

Transparent and Specular Object Reconstruction

Ivo Ihrke^{1,2}, Kiriakos N. Kutulakos³, Hendrik P. A. Lensch⁴, Marcus Magnor⁵ and Wolfgang Heidrich¹

¹University of British Columbia, Canada
ihrke@mci.uni-saarland.de,
heidrich@cs.ubc.ca

²Saarland University, MPI Informatik, Germany

³University of Toronto, Canada
kyros@cs.toronto.edu

⁴Ulm University, Germany
hendrik.lensch@uni-ulm.de

⁵TU Braunschweig, Germany
magnor@cg.tu-bs.de

Abstract

This state of the art report covers reconstruction methods for transparent and specular objects or phenomena. While the 3D acquisition of opaque surfaces with Lambertian reflectance is a well-studied problem, transparent, refractive, specular and potentially dynamic scenes pose challenging problems for acquisition systems. This report reviews and categorizes the literature in this field.

Despite tremendous interest in object digitization, the acquisition of digital models of transparent or specular objects is far from being a solved problem. On the other hand, real-world data is in high demand for applications such as object modelling, preservation of historic artefacts and as input to data-driven modelling techniques. With this report we aim at providing a reference for and an introduction to the field of transparent and specular object reconstruction.

We describe acquisition approaches for different classes of objects. Transparent objects/phenomena that do not change the straight ray geometry can be found foremost in natural phenomena. Refraction effects are usually small and can be considered negligible for these objects. Phenomena as diverse as fire, smoke, and interstellar nebulae can be modelled using a straight ray model of image formation. Refractive and specular surfaces on the other hand change the straight rays into usually piecewise linear ray paths, adding additional complexity to the reconstruction problem. Translucent objects exhibit significant sub-surface scattering effects rendering traditional acquisition approaches unstable. Different classes of techniques have been developed to deal with these problems and good reconstruction results can be achieved with current state-of-the-art techniques. However, the approaches are still specialized and targeted at very specific object classes. We classify the existing literature and hope to provide an entry point to this exiting field.

Keywords: range scanning, transparent, specular, and volumetric objects

ACM CCS: I.4.8 Scene Analysis, Range Data, Shape I.2.10 Vision and Scene Understanding, 3D Scene Analysis.

1. Introduction

The acquisition of three-dimensional (3D) real world objects or phenomena is an important topic in computer graphics as well as in computer vision. Current rendering tech-

niques achieve a high degree of realism once suitable computer models are available. However, manual generation of digital content is a labour-intensive task. Therefore, object









object type	surface / volume type	class	image formation
opaque	surface, rough	①	diffuse or near diffuse reflectance 
	surface, glossy	②	mixed diffuse and specular reflectance 
translucent	surface, smooth	③	ideal or near ideal specular reflectance 
	surface, sub-surface scattering	④	multiple scattering underneath surface 
transparent	surface, smooth	⑤	ideal or near ideal specular refraction 
	volume, emission / absorption	⑥	integration along viewing ray 
	volume, single scattering	⑦	integration along viewing ray 
	volume, multiple scattering	⑧	full global light transport without occluders 
inhomogeneous	mixed scenes, containing many / all of the above	⑨	full global light transport

Figure 1: A taxonomy of object classes based on increasing complexity in light transport. In this report we focus on reconstruction approaches for object classes 3–7 (yellow box).

digitization techniques that automatically generate digital models from real-world objects have been of considerable interest, both in research and in the industry.

Most techniques that have been developed over the past two decades have focused on opaque objects with Lambertian reflectance though this has changed in the last 5 years [ZBK02, DYW05]. There has been tremendous progress in this area, however, large classes of objects/phenomena still pose difficulties for traditional acquisition techniques. Because the majority of object acquisition approaches rely on observing light reflected off a surface, objects made of materials that exhibit significant effects of global light transport or that are simply too dark are difficult to handle.

In Figure 1, we show a taxonomy of object classes with different material properties giving rise to different modes of light transport. While methods for the acquisition of diffuse and glossy surfaces (green box) have been extensively studied, there are objects causing more complex light transport effects (yellow box). This report reviews state-of-the-art reconstruction techniques for object classes 3–7. Examples of objects or phenomena that can be modelled using these techniques include mirror-like objects, glass objects, water surfaces and natural phenomena like fire, smoke and interstellar nebulae. We focus on techniques that result in 3D models of objects' surfaces or volumetric descriptions, excluding purely image based approaches such as, for example, environment matting [ZWCS99, CZH*00, WFZ02, PD03, AMKB04] and approaches that enhance coarse geometry models, for example, opacity hulls [MPN*02, MPZ*02]. Image-based techniques are useful for rendering acquired objects with good quality, however it is not clear how to analyse and modify the data which is much simpler when surface geometry or volumetric object descriptions are available.

For object classes 8 and 9, full global illumination effects have to be taken into account. Participating media can be made more transparent by employing controlled illumination [GNS08]. Opaque objects immersed in participating media can be acquired using structured light techniques [NNSK05, FHL*08], and the scattering parameters of a homogeneous participating medium can be measured [NGD*06, JDJ06]. However, there are, with the exception of [FHL*08], no reconstruction techniques for inhomogeneous participating media exhibiting multiple-scattering effects in the computer graphics and computer vision literature. The problem is studied in the context of medical imaging, see for example [COW*96], having applications in ultrasound imaging of tissue.

The report covers principles and practice of automated acquisition techniques for transparent, refractive, specular and translucent objects. We review the major experimental setups, principles of surface or volume acquisition and experimental results for these types of objects. We discuss advantages and drawbacks of different methods with respect to each other and try to assess the current state-of-the-art in the field.

1.1. Overview of traditional diffuse object acquisition and its extensions

3D geometry acquisition is one of the major research directions of computer vision and related engineering disciplines. Several decades of development have led to reliable acquisition techniques for diffuse objects (Figure 1, class 1). A wide range of methods have been proposed, which can be coarsely divided into active and passive range sensing. Active range scanning techniques actively control the lighting in the scene, for example by projecting patterns of light, making feature

detection more reliable than in the uncontrolled case of passive sensing. Davis *et al.* [DNRR05] present a framework that unifies and generalizes active and passive range sensing approaches. Sensor fusion of active and passive range sensing techniques is discussed by Beraldin [Ber04].

1.1.1. Active structured light scanning

Examples of active range scanning include laser stripe projection, various structured light projection systems and time-of-flight scanners. For an overview of current state-of-the-art techniques in active range scanning we refer the interested reader to [Bla04]. Active light range scanning techniques belong to the most accurate object acquisition approaches known today. However, most of them rely on a clearly detectable pattern of light being reflected off the object's surface. Objects exhibiting significant effects of global light transport such as specular, refractive and translucent objects pose major difficulties for the proper analysis of the sensed reflection pattern. Another class of objects that is difficult to handle are objects with a very low surface albedo.

Active light range scanning techniques have been made more robust with respect to variations in surface reflectance by analysing space-time properties of the reflected light [CL95]. Curless and Levoy [CL95] show that varying surface reflectance as well as self-occlusions and object edges result in a systematic error in depth estimation. The proposed space-time analysis significantly improves range scanning results for glossy surfaces (Figure 1, class 2).

Trucco and Fisher [TF94] investigate the use of a second CCD sensor to disambiguate detected range data. They propose a number of consistency checks to exclude false measurements from the processing pipeline. Park and Kak [PK04, PK08] consider more extreme cases of non-Lambertian surfaces in systems based on laser stripe projection. They observe that for these types of surfaces, the reflected light often results in multiple peaks per scan-line of the imaging sensor, contrary to the assumption of a single peak being made in standard laser range scanning approaches. They suggest filtering methods based on local smoothness and global consistency and visibility constraints to clean up the recovered point clouds and achieve good results even for specular surfaces (Figure 1, class 3).

Another approach, based on polarization analysis of the reflected light is presented by Clark *et al.* [CTW97]. The laser stripe projector is equipped with a polarization filter. Three different measurements are taken with differently polarized laser light and the polarization state of the reflected light patterns is analysed. The recovered range scans are shown to be significantly more robust towards specular reflections than standard laser stripe projection techniques. Results are shown on industrial pieces, made of polished aluminum.

A different approach for range scanning of optically challenging objects has been proposed recently by Hullin *et al.*



Figure 2: Left panel: surface scanning of glass objects [HFI*08] by detecting laser sheet cut-off instead of the laser reflection off the object's surface. Right panel: resulting range scan.

[HFI*08]. The object is immersed in a fluorescent liquid and illuminated by laser planes as in standard laser range scanning. However, the fluorescent solution renders the light sheets visible, whereas the surface now disrupts the propagation of light and appears dark. Thus, the intersection between the plane of laser light and the surface can be identified, and 3D points can be triangulated. An example result for range scanning of a glass object is shown in Figure 2. Note however that this approach cannot be used with specular surfaces since the visible light sheets are now reflected by the object and are thus visible as a mirror image. This complicates the detection of the transition between free space and object.

1.1.2. Passive range sensing

Passive techniques in comparison do not influence the scene lighting and are thus more applicable in remote sensing applications. A variety of approaches, exploiting different properties of light reflection, have been proposed in the literature. These approaches include stereo and multi-view stereo techniques. A recent review article covering the majority of approaches is [REH06]. A performance evaluation of multi-view stereo techniques has been performed by Seitz *et al.* [SCD*06].

Passive range sensing usually makes assumptions about the material properties of the scene, the most common being Lambertian surface reflectance. However, recent research has aimed to relax this constraint. There are a variety of approaches targeting objects with non-Lambertian reflectance properties. For objects exhibiting Lambertian reflectance, a surface point can be assumed to have a similar colour in images taken from different view-points. This is no longer true for non-Lambertian surfaces. Large parts of the literature on non-Lambertian (class 2) surface reconstruction is based on the extension of (multi-view) stereo techniques.

One approach to extend multi-view stereo techniques to handle glossy surfaces is based on detecting specular highlights in the data and treating them as outliers. Bhat and Nayar [BN95] use a trinocular stereo system and analyse pairs of camera images to identify highlights. Nayar *et al.* [NFB93] employ polarization filters to identify and discard specular highlights, whereas Mallick *et al.* [MZKB05] propose a colour space transformation that is invariant to changes due to highlights. Brelstaff and Blake [BB88b] also identify highlight regions in a pre-processing step using *ad hoc* constraints describing deviations from Lambertian reflectance. Li *et al.* [LLL*02] mark specular image regions based on the uncertainty of depth estimates resulting from reconstructions performed in a multi-view stereo setup.

Another approach is based on generalizing the multi-view matching constraint. Instead of assuming Lambertian reflectance and thus colour constancy of a common feature between view-points, a more sophisticated model of colour variation is used. Stich *et al.* [STM06] propose to detect discontinuities in epipolar plane images using a constant baseline multi-view stereo setup. Similarly, Yang *et al.* [YPW03] propose the use of a linear colour variation model. Jin *et al.* [JSY03, JSY05] encode the colour variation in a tensor constraint by considering the local reflectance variation around each surface point. They show that for surface materials exhibiting a ‘diffuse+specular’ reflectance the radiance tensor is of rank two. Based on this constraint they derive a multi-view surface reconstruction algorithm while simultaneously estimating reflectance properties of the surface.

1.1.3. Photometric methods

Methods that employ a static view-point and observe changes in illumination are referred to as photometric stereo techniques [Woo80]. Using the observed radiance under changing, calibrated illumination, a normal map is recovered which can be integrated to obtain surface shape. Traditionally, photometric stereo methods have assumed distant illumination, an orthographic camera view and diffuse surface reflectance. Goldman *et al.* [GCHS05] present a photometric technique that is applicable to class 2 objects (Figure 1). They simultaneously recover BRDF parameters and surface normals by representing the surface BRDF as a linear combination of two to three basis BRDFs with unknown coefficients.

The fusion of photometric stereo with multi-view stereo approaches is another direction of research that enables the BRDF-invariant reconstruction of surfaces. One constraint that can be exploited for 3D reconstruction is Helmholtz reciprocity, that is, that viewing rays and light rays can be exchanged without altering the surface reflectance. Zickler *et al.* [MKZB01, ZBK02, ZHK*03] investigate the use of stereo images, where light and camera positions are exchanged during data acquisition. This way, the reflectance does not change even in the presence of glossy materials and surface highlights become features that can be used for

reconstruction purposes. Davis *et al.* [DYW05] consider another approach to fuse photometric information with multi-view stereo. Employing a static camera setup and static light source positions, they develop a constraint based on light transport constancy. The incident radiance at every scene point is varied, but the light’s incident direction remains constant. Therefore the reflected light observed by the cameras varies by the same amount. It is shown that the light sources do not have to be calibrated and that varying illumination intensity results in a robust (multi-view) stereo matching constraint. This constraint can be used as a matching metric in standard stereo algorithms.

1.2. Definition of scope

The previous discussion provides a brief overview of the state-of-the-art techniques in 3D range sensing for objects with Lambertian or glossy surface reflectance properties (class 1 and 2 in Figure 1). However, these techniques are not applicable in the case of global light transport like found in refractive or sub-surface scattering objects. Specular objects also pose challenges to the aforementioned techniques. Furthermore, there are classes of phenomena, that do not have a proper surface and need to be described as volumetric phenomena.

In this report, we focus on techniques specialized towards these kinds of objects and phenomena. We have structured the report such that each of the object classes is dealt with in a separate section to facilitate an easy localization of the required information for a certain application, for example, for glass scanning. This decision results in a separate description of similar technical approaches to solving similar problems encountered in each of the object classes. For this reason, and to aid in the identification of similarities between reconstruction approaches for different object classes, we have included an overview of the major experimental techniques as used in practical implementations, Figure 3.

