

State-of-the-art in Multi-Light Image Collections for Surface Visualization and Analysis

R. Pintus¹, T. G. Dulecha², I. Ciortan³, E. Gobbetti¹, and A. Giachetti²

¹CRS4, Italy

² Department of Computer Science, University of Verona, Italy

³ Norwegian University of Science and Technology, Department of Computer Science, Norway

Abstract

Multi-Light Image Collections (MLICs), i.e., stacks of photos of a scene acquired with a fixed viewpoint and a varying surface illumination, provide large amounts of visual and geometric information. In this survey, we provide an up-to-date integrative view of MLICs as a mean to gain insight on objects through the analysis and visualization of the acquired data. After a general overview of MLICs capturing and storage, we focus on the main approaches to produce representations usable for visualization and analysis. In this context, we first discuss methods for direct exploration of the raw data. We then summarize approaches that strive to emphasize shape and material details by fusing all acquisitions in a single enhanced image. Subsequently, we focus on approaches that produce relightable images through intermediate representations. This can be done both by fitting various analytic forms of the light transform function, or by locally estimating the parameters of physically plausible models of shape and reflectance and using them for visualization and analysis. We finally review techniques that improve object understanding by using illustrative approaches to enhance relightable models, or by extracting features and derived maps. We also review how these methods are applied in several, main application domains, and what are the available tools to perform MLIC visualization and analysis. We finally point out relevant research issues, analyze research trends, and offer guidelines for practical applications.

CCS Concepts

• **General and reference** → *Surveys and overviews*; • **Human-centered computing** → *Visualization systems and tools*; • **Computing methodologies** → *Computer vision representations*; *Reflectance modeling*;

1 Introduction

Multi-Light Image Collections (MLICs) are sets of digital images of a scene acquired from the same fixed viewpoint but with varying lighting conditions. The most general and common setup assumes that the color of the same scene point is captured with different incident light directions to gather information on the reflectance field. Such a collection of samples, typically arranged in image stacks, provides massive amounts of visual data that can be analyzed to extract information and knowledge on shape and appearance.

The continuous improvement of controllable lighting and digital photography has made the acquisition of high-resolution MLICs practical and affordable using many different physical setups (Sec. 3). As a result, a wide variety of applications are using MLICs as a major mean to non-destructively gather information on scenes and objects at many scales, as well as to provide users with useful visualization tools for object analysis (Sec. 6). These applications benefit from an array of computational tools that have been devised in different domains, such as feature detection and enhancement [FAR07], reconstruction of normals and 3D

shapes [Woo80, AG15], extraction of surface appearance and material behavior [WK15, DRS10], and creation of relightable images [MGW01]. The research community working on these problems appears, however, fragmented, as highlighted by the fact that even the input data are called in several different ways, including MLIC [FAR07], Reflectance Transformation Imaging (RTI) Stacks [MMSL06], or Photometric Stereo Datasets [SWM*16], depending on the type of processing that is applied.

In the recent past, extensive surveys have been presented for several aspects of MLICs, mainly focusing, however, on very specific acquisition and modeling aspects (Sec. 2). In this survey, we provide, instead, an up-to-date integrative view of MLICs as a mean to gain insight on scenes and objects through the analysis and visualization of the acquired image stacks. Our aim is to provide a fresh view of the subject, and help readers navigate through the constantly expanding literature on MLIC analysis and visualization. The target audience of our report includes researchers in analysis and visualization of imaging data, as well as scholars and practitioners in the relevant application fields. Researchers will find a structured overview of the

field, which organizes the various methods and their applications, and indicates challenging open problems. Domain experts will, in turn, find a presentation of the areas where MLICs are helping, or have the potential to improve, analysis work and scholarly research. After providing background information on related survey literature (Sec. 2) and MLIC capturing and representation techniques (Sec. 3), we introduce a taxonomy of the surveyed methods (Sec. 4) and provide a review of available techniques (Sec. 5). We then discuss how these methods are applied in practice (Sec. 6.1), and what are the available tools to perform MLIC analysis (Sec. 6.2). We finally point out relevant research issues, analyze research trends, and offer guidelines for practical applications (Sec. 7).

2 Related surveys

The areas of MLIC acquisition, processing, and analysis are vast and have many applications (Sec. 6.1). Available surveys have focused, so far, on specific acquisition and modeling aspects, such as the geometry computation [AG15] and the extraction of surface appearance and material behavior [WK15, DRS10], including Bi-directional Reflectance Distribution Function (BRDF) fitting [WLL*09, GGG*16] and Bidirectional Texture Function (BTF) extraction [FH09]. We refer the reader to these prior surveys for a more in-depth analysis of these vertical subjects.

Our analysis is focused on MLICs obtained with varying incident illumination, which is the wider class. We do not include capture and processing methods involving only frequency variations (multi-/hyper-spectral imaging), or only polarization variations, but we consider the literature coupling frequency/polarization and direction changes [VHW*18]. We refer the reader to a recent survey [YGC17] for coverage of multi-/hyper-spectral-only methods.

Since the ratio between the largest and smallest values acquired is typically extremely high, especially for glossy materials, most of the methods discussed here will benefit from being applied in a high-dynamic-range (HDR) setting. Many of the standard HDR techniques can be applied to MLICs at the acquisition stage to fuse multiple captures, and at the presentation stage to deal with tone mapping issues. This is an orthogonal aspect, and we thus refer the reader to a comprehensive book on the subject [RHD*10].

3 Background

MLICs are series of digital images arranged in a stack. They are acquired from the same fixed viewpoint but varying the lighting condition at each image (Fig. 1). Thus, in the image stack, a single view ray is associated with each pixel position, while different light rays are associated with it (typically, but not necessarily, one per image in the stack). In this survey, we consider that a variation of light direction is always present, while variations in other parameters (e.g., intensity, wavelength, polarization) are optional. In order to facilitate visual analysis, the image stack is typically transformed into intermediate representations, from which photorealistic or illustrative images can be generated to support visualization and analysis tasks, often in interactive settings. Before presenting a classification of the various approaches (Sec. 4) and discussing them in detail (Sec. 5), we provide relevant background information on light configurations, and on input and output data representations.




Figure 1: General concept. An imaging device acquires a stack of images while varying the illumination to generate a MLIC, which is then transformed into an intermediate representation for visualization and analysis.

3.1 Light configurations

Many physical realizations exist for MLIC capture, and full coverage is out of the scope of this survey. We summarize here only the main characteristics which have an impact on the subsequent processing and visualization steps.

The most common approach is to capture a MLIC by using a differently positioned small light source per image (often approximated by a directional, point, or spotlight model). This can be achieved by static setups, which sequentially fire a different light, by dynamic setups, which move a single light to a different location at each shot, or by a combination of both approaches. In terms of impact on visualization and analysis, the most important aspects are the number of light sources and the uniformity of their distribution.

The smallest static setup is the 4-source *light-ring* [RTF*04]. Since it is simple and easy to miniaturize, a version of this design has been tailored to medical applications (Sec. 6.1). However, this minimal setup leads to undersampling of the reflectance function and missing data due to shadows, so most of the MLIC processing is done with a much larger number of lights uniformly sampling the entire sphere or hemisphere around the scene. Implementations of this approach include large static *light-domes* approximating a far-light condition [MGW01, PSM05, CHI19, HP15, EBB*11, AIK13], smaller micro-domes [EBB*11, FBKR17, Ham15, WVM*05, PBFS14, VVP*18, HBMG02] requiring the handling of non-collimated light-rays through near-light or spot-light models [AG15, PCGG16, HWBC15], or virtual domes made with moving light arcs or robotic arms with a small number of light sources to span the entire hemisphere [MBW*14, SOSI03, DCCS06]. The known regular arrangement of lights in these solutions, using tens to hundreds of lights, allows methods, user interfaces, and tools to exploit a parameterization of the light space for navigation and/or interpolation (Sec. 5 and Sec. 6.2).

