIFONOV B., BRADLEY D., HEIDRICH W.: Tomographic reconstruction of transparent objects. In *Proceedings of the Eurographics Symposium on Rendering (EGSR)* (2006), pp. 51–60. 29

[TFG*13] TUNWATTANAPONG B., FYFFE G., GRAHAM P., BUSCH J., YU X., GHOSH A., DEBEVEC P.: Acquiring reflectance and shape from continuous spherical harmonic illumination. *ACM Transactions on Graphics 32*, 4 (2013), 109:1–109:12. 27, 42

[TGL*06] TARIQ S., GARDNER A., LLAMAS I., JONES A., DEBEVEC P., TURK G.: Efficient estimation of spatially varying subsurface scattering parameters. In *Proceedings of the International Fall Workshop on Vision, Modeling and Visualization* (2006), pp. 129–136. 49

[TI03] TAN R., IKEUCHI K.: Separating reflection components of textured surfaces using a single image. In *Proceedings of the International Conference on Computer Vision (ICCV)* (Oct 2003), vol. 2, pp. 870–877. 25

[TI05] TAN R., IKEUCHI K.: Separating reflection components of textured surfaces using a single image. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 27*, 2 (2005), 178–193. 25

[TLGS05] TARINI M., LENSCH H. P. A., GOESELE M., SEIDEL H.-P.: 3d acquisition of mirroring objects using striped patterns. *Graphical Models 67*, 4 (2005), 233–259. 26, 27

[TLQS03] TAN P., LIN S., QUAN L., SHUM H.-Y.: Highlight removal by illumination-constrained inpainting. In *Proceedings of the International Conference on Computer Vision (ICCV)* (Oct 2003), vol. 1, pp. 164–169. 25

[TM08] TUYTELAARS T., MIKOLAJCZYK K.: Local invariant feature detectors: A survey. *Foundations and Trends in Computer Graphics and Vision 3*, 3 (2008), 177–280. 18

[TMHF00] TRIGGS B., MCLAUCHLAN P. F., HARTLEY R. I., FITZGIBBON A. W.: Bundle adjustment - a modern synthesis. In *Proceedings of the International Workshop on Vision Algorithms: Theory and Practice* (2000), pp. 298–372. 17

[TMY12] TAGAWA S., MUKAIGAWA Y., YAGI Y.: 8-d reflectance field for computational photography. In *Proceedings of the International Conference on Pattern Recognition (ICPR)*. 2012, pp. 2181–2185. 44

[TS67] TORRANCE K. E., SPARROW E. M.: Theory for off-specular reflection from roughened surfaces. *Journal of the Optical Society of America 57*, 9 (1967), 1105–1112. 24, 25, 37

[TS91] THOMAS D., SUGIMOTO A.: Range image registration of specular objects under complex illumination. *Proceedings of the International Conference on 3D Data Processing, Visualization and Transmission 13*, 7 (1991), 635–657. 25

[TSA*05] TSUCHIDA M., SAKAGUCHI Y., ARAI H., NISHIKO M., FUJIKAWA N., YAMAGUCHI M., HANEISHI H., OHYAMA N.: High-fidelity color reproduction based on multi-channel BTF/BRDF acquisition, rendering and display. In *ACM SIGGRAPH 2005 Sketches*. 2005. 43

[TWL*05] TONG X., WANG J., LIN S., GUO B., SHUM H.-Y.: Modeling and rendering of quasi-homogeneous materials. *ACM Transactions on Graphics 24*, 3 (2005), 1054–1061. 45, 47, 49

[UB15] UMMENHOFER B., BROX T.: Global, dense multiscale reconstruction for a billion points. In *Proceedings of the International Conference on Computer Vision (ICCV)* (2015). 34

[VHTC07] VOGIATZIS G., HERNÁNDEZ C., TORR P. H. S., CIPOLLA R.: Multi-view stereo via volumetric graph-cuts and occlusion robust photo-consistency. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 29*, 12 (2007), 2241–2246. 23

[VSJ07] VELHO L., SOSSAI JR. J.: Projective texture atlas construction for 3d photography. *The Visual Computer 23*, 9-11 (2007), 621–629. 35, 36

[WAA*00] WOOD D. N., AZUMA D. I., ALDINGER K., CURLESS B., DUCHAMP T., SALESIN D. H., STUETZLE W.: Surface light fields for 3D photography. In *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)* (2000), pp. 287–296. 50

[War92] WARD G. J.: Measuring and modeling anisotropic reflection. *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH) 26*, 2 (1992), 265–272. 24, 37, 38, 41, 42, 60

[WAT92] WESTIN S. H., ARVO J. R., TORRANCE K. E.: Predicting reflectance functions from complex surfaces. *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH) 26*, 2 (1992). 37