Acquisition approaches for reflective objects (class 3) and approaches based on exploiting surface reflection for shape recovery are covered in Section 2. Translucent object reconstruction techniques (class 4) are discussed in Section 3. In Section 4, we review related work regarding the acquisition of refractive objects (class 5), whereas in Section 5 we describe approaches for the acquisition of volumetric, light emitting or scattering phenomena (class 6 and 7).

Finally, we discuss the merits and drawbacks of the presented methods and try to identify future directions of research in Section 6.

2. Specular Surface Acquisition

In this section, we discuss acquisition approaches for specular surfaces (Figure 1, class 3). The reconstruction of surface

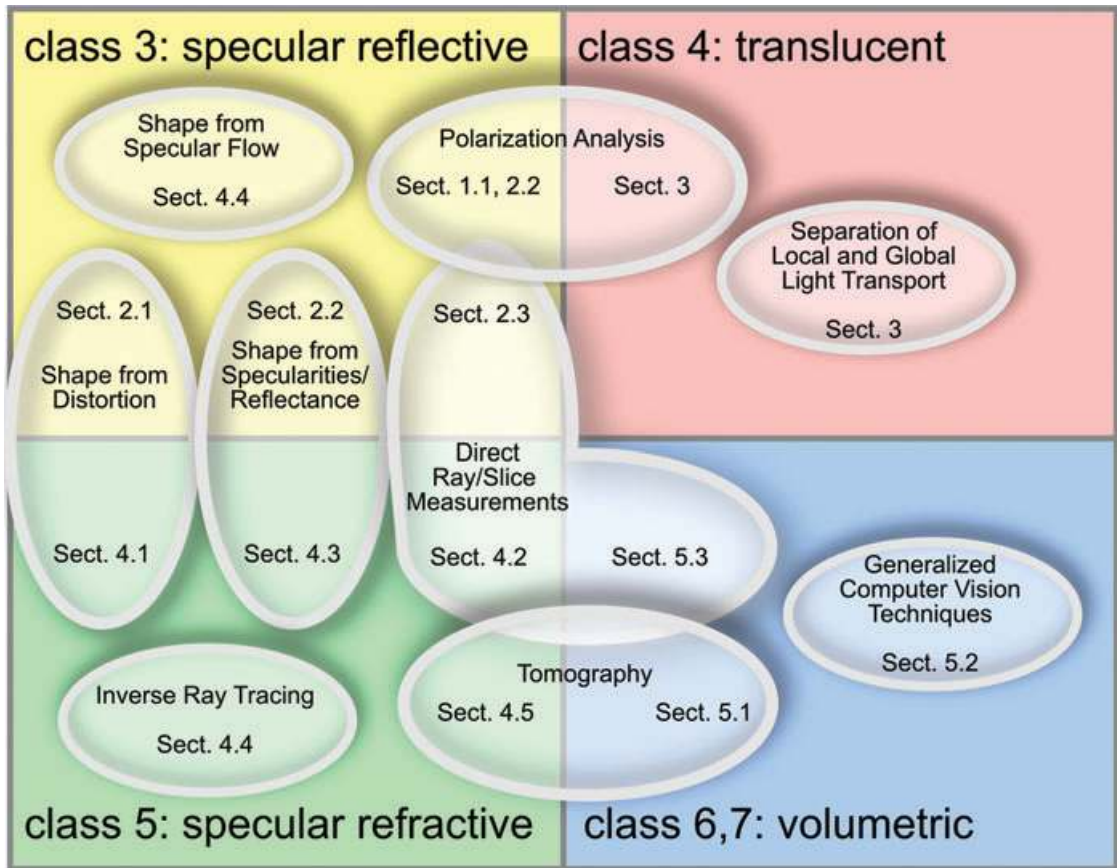


Figure 3: An overview of classes of algorithms used in the context of this report and their classification with respect to application area. The references indicate the sections where a particular technique is discussed.

geometry for specular objects is complicated by the fact that light is reflected off the surface. Therefore, there are no surface features that can be observed directly. When changing the view point, features appear to move on the surface and the law of reflection has to be taken into account. Fixing the viewing ray and a 3D world position on the incident light ray determines the depth and surface normal only up to a one-dimensional (1D) family of solutions [SWN88, SP01]. This ambiguity can be resolved by assuming distant illumination or by measuring an additional point on the incident light ray.

One group of methods for specular surface reconstruction makes use of known or unknown patterns that are distorted by specular reflection. These techniques usually assume perfect, mirror-like surface reflectance and are known as shape-from-distortion approaches, Section 2.1.

Another class of algorithms exploits surface reflectivity differently. While directly reflected light is very hard to detect for a ray-like light source, for example, a laser beam, light from point or extended light sources is usually reflected towards the imaging sensor at some points of the surface.

At these surface points highlights occur which are very disturbing for traditional passive range scanning approaches, see Section 1.1.2. However, the highlight information can be directly used to reconstruct these surface points. Techniques using this observation are termed shape-from-specularity approaches and are discussed in Section 2.2.

We refer to methods that measure two points on the light ray from the light source to a surface point as techniques based on direct ray measurements. For mirror-like objects with a single reflection event per ray, the surface can be reconstructed very accurately using this approach, Section 2.3.

2.1. Shape from distortion

Shape from distortion techniques are based on the observation of a known or unknown pattern that is distorted by a single specular surface. Multiple reflections are ignored by current techniques. The pattern is either given as a radiance map, assuming distant illumination, or placed close to the




Figure 4: The principle of shape-from-distortion based measurements. (a) The setup consists of a single or a number of patterns in fixed world positions. The pattern is illuminating a specular object diffusely while a camera takes images. (b) Captured example images. (c) The observed patterns encode one world position for every pixel of the camera. (d) From this information depth and normal can be extracted. (e) A resulting surface scan [TLGS05].

object, resulting in the depth-normal ambiguity mentioned before. The principle of the experimental setup for shape-from-distortion approaches is shown in Figure 4.

2.1.1. Calibrated patterns

One of the earliest approaches of specular surface acquisition based on shape-from-distortion was proposed by Schultz [Sch94]. The pattern is assumed to be known and consists of a partial radiance map of the sky-sphere. The author develops an algorithm based on information propagation from known seed points. Reflections on refractive and mirror-like surfaces are simulated from four viewpoints and the algorithm is evaluated on synthetic data.

Halstead *et al.* [HBKM96] present a shape-from-distortion approach for the measurement of the human cornea. They describe a one-view setup where the camera is placed at the tip of a conically shaped pattern. By observing the reflections in the human eye, and employing an inverse ray-tracing approach the authors reconstruct 3D surface models of the human eye. The reconstruction approach is iterative and performs normal fitting using a spline representation of the surface followed by a refinement step.

Bonfort and Sturm [BS03] develop a multi-view technique based on specularly reflected observations of a calibrated world pattern. The method is voxel-based and is similar to space carving techniques [KS00]. The algorithm first computes a normal for every voxel in every view. This is possible because of the depth-normal ambiguity in the one-view case where only one reflected world point is measured. In the second phase the algorithm determines the object surface by voxel colouring, the voxels with the most consistent normals for different views are considered to be surface voxels.

Nehab *et al.* [NWR08] also define a consistency measure for normal directions. They use this measure to replace the matching cost in standard stereo algorithms, an approach introduced by Sanderson *et al.* [SWN88]. However, because

the matching lacks robustness, the authors anisotropically diffuse the error landscape to remove spurious mismatches. Additionally, an analysis of reconstruction ambiguities is presented.

Tarini *et al.* [TLGS05] present a one-view approach where different patterns at the same world location are used to compute pixel to world plane correspondences with sub-pixel accuracy. The patterns are generated using a computer monitor. Since the monitor is placed in close proximity of the object the inherent depth-normal ambiguity has to be considered. The authors resolve it using an iterative approach. An initial guess for the depth value is propagated and corresponding normals are computed. The normal field is then integrated to obtain an updated depth estimate from which updated normals are computed. The process is iterated until the surface shape converges. The approach is unique in that it includes attenuation of the reflected pattern as is, for example the case in coloured metals like copper and gold. An overview of the approach which exemplifies the shape-from-distortion framework is shown in Figure 4.

2.1.2. Theoretical analysis

A theoretical analysis of shape-from-distortion for specular surfaces has been presented by Oren and Nayar [ON96] and Savarese *et al.* [SP01, SP02, SCP05].

Oren and Nayar [ON96] consider specular surface reconstruction in a structure-from-motion [HZ00] setting. The apparent motion of features in the image plane of a moving camera is analysed. The authors develop a classification between ‘real’ features, that is, world points not reflected by a specular object and ‘virtual’ features, that is, features influenced by specular reflection. The theory is based on envelopes of reflected rays, that is, caustic curves. It is shown that in the case of co-planar camera movement with respect to the surface, a profile can be computed from just two specularly reflected features. For 3D profiles, tracking of a single specular feature from the occluding boundary of the object is

sufficient to reconstruct a 3D curve on its surface. The point on the occluding boundary serves as a boundary condition since the normal is known at this position.

Savarese *et al.* [SP01, SP02, SCP05] theoretically analyse shape-from-distortion using a single, calibrated view and a known pattern with tangential information in a calibrated world position. Under these conditions, the authors analyse the differential relationship between the local geometry of a known, planar world pattern, the specular surface and the local geometry in the image plane of a camera observing it. This relationship is then inverted and necessary and sufficient conditions for the inverse mapping to exist are given. It is shown that known position and tangential information in the world plane in conjunction with second-order curve measurements in the image plane determine the position and the normal of a specular surface. In general, third-order surface information can be extracted from a single view setup with a calibrated planar scene, given the reflections of six or more scene points.

Links between differential surface properties and images of specular and refractive objects have also been analysed by Ding and Yu [YY08]. The authors employ the concept of general linear cameras [YM04] to classify surface points into one of eight classes of these types of cameras. Assuming an orthographic observing camera and height field-like surfaces without interreflections, it is shown that Gaussian and mean curvature of object surfaces can be extracted from reflection or refraction images.

The multi-perspective imaging concept employed above has also been analysed in a different context. Treibitz *et al.* [TSS08], for example show that a projective camera when used in a multi-medium setting, such as in the presence of air/water interfaces, in general changes its focusing properties. The camera is no longer focused in a single point, but the focal point is expanded into a surface; the locus of focal points or viewpoint caustic. Caustic surfaces due to catadioptric imaging systems have been analysed by Swaminathan *et al.* [SGN01, SGN06] for specular reflectors that are conic sections. A similar analysis has been performed by the same authors to explain the movement of highlights on specular surfaces [SKS*02], see Section 2.2.

2.1.3. Shape from specular flow

Instead of relying on distant, calibrated patterns, lately researchers have investigated the dense tracking of specularly moving features reflected from a distant, unknown environment map. Roth and Black [RB06] introduced the notion of specular flow, similar to optical flow [HS81, LK81] for image movement due to diffuse surfaces. The authors consider a surface composed of a mixture of diffuse and specular regions. The camera motion is assumed to be known and distant illumination by an unknown environment map is modelled. A vector field describing the standard optical flow between

an image pair is used as input to the algorithm. The material distribution is modelled in a probabilistic way and an expectation–maximization algorithm is employed to infer a segmentation between regions moving due to diffuse optical flow and regions with apparent movement due to specular reflection. Simultaneously, a parametric surface model (a sphere) is optimized. The authors present synthetic and real world evaluations using spheres with varying surface properties. It is shown that the incorporation of specular information yields a notably better reconstruction than in the case of only using the diffuse model.

Adato *et al.* [AVBSZ07] also use specular flow, but reconstruct general surface shape under distant, unknown illumination by an environment map and a static observer, assuming orthographic projection. The relative positions between camera and object must remain static, that is, only the environment map is allowed to move. The focus of the paper is the theoretical analysis of this setup. The authors show that in two dimensions an analytic solution is possible if an analytical description of the specular flow is available. Extending their results to 3D, they develop a coupled second-order non-linear system of PDEs which they solve for the special case of constant angular velocity rotation around the optical axis of the camera. In this case the equations uncouple and can be solved by the method of characteristics. An example on real world data validates the approach.

Vasilyev *et al.* [VAZBS08] extend this approach by observing that the specular flows induced by two equi-azimuthal rotations, or more general, by three rotations around arbitrary axes, can be linearly combined to form the equivalent of one rotation around the camera's optical axis. The problem is thus reduced to the solvable case described by Adato *et al.* [AVBSZ07]. The assumption of a constant, known, angular velocity is still implied. However, the authors show that the angular velocity can be recovered in conjunction with the surface properties without having to include additional constraints. The constancy assumption on the angular velocity is, however, still required.

2.2. Shape from specularity

Shape from specularity approaches rely on the observation of surface highlights caused by specular reflection at some surface points, see for example Figure 5 (left panel). If standard stereo techniques are applied to such features, the depth estimate results in a point in front of the surface for concave surfaces and in its back when the surface shape is convex [Bla85, BB88a] since specular highlights do not remain stationary on a surface when the viewpoint is changed. Apart from that, the situation is similar to shape from distortion since the light sources causing the highlight specify a 3D position in space that is usually calibrated. This results again in a 1D ambiguity for depth and normals of the surface [SP01]. The depth-normal ambiguity can be avoided if the




Figure 5: The photo of a jelly candy, exhibiting sub-surface scattering (left panel), a normal map acquired with the shape-from-specularity approach of Chen *et al.* [CGS06] (middle panel), and three-dimensional surface shape obtained by integrating the normal field. Image courtesy of Tongbo Chen, Michael Goesele and Hans-Peter Seidel.

illumination is distant with respect to the object size, for example [Ike81, SWN88, RB06, AVBSZ07] or if polarization measurements are being used [SSIK99].

2.2.1. Direct measurement of highlights

One of the earliest examples of shape recovery from specular information is given by Ikeuchi [Ike81]. The author considers one-view acquisition under changing illumination using an extended, distant light source. By employing three light distributions, surface orientations of a specular object are recovered.