In free-form setups, instead, the light source is freely moved in space with hand-held or robotics setups [MDA02, CGS06, MVSL05, KK13]. This approach is more flexible but requires more calibration efforts to deal with non-uniformity of light intensity and the variable spacing of positioning [GCD*18], and more effort to design interfaces to navigate in the data. For these reasons, physical constraints are often applied to approximate a uniform sampling of the hemisphere in a free-form setting [CHI19], or, alternatively, a regular light distribution is simulated through interpolation before further processing or interaction [PCS18].

As an alternative to moving point lights, several specialized solu-

tions have been introduced to control incident illumination through *extended light sources*, e.g., linear lights [GTHD03, RDL*15, RWS*11], or more complex *structured illumination models*, such as spherical gradient illumination [MHP*07, GCP*09, GCP*10] or spherical harmonic illumination [TFG*13]. However, these approaches are mostly used to support specialized processing methods, especially for BRDF extraction (Sec. 5.4), and, so far, have not been directly exploited for visualization and analysis tasks, since the generated model is oblivious of the illumination pattern.

3.2 Captured data representation

Some of the techniques for visualizing and analyzing MLICs can work without any other information than the captured images, eventually after performing undistortion as well as a wavelength-dependent correction to take into account the camera response (Sec. 5.2). However, most of the visualization and analysis methods also require additional information on camera parameters and lights to reason on the relation between the measured light, the view direction, and the incoming light direction and intensity. The images are thus augmented with metadata coming out of calibration.

A large variety of solutions have been proposed, including assuming mechanically calibrated arrangements for domes and rings [RTF*04], using known calibration objects [GCD*18, Mac15, CCC08, MDA02, GCD*17], and/or autocalibration techniques that reconstruct illumination parameters by making assumptions on the image reflectance field [WMTG05, XDW15, PF14, GCD*18] or on the light form factor [ASSS14, USTY15, PCGG16, XSJ*15]. Ackermann et al. [AFG13] present a validation through error statistics of both a forward and backward geometric point light source calibration by using sets of different numbers of reflective spheres, showing that significant errors are present even in common use cases just for the identification of light position. However, while MLIC calibration strongly affects tasks such as 3D reconstruction, many of the MLIC visualization and analysis techniques proved to be reasonably robust to inaccuracies in a fine calibration, since the incident light field is smooth, and small global shifts have reduced effects in local details.

In the RTI settings [CHI19, MGW01, MMC*08, PLGF*15, PCS18], the classical approximation is a directional light per shot. For domes or planar arrangements, the recovered information stored together with the MLIC can be light positions, so that the classical encoding for planar lights is a discretized light matrix [RDL*15]. In order to cope with more complex real-world lights, some methods virtually use a different per-pixel light direction. Since the field is smooth and slowly varying, this per-pixel lighting information is not usually stored at image resolution, but approximated by a low-degree polynomial function [GCD*18]. It should be noted that in these methods the incident lighting direction is expressed relative to a particular reference object (e.g., a constant-depth plane), and not to the imaged object, whose actual depth is generally a result of further processing. More complex lighting patterns using light bases [GTHD03, MHP*07] or polarized light [GCP*10] also store extra information related to the specific light configuration. Finally, note that different terms are often used to refer to the same light parameterization. For instance, the measured color as a function of incoming light direction is sometimes called Light transport function [HWBC15] or Appearance profile [GCD*18] and its discretization on a grid has been called Observation map [Ike18].

3.3 Output data representation

While the processing of MLIC data can result in a large variety of information, e.g., 3D surface meshes from photometric stereo [XDW15] or single-material BRDFs from imaging of planar samples [NDM05], the most widespread representation in the context of visualization and analysis is image-based. Thus, visualization and analysis-friendly representations are associated with each pixel (or pixel group), by using the same view of the input MLIC. Per-pixel representations range from rearranged/compressed/resampled input data samples (Sec. 5.1), to single colors resulting from data fusion (Sec. 5.2), to a variety of learned parameters for computing novel relighting (Sec. 5.3), to shape and material samples usable in a deferred shading settings (Sec. 5.4), to data usable for illustrative visualization (Sec. 5.5), to a variety of derived maps, typically presented as overlays over original images (Sec. 5.6). In most cases, the processing leads to visualization methods and user interfaces that work in image space (Sec. 6.2).

4 Overview and taxonomy

A wide variety of methods have been devised over the past decades to extract information from MLICs. In this survey, we focus on techniques from the point of view of the data and interaction means that they make available to users for analysis and visualization. Our proposed taxonomy is summarized in Table 1.

A first set of methods strives to support the *direct exploration* of the captured MLIC (Sec. 5.1). These techniques typically exploit a light-space parameterization to offer a user-interface to browse the input dataset, eventually performing interpolation to provide continuity of motion, or to simulate in-between lighting configurations. Little processing or enhancements are applied in these techniques, which are mostly used as components of larger data exploration systems. Minimal calibration is required here, most of the time limited to an approximation of the spatial configuration of lights.

The second group of techniques generates a *fused single-image* (Sec. 5.2) that maximizes for each output pixel the amount of information present in the input MLIC. This is generally done by an optimization process, and the resulting image is static, and not used in interactive settings. User interaction is typically limited to changing parameters to guide the detail enhancement process. Moreover, many of them can be applied without the knowledge of the illumination pattern, thus limiting calibration needs.

User interaction is, instead, at the core of the techniques used to generate *relightable models*. These techniques typically rely on, at least approximately, calibrated data, and compute per-pixel representations that can be exploited at visualization time to generate novel illumination patterns. A first set of methods in this class, possibly the most popular one, constructs fitted representations that approximate the light transport function, or perform a light matrix domain fitting (Sec. 5.3). A second set of methods tackles the more complex problem of separating shape and material contributions, typically producing Spatially-Varying BRDF (SVBRDF), models coupled with mesostructure descriptors (Sec. 5.4). While the latter methods are, in general, usable also to solve the 3D reconstruction problem, in our context they are employed to produce relightable models, which can be evaluated in a deferred shading context. A