[WB91] WOLFF L. B., BOULT T. E.: Constraining object features using a polarization reflectance model. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 13*, 7 (1991), 635–657. 25, 48

[WD06] WANG J., DANA K.: Relief texture from specularities. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 28*, 3 (2006), 446–457. 44

[Wei13] WEINMANN M.: Visual features – from early concepts to modern computer vision. In *Advanced Topics in Computer Vision*, Advances in Computer Vision and Pattern Recognition. Springer London, 2013, pp. 1–34. 18

[WGS*11] WU L., GANESH A., SHI B., MATSUSHITA Y., WANG Y., MA Y.: Robust photometric stereo via low-rank matrix completion and recovery. In *Proceedings of the Asian Conference on Computer Vision (ACCV)* (2011), pp. 703–717. 19

[WGT*05] WENGER A., GARDNER A., TCHOU C., UNGER J., HAWKINS T., DEBEVEC P.: Performance relighting and reflectance transformation with time-multiplexed illumination. *ACM Transactions on Graphics 24*, 3 (2005), 756–764. 45, 50

[WI93] WANG Z., INOKUCHI S.: Determining shape of specular surfaces. In *Proceedings of the Scandinavian Conference on Image Analysis (SCIA)* (1993), pp. 1187–1194. 27

[Wit80] WITKIN A. P.: *Shape from contour*. Tech. rep., Massachusetts Institute of Technology, Cambridge, MA, USA, 1980. 16

[Wit81] WITKIN A. P.: Recovering surface shape and orientation from texture. *Artificial Intelligence 17* (1981), 17–45. 16

[WJV*05] WILBURN B., JOSHI N., VAISH V., TALVALA E.-V., ANTUNEZ E., BARTH A., ADAMS A., HOROWITZ M., LEVOY M.: High performance imaging using large camera arrays. In *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)* (2005), pp. 765–776. 50

[WK15a] WEINMANN M., KLEIN R.: Material recognition for efficient acquisition of geometry and reflectance. In *Proceedings of the ECCV 2014 Workshops* (2015), pp. 321–333. 56

[WK15b] WEINMANN M., KLEIN R.: A short survey on optical material recognition. In *Proceedings of the Eurographics Workshop on Material Appearance Modeling* (2015), pp. 35–42. 56

[WLDW11] WU C., LIU Y., DAI Q., WILBURN B.: Fusing multiview and photometric stereo for 3d reconstruction under uncalibrated illumination. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 17*, 8 (2011), 1082–1095. 19, 23, 45

[WLH*14] WEINMANN M., LEITLOFF J., HOEGNER L., JUTZI B., STILLA U., HINZ S.: Thermal 3d mapping for object detection in dynamic scenes. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences II-1* (2014), 53–60. 20

[WLL*07] WEYRICH T., LAWRENCE J., LENSCH H., RUSINKIEWICZ S., ZICKLER T.: Principles of appearance acquisition and representation. In *ICCV 2007 Short Course* (2007). 41, 48

[WLZ*09] WANG H., LIAO M., ZHANG Q., YANG R., TURK G.: Physically guided liquid surface modeling from videos. *ACM Transactions on Graphics 28*, 3 (2009), 90:1–90:11. 28

[WMG14] WAECHTER M., MOEHRLE N., GOESELE M.: Let there be color! Large-scale texturing of 3d reconstructions. In *Proceedings of the European Conference on Computer Vision (ECCV)* (2014), pp. 836–850. 35, 36, 37, 60

[WMP*05] WEYRICH T., MATUSIK W., PFISTER H., NGAN A., GROSS M.: *Measuring Skin Reflectance and Subsurface Scattering*. Tech. rep., Mitsubishi Electric Research Laboratories, TR2005-046, 2005. 45, 46, 47

[WMP*06] WEYRICH T., MATUSIK W., PFISTER H., BICKEL B., DONNER C., TU C., MCANDLESS J., LEE J., NGAN A., JENSEN H. W., GROSS M.: Analysis of human faces using a measurement-based skin reflectance model. *ACM Transactions on Graphics 25* (2006), 1013–1024. 41, 45, 46, 47, 49

[Wol89] WOLFF L.: Using polarization to separate reflection components. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (1989), pp. 363–369. 25

[Woo80] WOODHAM R. J.: Photometric method for determining surface orientation from multiple images. *Optical Engineering 19*, 1 (1980), 139–144. 18, 26

[WORK13] WEINMANN M., OSEP A., RUITERS R., KLEIN R.: Multi-view normal field integration for 3d reconstruction of mirroring objects. In *Proceedings of the International Conference on Computer Vision (ICCV)* (2013), pp. 2504–2511. 19, 27, 28, 30, 31, 55, 57