Healy and Binford [HB88] investigate the information that is inherent in specular highlights. Using the physical Torrance–Sparrow BRDF model [TS67], the radiance fall-off in extended specular highlights is analysed and it is shown that second-order surface information, that is, the directions and magnitudes of principal curvature can be extracted from a single highlight. The authors also investigate degenerate cases and propose detection and interpolation methods for surface edges and corners.

Zisserman *et al.* [ZGB89] study the movement of specular highlights due to known movement of the imaging sensor. The authors show that a tracked specular highlight contains information about a 1D path on the object surface, although a 1D ambiguity remains. This ambiguity can be removed by specifying one point on the object's surface through which the family of curves passes.

Sanderson *et al.* [SWN88] propose an approach to scan specular surfaces termed structured highlight scanning. The approach is based on the distant source assumption and uses an array of point light sources distributed around the object at a distance meeting this assumption. By sequentially activating the light sources and observing the corresponding highlights, a normal field of the surface can be reconstructed. Nayar *et al.* [NSWS90] improve on this method by binary coding the array of point light sources. They employ 127 light sources, the resulting highlights of which can be scanned in

$\log_2 N$ passes. This is possible if the highlights do not overlap in the image domain. The techniques are applied to the quality inspection of industrial parts. Graves *et al.* [GNS07] investigate the accuracy and limitations of structured highlight approaches. They find that the accuracy of this approach diminishes with increasing surface curvature.

In a series of papers Zheng *et al.* [ZMFA96, ZFA97, ZM98, ZM00] develop a surface recovery algorithm based on extended radial light sources illuminating a glossy or specular object. The light sources surround the object and are observed by a static camera. By rotating the object, moving connected highlight regions (stripes) are observed by the camera. The images are accumulated in a space–time stack of images. Since orthographic projection is employed, epipolar plane images can be analysed to recover the apparent motion of the highlight on the object surface. This information enables the extraction of the surface geometry of an entire object. The authors observe that sharper specular highlights, that is, for objects of class 3 (Figure 1), result in better reconstruction accuracy. Another work, analysing the geometry of specular highlights in epipolar plane images is [CKS*05]. Instead of reconstructing the geometry from this information, the authors concentrate on highlight removal from image sequences.

Tracking specularities in a structure-from-motion setting [HZ00] is investigated by Solem *et al.* [SAH04]. The camera path and its internal parameters are estimated from tracked diffuse features in the scene. Additionally, specular highlights are tracked through the image sequence and a variational framework for shape recovery from these sparse features is developed. To regularize the solution and make the problem tractable, the authors include smoothness constraints for the surface. The variational problem is solved using a level-set formulation [Set99, OF03] with diffuse features as boundary conditions. The application is a generalization of structure-from-motion approaches, where specular surfaces like windows and metallic surfaces are permitted to be present in the scene.

A completely different approach to exploit highlight information from surface highlights is presented by Saito *et al.* [SSIK99]. Their technique is based on partial polarization of light due to reflection off non-metallic surfaces. Examples for such surfaces include asphalt, snow, water or glass. If light is polarized by reflection, the polarization is minimal in the plane of reflection, that is, in the plane containing the incident light ray, the surface normal, and the viewing ray. Saito *et al.* exploit this effect by measuring the polarization state of light with a rotating linear polarizer in front of a camera. The minimum intensity response of the surface highlight is assumed to correspond to the linear polarizer being parallel to the plane of reflection. The angle of the incident light ray with respect to the surface normal is then inferred from the degree of polarization of the light measured by the imaging sensor.

2.2.2. Surface detail from specularities

Recently, several methods have been proposed to recover geometric surface detail from specular measurements. These details are also referred to as surface meso-structure.

Wang *et al.* [WD06] propose to use a BRDF/BTF measurement device [Dan01] to also recover surface normal information along with spatially varying BRDFs. The device uses a double optical path. A parabolic mirror section is placed above the surface such that its focal point is incident on the surface. Using parallel light that is shifted by a movable aperture, different incident light directions can be achieved. Simultaneously observing an orthographic projection of the surface through a beam splitter enables the observation of dense two-dimensional (2D) BRDF slices. By detecting the highest intensity point in these slices, the major reflection direction of the surface under the incident illumination direction can be recovered, allowing for the extraction of the surface normal. By moving the mirror across the planar object surface, a 2D sampling of the spatially varying BRDF and surface normals can be achieved. The surface is then obtained by integrating the normal information.

Another interesting approach to exploit surface reflectance for object measurement has recently been proposed by Holroyd *et al.* [HLHZ08]. It is based on static viewpoint sampling of different incident illumination directions due to a moving point light source. Per-pixel, the measurements correspond to a 2D slice of the BRDF. The authors continue to establish a symmetry measure on the recovered BRDF slice. This measure enables the extraction of a local coordinate frame, that is, normal, bi-normal and tangent vectors, for every pixel in the image plane in addition to recovering a 2D BRDF slice at each pixel location.

A simpler approach using a hand-held sampling device is proposed by Chen *et al.* [CGS06]. An object with small variation in surface height is observed under an approximately orthographic view. A hand-held point light source is moved around the object in a large distance compared to the size of the object. The illumination direction is recovered from four specular spheres placed in the field-of-view of the camera. By online-thresholding the video frames, specular highlights are identified and used to recover the surface normal. The sampling stage has a user feedback, showing the sampling density at every point in time. Thus, sufficient data can be accumulated to allow for a dense reconstruction of the surface normal field. Again, the normal field is integrated to obtain the final surface shape. An example object along with the recovered normal map and shape reconstruction is shown in Figure 5.

Francken *et al.* [FCB08, FCM*08] propose an extension to this scheme by using coded highlights as in [NSWS90]. The light source used for producing coded illumination is an LCD display. However, since the display has to be placed close to the object in order to cover sufficiently many incident light

directions, the distant illumination assumption is violated and the 1D ambiguity between surface normal and depth would have to be considered. It is unclear, how this affects practical results obtained with this setup.

Whereas the previous techniques assumed approximately planar surfaces, Ma *et al.* [MHP*07] consider high resolution recovery of normals for geometrically arbitrary surfaces. The authors employ gradient illumination over the incident light sphere. By linking the measured radiance to the surface normals they show that three images taken under gradient illumination and one under constant illumination suffice to recover surface normals of arbitrary objects. However, the derivation is different for diffusely and specularly reflecting surfaces. Since most surfaces found in nature exhibit combined diffuse and specular reflection, the authors propose to separate the two components using linear or spherical polarization. A separate set of surface normals is computed for the diffuse and the specular reflection components. Using a low resolution structured light scan and the highly detailed specular normals, they recover high resolution surface models.

Francken *et al.* [FHC08] propose a similar setup as previously discussed. However, instead of generating the distant gradient illumination by employing spherically distributed LED's like Ma *et al.* [MHP*07], they use an LCD screen to display the pattern. The LCD screen simultaneously serves as a polarized light source, enabling the separation of diffuse and specular reflection components as in [MHP*07].

2.3. Direct ray measurements

To avoid the depth-normal ambiguity in the case of near-field illumination, viewing rays reflected by the specular object can be measured. A calibrated planar target is positioned in different locations with respect to the object and the distorted pattern is observed by a camera. By decoding at least two world positions for every pixel the reflected viewing ray can be measured.

A practical algorithm for specular surface reconstruction based on direct ray measurements is developed by Kutulakos and Steger [KS05, KS07]. They assume that exactly one reflection event occurs along the ray. Using the reflected ray and the viewing ray, a surface position and an associated normal direction are recovered independently for every pixel. The authors report very precise measurements for a planar front-surface mirror. Bonfort *et al.* [BSG06] present a more detailed description of the approach and show its applicability to arbitrary surface shapes. An example of the reconstruction results achievable with this technique is shown in Figure 6.

3. Translucent Objects

Translucent objects (Figure 1, class 4) are difficult to acquire for traditional range scanning techniques due to the




Figure 6: Shape reconstruction of specular objects using direct ray measurements. An input photograph with three curved and one planar specular object (left panel) and reconstruction result seen from a different view-point (right panel). The large object to the left has only been partially reconstructed due to missing pattern information. Image courtesy of Thomas Bonfort, Peter Sturm and Pau Gargallo [BSG06].

non-locality of light transport introduced by multiple scattering just beneath the object surface. Active light techniques often observe blurred impulse responses and the position of the highest intensity measurement might not coincide with the actual surface position that was illuminated [CLFS07]. A notable bias in surface measurements by laser range scanning is reported by Godin *et al.* [GBR*01].

Techniques applicable to surface detail acquisition have already been discussed in the context of shape-from-specularity approaches, Section 2.2. Since specular reflection is not influenced by sub-surface light transport, specular highlights appear in the same positions as they would for a surface not exhibiting global light transport within the object. This property has been used by Chen *et al.* [CGS06] to acquire the mesostructure of sub-surface scattering objects and by Ma *et al.* [MHP*07] to obtain detailed surface normals for translucent materials like human skin.

Chen *et al.* [CLFS07] present a structured light scanning approach directly targeted at surface scanning of translucent objects. The authors employ an approach based on a combination of polarization and phase-shifting structured light measurements. Phase-shifting of high-frequency light patterns has been shown to enable the separation of specular highlights and diffuse reflection components [NKGR06]. Chen *et al.* combine this observation with polarization based separation of surface highlights to robustly scan translucent objects. Since light gets unpolarized by global light transport effects, the authors equip a light source and the camera with a polarizer. By observing two orthogonally polarized images, multiple-scattering effects can be removed from the structured light images and improved geometry is recovered. A photo and a recovered surface scan for a translucent object are shown in Figure 7.

The phase-shifting technique [CLFS07] discussed previously relies on polarization-based removal of surface highlights. In [CSL08], Chen *et al.* improve their separation strat-




Figure 7: Photo of a star fruit (left panel) and a structured light scan (right panel) acquired with the method of Chen *et al.* [CLFS07].

egy by modulating the low-frequency pattern used for phase shifting with a 2D high frequency pattern that allows for the separation of local direct illumination effects and light contributions due to global light transport [NKGR06]. This removes the need for polarization filters and is reported to yield more robust results than [CLFS07].

The use of high-frequency patterns for the separation of local and global light transport effects is not restricted to the improvement of surface scans, but can also be applied in a volumetric fashion. Fuchs *et al.* [FHL*08] consider the acquisition of object models in the presence of participating media, or, alternatively, the volumetric acquisition of translucent, multiple-scattering objects. The authors combine high-frequency de-scattering [NKGR06] with confocal illumination and confocal imaging. By selectively illuminating and observing different confocal planes inside the scattering volume, and de-scattering the observations using high frequency patterns, a complex scene can be acquired volumetrically. Results are shown on object reconstruction in scattering media similar to Narasimhan *et al.* [NNSK05], and for the volumetric object reconstruction of a frosted glass light bulb.

An approach not directly related to surface acquisition is presented by Goesle *et al.* [GLL*04]. The authors acquire the object geometry by covering the object with removable, diffuse dust and employing a standard laser range scan. They then proceed to capture the point response of the translucent object for the whole object surface and different wavelengths, enabling the photo-realistic rendering of sub-surface scattering objects. This approach recovers a representation of sub-surface light transport within the object, but not its geometry.

4. Refractive Surface Acquisition

In this section we consider reconstruction approaches for refractive objects (Figure 1, class 5). The problem of acquiring complete surface descriptions of refractive objects with

possibly inhomogeneous material properties is very complex. In its most general form inclusions like air bubbles, cracks or even opaque or specular materials would have to be considered. The image formation for such objects is non-trivial and to date no reconstruction approaches exist for the general problem. Researchers have so far restricted themselves to sub-problems like single surface reconstruction where a well defined surface represents the transition from one medium to the other. Often the refractive index of the object needs to be known. Almost all methods assume that the refractive material is homogeneous. The earliest approaches considering refractive media can be found in the photogrammetry literature, for example [Höh71, Maa95]. However, these approaches consider underwater opaque object reconstruction, that is, a camera positioned in or outside water, the imaging sensor being separated by a planar layer of glass from the water in which the object is immersed. Photogrammetry solutions are based on the bundle adjustment technique [TMHF00].

In the following, we cover the main categories of algorithms for refractive surface acquisition. Similar to specular surface reconstruction, shape from distortion approaches, Section 4.1, and methods based on direct ray measurements, Section 4.2, have been proposed. Additionally, it is possible to sample parts of the surface reflectance field densely, Section 4.3. This approach allows for the acquisition of refractive objects with complex, inhomogeneous interior. Another class of methods is based on indirect measurements like optical thickness or measurements of the polarization state of the observed light, Section 4.4. These methods employ inverse ray-tracing based on physical image formation models to recover surface shape. Finally, light paths can be linearized by physical or chemical means, enabling the application of tomographic reconstruction, Section 4.5, or direct sampling techniques, Section 4.6.

4.1. Shape from distortion

The basics of shape-from-distortion techniques have already been discussed in Section 2.1. Here, we discuss techniques dealing explicitly with refractive surfaces. The acquisition of refractive surfaces is more complex than the corresponding specular surface case because the ray path depends on the refractive index in addition to the dependence on the surface normal.

Shape from distortion approaches are limited to the recovery of a general single refractive surface or the reconstruction of parametric surface models of simple shapes and thus are not suitable for general object acquisition.

4.1.1. Water surfaces

In computer vision, the problem of refractive surface reconstruction was introduced by Murase [Mur90, Mur92]. The

author considers the problem of reconstructing a water surface using an orthographic one-view setup where the camera is placed normal to the average water surface. An unknown pattern is placed at the bottom of a water tank. A sequence of distorted images due to water movement is recorded by the camera and analysed using optical flow [HS81, LK81]. The mean value of the pixel trajectories is used as an approximation to the average water surface, enabling the extraction of the undistorted (undistorted in the sense that refraction is taking place at a planar interface only) background pattern. Using the orthographic view assumption, a relationship between the distortion vectors with respect to the medium point of the trajectory and the surface gradient can be established for every frame of the video sequence. The gradient vectors are then integrated to obtain the final surface up to scale. The scale of the surface is influenced by the refractive index and the distance between the water surface and the bottom of the water tank.