Class		Approaches	Features
Direct exploration (Sec. 5.1) <i>Visualization of input data, possibly with interpolated light positions</i>		Input browsing [Mac15]; Resampling & interpolation [Buh03, GCD*17, PCS18].	Good fidelity but little or no relighting possibilities; Interaction limited to moving in light space (best for dense sampling); Static visualization limits depth cues; Enhancements only on single images; No calibration or just light position approximation required for in-between light generation; Data size is an issue.
Single-image data fusion (Sec. 5.2) <i>Static presentation of all data in a single image</i>		Albedo/Normal [CDG*18, Woo80, AG15]; Multi-scale analysis [FAR07, ZLR*10]; Discontinuity detection [RTF*04]; Multi-spectral fusion [CSS15, VHW*18, VdPHB*16, WVP15].	Static visualization; Interaction limited to panning/zooming and parameter tuning; Enhanced perception of visible and invisible details; Automatic illustration and false color visualization; No calibration or just light direction approximation required; High per-image computation costs; Low memory footprint.
Relightable models <i>Interactive inspection with novel illumination</i>	Relightable images (Sec. 5.3) <i>Reflectance field models</i>	Polynomial Texture Map [DHOMH12, MGW01, PGPG17, ZD14]; (Hemi-)Spherical harmonics [BJK07, MMC*08, ZD14]; Discrete Modal Decomposition [PLGF*15, PLGM*17a]; Layered Models [DHOMH12, Mac15, Mac18, FBKR17, ZD14]; Deep Learning [RDL*15, XSHR18]; Light Transport Matrix [TBF*17].	Interactive relighting; Mostly low-frequency modeling and no global illumination effects; High-frequency components with increased computational complexity and less portable models; Good light/camera calibration needed; Compact and fast analytic models available.
	(SV)BRDF fitting (Sec. 5.4) <i>Decoupled shape and appearance models</i>	Example-based [HS17]; Analytic BRDF [AWL13, AWL15, GTHD03, HNI08, RPG16, RMS*08]; Polarization based [GCP*09, GCP*10, MHP*07]; Gradient illumination [GCP*09, MHP*07]; Spherical harmonic lighting [TFG*13].	Interactive relighting; Supports novel viewpoints; Interactive pan/zooming/rotate and BRDF parameter tuning and editing; Hard to compute in practice from small samples; Supports only specific classes of objects; Can be extremely photorealistic; Good light/camera calibration needed; Good trade-off memory/visualization quality.
Non-photorealistic enhancement (Sec. 5.5) <i>Illustrative presentation</i>		Diffuse gain [CM11, EBB*11, HBMG02, MVSL05, MGW01, New15, SSM*17, UW13]; Specularity Enhancement [AIK13, CM11, CSS15, EBB*11, GCD*17, HBMG02, HP15, KK13, MGW01, MWGA06, MVSL05, MMSL06, MCPC17, New15, SSM*17, UW13, WVM*05]; Exaggerated shading [GCD*17, MWGA06, PCC*10, VdPHB*16, VHW*18, WVM*05]; Unsharp masking [CM11, HP15, PCC*10, SSM*17]; Light extrapolation [HBMG02, MGW01, PSM05]; Multi-light enhancement [GCD*17, HP15, PCC*10].	Interactive relighting; Interactive pan/zoom/parameter tuning; Enhanced tiny details and global shape visualization; May be used for real-time enhancement during capture at video rates; Compact representation.
Features and derived maps (Sec. 5.6) <i>Extraction of shape and appearance characteristics</i>		Edge extraction [BCDG13, EBB*11, Pan16, RTF*04, VVP*18]; Saliency maps [PLGM*17b, PLGM*17c]; Sketch-like visualization [EBB*11, Ham15, WVM*05, ZHS*18]; Slope exaggeration [MWGA06, VVP*18, WVM*05]; Normal/gradient/curvature visualization [AG15, BCDG13, DGL*14, EBB*11, GCD*17, CDG*18, MHP*07, MBW*14, MCPC17, PSM05, SSM*17, TML*13, VdPHB*16, VHW*18, VVP*18, WLW17, WVM*05]; Normal maps [GCP*09, MBW*14]; Depth maps [CB12, DGL*14, CDG*18, TLQ06, WVM*05]; Other derived maps [GCD*17, GPDG12, CDG*18, PSM05].	Many visualization modes (edges, sketches, toon-like, normal maps, saliency); Interaction limited to panning/zooming and parameter setting unless mixed with other relighting techniques; Increased detail perception; Small memory footprint.

Table 1: Taxonomy. Major approaches and their features for each MLIC method class.

user interface is generally provided to control the simulation of novel lighting, which is used in exploration applications (e.g., using simulated raking lights to enhance particular features). In contrast to the *direct exploration* methods, the lighting used during the interactive task is not limited to be an interpolation of the ones used for the acquisition.

While these relighting techniques are typically employed in photorealistic settings, *illustrative techniques* help in gaining insight on objects through visual effects (Sec. 5.5). These methods are rarely applied for enhancing specific aspects of the raw images (e.g., color discontinuities and color frequencies at different levels of resolution), while, most often, they are combined with relightable image methods to support an improved visual interpretation of the scene of interest in an interactive setting.

Finally, *feature maps* and *derived maps* (Sec. 5.6) can be extracted from the data to identify interesting features concerning overall

shape, discontinuities, roughness, material changes, and so on. In the visualization and analysis setting, these maps are directly visualized (often as overlays) to highlight particular aspects of the imaged object, or are used to drive more complex visualization techniques (e.g., for illustrative rendering).

5 Methods

In the following sections, we build on our taxonomy to provide an in-depth analysis of the various methods that have been proposed in the literature. For each class of methods, we provide a general description of the overall goals of the methods, comparatively discuss the various techniques that have been proposed to reach those goals, and discuss the related visualization and analysis issues.

5.1 Direct exploration of acquired data

Direct exploration approaches aim at providing a framework to browse original data, by performing little processing except, eventu-

ally, for the computation of in-between images to simulate a continuous sampling of the light space. One of the main challenges is the need for dealing with large amounts of data at exploration time.

5.1.1 Approaches

In the case of original image navigation, no processing is needed on the image stack, so the only issue is how to offer interactive access to the displayed images, which is done either by a simple discrete selector, or by finding the nearest neighbor in a continuous interface for specifying the light position [Mac15].

In the direct interpolation scenario, Radial Basis Functions (RBF) [Buh03] interpolation has been typically used to produce a relighted image with a new virtual light from the MLIC stack [GCD*17]. Given the N images corresponding to the closest neighbors of the new light direction, RBF interpolation is achieved by computing the parameters that define a sum of N radial functions.

The reflectance I is expressed as $I(\vec{l}) = \sum_{i=1}^N \alpha_i e^{-\frac{\|\vec{l}-\vec{l}_i\|^2}{R^2}}$, where α_i is estimated per pixel by solving a least squares problem. A recent objective and subjective evaluation has shown, in particular, how RBF interpolation is capable of better visualizing the behaviour of complex materials for in-between views with respect to classic fitting approaches [PDJ*18] (Fig. 2). However, the need to access large amounts of data makes the method very computation and memory-intensive. For this reason, it has been combined with compression approaches. Since the extraction of a set of basis has been demonstrated to work well to approximate the appearance [SOSI03], Ponchio et al. [PCS18], in particular, propose a relighting method based on a simple bilinear interpolation. In the pre-processing step, they create a novel relighting with RBF by using novel light directions that form a regular grid in the light direction space. Then, they combine this information with a per-pixel compression strategy based on Principal Component Analysis (PCA). At exploration time, they perform relighting from a novel light direction by applying bilinear interpolation to the pre-computed grid, and by selecting the final color from the compact PCA basis. Their data encoding is extremely suitable for online relighting, and a WebGL based interactive viewer has been realized that exploits the combination of those two techniques. Finally, another interesting thing is that PCA components can be stored in byte planes, and compressed as JPEG images to support compact relightable data exchange [PCS18].