[WRO*12] WEINMANN M., RUITERS R., OSEP A., SCHWARTZ C., KLEIN R.: Fusing structured light consistency and Helmholtz normals for 3d reconstruction. In *Proceedings of the British Machine Vision Conference (BMVC)* (2012), pp. 1–12. 13, 23, 25, 26, 27, 30, 31, 55, 56, 57

[WSB*98] WHITE D. R., SAUNDERS P., BONSEY S. J., VAN DE VEN J., EDGAR H.: Reflectometer for measuring the bidirectional reflectance of rough surfaces. *Applied Optics 37*, 16 (1998), 3450–3454. 38

[WSRK11] WEINMANN M., SCHWARTZ C., RUITERS R., KLEIN R.: A multi-camera, multi-projector super-resolution framework for structured light. In *Proceedings of the International Conference on 3D Imaging, Modeling, Processing, Visualization and Transmission (3DIMPVT)* (2011), pp. 397–404. 13, 21, 24, 26, 47, 52, 56, 57

[WWHL07] WEISTROFFER R. P., WALCOTT K. R., HUMPHREYS G., LAWRENCE J.: Efficient basis decomposition for scattered reflectance data. In *Proceedings of the Eurographics Symposium on Rendering (EGSR)* (2007), pp. 207–218. 37

[WZT*08] WANG J., ZHAO S., TONG X., SNYDER J., GUO B.: Modeling anisotropic surface reflectance with example-based microfacet synthesis. *ACM Transactions on Graphics 27*, 3 (2008), 41:1–41:9. 50, 51

[XS97] XIONG Y., SHAFER S. A.: Moment and hypergeometric filters for high precision computation of focus, stereo and optical flow. *International Journal of Computer Vision (IJCV) 22*, 1 (1997), 25–59. 16

[YAC06] YU T., AHUJA N., CHEN W. C.: Sdg cut: 3d reconstruction of non-lambertian objects using graph cuts on surface distance grid. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2006), vol. 2, pp. 2269–2276. 23

[YBT10] YUAN J., BAE E., TAI X.-C.: A study on continuous max-flow and min-cut approaches. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2010), pp. 2217–2224. 23, 31

[YIX07] YAMAZAKI M., IWATA S., XU G.: Dense 3d reconstruction of specular and transparent objects using stereo cameras and phase-shift method. In *Proceedings of the Asian Conference on Computer Vision (ACCV)* (2007), pp. 570–579. 26, 55

[YWT*11] YEUNG S.-K., WU T.-P., TANG C.-K., CHAN T. F., OSHER S.: Adequate reconstruction of transparent objects on a shoestring budget. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2011), pp. 2513–2520. 27, 55

[YY11] YOSHIYASU Y., YAMAZAKI N.: Topology-adaptive multi-view photometric stereo. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2011), pp. 1001–1008. 23

[ZBK02] ZICKLER T. E., BELHUMEUR P. N., KRIEGMAN D. J.: Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction. *International Journal of Computer Vision (IJCV) 49*, 2-3 (2002), 215–227. 23, 25, 30, 50

[ZGB89] ZISSERMAN A., GIBLIN P., BLAKE A.: The information available to a moving observer from specularities. *Image and Vision Computing 7*, 1 (1989), 38–42. 21, 26

[ZGHG11] ZHOU K., GONG M., HUANG X., GUO X.: Data-parallel octrees for surface reconstruction. *IEEE Transactions on Visualization and Computer Graphics 17*, 5 (2011), 669–681. 32

[ZHK*03] ZICKLER T. E., HO J., KRIEGMAN D. J., PONCE J., BELHUMEUR P. N.: Binocular Helmholtz stereopsis. In *Proceedings of the International Conference on Computer Vision (ICCV)* (2003), vol. 2, pp. 1411–1417. 25

[ZM00] ZHENG J. Y., MURATA A.: Acquiring a complete 3d model from specular motion under the illumination of circular-shaped light sources. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 22*, 8 (2000), 913–920. 26

[ZMLC10] ZHENG Z., MA L., LI Z., CHEN Z.: An extended photometric stereo algorithm for recovering specular object shape and its reflectance properties. *Computer Science and Information Systems 7*, 1 (2010), 1–12. 24

[ZREB06] ZICKLER T., RAMAMOORTHI R., ENRIQUE S., BELHUMEUR P. N.: Reflectance sharing: Predicting appearance from a sparse set of images of a known shape. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 28*, 8 (2006), 1287–1302. 37, 41

[ZTCS99] ZHANG R., TSAI P., CRYER J. E., SHAH M.: Shape from shading: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 21*, 8 (1999), 690–706. 16

[ZWWW10] ZHANG W., WANG F., WANG Z., WANG H.: Measuring of spectral BRDF using fiber optic spectrometer. In *Proceedings of SPIE* (2010), vol. 7658, pp. 76582P–1 – 76582P–8. 38