The problem of time-varying water surface reconstruction is also considered by Morris and Kutulakos [MK05]. The authors lift several restrictions of Murase's work by employing a stereo setup and using a known background pattern. With this extended setup it is shown that an unknown refractive index can be recovered in conjunction with accurate per-pixel depth and normal estimates. Furthermore, the method does not rely on an average surface shape and is also robust against disappearing surfaces, as in the case of an empty tank that is being filled with water. The algorithm is a special case study for a more general analysis of reconstructing piecewise linear light paths conducted by Kutulakos and Steger [KS05, KS07]. The method shares similarities with the normal matching methods [SWN88, BS03, NWR08] discussed in Section 2.1 for specular object reconstruction, for example the normal matching cost function used by [NWR08] is very similar to the one used here. Some results of this technique are shown in Figure 8.

4.1.2. Glass objects

Shape from distortion techniques have also been applied to recover the surface shape of glass objects. Hata *et al.* [HSKK96] consider glass objects and drop-like structures with one planar surface, resting on a diffuse base. The




Figure 8: Experimental setup (left) and result of reconstructing a time-varying water surface (right) using the method of Morris and Kutulakos [MK05].

authors use a structured light setup to project stripe patterns into the object, the distorted patterns of which are observed by an imaging sensor. Since there are two refracted light paths, one from the projector and one for the observing camera, the problem is more complex than the methods discussed previously and no analytic solution is known. Instead, the authors employ a genetic algorithm to recover the surface shape.

Another, model-based, approach to surface reconstruction of glass objects is proposed by Ben-Ezra and Nayar [BEN03]. The authors assume an unknown, distant background pattern and a known parametric model for the object as well as its refractive index. The method differs from the previous techniques in that it can recover the surface shape of complete objects and not just single surfaces. The authors use a single-view setup and track refracted scene features over a motion sequence of the object, similar to Murase [Mur90, Mur92]. Using a steepest descent method, they solve for the shape and pose of the object. A-priori geometric models can also be used to perform camera pose refinement based on specular reflections [LSLF08].

The tracking of refracted scene features might be complicated by the fact that refraction results in sometimes severe magnification or minification of the background pattern. Additionally, if the object is not completely transparent, absorption might change the intensity of the observed features, complicating feature tracking. A solution to this problem, an extension to standard optical flow formulations, has been presented by Agarwal *et al.* [AMKB04].

4.2. Direct ray measurements

Direct ray measurement approaches, c.f. Section 2.3, have also been used for refractive surface reconstruction. Rays are measured after having passed through the refractive object. Ray measurements are either based on the measurement of calibrated planar targets imaged in several positions with respect to the object [KS05, KS07] or approximated from optical flow data [AIH*08]. The measurement-based approach allows for the recovery of several 3D world points per camera pixel to which a line is fit that describes the ray exitant from the object. Optical flow based techniques on the other hand are approximate and assume the size of the object to be small compared to the distance between the object and the background pattern.

Kutulakos and Steger [KS05, KS07] investigate several applications of direct ray measurements. The authors provide a thorough theoretical analysis of reconstruction possibilities based on pixel-independent ray measurements. They categorize reconstruction problems involving refractive and specular surfaces as pairs $\langle N, M, K \rangle$, where N is the number of view-points that are necessary for reconstruction, M is the number of specular or refractive surface points on a piecewise linear light path and K is the number of calibrated reference points on a ray exitant from the object. Two




Figure 9: Faceted glass object with refractive index $n \approx 1.55$ (left panel) and pixel-independent reconstruction result (right panel) using the method of Kutulakos and Steger [KS05, KS07].

practical examples, $\langle 1, 1, 2 \rangle$ reconstruction (1 viewpoint, 1 specular interaction and 2 reference points) of specular surfaces, Section 2.3, and $\langle 2, 1, 1 \rangle$ -reconstruction (two viewpoints, one refractive interaction and one reference point on each refracted ray) [MK05], Section 4.1.1, have already been discussed. The authors investigate the tractability of general $\langle N, M, K \rangle$ -reconstruction algorithms and show that a pixelwise independent reconstruction is not possible for more than two specular or refractive surface intersections, regardless of the number of input views and the number of reference points on each exitant ray. It is also shown that more than two known points on an exitant ray do not contribute information to the reconstruction problem.

For the purpose of this section, the $\langle 3, 2, 2 \rangle$ reconstruction problem is of interest. Kutulakos and Steger [KS05, KS07] develop a practical algorithm for the reconstruction of two interface refractive light interaction using a three viewpoint setup and measuring the exitant ray directions. They recover four surface points and four normal estimates per pixel of the imaging sensor. One point and corresponding normal are situated at the front surface of the object. The other three points and normals are found on the back surface separately for each of the differently refracted viewing rays. Results of this algorithm on a faceted glass object with refractive index of $n \approx 1.55$ are shown in Figure 9. The refractive index is recovered along with the surface points and normals.

Another approach that is based on the measurement of (approximate) exitant rays is presented by Atcheson *et al.* [AIH*08]. The authors focus on the reconstruction of gas flows, more specifically the reconstruction of refractive index variation due to temperature changes within such flows. Since the refractive index variations due to temperature changes are very low (in the range of 10^{-4} to 10^{-3}), the exitant rays exhibit only minor changes in direction. Due to these constraints, their apparent deflection in image space can be computed by optical flow methods [HS81, LK81]. Suitable choices of background patterns and optical flow algorithms are discussed in [AHI09]. The image space deflections are then converted into exitant ray measurements by centring the ray




Figure 10: Optical flow detection in gas flows (left panel) and an iso-surface rendering of a 3D reconstruction of a time-varying, inhomogeneous refractive index field (right panel) using the approach of Atcheson et al. [AIH*08].

at the midpoint of the gas flow, the region of space occupied by it being small compared to the distance to the background pattern. Using the exitant ray measurements, the authors set up a linear system that describes the differential change in the ray directions which is related to the refractive index gradient. The linear system is then inverted in a least-squares sense to yield a volumetric description of the refractive index gradients, which is integrated to obtain volumetric refractive index measurements. The method uses a multi-view setup and is suitable for the reconstruction of time-varying inhomogeneous refractive index distributions. An example of the refractive index distribution above a gas burner is shown in Figure 10.

4.3. Reflectance-based reconstruction

Reflectance-based reconstruction of refractive objects has recently been introduced by Morris and Kutulakos [MK07]. The authors employ a static one-view setup with a moving near-field light source. By moving the light source to a 2D set of positions on a regular grid while taking images with the camera, they acquire a dense set of reflectance measurements for each pixel of the imaging sensor. The reflectance measurements are influenced by direct surface reflection and additional global light transport effects.

Since the positions of the imaging sensor and the object are static throughout the measurement process, the reflectance response of the object stays static with respect to the viewing direction and a 2D slice of the surface BRDF is measured. These measurements are corrupted by the indirect light transport within the object, however, the authors show that it is possible to separate the direct reflection component from the indirect lighting effects by exploiting the physical properties of light transport, that is, light travels linearly before hitting the object and there is a radial fall-off of the incident irradiance. This way, it is possible to detect incident light rays in the measurements corrupted by additional global light transport effects. The incident light rays converge towards the surface point that reflects light towards the camera. An additional constraint is that the reflection point must lie on the




Figure 11: A refractive object with complex inhomogeneous interior (left panel), reconstructed normal map (middle panel), and depth map (right panel). Reconstructions are obtained with the method of Morris and Kutulakos [MK07].

imaging sensor's viewing ray. Based on these constraints it is possible to reconstruct very detailed depth and normal maps of refractive objects with complex, inhomogeneous interior, see Figure 11.

4.4. Inverse ray-tracing

Inverse ray-tracing relies on the comparison of suitably chosen input data with synthetically generated images. The experimental setup has to be chosen carefully to enable the formulation of a proper image formation model. Starting with an initial guess for the surface shape, the forward ray-tracing problem is solved. By relating the residual error in the image plane to surface deformations the surface shape is optimized, usually in a non-linear way.

One possibility is an experimental setup that is based on the effect of fluorescence [IGM05] or chemiluminescence [GILM07]. Ihrke et al. [IGM05] and Goldluecke et al. [GILM07] consider the reconstruction of 3D, time-varying surfaces of free-flowing water, such as water columns that splash into a glass when being filled. The method is based on mixing the water with either a fluorescent dye or a chemiluminescent chemical. This measure makes the water self-emissive when illuminated by UV-light in the case of fluorescence or by a chemical process that lasts for several minutes in the case of chemiluminescence. Self-emission is assumed to be homogeneous throughout the water, resulting effectively in optical path length measurements of the (multiple) refracted rays. The authors employ an image formation model based on constant self-emissivity and perform a level-set optimization [Set99, OF03] of the water surface to match the input video frames acquired using a multi-view setup with synthetically generated images. The surface is initialized with the visual hull [Lau94]. Synthetic simulations show the capability of the approach to recover even major concavities. An input image and reconstruction results on real world data for this technique are shown in Figure 12. The method can be considered a binary tomographic approach since a binary volumetric reconstruction is performed.

Recently, water columns have been reconstructed by dyeing the fluid with opaque white paint and projecting a pattern onto its surface [WLZ*09]. This measure allows for the



Figure 12: Input video frame with chemiluminescent water column (left panel), the reconstructed geometry (middle panel) and another acquired surface rendered into a virtual environment with modified material properties (right panel). Surface geometry was reconstructed using the technique of Ihrke et al. [IGM05, GILM07].

application of standard stereo reconstruction methods to recover surface geometry. In addition, the authors regularize the surface reconstruction with a quasi-physical prior mimicking the behavior of the Navier–Stokes equations. Because the recorded data is sparse, only physically plausible surfaces are generated. They do not necessarily correspond to the actual physical surface.

A different approach based on inverse ray-tracing taking polarization into account is presented by Miyazaki and Ikeuchi [MI05]. The measurement setup consists of a single camera equipped with a linear polarizer. The refractive object is mounted inside a geodesic dome of light sources that are diffused by a plastic sphere surrounding the object. The shape of the object's back surface as well as its refractive index and the illumination distribution are assumed to be known. The measurement process consists of acquiring four differently polarized images by rotating the linear polarizer in front of the camera. The reconstruction is then performed using an iterative scheme that minimizes the difference between the measured polarization state and the polarization ray-traced image assuming a specific surface configuration.

4.5. Reduction to tomography

Under certain circumstances light is not refracted by refractive objects. This is the case if the wavelength of the illumination is sufficiently high, that is, in the case of X-ray illumination, and when the refractive index of the medium surrounding the refractive object is the same as the object's refractive index. X-ray scanning of refractive objects is straight forward [KTM*02]. Although the authors do not concentrate on refractive object scanning, computed tomography reconstruction of glass objects is possible as well, as long as no metal inclusions are present inside the object.

A method that operates in the visible wavelengths and does not resort to expensive equipment is presented by Trifonov et al. [TBH06]. Volumetric descriptions of glass objects are acquired by immersing them into a refractive index matched



Figure 13: A photograph of a refractive object with absorptive properties (left), a tomographic projection obtained by matching the refractive index of a surrounding medium to the one of the object (middle) and the object's surface generated by iso-surface extraction on a volumetric representation of the absorption density of the object (right). The absorption density is reconstructed by tomographic means [TBH06].

fluid to 'straighten' the light paths. Refractive index matching is achieved by mixing water with, usually toxic, chemicals. In [TBH06] potassium thiocyanate is used, solutions of which in water can achieve refractive indices of $n \approx 1.55$. If the refractive object is completely transparent, it ideally disappears in a refractive index matched immersing medium. Therefore, it is necessary to dye the surrounding medium in this case. However, if the refractive object is itself absorptive dyeing the surrounding medium can be omitted. The authors acquire 360 images spaced evenly around the object and solve a standard tomographic reconstruction problem. Results of this approach are shown in Figure 13.

4.6. Direct sampling

Another technique avoids the computationally expensive and ill-posed tomographic reconstruction process by altering the immersing medium. Hullin et al. [HFI*08] propose to dye the refractive index matched liquid with a fluorescent chemical agent like eosin Y. For measurement, a laser sheet scanning approach similar to the techniques discussed in Section 5.3, is adopted by the authors. The laser sheet is rendered visible by the fluorescing medium while the glass object remains dark. Since light is not bent inside the tank used for measurement purposes, the laser sheets remain surface-like even though they are slightly curved due to refraction at the tank boundary. The curved light sheets can, however, be calibrated prior to acquisition. Scanning results for an acrylic glass bowl filled with cylinders of the same material are shown in Figure 14. Note that the method can also be used for first-surface scanning of refractive, translucent, and low albedo objects, see Section 1.1.1.

A related technique that also tries to avoid the refractive properties of transparent objects is investigated by Eren et al. [EAM*09]: the infrared spectrum is not refracted by glass objects, therefore using an IR laser instead of visible light allows for laser range scanning of glass objects. The method is termed 'scanning from heating' because the glass is heated up by the incident IR radiation which is then recorded using an IR-sensitive camera. The resolution of this technique is, however, restricted since the wavelength of the




Figure 14: Fluorescent immersion range scanning [HFI*08] can recover all surfaces of glass objects of homogeneous refractive index. A photograph of an acrylic glass object (left panel), a direct volume rendering of a recovered voxel model (middle, left panel), A cut-away iso-surface view (middle, right panel), and a realistic rendering using a ray-tracer with photon mapping (right panel).

incident illumination is much larger than for visible light and thus cannot be focussed as well. In addition, glass dissipates the heat quite well and blurry impulse responses result.