5.1.2 Visualization and Analysis

Direct exploration approaches are the least sensitive to calibration and raw data quality (e.g., for original data viewing without interpolation, the visual quality is simply the same as the capture sensor). For this reason, these techniques have been successfully employed in many application domains (e.g., Cultural Heritage - CH, Sec. 6.1), and they prove very suitable to directly treat complex materials with behaviour that is difficult to model [GCD*17]. Original data visualization, and to a lesser extent interpolated data, is capable of reproducing both high-frequency components (e.g., highlights) of the reflectance field, and global effects, such as cast shadows and interreflections, that allow users to appreciate and understand the geometry at a meso/macro level. With an extremely dense acquisition, and a continuous light selection interface coupled with nearest neighbor interpolation, the user can naturally select the




Figure 2: Direct interpolation. For each row, here we show (from left to right): one of the original images in the MLIC; PTM-based rendering (Sec. 5.3); HSH-based rendering (Sec. 5.3); RBF rendering. RBF interpolation is capable of visualizing better novel virtual light directions, compared to classic methods. Courtesy of Pintus et al. [PDJ*18].

original image to display, and interactively simulate controlled illumination. Direct exploration of the original images may be, however, difficult for visualizing massive image collections in remote settings. Moreover, the Graphical User Interface (GUI) for selecting a light is efficient only for regular acquisitions (e.g., domes), which offer a natural 2D parameterization. Further, the lack of continuity may make the perception of details difficult, as it is not possible, in most practical cases, to exploit motion parallax. Compression techniques combined with interpolation [PCS18] allow for simulating more flexible scenarios, including, for instance, moving a raking light, which is one of the classical methods used in real-world surface analysis and inspection. While the simulated continuous light motion may improve perception through motion parallax cues, it may also produce artifacts for cast shadows and highlights, especially if the sampling is not dense. Thus, some authors [GCD*17] have proposed to let users interactively tune the locality of the interpolation, but this approach is far from optimal.

5.2 Single-image data fusion

While direct exploration strives to provide direct access to all the individual images composing the MLIC, single-image fusion approaches have the goal to maximize the information that can be conveyed in a single static image. Often the performed data fusion produces physically inconsistent or impossible images, but nevertheless serves to increase the amount of shape information (less often appearance) presented in a single view of the object.

5.2.1 Approaches

The most straightforward way, and classic way, to extract a single image that conveys object surface information from a MLIC is the albedo visualization. This is a light-independent view of the low-frequency (typically constant), spatially-varying BRDF across




Figure 3: Single-image data fusion. The input data consists in a 3 image MLIC (a). The multi-scale approach of Fattal et al. [FAR07] has been applied and two possible results have been obtained with different parameter configurations (b). The enhanced view on the left locally highlights details and flattens the global appearance of the object, while the result on the right implicitly uses some shadows to increase detail perception. Courtesy of Fattal et al. [FAR07].



(a) A car engine (original vs enhanced).



(b) A flower plant (original vs enhanced)

Figure 4: Single-image data fusion for illustrative visualization. Some MLIC-based fusion methods are capable of enhancing visual information in complex shapes, e.g., mechanical parts (a), and of removing texture data through a digital visual abstraction that still preserves important features across the viewed surface (a)(b). Courtesy of Raskar et al. [RTF*04].

the surface [CDG*18, Woo80, AG15]. Since the albedo removes highlights or other high-frequency effects that might disturb the underlying content in the visualization, this approach is often used for objects where pigment information is important (e.g., think of a painting inspection). However, a large amount of information about small, fine surface details is contained into the high-frequency part of that signal and is completely lost.

Motivated by artistic techniques, which bring-out fine-scale details, while at the same time preserving the global object shape, several methods perform a multi-scale analysis of the MLIC signal at different spatial resolutions and frequencies, and then reconstruct an enhanced image that combines detail information at each scale across all the input images. Fattal et al. [FAR07] decompose each image in a MLIC into a series of detail images plus a base image. The detail images are the differences between two bilaterally filtered versions at consecutive scales, while the base image encodes the low-frequency shading. The final visualization is built by a weighted sum of those contributions. Fig. 3 shows both an input 3 image MLIC and two possible results given different parameters,

which drive the image fusion. One enhanced version makes the local finest details pop up, but flattens the global appearance of the object, while the other output implicitly uses some shadows (no shadow map has been explicitly extracted though) to augment the global perception of the object geometry. Since some high-frequency signals are generated by phenomena that are not related to local surface topology, or, in other terms, to real details or surface features (e.g., globally cast shadows), some algorithms integrate this approach with an additional step of shadow detection and removal [ZLR*10]. Other works exploit shadows cast by multi-light acquisitions in order to classify spatial changes among those caused by spatially-varying materials across the surface, rather than depth discontinuities [RTF*04]. They use that information to build a single image visualization, which superimposes a stylized rendering on a processed version of the original input MLIC that removes unnecessary details while increasing the three-dimensional perception. This results in a final toon-like visualization, which highlights boundaries between the various shapes in the framed scene. Fig. 4 shows how this kind of illustrative techniques has been employed to highlight boundaries of complex shapes (Fig. 4a), and to remove texture data that contains less information through an abstraction strategy that preserves important features across the viewed surface. In this approach, an interactively controllable "degree of abstraction" parameter controls the amount of local details that can be removed (Fig. 4). When MLIC acquisition is performed also in a colored or even multi- or hyper-spectral framework [CSS15], it is possible to exploit not only detail information from the varying light position/direction but also from shading-independent wavelength response across visible and invisible spectra. Once light-position-dependent components have been fused, the multi-spectral signals can be combined to produce false color visualizations, which fuse per-pixel contributions into a single image with highlighted details and main polychrome regions [VHW*18]. In this context, a common strategy to visualize high-dimensional chromatic data is to perform false color PCA visualization. This data-reduction approach computed on a multi-spectral MLIC proved to be very powerful in significantly increasing the legibility of almost-faded details on the surface [VdPHB*16] and subsurface layers (e.g., underdrawings in paintings or manuscripts [WVP15]), which are completely invisible in the original photographs.

5.2.2 Visualization and analysis

The combination, through a data fusion process, of all the MLIC information in a single image has the inherent disadvantage that some information is lost, interactive browsing is no more possible, and motion parallax cannot be exploited. On the other hand, the data-fusion process produces images that convey much more readable information than any single image in the MLIC. Compared to a photorealistic visualization by using direct exploration (Sec. 5.1), single image methods are capable of revealing both tiny surface texture elements and patterns, and global shape features (e.g., depth discontinuities in cluttered/complex environments), which are otherwise poorly visible or completely invisible in the photorealistic version or in the original images. Since they try to emphasize geometrical features and to convey a visualization with a high level of abstraction, these techniques can provide useful insights into the nature of the real material appearance. For these depicting capabilities and their efficiency, these data fusion methods are employed

for several tasks. Examples include visualization of surface texture and cracks in rock [FAR07], fur visualization [ZLR*10, FAR07], as well as venation and texture enhancement in botanic visual inspection [ZLR*10, FAR07, RTF*04]. When multi- or hyper-spectral capture is available, the shading-independent response to several different wavelengths can be used to enhance small poorly visible inscriptions (e.g., incisions on stones) [CSS15, VdPHB*16], to visualize abundance maps of pigments across the acquired object [CSS15, VHW*18, WVP15], or to improve the visibility of faded inks in old written text [VHW*18]; moreover geometrical structures and irregularities under the surface layer (e.g., pictorial layer visualization) can be made visible together with geometrical gradients. Moreover, these methods may produce a meaningful visualization even with a small amount of input data (e.g., 3–4 images). For this reason, they have also been used for the enhancement of images during dynamic capture. Most of the approaches, moreover, do not require camera and light calibration, since they perform the analysis just starting from the intensity changes in the image stack. The fact that they analyze multiple images at once makes them more effective than single-image enhancements or exaggerated shading techniques directly applied to 3D geometry [FAR07]. They are employed to display scenes with low-contrast depth edges, and complex geometries such as mechanical parts (Fig. 4) or medical data (e.g., in endoscopy [RTF*04]). In addition, during the visualization, the user has only to tune a small number of parameters to interactively weight the contribution of various surface behaviour components (e.g., weights of base and detail signals [FAR07], amount of detail enhancement [ZLR*10], or degree of abstraction). Single image approaches are also very suitable for storage and communication purposes in the context of web-based visualization platforms, since they exhibit a really compact and compressed representation, and are very appropriate for printing. The limitation to static visualization may be an interesting area of future work, since multi-image data fusion might be used in the context of real-time exploration methods, e.g., by fusing only subsets of images near the currently, interactively determined, focus light, or by exploiting a globally fused field in combination with other methods to improve shape and detail perception while performing relighting or other interactive tasks.