5. Volumetric Phenomena

In this section, we review acquisition techniques related to 3D sensing of volumetric phenomena (Figure 1 object classes 6 and 7). The methods presented here assume a volumetric description of the scene content. Unlike in space carving approaches [KS00] the scene is assumed to be either completely or partially transparent. Furthermore, all methods presented in this section assume that light rays pass straight through the scene and that refractive effects can be neglected. The main application of these techniques is the acquisition of transparent, volumetric phenomena such as fire and smoke, but also 3D descriptions of plasma effects like planetary and reflection nebulae have been recovered this way. Fire and smoke are inherently dynamic phenomena whereas interstellar object reconstruction suffers from the availability of only a single view-point. Thus, the methods covered in this section typically cannot employ multiple measurement passes to stabilize the reconstruction.

We classify the approaches to volumetric phenomena acquisition into tomographic approaches Section 5.1, computer vision techniques that assume partial scene transparency, Section 5.2 and techniques based on direct measurements, Section 5.3.

5.1. Tomographic approaches

Observing volumetric phenomena with an imaging sensor results in integral measurements of the volumetric light distribution¹ over the line of sight for every sensor element.

¹The light reaching a sensor element is usually a combination of emitted light, and light that is scattered into the direction of the observer. On its way through the volume it is generally subject to attenuation due to out-scatter and extinction.

Integral measurements are usually called projections and the task of recovering an n -dimensional function from its $(n - 1)$ -dimensional projections is known as tomography. The mathematical foundations and the existence of a unique solution for infinitely many measurements have been shown by Radon [Rad17]. The major difficulties in computed tomography, that is, in the numerical inversion of the projection operator for real measurements, are the finite number of measurements that are usually available and the instability of the inversion with respect to noise. A classical text on numerical inversion techniques for the computed tomography problem is Kak and Slaney [KS01]. Tomographic reconstruction techniques have lately been used for the acquisition of (time-varying) volumetric phenomena like fire, smoke, astronomical objects and biological specimen.

5.1.1. Fire and smoke

In computer vision, the sparse-view tomographic reconstruction of fire was introduced by Hasinoff and Kutulakos [Has02, HK03]. In [Has02] a simplified image formation model based on self-emission is introduced. A collection of Gaussian blobs with varying standard deviation is used as a reconstruction basis for the tomographic problem. The blobs are initially evenly distributed. Their positions and standard deviations are then optimized in an iterative manner. The results of this technique, however, suffer from over-fitting [HK07].

This short-coming was addressed in subsequent work [HK03, HK07]. The proposed algorithm maintains the majority of high frequency detail of the input images in interpolated views while simultaneously keeping the number of input images as low as possible. To achieve this, the authors develop a basis for the reconstructed density fields that is spatially compact and simultaneously allows for a convex representation of the density field. They term this basis *decomposed density sheets*. The basis consists of sheet-like spatial structures and is proved to be complete [HK07]. The convexity constraint on the basis functions' coefficients is the major difference to standard tomographic approaches. It




Figure 15: Density sheet basis (left panel), reconstruction result for a two-view setup (middle panel) and reconstruction result from two views for data set in Figure 16 (left panel) computed with the algorithm of Hasinoff and Kutulakos [HK03, HK07].

enables the recovery of a globally minimal solution at the cost of employing a quadratic programming solver. The rendering of a simplified version of the basis functions as well as reconstruction results are shown in Figure 15. The computational cost currently limits the acquisition setup to camera configurations that lie in the same plane and allow for epipolar slicing of the reconstruction volume, essentially reducing the 3D reconstruction problem to a 2D one. Due to the spatial compactness of the basis functions, view-generation from view-points significantly above or below the plane containing the cameras' optical axes results in sheet-like structures being noticeable in the renderings.

Tomographic 3D reconstruction of time-varying fire and smoke volumes from sparse view input data has also been investigated by Ihrke and Magnor [IM04, IM05, IM06]. Compared to [HK03, HK07] the reconstruction is performed with standard basis functions, resulting in a better representation of inner structure. This is helpful when synthesizing views from atop or below the flame. The experimental setup involves a multi-camera acquisition setup arranged in an approximately circular ring around the phenomenon. In the case of fire reconstruction [IM04], recording takes place in a dark environment, while smoke acquisition [IM05, IM06] is performed in a homogeneously and diffusely lit room. The diffuse lighting is a measure to make scattering effects in the smoke volume approximately homogeneous. This way, the smoke volume can be treated as a self-emissive medium as in the case of fire. The authors then set up a system of linear equations that describes the tomographic projection operation into all views simultaneously. By inverting the linear system in a least squares sense, a volumetric description of the phenomenon under observation is recovered. However, the low number of eight cameras in the multi-view setup leads to ghosting artefacts in the reconstruction [IM04]. Photo-realistic results are achieved by constraining the reconstruction to the visual hull [Lau94] of the phenomenon. In [IM05, IM06] it is shown that the visual hull restriction can be performed by analysing the linear system only. Based on this observation, and a method



Figure 16: Results of tomographic reconstruction on fire [IM04] (left panel) and smoke data sets [IM05, IM06] (right panel). The multi-view setup used to capture the input data consists of 8 cameras, Reconstruction resolution is 128^3 (left panel) and an octree-representation with effective resolution of 256^3 (right panel).

to project image space residual errors into the volumetric reconstruction domain, an adaptive reconstruction scheme is proposed. This allows for higher effective reconstruction resolutions and a better representation of fine detail. Reconstruction results achieved with this method are shown in Figure 16.

5.1.2. Astronomical objects

Tomographic reconstruction has also been used in the context of model acquisition for emissive or scattering astronomical objects like planetary [MKHD04, LLM*07a, LLM*07b] and reflection nebulae [LHM*07].

Magnor *et al.* [MKHD04] describe an inverse rendering approach for planetary nebulae that is applicable to the reconstruction of these objects in a purely emissive setting (Figure 1, class 6). Planetary nebulae exhibit axial symmetry. Therefore, the single view point that is available by observations from earth is sufficient to recover the 2D emission map of the nebula. The reconstruction is performed independently at three different wavelengths to enable realistic rendering of the recovered objects. A non-linear optimization method is employed to solve for the emission maps and inclination angle of the nebula. Reconstruction results from [MKHD04] are shown in Figure 17.

Lințu *et al.* [LLM*07a, LLM*07b] extend this scheme to include effects of emission, absorption and scattering (Figure 1, class 7). The authors reconstruct a gas map (emissive plasma) and a dust map (absorption). The reconstruction of the two different maps is performed by using the physical properties of different wavelength measurements. Radio imaging is nearly unaffected by absorption and thus allows for the reconstruction of the emissive gas map unobstructed by dust particles. In a second step, an image at visible wavelengths is used to recover the dust density, again employing




Figure 17: Results of axisymmetric tomography on planetary nebulae [MKHD04]. Realistic rendering of reconstructed nebulae (left column) and iso-surface rendering (right column).

an inverse rendering scheme. The same can be performed for infrared/visible image pairs because infrared components are emitted mainly by the dust distribution, whereas visible wavelength images include both emission and absorption effects.

In another study, Lințu *et al.* [LHM*07] show a reconstruction approach for reflection nebulae. Reflection nebulae do not exhibit symmetry and the reconstruction problem is ill-posed. The major component in the appearance of these objects are scattering and absorption of light emitted by the central star (class 7). The authors again employ an inverse rendering approach with an additional regularization component to acquire plausible reconstructions of reflection nebulae. However, it has to be emphasized that the reconstruction results have no physical basis in the direction parallel to the optical axis of the camera.

5.1.3. Biological specimen

Levoy *et al.* [LNA*06] describe a light field [LH96] microscope. A standard microscope is modified by inserting a lenslet array into the optical path. This measure allows for the simultaneous acquisition of multiple orthographic integral projections of the specimen which in turn can be used to reconstruct a 3D description of the object. Since the viewpoint varies only over the area of the main lens, the reconstruction is equivalent to limited angle tomography, that is, tomographic reconstruction with large parts of missing data. The reconstruction problem is solved by deconvolution with the point spread function of the microscope, the equivalence of which to tomographic reconstruction is proofed in the paper.

A non-tomographic way of reconstructing 3D descriptions of biological specimen is the use of a light sheet scanning microscope as developed by Fuchs *et al.* [FJLA02]. A laser sheet is scanned through a volumetric specimen placed under the microscope. This allows for the direct sampling of object layers provided sufficient amounts of light are scattered towards the observer and attenuation and out-scatter within the object are negligible. This is similar to the smoke scanning techniques discussed in Section 5.3. The technique has been improved for dynamic high-speed fluorescence imaging by Keller *et al.* [KSW08]. The method was employed to observe the growth of embryonic cells.

5.2. Transparency in multi-view stereo

In this subsection we discuss volumetric reconstruction approaches that are deviating from a classical tomographic reconstruction formulation. Most algorithms can be considered as specialized tomographic approaches though. The distinction is thus not strict. Future research could establish links between the methods presented here and classical tomography approaches, potentially leading to more efficient or more accurate reconstruction algorithms in both domains.

Linear transparency models have also been considered in the stereo literature. Mixed representations of opaque and transparent scene parts have been developed by Szeliski and Golland [SG99] in the context of better boundary and occlusion treatment in passive stereo vision. De Bonet and Viola [BV99] describe a reconstruction method that assumes partially transparent objects to be present in the scene. The authors formulate the image formation in such scenes by using a formulation similar to environment matting approaches [ZWCS99]. However, their formulation confounds transparency and uncertainty and thus does not have a clear physical basis. The results produced by this method are not suitable for photo-realistic view synthesis. A formal probabilistic treatment of occupancy uncertainty in multi-view stereo is presented in [BFK02]. This approach, however, is focused on scenes that contain only opaque surfaces.

An approach similar to [BV99] is used by Yamazaki *et al.* [YMK06] to model opaque objects with intricate surface details such as fur and hair. The reconstruction volume is separated into opaque and transparent regions based on the visual hulls of environment matted input images. The visual hull of completely opaque pixels is modelled as opaque voxels and the surrounding region that projects to partially transparent pixels in at least one of the input views is considered transparent. A solution is then computed based on an expectation-maximization algorithm. The estimation problem can be interpreted as a tomographic approach in the presence of occlusions such as metal implants in medical computed tomography. A toy animal reconstructed using this method is shown in Figure 18 (right panel).




Figure 18: Results for the reconstruction of partially transparent objects. Image-based tree modelling Image courtesy of Alex Reche, Ignacio Martin and George Drettakis [RMD04] (left panel) and reconstruction of a furry toy animal Image courtesy of Shuntaro Yamazaki, Masaaki Mochimaru and Takeo Kanade [YMK06] (right panel).

Reche *et al.* [RMD04] also use an essentially tomographic approach to acquire image-based representations of trees. Similar to [YMK06] they model the tree as a transparent volume and use an image-based texturing technique to render the foliage of the tree in a photo-realistic way, see Figure 18 (left panel).

5.3. Direct measurement interpolation

The previous subsections concentrated on inverse problem formulations of the acquisition problem for volumetric phenomena. Inverse problems are often ill-conditioned which makes them susceptible to instability due to noise and missing data. It is therefore advisable to consider direct measurement approaches if this is feasible.

Direct volumetric measurements of time-resolved phenomena have only recently been performed in the computer graphics community [HED05, FCG*06, FCG*07, GNG*08]. The only phenomenon that has been tackled so far is smoke. However, smoke as a participating medium exhibits strong

global illumination effects such as single and multiple scattering. Direct measurements of participating media are performed using structured illumination techniques.

A laser ray is either transformed into a sheet of light [HED05] and swept through the acquisition volume or split up into multiple static laser lines [FCG*06, FCG*07]. Alternatively, a digital projector can be used to generate more general basis illumination [GNG*08]. The approaches covered in this section assume negligible amounts of multiple scattering to be present in the scene. They are applicable for the reconstruction of class 7 objects, Figure 1.

5.3.1. Laser sheet scanning

The laser sheet scanning approach by Hawkins *et al.* [HED05] is similar to a technique termed laser-induced fluorescence (LIF) in the fluid imaging community, see for example [DD01, VVD*04]. The measurement setup consists of a laser light source, a mirror galvanometer and a cylindrical lens. This way, a laser sheet is created that can be swept through the acquisition volume, illuminating only a 2D slice of it at every point in time. The laser plane is observed with an imaging sensor placed approximately orthogonal to the laser illumination direction. By scanning the volume very quickly and observing the illuminated slice images with a high speed camera synchronized to the mirror galvanometer, 200 slices of the volume can be acquired at 25 frames per second [HED05]. To achieve a good signal-to-noise ratio a rather powerful 3W ion-laser has to be employed. A sketch of the measurement setup is shown on the left-hand side of Figure 19.

The measured image intensities are directly interpreted as volume densities after radiometric compensation for non-uniform laser illumination intensity throughout the volume. Multiple scattering and extinction effects are ignored, restricting the method to the acquisition of optically thin participating media. Furthermore, the different slices of the volume are captured at different instances in time, causing a




Figure 19: Left: Laser sheet scanning [HED05] sweeps a sheet of laser light through the volume. Middle: Structured light scanning employs various coded patterns, compressive sensing strategies can be employed [GNG*08]. Right: Laser line scanning [FCG*06] uses static patterns of laser lines to increase temporal resolution while trading spatial resolution. In all setups, the scattered radiance is observed by an off-axis camera, absorption and multiple scattering effects are assumed to be negligible.

shear in the acquired data if the advection speed of the underlying fluid flow is too fast [VVD*04]. However, as shown by Van Vliet *et al.* [VVD*04] this effect can be compensated for by estimating the velocity vectors of the fluid flow using 3D optical flow techniques.

In addition to directly measuring volumetric smoke density distributions, Hawkins *et al.* [HED05] propose a setup to measure the scattering phase function and the albedo of a participating medium. Measuring these data allows for photo-realistic rendering of the acquired time-varying smoke volumes. To measure the scattering phase function of the smoke, a laser line is projected into a spherical chamber filled with a homogeneous smoke distribution. The laser response is then observed in a single image using a conical mirror surrounding the chamber. For albedo estimation, a similar setup is employed, this time using a glass tank with a planar wall and a high, homogeneous smoke concentration within the measurement volume. The observed attenuation of the laser ray is then related to the albedo of the smoke.

5.3.2. Laser line scanning

It is possible to modify the scanning setup in order to achieve equi-temporal volume measurements. One such approach has been described by Fuchs *et al.* [FCG*06, FCG*07]. The geometry of the illuminating laser is changed from a 2D plane sweeping through space to a static set of laser lines, see Figure 19 (right). The elimination of time-varying illumination enables the acquisition of equi-temporal data sets, thus allowing for the capture of fast moving smoke or longer integration times if necessary. However, this comes at the cost of reduced spatial resolution. The participating medium is probed densely on the illuminating laser lines but in order to separate spatial samples in the image plane, the lines have to be coarsely spaced throughout the volume, such that no two lines overlap on the image sensor. To increase the sampling resolution in the image plane, lasers of different wavelength (red and blue) are used and sensed in the different colour channels of the CCD camera. Similar to Hawkins *et al.* [HED05], the pixel values are interpreted as a measure of smoke density. However, Fuchs *et al.* [FCG*06, FCG*07] subtract median values of background pixels, not occupied by laser lines in image space to account for multiple scattering effects.

The data acquired with this approach is densely measured along the laser lines projected into the medium, but coarsely throughout the medium itself. The biased resolution along the laser lines poses a problem for the interpolation of the data into the surrounding space. Methods using the full data are shown to have unsatisfactory performance. Instead, a radial basis function (RBF) style interpolation is found to yield adequate results. Using a simulation, the resulting interpolated data is shown to resemble a low pass filtered version




Figure 20: Result of laser sheet scanning of a complex smoke column (left panel). Image courtesy of Tim Hawkins, Per Einarsson and Paul Debevec [HED05]. Photographs of laser lines illuminating a smoke volume and resulting reconstructions (right). Image courtesy of Christian Fuchs, Tongbo Chen, Michael Goesele and Hans-Peter Seidel [FCG*07].

of the ground truth volumes. Results of laser sheet and laser line scanning are shown in Figure 20.

5.3.3. Arbitrary basis scanning

Gu *et al.* [GNG*08] generalize the laser sheet and laser line scanning approaches by using a digital projector in an off-axis camera–projector setup, see Figure 19 (middle). The use of the projector enables the projection of arbitrary basis illumination into the scattering volume. Because light is scattered into the direction of the observer from every spatial position, the image intensities super-impose, similar to the tomographic problems discussed previously in Section 5.1. However, because the line of sight stays constant, and the authors assume the smoke to move sufficiently slow for a constant smoke volume assumption, the volume densities are expanded in the dual of the projected basis. The integration performed by the sensor computes the scalar product between the projected basis and the volume densities, resulting in the measurement of the dual basis coefficients in image space. This way, the 3D tomographic problem is reduced to a much simpler per-pixel 1D basis transformation. In practice, the authors employ a noise basis and use compressive sensing [CRT06] strategies to reduce the number of basis functions, and thus the acquisition time interval needed for dynamic smoke capture. A fast way of imaging the required basis representation could employ digital light processing (DLP) projector technology and high-speed cameras similar to the techniques described by Narasimhan *et al.* [NKY08].

6. Conclusions

We have reviewed and classified methods for the acquisition of surface geometry or volumetric descriptions of objects with complex optical characteristics. Currently there exist approaches that can deal relatively well with different subclasses of objects. However, the algorithms are still very

specific and not generally applicable. Furthermore, many techniques require considerable acquisition effort and careful calibration.

Except for X-ray tomography, there are as of yet no commercially available scanners for refractive, specular, sub-surface scattering or volumetric objects. However, considerable progress has been made in the last couple of years and today it is possible to recover accurate depth and normal maps for reflective and refractive objects, a task that seemed close to impossible only a few years ago. Current state-of-the-art approaches hint at the possibility of robustly recovering the surface shape of these challenging objects. Still, the acquisition of general, mixed objects remains a daunting task (Figure 1, class 9).

Especially in the case of refractive objects, it is usually not sufficient to recover the first surface of an object. Glass objects, for example, are seldom solid, at least for objects that are interesting from a rendering perspective. They often have holes or colourful inclusions, complicating the acquisition process [MK07]. While it would be desirable from a computer graphics point of view to acquire such objects, at the moment it is not clear how to achieve this goal. Rendering techniques can handle much more complex objects than can currently be acquired. Apart from the problem of geometry acquisition, surface reflectance properties have to be estimated simultaneously to achieve high quality renderings. There exist approaches to capture the surface properties of objects once the geometry is known [LKG*03, GLL*04], but the simultaneous acquisition is still challenging even for objects that do not give rise to major global illumination effects [TAL*07]. This is another large body of work and we have barely scraped the surface. In the long run we would wish for a flexible reconstruction paradigm that would enable the acquisition of arbitrary objects made from arbitrary mixtures of materials.

The results so far indicate that a convergence of reconstruction methods for specularly reflective and specularly refractive objects might be achieved. Kutulakos and Steger, for example, have studied the problem in a common framework. Even though they arrive at a theoretical result that shows that for light paths intersecting more than two surfaces of the reflective/refractive type a reconstruction of the surface shape is impossible *in principle* [KS07], this result only holds for per-pixel independent reconstructions. Additional assumptions on surface shape might allow for more relaxed conditions. Furthermore, the diagram in Figure 3 shows that the basic ideas underlying reconstruction algorithms in these areas are very similar even though their implementation differs in applications. It should therefore be possible to derive unifying representations and reconstruction algorithms for class 3 and class 5 objects, Figure 1.

For inherently volumetric phenomena, model acquisition of time-varying data suitable for photo-realistic view-

synthesis is within reach. While the acquisition effort is still considerable, and only sub-classes of phenomena have been covered, it seems to be possible to converge to a generally applicable reconstruction paradigm. The major task here is the improvement of existing techniques, using, for example, temporal coherence of the data. A major drawback of current techniques is that all of them recover appearance parameters of these phenomena only. It is desirable to also measure the parameters of the underlying physical process. This would allow for true modification of the measured data. Combining, for example, fluid simulations and measurement approaches seems to be a promising avenue for future work. A first step into this direction has been attempted recently by Wang *et al.* [WLZ*09]. While they recover physically plausible fluid motion from videos, the results are not reconstructions in the sense that they interpolate and fill large parts of missing data with physically interpolated results. Being able to measure physical processes in all their aspects would yield better insight into the underlying phenomenon and would enable the development of improved models for simulation purposes.

The aforementioned challenges require work on the theoretical foundations of image formation and especially its inversion. However, in most of the existing work a traditional camera model is assumed. Recently, dramatic progress has been made in the area of Computational Photography suggesting that improved computational sensors will become available also for measurement purposes. Light field cameras, for example, allow for the capture of a high angular resolution while simultaneously spacing view points very closely. This qualitatively new amount of data would be suitable to analyse difficult objects exhibiting global light transport effects as discussed in this report. Another promising example, multi-spectral imaging, could potentially aid in detecting surfaces that are difficult to acquire otherwise, as demonstrated by the infrared range scanning technique of Eren *et al.* [EAM*09]. In a similar manner, ultraviolet radiation might be used for object detection in future range scanning systems. Thus extending the spectral capabilities of our sensing devices could improve detection and reconstruction accuracy of otherwise challenging objects.

In summary, there are many applications for general object acquisition techniques. In the area of computer graphics, model acquisition can be automated, simplifying the laborious task of object modelling. Additionally, the acquired real-world objects can be used for improved rendering fidelity. An analysis of the acquired objects or object parameters may facilitate the creation of improved rendering models. Practical applications include the preservation of cultural heritage, more life-like appearance of virtually generated imagery and content production for digital entertainment.

From a computer vision perspective, the investigation of general object acquisition approaches can yield more robust algorithms for the recovery of surface shape in the presence of

global illumination effects. Computer graphics approaches, describing the forward problem of light transport can be used to yield additional insight for the inverse problem of recovering scene parameters. Applications include quality control of industrial parts which are often manufactured from non-diffuse materials as well as scientific applications in areas as diverse as oceanography, photogrammetry, applied optics and experimental physics.

Acknowledgments

The authors are grateful to the anonymous reviewers for their helpful suggestions, and to the following researchers for granting us permission to use their images: Christian Fuchs, Tongbo Chen, Michael Goesele, Holger Theisel, Hans-Peter Seidel, Tim Hawkins, Per Einarsson, Paul Debevec, Alex Reche, Ignacio Martin, George Drettakis, Shuntaro Yamazaki, Masaaki Mochimaru, Takeo Kanade, Thomas Bonfort, Peter Sturm, Pau Gargallo and Marco Tarini. Ivo Ihrke was supported by a Feodor Lynen Fellowship granted by the Alexander von Humboldt Foundation, Germany and by the German Research Foundation (DFG) through the Cluster of Excellence 'Multimodal Computing and Interaction'.

References

- [AHI09] ATCHESON B., HEIDRICH W., IHRKE I.: An evaluation of optical flow algorithms for background oriented schlieren imaging. *Experiments in Fluids* 46, 3 (2009), 467–476.
- [AIH*08] ATCHESON B., IHRKE I., HEIDRICH W., TEVS A., BRADLEY D., MAGNOR M., SEIDEL H.-P.: Time-resolved 3D capture of non-stationary gas flows. *ACM Transactions on Graphics (SIGGRAPH Asia 2008)* 27, 5 (2008), 132:1–132:9.
- [AMKB04] AGARWAL S., MALLICK S., KRIEGMAN D., BELONGIE S.: On refractive optical flow. In *Proceedings of European Conference on Computer Vision (ECCV)* (Prague, Czech Republic, 2004), pp. 279–290.
- [AVBSZ07] ADATO Y., VASILYEV Y., Ben-Shahar O., ZICKLER T.: Towards a theory of shape from specular flow. In *Proceedings of IEEE International Conference on Computer Vision (ICCV)* (Rio de Janeiro, 2007), pp. 1–8.
- [BB88a] BLAKE A., BRELSTAFF G.: Geometry from specularly. In *Proceedings of IEEE International Conference on Computer Vision (ICCV)* (Tarpon Springs, FL, 1988), pp. 297–302.
- [BB88b] BRELSTAFF G., BLAKE A.: Detecting specular reflections using Lambertian constraints. In *Proceedings of IEEE International Conference on Computer Vision (ICCV)* (Tarpon Springs, FL, 1988), pp. 297–302.
- [BEN03] BEN-EZRA M., NAYAR S.: What does motion reveal about transparency? In *Proceedings of IEEE International Conference on Computer Vision (ICCV)* (Nice, France, 2003), vol. 2, pp. 1025–1032.
- [Ber04] BERARDIN J.-A.: Integration of laser scanning and close-range photogrammetry—the last decade and beyond. In *Proceedings of the XXth ISPRS Congress* (Istanbul, 2004), pp. 972–983.
- [BFBK02] BHOTIKA R., FLEET D. J., KUTULAKOS K. N.: A probabilistic theory of occupancy and emptiness. In *Proceedings of European Conference on Computer Vision (ECCV)* 3 (Copenhagen, Denmark, 2002), pp. 112–132.
- [Bla85] BLAKE A.: Specular stereo. In *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI)* (Los Angeles, CA, USA, 1985), Morgan Kaufmann, pp. 973–976.
- [Bla04] BLAIS F.: Review of 20 years of range sensor development. *Journal of Electronic Imaging* 13, 1 (2004), 231–243.
- [BN95] BHAT D., NAYAR S.: Stereo in the presence of specular reflection. In *Proceedings of IEEE International Conference on Computer Vision (ICCV)* (Boston, MA, 1995), pp. 1086–1092.
- [BS03] BONFORT T., STURM P.: Voxel carving for specular surfaces. In *Proceedings of IEEE International Conference on Computer Vision (ICCV)* (Nice, France, 2003), pp. 591–596.
- [BSG06] BONFORT T., STURM P., GARGALLO P.: General specular surface triangulation. In *Proceedings of the Asian Conference on Computer Vision* (Hyderabad, India, 2006), vol. 2, pp. 872–881.
- [BV99] BONET J. S. D., VIOLA P. A.: Roxels: responsibility weighted 3D volume reconstruction. In *Proceedings of IEEE International Conference on Computer Vision (ICCV)* (Corfu, 1999), pp. 418–425.
- [CGS06] CHEN T., GOESELE M., SEIDEL H.-P.: Mesostructure from specularly. In *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)* (New York, NY, USA, 2006), pp. 17–22.
- [CKS*05] CRIMINISI A., KANG S. B., SWAMINATHAN R., SZELISKI R., ANANDAN P.: Extracting layers and analyzing their specular properties using epipolar-plane-image analysis. *Computer Vision and Image Understanding* 97, 1 (2005), 51–85.
- [CL95] CURLESS B., LEVOY M.: Better optical triangulation through spacetime analysis. In *Proceedings of IEEE*