5.3 Relightable images

While single-image data-fusion mostly focuses on static visualization, offering the possibility of dynamically relighting the models is at the core of relightable model based approaches, which analyze the MLIC in order to learn the relation between measured images and incident illumination, and synthesize it in a relighting model. The first set of methods, *relightable images*, surveyed in this section, strives to directly model the reflectance field, without separating shape and material contribution. The second set of methods, that we will analyze in the next section, performs, instead, that separation. In both cases, the techniques require, in addition to the MLIC, metadata with information on camera and light parameters to reason on the relation between measured reflectance, view direction, and incoming light direction and intensity.

5.3.1 Approaches

The seminal insight behind relightable images, introduced by Malzbender et al. [MGW01] in terms of Polynomial Texture Maps




Figure 5: Layered models. This class of models combine a matte model with a specular/shadow layer to better render complex materials. In this figure we show several layers that are modeled from a MLIC data: (1st row, left to right) surface normals, matte albedo; (2nd row, left to right), specular quotient (the ratio between specular and diffuse intensities), and the specular color. Courtesy of MacDonald [Mac15].

(PTMs), was to replace classical multi-light processing based on Photometric Stereo [Woo80] with a more flexible interpolation-based approach, which transforms the high amount of data in RTI image stacks into a compact multi-layered representation focused on approximating the measured appearance as a function of light direction. The method is based on a per-pixel bi-quadratic polynomial interpolation of the image stack producing a spatially varying reflectance I as a function of the light direction. At each pixel, given the six parameters α_i computed from data fitting, and a light direction \vec{l} , the reflectance is computed as $I(\vec{l}) = \alpha_0 \vec{l}_x^2 + \alpha_1 \vec{l}_y^2 + \alpha_2 \vec{l}_x \vec{l}_y + \alpha_3 \vec{l}_x + \alpha_4 \vec{l}_y + \alpha_5$. This method was the basis of the development of the RTI framework. A drawback of the quadratic model is that it captures only the low-frequency reflectance behaviour, and it is not capable of reproducing high-frequency effects like specularities and shadows. The use of high order polynomials has been also tested [ZD14], where third order coefficients have been used, increasing the quality of the relighted images at the expense of increased oscillations in case of unreliable data. High-frequency effects may also act as outliers in the fitting process, damaging the accuracy of the reconstruction. For this reason, robust fitting methods have been successfully proposed. For instance, Drew et al. [DHOMH12] propose a Least Median of Squares fitting, while Pintus et al. [PGPG17] proposed an accelerated guided robust matte-model fitting method that exploits similarity of nearby pixels to speed-up computation.

Similarly to the Fourier series, spherical harmonics are a good set of basis to describe functions on the surface of a sphere, and low order sets of those bases are typically used in modeling low-frequency reflectance, e.g., applied in photometric stereo scenarios [BJK07]. In typical MLIC acquisition setups, the incident light at any surface point is defined on the upper hemisphere only, and a full spherical

representation is not needed. For this reason, the hemispherical basis defined from the shifted associated Legendre polynomials [ERF11] has been introduced to represent image irradiance [MMC*08]. To create relightable images from a captured dataset, the projection of the images onto the elements of the basis up to a maximal order is estimated. Typically, the maximal order used is 1 (4 coefficients), 2 (9 coefficients) or 3 (16 coefficients). The per-pixel reflectance is computed as $I(\vec{l}) = \sum_{l=0}^{n-1} \sum_{m=-l}^l \alpha_l^m H_l^m(\vec{l})$, where H_l^m are the hemispherical harmonic (HSH) basis functions. Relightable images encoded with HSH coefficients are also currently used in application domains, and are supported by RTI viewers (Sec. 6.2). HSH fitting is generally preferred to PTM since it captures better high-frequency reflectance behaviours. However, the analysis performed by Zhang and Drew [ZD14] shows that, with a similar number of coefficients, polynomial and HSH coefficients provide similar results.

Discrete Modal Decomposition (DMD) has also been proposed to generate a continuous model of the local reflection. The DMD method is based on a projection on the modal bases, which are composed of elementary forms that take their origin from a structural dynamic problem. Pitard et al. [PLGF*15, PLGM*17a] adopt this technique for modeling angular components of the local reflectance by devising new modal shapes called Reflectance Modal Basis (RMB) from the set of images based on the Discrete Modal decomposition. They model the surface reflectance function, f , as the sum of a linear combination of the modal vectors and the residue of the decomposition R_n , which can be expressed as $f(\theta_v, \varphi_u)(\theta_i, \varphi_i) = \sum_{k=1}^n \lambda_k(\theta_v, \varphi_u) Q_k(\theta_i, \varphi_i) + R_n$, where the modal coefficients λ_k can be obtained as the projection of the vector of measured luminance \mathbf{L} onto each modes (Q_k , $k=1, \dots, n$, where n is the number of modes) of the non-orthonormal basis. A comparative analysis with other fitting techniques has shown that the DMD is well suited for approximating the complex physical behavior of light reflections, especially for shiny reflective surfaces [PLGF*15, PLGM*17a].

Simple parametric functions are, however, usually well suited to model the matte reflection and some shininess, but fail to model higher-degree phenomena such as transparency, interreflections, specularities, and shadows. To tackle this problem, a possible way is to use a layered model that composes a matte model and specular/shadow models to obtain the relighted image. Some proposed methods [DHOMH12, ZD14] are based on the assumption that the matte component can be described by the inliers in a robust regression of a PTM or an HSH model. The difference between the original images and the matte interpolation is then interpolated with a RBF, in order to model the high-frequency component of the original signal, and to obtain a more photorealistic rendering. Fornaro et al. [FBKR17] employ a PTM to model the matte signal and compute the distance between that matte function and the intensity of each pixel for all lighting directions. The average of those distances is considered a measure of the pixel glossiness, and used to add a Phong based shading component to the rendering. This method is less general than the previous approaches, since it makes the implicit assumption that the object is a reflective surface. Similarly, in the method proposed by Macdonald [Mac15, Mac18] the matte model is used to extract the albedo and the normals (Fig. 5). Specularities are then obtained in the following way: first, the specular quotients are calculated as the ratio of the actual intensity divided by the Lamber-

tian intensity, then a modified Lorentzian function is fitted on them (Fig. 5). The method, however, has not been used to interactively render images with arbitrary light directions.

In contrast to the methods that recover the relighted image independently per pixel, Xu et al. [XSHR18] have proposed a method to learn the behaviour for entire image patches using specialized Convolutional Neural Networks (CNNs). Two networks are trained, one to select optimal light sampling (that, therefore, is then fixed with a certain tolerance), and the other to relight the scene given a novel light direction. The method can reasonably model global effects (e.g., shadows and interreflections) and complex materials (general isotropic) with an extremely small number of images compared to state-of-the-art techniques (e.g., only 5 images). However, the rendering of global high-frequency behaviours, such as sharp cast shadow boundaries, still remains a challenging issue. That is the reason why that approach is more suitable (and can obtain good error metrics) when environmental light is provided rather than a directional light.

Among the methods that model appearance as a function of light direction domain, without explicitly recovering shape/material descriptions, a few of them do not focus on the light transport function but they try to interpolate the measured signal in the light matrix domain. Specifically, they interpolate a discretized matrix $I(i, j)$ called Light Transport Matrix (LTM), which represents the light reflected at each pixel location as a function of the illumination on a planar surface parallel to the image plane. Thanikachalam et al. [TBF*17] sample the matrix entries by acquiring the object under varying illumination with a mobile phone LED that moves on a plane, and the relighting is performed with a compressive-sensing based reconstruction that has been proved to interpolate better than bicubic interpolation or kernel regression. Similarly, the LTM can be also recovered from a sparse sampling with the use of a Neural Network ensemble, by exploiting spatial coherence [RDL*15]. However, methods interpolating the LTM usually perform relighting based on a larger number of input images (e.g. 200-300) with respect to methods sampling input light in the light direction space (50-100).

5.3.2 Visualization and analysis

This class of relightable models is very general and easy to apply. By avoiding to solve the complex problems of separating shape from material contributions, and of modeling complex light transports, these methods can, in principle, provide visually realistic approximations while interactively controlling a virtual light. Similarly to the local interpolation methods (Sec. 5.1), this class of fitting techniques can be used to simulate interactive lighting or raking light. The user is provided with a simple interface element to select a virtual light source (typically a Gaussian sphere projected in 2D), and the visualization algorithm uses the parameters of the fitted analytic model to transform that light direction input in the proper relighted image; in addition, the user might pan and zoom, and select regions of interest within the relightable image. In addition to simple interpolation, moreover, more complex modifications can be made to the illumination, including spatially-varying intensity and direction control.

Due to the rich viewing experience that this approach provides,




Figure 6: Light Transport Matrix Reconstruction. The first row shows a series of ground truth images of a light transport matrix column which is not included in the input data used to produce the relightable image model. The second row are the same matrix columns reconstructed by using a trained relighting Neural Network. Courtesy of Ren et al. [RDL*15].

this is the most used approach in a wide range of visualization and analysis tasks. Interactive relighting controls allow viewers to better understand object morphology in terms of meso- (e.g., local surface bumps, roughness, superficial patterns) and macro-structures (e.g., global curvature, convexity and concavity cues) by using an interface that mimics the typical real-world inspection (e.g., with a raking light). This fact makes them a good virtual alternative to real-world inspection under controlled light for surface quality control in many fields from CH study and inspection, criminal forensics, art conservation, and paleontology [MMC*08], to micro-mechanics, biomedicine, and horology.

In most cases, however, only the low-frequency behavior is sufficiently accurate for typical acquisition sampling rates (i.e., hundreds of images). A full model in terms of geometry and material is more amenable to produce physically-valid relighting; such methods are only computable for specific classes of objects (Sec. 5.4), but relightable images are more generally applicable, and have been shown to help in understanding how different optical features as matte nature or glossiness, and generally heterogeneous materials, are distributed across the object, as well as to visually highlight surface grain/granularity, damages, weathering patterns, erosion, and general decay.

All the various methods, based on attempting to directly map, through a small set of parameters, the changes in perceived color to change in illumination direction, tend to exhibit visual artifacts due to the global interpolation in the light space. To overcome this issue, LTM-based methods and CNN-based methods strive to learn more information on global transport. In principle, LTM-based methods are capable of producing an extremely accurate relighting, with anisotropic, transparent and highly complex local behaviour of materials and global illumination effects (Fig. 6). Unfortunately, while they meet these goals, LTM approaches pose several problems in terms of capture complexity, and memory occupancy for both storage and communication issues. In terms of capture, to avoid hallucinations, very high-frequency MLIC should be employed. Moreover, the size of typical reconstructed light transport function for each pixel is about 200k samples [TBF*17], so, even for small images, the representation tends to pose very hard problems in terms of storage, transmission, and computation required to meet interactive rates during relighting. The use of neural networks [RDL*15],

trained for each specific dataset, tries to cope with this problem, by producing a relatively small model with tractable memory footprint and fast relighting performance. However, the methods proposed so far tend to produce hallucinations, and are, therefore, not very appropriate to visual inspection tasks.

For these reasons, the most employed methods in practice are still those based on low-frequency fitting (e.g., PTM, HSH). These representations are easy to compute and to compress, eventually exploiting available image-based encoding (e.g. JPEG) to compress the various layers, leading to straightforward incorporation in web tools and mobile devices (Sec. 6.2. In these tools, visualization is often integrated with more complex knowledge organization and representation software, and within systems for the annotation of specific sets of objects, e.g., in the arts and humanities fields. Some of those tools allow the user to link comments and annotate the relightable images during the visualization experience [FBKR17], and export or store this information as metadata. This opens the way for collaborative visualization environments.

5.4 (SV)BRDF fitting and meso-structure visualization

A clear limitation of the reflectance fitting/interpolation approach is the lack of proper decoupling of shape and material information. For this reason, a number of methods try to extract a geometric model from the MLIC and associate it to a material model. This problem is very complex, and feasible only by limiting the class of scenes to which the approaches are applied. Even in the simplest case of opaque surfaces, the key challenge is that the reflectance, characterized in terms of SVBRDF, and the shape, characterized in terms of surface normals, are inherently coupled and need to be estimated jointly. Furthermore, the SVBRDF is a 6D function of space and incident/outgoing angles, and, in the absence of additional assumptions, its estimation requires a large number of input images captured with precisely calibrated, and often prohibitively expensive, setups. Moreover, in the typical MLIC setting, modification of the view direction is not available. The problem of estimating a full representation of the shape and the material from a single viewpoint data is thus heavily ill-posed and requires the introduction of extra knowledge. To this end, a variety of approaches have been proposed.




Figure 7: SVBRDF from MLICs. Comparison between the real photograph (first row) and a virtually rendered surface under novel lighting and viewing conditions (second row). Courtesy of Aittala et al. [AWL13].

5.4.1 Approaches

The classic approach to shape and BRDF recovery is to first extract surface normals by using a photometric stereo approach, and, then, perform a fitting assuming that the geometry is known. The underlying assumption is that the surface and the material are sufficiently simple so that second bounce effects are negligible. Shi et al. [SWM*16] have benchmarked state-of-the-art photometric stereo solutions, showing that reasonable accuracy can be obtained for non-Lambertian objects when using calibrated data. Once the normal field is known, the most straightforward approach is to fit the data to an analytical model by using a non-linear optimizer that minimizes the fitting error [NDM05]. The fitting, however, is heavily underconstrained, since only a very small portion of the 4D space and incident/outgoing angles are measured at each point.

A common assumption for enabling computationally tractable models for SVBRDF is that the BRDF at each pixel is a weighted combination of a few, unknown reference BRDFs [LBAD*06]. Based on this concept, Hui and Sankaranarayanan [HS17] present a method that recovers normals and BRDF parameters for complex materials from MLICs. The method is based on the assumption that the BRDF at each pixel lies in the non-negative span of a known BRDF dictionary. First, the normals are estimated using a variant of example-based photometric stereo. Given the surface normals, the SVBRDF is obtained by constraining the BRDF at each pixel to be in the span of the dictionary while leveraging additional priors to regularize the solution.