- International Conference on Computer Vision (ICCV)* (Boston, MA, 1995), pp. 987–994.
- [CLFS07] CHEN T., LENSCH H. P. A., FUCHS C., SEIDEL H.-P.: Polarization and phase-shifting for 3d scanning of translucent objects. In *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)* (Minneapolis, MN, 2007), pp. 1–8.
- [COW*96] CHEW W. C., OTTO G. P., WEEDON W. H., LIN J. H., LU C. C., WANG Y. M., MOGHADDAM M.: Nonlinear diffraction tomography—the use of inverse scattering for imaging. *International Journal of Imaging Systems and Technology* 7 (1996), 16–24.
- [CRT06] CANDÈS E., ROMBERG J., TAO T.: Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory* 52, 2 (2006), 489–509.
- [CSL08] CHEN T., SEIDEL H.-P., LENSCH H. P. A.: Modulated phase-shifting for 3D scanning. In *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)* (Anchorage, 2008), pp. 1–8.
- [CTW97] CLARK J., TRUCCO E., WOLFF L. B.: Using light polarization in laser scanning. *Image and Vision Computing* 15, 1 (1997), 107–117.
- [CZH*00] CHUANG Y.-Y., ZONGKER D. E., HINDORFF J., CURLESS B., SALESIN D. H., SZELISKI R.: Environment matting extensions: Towards higher accuracy and real-time capture. In *Proceedings of ACM SIGGRAPH* (New Orleans, 2000), pp. 121–130.
- [Dan01] DANA K. J.: BRDF/BTF Measurement device. In *Proceedings of International Conference on Computer Vision (ICCV)* (Vancouver, BC, 2001), pp. 460–466.
- [DD01] DEUSCH S., DRACOS T.: Time resolved 3D passive scalar concentration-field imaging by laser-induced fluorescence (LIF) in moving liquids. *Measurement Science and Technology* 12 (2001), 188–200.
- [DNRR05] DAVIS J., NEHAB D., RAMAMOORTHY R., RUSINKIEWICZ S.: Spacetime stereo: A unifying framework for depth from triangulation. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)* 27, 2 (2005), 296–302.
- [DYW05] DAVIS J., YANG R., WANG L.: BRDF Invariant stereo using light transport constancy. In *Proceedings of IEEE International Conference on Computer Vision (ICCV)* (Beijing, China, 2005), pp. 436–443.
- [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.
- [FCB08] FRANCKEN Y., CUYPERS T., BEKAERT P.: Mesostucture from specularity using gradient illumination. In *Proceedings of ACM/IEEE International Workshop on Projector Camera Systems* (Los Angeles, CA, USA, 2008), ACM, pp. 11:1–11:7.
- [FCG*06] FUCHS C., CHEN T., GOESELE M., THEISEL H., SEIDEL H.-P.: Volumetric density capture from a single image. In *Proceedings of International Workshop on Volume Graphics* (Boston, MA, USA, 2006), IEEE, pp. 17–22.
- [FCG*07] FUCHS C., CHEN T., GOESELE M., THEISEL H., SEIDEL H.-P.: Density estimation for dynamic volumes. *Computer and Graphics* 31, 2 (April 2007), 205–211.
- [FCM*08] FRANCKEN Y., CUYPERS T., MERTENS T., GIELIS J., BEKAERT P.: High quality mesostructure acquisition using specularities. In *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)* (Anchorage, 2008), pp. 1–7.
- [FHCB08] FRANCKEN Y., HERMANS C., CUYPERS T., BEKAERT P.: Fast normal map acquisition using an LCD screen emitting gradient patterns. In *Proceedings of Canadian Conference on Computer and Robot Vision* (Windsor, ON, CA, 2008), pp. 189–195.
- [FHL*08] FUCHS C., HEINZ M., LEVOY M., SEIDEL H.-P., LENSCH H. P. A.: Combining confocal imaging and descattering. *Computer Graphics Forum, Special Issue for the Eurographics Symposium on Rendering (EGSR)* 27, 4 (Jun. 2008), 1245–1253.
- [FJLA02] FUCHS E., JAFFE J., LONG R., AZAM F.: Thin laser light sheet microscope for microbial oceanography. *Optics Express* 10, 2 (2002), 145–154.
- [GBR*01] GODIN G., BERARDIN J.-A., RIOUX M., LEVOY M., COURNOYER L., BLAIS F.: An assessment of laser range measurement of marble surfaces. In *Proceedings of Fifth Conference on Optical 3-D Measurement Techniques* (Wien, Austria, 2001), pp. 49–56.
- [GCHS05] GOLDMAN D. B., CURLESS B., HERTZMANN A., SEITZ S. M.: Shape and spatially-varying BRDFs from photometric stereo. In *Proceedings of IEEE International Conference on Computer Vision (ICCV)* (Beijing, China, 2005), pp. 341–348.
- [GILM07] GOLDLÜCKE 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.

- [GLL*04] GOESELE M., LENSCH H. P. A., LANG J., FUCHS C., SEIDEL H.-P.: Disco—acquisition of translucent objects. In *Proceedings of ACM SIGGRAPH* (Los Angeles, CA, USA, 2004), pp. 835–844.
- [GNG*08] GU J., NAYAR S. K., GRINSUN E., BELHUMEUR P. N., RAMAMOORTHY R.: Compressive structured light for recovering inhomogeneous participating media. In *Proceedings of European Conference on Computer Vision (ECCV)* (Marseille, France, 2008), pp. 845–858.
- [GNS07] GRAVES K., NAGARAJAH R., STODDART P. R.: Analysis of structured highlight stereo imaging for shape measurement of specular objects. *Optical Engineering* 46, 8 (Aug. 2007), published online, 30, August 2007.
- [GNS08] GUPTA M., NARASIMHAN S. G., SCHECHNER Y. Y.: On controlling light transport in poor visibility environments. In *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)* (Anchorage, 2008), pp. 1–8.
- [Has02] HASINOFF S. W.: Three-dimensional reconstruction of fire from images. MSc Thesis, University of Toronto, Department of Computer Science, 2002.
- [HB88] HEALEY G., BINFORD T. O.: Local shape from specularly. *Computer Vision, Graphics, and Image Processing* 42, 1 (1988), 62–86.
- [HBKM96] HALSTEAD M. A., BARSKY B. A., KLEIN S. A., MANDELL R. B.: Reconstructing curved surfaces from specular reflection patterns using spline surface fitting of normals. In *Proceedings of ACM SIGGRAPH* (New Orleans, 1996), pp. 335–342.
- [HED05] HAWKINS T., EINARSSON P., DEBEVEC P.: Acquisition of time-varying participating media. In *Proceedings of ACM SIGGRAPH* (Los Angeles, CA, USA, 2005), pp. 812–815.
- [HFI*08] HULLIN M. B., FUCHS M., IHRKE I., SEIDEL H.-P., LENSCH H. P. A.: Fluorescent immersion range scanning. *ACM Transactions on Graphics (SIGGRAPH'08)* 27, 3 (2008), 87:1–87:10.
- [HK03] HASINOFF S. W., KUTULAKOS K. N.: Photo-consistent 3D fire by flame-sheet decomposition. In *Proceedings of IEEE International Conference on Computer Vision (ICCV)* (Nice, France, 2003), pp. 1184–1191.
- [HK07] HASINOFF S. W., KUTULAKOS K. N.: Photo-consistent reconstruction of semi-transparent scenes by density sheet decomposition. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)* 29, 5 (2007), 870–885.
- [HLHZ08] HOLROYD M., LAWRENCE J., HUMPHREYS G., ZICKLER T.: A photometric approach for estimating normals and tangents. *ACM Transactions on Graphics (SIGGRAPH Asia 2008)* 27, 5 (2008), 133:1–133:9.
- [Höh71] HÖHLE J.: Reconstruction of the underwater object. *Photogrammetric Engineering* (1971), 948–954.
- [HS81] HORN B. K. P., SCHUNCK B. G.: Determining optical flow. *Artificial Intelligence* 17 (1981), 185–203.
- [HSKK96] HATA S., SAITO Y., KUMAMURA S., KAIDA K.: Shape extraction of transparent object using genetic algorithm. In *Proceedings of International Conference on Pattern Recognition (ICPR)* (Vienna, 1996), vol. 4, pp. 684–688.
- [HZ00] HARTLEY R., ZISSERMAN A.: *Multiple View Geometry*. Cambridge University Press, Cambridge, UK, 2000.
- [IGM05] IHRKE I., GOLDLUECKE B., MAGNOR M.: Reconstructing the geometry of flowing water. In *Proceedings of IEEE International Conference on Computer Vision (ICCV)* (Beijing, China, 2005), pp. 1055–1060.
- [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), 661–670.
- [IM04] IHRKE I., MAGNOR M.: Image-based tomographic reconstruction of flames. In *Proceedings of ACM SIGGRAPH Symposium on Computer Animation (SCA)* (Los Angeles, CA, USA, 2004), pp. 367–375.
- [IM05] IHRKE I., MAGNOR M.: Adaptive grid optical tomography. In *Proceedings of Vision, Video, and Graphics (VVG)* (Edinburgh, UK, 2005), TRUCCO E., CHANTLER M., (Eds.), *Eurographics*, pp. 141–148.
- [IM06] IHRKE I., MAGNOR M.: Adaptive grid optical tomography. *Graphical Models* 68 (2006), 484–495.
- [JDJ06] JOSHI N., DONNER C., JENSEN H. W.: Non-invasive scattering anisotropy measurement in turbid materials using non-normal incident illumination. *Optics Letters* 31 (2006), 936–938.
- [JSY03] JIN H., SOATTO S., YEZZI A. J.: Multi-view stereo beyond Lambert. In *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)* (Madison, WI, 2003), pp. 171–178.
- [JSY05] JIN H., SOATTO S., YEZZI A. J.: Multi-view stereo reconstruction of dense shape and complex appearance. *International Journal of Computer Vision (IJCV)* 63, 3 (2005), 175–189.
- [KS00] KUTULAKOS K. N., SEITZ S.: A theory of space by space carving. *International Journal of Computer Vision (IJCV)* 38, 3 (2000), 199–218.

- [KS01] KAK A. C., SLANEY M.: *Principles of Computerized Tomographic Imaging*. Society of Industrial and Applied Mathematics, Philadelphia, USA, 2001.
- [KS05] KUTULAKOS K. N., STEGER E.: A theory of refractive and specular 3D shape by light-path triangulation. In *Proceedings of IEEE International Conference on Computer Vision (ICCV)* (Beijing, China, 2005), pp. 1448–1455.
- [KS07] 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.
- [KSW08] KELLER P. J., SCHMIDT A. D., WITTBRODT J., STELZER E. H. K.: Reconstruction of zebrafish early embryonic development by scanned light sheet microscopy. *Science* 322, 5904 (2008), 1065–1069.
- [KTM*02] KANITSAR A., THEUSSL T., MROZ L., ŠRÁMEK M., BARTROLÍ A. V., CSÉBFALVI B., HLÁDUVKA J., GUTHE S., KNAPP M., WEGENKITTL R., FELKEL P., RÖTTGER S., FLEISCHMANN D., PURGATHOFER W., GRÖLLER 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 IEEE Visualization* (Boston, 2002), IEEE, pp. 489–492.
- [Lau94] LAURENTINI A.: The visual Hull concept for silhouette-based image understanding. *IEEE Transactions on Pattern Analysis and Machine Recognition (PAMI)* 16, 2 (1994), 150–162.
- [LH96] LEVOY M., HANRAHAN P.: Light field rendering. In *Proceedings of SIGGRAPH'96* (New Orleans, 1996), ACM, pp. 31–42.
- [LHM*07] LINȚU A., HOFFMAN L., MAGNOR M., LENSCH H. P. A., SEIDEL H.-P.: 3D reconstruction of reflection nebulae from a single image. In *Proceedings of Vision, Modelling and Visualization (VMV)* (Saarbruecken, Germany, 2007), Aka GmbH, pp. 109–116.
- [LK81] LUCAS B., KANADE T.: An iterative image registration technique with an application to stereo vision. In *Proceedings of Seventh International Joint Conference on Artificial Intelligence* (Karlsruhe, West Germany, 1981), William Kaufmann, pp. 674–679.
- [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, 2 (2003), 234–257.
- [LLL*02] LI Y., LIN S., LU H., KANG S. B., SHUM H.-Y.: Multibaseline stereo in the presence of specular reflections. In *Proceedings of International Conference on Pattern Recognition (ICPR)* (Quebec City, Quebec, Canada, 2002), pp. 573–576.
- [LLM*07a] LINȚU A., LENSCH H. P. A., MAGNOR M., EL-ABED S., SEIDEL H.-P.: 3D reconstruction of emission and absorption in planetary Nebulae. In *Proceedings of IEEE/EG International Symposium on Volume Graphics* (Prague, Czech Republik, 2007), H.-C. Hege and R. Machiraju, (Eds.), pp. 9–16.
- [LLM*07b] LINȚU A., LENSCH H. P. A., MAGNOR M., LEE T.-H., EL-ABED S., SEIDEL H.-P.: A multi-wavelength-based method to de-project gas and dust distributions of several planetary Nebulae. In *Proceedings of the Asymmetrical Planetary Nebulae IV* (La Palma, Spain, 2007), pp. 1–6.
- [LNA*06] LEVOY M., NG R., ADAMS A., FOOTER M., HOROWITZ M.: Light field microscopy. In *Proceedings of ACM SIGGRAPH* (Boston, 2006), pp. 924–934.
- [LSLF08] LAGGER P., SALZMANN M., LEPETIT V., FUA P.: 3D pose refinement from reflections. In *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)* (Anchorage, 2008), pp. 1–8.
- [Maa95] MAAS H.-G.: New developments in multimedia photogrammetry. In *Optical 3D Measurement Techniques III*. A. GRÜN, and H. KAHMEN (Eds.), Wichmann Verlag, Berlin, Germany, 1995.
- [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 Eurographics Symposium on Rendering (EGSR)* (Grenoble, 2007), pp. 183–194.
- [MI05] MIYAZAKI D., IKEUCHI K.: Inverse polarization ray-tracing: estimating surface shapes of transparent objects. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (San Diego, 2005), vol. 2, pp. 910–917.
- [MK05] MORRIS N. J. W., KUTULAKOS K. N.: Dynamic refraction stereo. In *Proceedings of IEEE International Conference on Computer Vision (ICCV)* (Beijing, China, 2005), pp. 1573–1580.
- [MK07] MORRIS N. J. W., KUTULAKOS K. N.: Reconstructing the surface of inhomogeneous transparent scenes by scatter-trace photography. In *Proceedings of IEEE International Conference on Computer Vision (ICCV)* (Rio de Janeiro, 2007), pp. 1–8.
- [MKHD04] MAGNOR M., KINDLMANN G., HANSEN C., DURIC N.: Constrained inverse volume rendering for planetary nebulae. In *Proceedings of IEEE Visualization* (Austin, TX, USA, 2004), IEEE, pp. 83–90.