By imposing strong limitations on the shape and/or material distribution, the recovery of the shape and the material can also be obtained with a limited number of images. A common assumption is that objects are (mostly) planar. For this case, Gardner et al. [GTHD03] exploit a linear light source instead of a point or directional lights to create the MLIC data. They have developed a method capable of finding parameters of a Ward BRDF model (diffuse and specular reflectance, and roughness) for planar surfaces by building a reflectance table to determine which parameters most closely fit the captured data. The method is adapted to capture surfaces with perturbed normals, allowing for the rendering of images of the surfaces under novel lights. Aittala et al. [AWL13] recover SVBRDF parameters for near-planar surfaces; they use ad-hoc illumination patterns for the MLIC generation by creating point samples in the 2D Fourier space. When applied to surfaces in which repeated pat-

terns of the same materials and varying normals are known to be present, these methods are capable to obtain the SVBRDF from just two images [AWL15].

Fig. 7 shows a comparison between original photographs and virtual renderings under novel lighting and viewing conditions for the method of Aittala et al. [AWL13]. It proves how photorealistic visualization can be provided by these kinds of approaches if the assumptions on the imaged scene are met in practice. With this approach, moreover, the recovery of the per-pixel BRDF parameters allows for generating novel views in addition to novel illumination, even if the result is not satisfactory for all the materials.

The polarization of light can also be used to recover reflectance parameters for relighting. The creation of MLICs by using spherical gradient illumination patterns combined with linear or circular polarization [MHP*07] allows for the separation of spatially varying diffuse and specular optical behaviour across the surface, and it is used also to extract normals. Rendering is then performed with a hybrid normal mapping that combines the diffuse and specular shading. While Riviere et al. [RPG16] present a method to practically fit BRDF parameters from a MLIC capture performed with a fixed viewpoint and gradient lights, Ghosh et al. [GCP*09] extend the method of Ma et al. [MHP*07] by using polarizers and gradient lights to recover view-independent BRDF parameters. A more general isotropic BRDF extraction has been presented by Ghosh et al. [GCP*10], which recover per-pixel estimations of the diffuse and specular albedo, the index of refraction, and the specular surface roughness. It is based on the estimation of Stokes parameters of the reflected light based on a three image MLIC, acquired with differently oriented linear polarizers positioned in front of the camera, and one additional photograph captured with a circular polarizer. For the computation of normals, this method exploits a classic photometric stereo approach [Woo80]. Given the known object shape, the reflectance properties can be obtained from small MLICs (or even a single image) by representing the specular reflection as a mixture of probability distributions on the unit sphere, and then by using the Expectation Maximization (EM) algorithm to estimate parameters [HNI08]. Another technique to acquire and visualize accurate spatially-varying reflectance properties of an object with complex anisotropic materials is based on continuous spherical harmonic lighting conditions [TFG*13].

5.4.2 Visualization and analysis

With respect to all previous methods, the recovery of shape and material in terms of per-point normal and BRDF has major limitations in terms of managed materials and the computational complexity (involving, often, the introduction of several additional constraints to make the shape and material decoupling possible, and to cope with insufficient data). Moreover, joint computation of BRDF and normal fields requires highly precise acquisition setup calibration, and the resulting model is very sensitive to errors in camera, light and radiometric parameters. On the other hand, in the case where the recovery is possible, the BRDF and normals representation allows for the visualization of geometric topography at different scales, i.e., microfacets distribution, mesoscale structures, and global/macro geometry/curvature, and the extraction of a normal vector field can also lead to the computation of the depth (see also Sec. 5.6). Moreover, the representation is usable within realistic renderers, making it possible to generate faithful visual representations of the optical response of the surface to light and viewing conditions. The full shape/BRDF estimation not only allows for interactive re-lighting, but supports other different visualization modes up to rendering under environmental maps and even novel viewpoints. Due to the fact that BRDFs are extremely useful in a wide range of applications (e.g., CH, industry, automotive, medical imaging, etc.), tools to visualize, explore, and edit those BRDFs (and other related optical functions) are continuously evolving [BS12,BBP14,FPBP09]; those tools can handle analytic representations, data-driven BRDFs, and anisotropic materials. The user can both visualize the object with the corresponding BRDF properties, and the parameters that control the material appearance; by interactively tuning those parameters, and by visualizing the corresponding 2D or 3D graphs in real-time, those interfaces help in understanding and analyzing different surface optical behaviours (e.g., diffuse or specular albedo, index of refraction, specular roughness), and to see how they respond to several illumination conditions. Moreover, they serve as a workbench to compare expected BRDFs and sample measurements. Apart from virtual reality renderings, BRDF capture can serve to calibrate 3D printers, and to pre-process 3D models for rapid-prototyping, so that the visualization will happen in the physical real-world with photo-realistic replicas of the acquired objects. In this context, BRDFs are used to decrease the perceptual color error between digital copies and real objects created by multi-material 3D printers [TU16].

5.5 Non-photorealistic enhancements

Relightable image models are often combined with illustrative techniques, in order to improve the understanding of scenes beyond what is possible by sticking to photorealistic approaches. Differently from static methods, which extract single images that convey an enhanced view of the object (already seen in Sec. 5.2) or feature maps showing shape or material properties (that we will see in Sec. 5.6), the methods described here focus on interactive-rate solutions that modify the employed representations, or that post-process the generated images.

5.5.1 Approaches

The most popular strategies to improve surface detail readability in real-time relighting applications are the *Diffuse Gain* [UW13,




Figure 8: Non-photorealistic enhancement. Exaggerated shading helps in the visualization of the finest object details. Here, for instance, directions of smoothing marks across the surface have been highlighted by this non-photorealistic way of display MLIC-derived data. Courtesy of Van et al. [VdPHB*16].

CM11,SSM*17,EBB*11,HBMG02] and the *Specular enhancement* approaches [EBB*11,CM11,AIK13]. The *Diffuse Gain* has been presented for the first time by Malzbender et al. [MGW01]. The main idea is to apply a transformation to the fitting parameters of the matte model function (in their case, the six PTM parameters). By non-linearly changing those luminance coefficients, they are capable of increasing the visual perception of the curvature of the surface response to light direction variations. That transformation is designed in order to maintain the same per-pixel normal field. This reflectance-based contrast improvement, together with the dynamically moving light position, proved to be an efficient and fast way for aided morphology inspection and the visualization of surface subtleties not discernible to the naked eye [MVSLO5,UW13,SSM*17]. Having an underlying normal vector field, all types of materials can be visualized by adding a specular component to their analytical or data-driven BRDF function. This is an old and widely used technique, due to its simplicity, efficiency and effectiveness in enhancing meso and micro structures across the object surface [MMSLO6,MVSLO5], to better discern object content [MCPC17,SSM*17,UW13]; figure 8 is an example of how those exaggerated shading can reveal not only shape content but also semantic information about how the object has been manufactured [VdPHB*16]. The most simple specular model consists in a combination of a specular color k_s and a non-linear (typically exponential) function driven by a roughness parameter α . During the visualization, the user can easily tune those parameters in real-time [MWGA06] while dynamically changing the light source to even extreme conditions, up to raking light [KK13] (see also light extrapolation explained below), in order to obtain the desired effect and to highlight different shape properties and features [HBMG02]. The *Specular Enhancement* can be used together with the color information blended with it [WVM*05], or in a visualization that removes the albedo or generally the chromatic information, by displaying only the reflectivity component, and by focusing only on the amplification of slope perception [CSS15,UW13].