- [MKZB01] MAGDA S., KRIEGMAN D. J., ZICKLER T., BELHUMEUR P. N.: Beyond Lambert: Reconstructing surfaces with arbitrary BRDFs. In *Proceedings of IEEE International Conference on Computer Vision (ICCV)* (Vancouver, BC, 2001), pp. 391–398.
- [MPN*02] MATUSIK W., PFISTER H., NGAN A., BEARDSLEY P., ZIEGLER R., MCMILLAN L.: Image-based 3D photography using opacity hulls. In *Proceedings of ACM SIGGRAPH* (San Antonio, 2002), pp. 427–436.
- [MPZ*02] MATUSIK W., PFISTER H., ZIEGLER R., NGAN A., MCMILLAN L.: Acquisition and rendering of transparent and refractive objects. In *Proceedings of Eurographics Symposium on Rendering (EGSR)* (Pisa, 2002), pp. 267–278.
- [Mur90] MURASE H.: Surface shape reconstruction of an undulating transparent object. In *Proceedings of IEEE International Conference on Computer Vision (ICCV)* (Osaka, 1990), pp. 313–317.
- [Mur92] MURASE H.: Surface shape reconstruction of a non-rigid transparent object using refraction and motion. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)* 14, 10 (1992), 1045–1052.
- [MKZB05] MALLICK S. P., ZICKLER T., KRIEGMAN D. J., BELHUMEUR P. N.: Beyond Lambert: reconstructing specular surfaces using color. In *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)* (San Diego, 2005), pp. 619–626.
- [NFB93] NAYAR S. K., FANG X.-S., BOULT T.: Removal of Specularities using color and polarization. In *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)* (New York, NY, USA, 1993), pp. 583–590.
- [NGD*06] NARASIMHAN S., GUPTA M., DONNER C., RAMAMOORTHY R., NAYAR S., JENSEN H. W.: Acquiring scattering properties of participating media by dilution. In *Proceedings of ACM SIGGRAPH* (Boston, 2006), pp. 1003–1012.
- [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. In *Proceedings of ACM SIGGRAPH* (Boston, 2006), pp. 935–944.
- [NKY08] NARASIMHAN S. G., KOPPAL S. J., YAMAZAKI S.: Temporal dithering of illumination for fast active vision. In *Proceedings of the 10th European Conference on Computer Vision (ECCV)* (Marseille, France, 2008), pp. 830–844.
- [NNSK05] NARASIMHAN S. G., NAYAR S. K., SUN B., KOPPAL S. J.: Structured light in scattering media. In *Proceedings of IEEE International Conference on Computer Vision (ICCV)* (Beijing, China, 2005), pp. 420–427.
- [NSWS90] NAYAR S. K., SANDERSON A. C., WEISS L., SIMON D.: Specular surface inspection using structured highlight and Gaussian images. *IEEE Transactions on Robotics and Automation* 6, 2 (Apr. 1990), 208–218.
- [NWR08] NEHAB D., WEYRICH T., RUSINKIEWICZ S.: Dense 3D Reconstruction from specular consistency. In *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)* (Anchorage, 2008), pp. 1–8.
- [OF03] OSHER S., FEDKIW R.: *Level Set Methods and Dynamic Implicit Surfaces*, vol. 153 of *Applied Mathematical Sciences*. Springer-Verlag, New York, 2003.
- [ON96] OREN M., NAYAR S. K.: A theory of specular surface geometry. *International Journal of Computer Vision (IJCV)* 24, 2 (1996), 105–124.
- [PD03] PEERS P., DUTRÉ P.: Wavelet environment matting. In *Proceedings of Eurographics Symposium on Rendering (EGSR)* (Leuven, 2003), pp. 157–166.
- [PK04] PARK J., KAK A. C.: Specularity elimination in range sensing for accurate 3D modeling of specular objects. In *Proceedings of First International Symposium on 3D Data Processing, Visualization and Transmission (3DPVT)* (Thessaloniki, Greece, 2004), IEEE, pp. 707–714.
- [PK08] PARK J., KAK C.: 3D modeling of optically challenging objects. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 14, 2 (2008), 246–262.
- [Rad17] RADON J.: Über die Bestimmung von Funktionen durch ihre Integralwerte längs gewisser Mannigfaltigkeiten. *Ber. Ver. Sachs. Akad. Wiss. Leipzig, Math-Phys. Kl.* 69 (1917) 262–277.
- [RB06] ROTH S., BLACK M. J.: Specular flow and the recovery of surface structure. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (New York, NY, USA, 2006), pp. 1869–1876.
- [REH06] REMONDINO F., EL-HAKIM S.: Image based 3D modeling: A review. *The Photogrammetric Record* 21, 115 (2006), 269–291.
- [RMD04] RECHE A., MARTIN I., DRETTAKIS G.: Volumetric reconstruction and interactive rendering of trees from photographs. In *Proceedings of ACM SIGGRAPH* (Los Angeles, CA, USA, 2004), pp. 720–727.

- [SAH04] SOLEM J. E., AANAES H., HEYDEN A.: A Variational analysis of shape from specularities using sparse data. In *Proceedings of International Symposium on 3D Data Processing, Visualization and Transmission (3DPVT)* (Thessaloniki, Greece, 2004), IEEE, pp. 26–33.
- [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 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)* (New York, NY, USA, 2006), pp. 519–526.
- [Sch94] SCHULTZ H.: Retrieving shape information from multiple images of a specular surface. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)* 16, 2 (1994), 195–201.
- [SCP05] SARAVESE S., CHEN M., PERONA P.: Local shape from mirror reflections. *International Journal of Computer Vision (IJCV)* 64, 1 (2005), 31–67.
- [Set99] SETHIAN J. A.: *Level Set Methods and Fast Marching Methods*, 2nd ed. Monographs on Applied and Computational Mathematics. Cambridge University Press, Cambridge, UK, 1999.
- [SG99] SZELISKI R., GOLLAND P.: Stereo matching with transparency and matting. *International Journal of Computer Vision (IJCV)* 32, 1 (1999), 45–61.
- [SGN01] SWAMINATHAN R., GROSSBERG M. D., NAYAR S. K.: Caustics of catadioptric cameras. In *Proceedings of International Conference on Computer Vision (ICCV)* (Vancouver, BC, 2001), pp. 2–9.
- [SGN06] SWAMINATHAN R., GROSSBERG M. D., NAYAR S. K.: Non-single viewpoint catadioptric cameras: geometry and analysis. *International Journal of Computer Vision (IJCV)* 66, 3 (2006), 211–229.
- [SKS*02] SWAMINATHAN R., KANG S. B., SZELISKI R., CRIMINISI A., NAYAR S. K.: On the motion and appearance of specularities in image sequences. In *Proceedings of European Conference on Computer Vision (ECCV)* (Copenhagen, Denmark, 2002), pp. 508–523.
- [SP01] SAVARESE S., PERONA P.: Local analysis for 3d reconstruction of specular surfaces. In *Proceedings of IEEE International Conference on Computer Vision (ICCV)* (Vancouver, BC, 2001), pp. 594–603.
- [SP02] SAVARESE S., PERONA P.: Local analysis for 3D reconstruction of specular surfaces—part II. In *Proceedings of European Conference on Computer Vision (ECCV)* (Copenhagen, Denmark, 2002).
- [SSIK99] SAITO M., SATO Y., IKEUCHI K., KASHIWAGI H.: Measurement of surface orientations of transparent objects using polarization in highlight. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (Fort Collins, CO, 1999), vol. 1, pp. 381–386.
- [STM06] STICH T., TEVS A., MAGNOR M.: Global Depth from epipolar volumes—a general framework for reconstructing non-Lambertian surfaces. In *Proceedings of Third International Symposium on 3D Data Processing, Visualization and Transmission (3DPVT)* (Chapel Hill, NC, 2006), IEEE, pp. 1–8.
- [SWN88] SANDERSON A. S., 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.
- [TAL*07] THEOBALT C., AHMED N., LENSCH H. P. A., MAGNOR M., SEIDEL H.-P.: Seeing people in different light: Joint shape, motion, and reflectance capture. *IEEE Transactions on Visualization and Computer Graphics* 13, 4 (2007), 663–674.
- [TBH06] TRIFONOV B., BRADLEY D., HEIDRICH W.: Tomographic reconstruction of transparent objects. In *Proceedings of Eurographics Symposium on Rendering (EGSR)* (Cyprus, 2006), pp. 51–60.
- [TF94] TRUCCO E., FISHER R. B.: Acquisition of consistent range data using local calibration. In *Proceedings of IEEE International Conference on Robotics and Automation* (1994), pp. 3410–3415.
- [TLGS05] TARINI M., LENSCH H. P. A., GOESELE M., SEIDEL H.-P.: 3D Acquisition of mirroring objects. *Graphical Models* 67, 4 (2005), 233–259.
- [TMHF00] TRIGGS B., MCLAUCHLAN P., HARTLEY R., FITZGIBBON A.: Bundle adjustment—a modern synthesis. In *Vision Algorithms: Theory and Practice*. W. TRIGGS, A. ZISSERMAN, and R. SZELISKI, (Eds.), LNCS. Springer Verlag, London, UK, 2000, pp. 298–375.
- [TS67] TORRANCE K., SPARROW E.: Theory for off-specular reflection from rough surfaces. *Journal of the Optical Society of America* 57, 9 (Anchorage, 1967), 1105–1114.
- [TSS08] TREIBITZ T., SCHECHNER Y. Y., SINGH H.: Flat refractive geometry. In *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)* (Anchorage, 2008), pp. 1–8.
- [VAZBS08] VASILYEV Y., ADATO Y., ZICKLER T., BEN-SHAHAR O.: Dense specular shape from multiple specular flows. In *Proceedings of IEEE Computer Society Conference*

- on *Computer Vision and Pattern Recognition (CVPR)* (Anchorage, 2008), pp. 1–8.
- [VVD*04] VAN VLIET E., VAN BERGEN S. M., DERKSEN J. J., PORTELA L. M., VAN DEN AKKER H. E. A.: Time-resolved, 3D, laser-induced fluorescence measurements of fine-structure passive scalar mixing in a tubular reactor. *Experiments in Fluids* 37 (2004), 1–21.
- [WD06] WANG J., DANA K. J.: Relief texture from specularities. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)* 28, 3 (2006), 446–457.
- [WFZ02] WEXLER Y., FITZGIBBON A., ZISSERMAN A.: Image-based environment matting. In *Proceedings of Eurographics Symposium on Rendering (EGSR)* (Pisa, 2002), pp. 279–290.
- [WLZ*09] WANG H., LIAO M., ZHANG Q., YANG R., TURK G.: Physically guided liquid surface modeling from videos. *ACM Transactions on Graphics (SIGGRAPH '09)* 28, 3 (2009), 90:1–90:12.
- [Woo80] WOODMAN R. J.: Photometric method for determining surface orientation from multiple images. *Optical Engineering* 19, 1 (1980), 139–144.
- [YM04] YU J., MCMILLAN L.: General linear cameras. In *Proceedings of European Conference on Computer Vision (ECCV)* (Prague, Czech Republic, 2004), pp. 14–27.
- [YMK06] YAMAZAKI S., MOCHIMARU M., KANADE T.: Inverse volume rendering approach to 3D reconstruction from multiple images. In *Proceedings of Asian Conference on Computer Vision (ACCV)* (Hyderabad, India, 2006), pp. 409–418.
- [YPW03] YANG R., POLLEFEYS M., WELCH G.: Dealing with textureless regions and specular highlights—a progressive space carving scheme using a novel photo-consistency measure. In *Proceedings of IEEE International Conference on Computer Vision (ICCV)* (Nice, France, 2003), pp. 576–584.
- [YY08] YING Y., YU J.: Recovering shape characteristics on near-flat specular surfaces. In *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)* (Anchorage, 2008), pp. 1–8.
- [ZBK02] ZICKLER T., 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.
- [ZFA97] ZHENG J. Y., FUKAGAWA Y., ABE N.: 3D Surface estimation and model construction from specular motion in image sequences. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)* 19, 5 (May 1997), 513–520.
- [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.
- [ZHK*03] ZICKLER T. E., HO J., KRIEGMAN D. J., PONCE J., BELHUMEUR P. N.: Binocular Helmholtz stereopsis. In *Proceedings of IEEE International Conference on Computer Vision (ICCV)* (Nice, France, 2003), pp. 1411–1417.
- [ZM98] ZHENG J. Y., MURATA A.: Acquiring 3D object models from specular motion using circular lights illumination. In *Proceedings of IEEE International Conference on Computer Vision (ICCV)* (Mumbai (Bombay), India 1998), pp. 1101–1108.
- [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 (Aug. 2000), 913–920.
- [ZMFA96] ZHENG J. Y., MURATA A., FUKAGAWA Y., ABE N.: Reconstruction of 3D models from specular motion using circular lights. In *Proceedings of International Conference on Pattern Recognition (ICPR)* (Vienna, 1996), pp. 869–873.
- [ZWC99] ZONGKER D., WERNER D., CURLESS B., SALESIN D.: Environment matting and compositing. In *Proceedings of ACM SIGGRAPH* (Los Angeles, CA, USA, 1999), pp. 205–214.