Besides those two ways of enhancing relightable image visualizations, which are of paramount importance for visual details inspection, other techniques are less used, although equally effective for the task. One consists in applying unsharp masking to the luminance field [CM11,SSM*17]. Others exploit a light that is not constrained

to have a unit norm, and goes beyond that physically acceptable condition, and perform light extrapolation [HBMG02, PSM05]. This non-photorealistic representation provides a more "oblique", raking light, and a higher image contrast. Of course, by enabling virtual relighting, one can also increase the number of virtual lights. This can be used to increase object readability by automatically choosing the appropriate set of lights that augments details perception. Palma et al. [PCC*10] present a multi-lighting detail enhancement that, given a single light direction selected by the user, computes a virtual lighting environment that aims at increasing the local contrast across the image. This is a non-photorealistic method because that virtual lighting is not feasible and reproducible in the real world due to the high number of lights, the very localized contribution (or support volume) of each light (which is not physically plausible), and the real-world interreflections which would ruin the original enhancement purpose by smoothing out the desired sharpening effect.

In order to enrich the quality of the visualization enhancement, the aforementioned interactive techniques are typically combined with the real-time normal enhancement mode [MWGA06] to increase granularity and sharpness in the final visualization. Other approaches consist in directly applying the same rationale of unsharp masking [PCC*10], or a simple contrast enhancement operator [GCD*17], to the normal vector field. While in Sec. 5.6 normal maps are exploited to extract features for further analysis, here they are used within the visualization pipeline as a layer to produce the effect of a slope and shading exaggeration [WVM*05, VHW*18].

5.5.2 Visualization and analysis

With the same purpose of the methods seen in Sec. 5.2, non-photorealistic enhancement is aimed at producing a visualization that increases details perception. Similarly to single-image data-fusion methods, many of the approaches are specially devoted to highlight microfeatures (e.g., roughness) and mesostructures/patterns across the object surface, and mostly focus on shape rather than material enhancement. Unlike single-image data-fusion methods, however, most of the illustrative techniques are interactive and tightly coupled with both a dynamic relighting, and, in some cases (e.g., *Diffuse Gain*, *Specular enhancement*, and *Unsharp Masking*) even with real-time MLIC acquisition and enhancement [MWGA06]. Direct interactive control of parameters is generally used to tune various contributions of different appearance layers. Since most of the methods can be applied to the data stored in classic relightable image formats, the same advantages in terms of portability, compression and communication described in Sec. 5.3 and Sec. 5.4 are valid here as well.

Many of the non-photorealistic techniques focus more on the shape than on the appearance. For this reason, many of them rely on the computation of normals from MLIC in order to convey shape information, which, as discussed for BRDF extraction, is not applicable to all imaged scenes. When this is possible, those non-photorealistic enhancements have proved to be successful in highlighting hardly visible features (e.g., poorly visible inscriptions [WVM*05]). In some cases, to highlight geometrical details, the appearance is even changed on purpose, e.g., when applying specular enhancement on completely matte objects to emphasize surface finish, so optical material interpretation is not applicable

and misleading. Further, some visualization techniques consider the color or appearance information as a cluttering element, and they remove that to bring the wanted details to light [MCPC17]. This is true also in the cases where the appearance/color information is not removed but rather exploited; the multi-spectral signal is used to produce false-color images that, rather than containing information about the optical behaviour of the surface, make it visible underlying shapes and features [VHW*18]. Since the scope of non-photorealistic methods is the enhanced visualization of semantic structures and of details across the surface, visualization is not so dependent on an accurate light/camera calibration, since, generally, small scale calibration errors do not change much the overall information content of the imaged scene. Unlike the methods in Sec. 5.3, illustrative techniques are thus more robust to acquisition noise and inaccurate, relaxed computation of light, camera, and radiometric parameters. Due to the combination of a compact data format, the user-friendly interface, the robustness to noise or poor calibration, and the effectiveness of those illustrative techniques, this group of methods are the most used ones in a large amount of daily applications (Sec. 6.1).

5.6 Feature maps and derived maps

All the methods presented so far, whether photorealistic or illustrative, aim at visualizing the original object with all its chromatic and geometrical information, and, eventually, with some enhancement or some exaggerated representation of specific details or regions across the surface. Conversely, other techniques (visually) discard all the input data, and are more focused on the visualization of just the results of some particular MLIC processing, in one or multiple layers. More precisely, those pipelines have basically two steps, i.e., the extraction of a particular set of features or per-pixel attributes (≥ 1), and the presentation of them in a set of MLIC derived maps.

5.6.1 Approaches

By exploiting the high-correlation between appearance changes due to moving light direction and depth discontinuities in the observed scene, a group of works directly visualize a per-pixel attribute that maps edges and contours on the object, without computing any geometrical representation of the object shape. With this rationale, Raskar et al. [RTF*04] are able to produce edge-rendering, sketches and toon-based visualizations of a scene captured with just four lights, by exploiting depth discontinuities retrieved without the estimation of an explicit normal map (as Photometric Stereo methods do). They combine an ambient signal and the four photographs lit by different light sources by using several differences and ratio images. Other techniques exploit the variation of the parameters of a particular MLIC fitting model; for instance, Pan [Pan16] visualizes surface edges by computing them directly from Polynomial Texture Map coefficients (Sec. 5.3); the probability of having an edge is proportional to the largest variation of one of its six parameters in terms of direction and magnitude. The amount of edges or discontinuities in the original MLIC signal is also a cue for spotting details or important regions with significant features. Some papers use this idea to extract saliency maps out of MLICs [PLGM*17b]. Pitard et al. [PLGM*17c] use DMD (Sec. 5.3) to extract saliency maps that highlight the local changes of reflectance; they develop a method which can split the modes into rotation-dependent and




Figure 9: Sketch-like extraction. We present here two images: a color representation of cones with inscriptions, and a sketch-like visualization obtained by applying a line drawing algorithm to the captured MLIC. Courtesy of Hameeuw [Ham15]

rotation invariant ones, so that their robust computation of surface saliency exhibits an invariance to object rotation.


Unlike the previous “implicit” methods, another group of works produces visual representations of feature maps, such as edges, by explicitly relying on normal map computation from MLIC data. Brognara et al. [BCDG13] take the normals computed from PTM coefficients and extract surface discontinuities by the sum of the squared directional derivatives of the normal vector field. The final edge, or sketch-like, visualization is obtained by simply applying a Canny filter to the discontinuity signal. They apply this process at different scales and merge them into a single representation, in order to smooth out the final visualization and to make it more meaningful, clean, and more artifact/outlier free. Similarly, Earl et al. [EBB*11] directly apply a Sobel filter on the normal map to produce a per-pixel edge map visualization of the object under study. In addition, very accurate and extremely readable sketch-like views can be extracted by feeding automated line drawing algorithms with the normal map computed by MLICs [Ham15]. Fig. 9 presents the result of a typical and very useful transformation from the MLIC to a more effective sketch-like view of the object under study; in this case, it is very clear how, by removing chromatic information, low-frequency luminance signals, and unimportant details, all the inscriptions on ancient cones are presented in a more abstract way, and ultimately are more readable than in the original photo. Of course, with the normal maps at our disposal, it stands to reason that other measurements can be derived, whose values represent better edges and surface features. Normal unsharp masking or gradient maps, as well as the computation of curvature from normal vector fields, are able to display exaggerated surface features [VVP*18]; they can reveal fine structures across the object that only non-photorealistic renderings can bring to light, such as curvature coloring, slope exaggeration or a combination of them [WVM*05].

Edges, sketch maps, and the curvature-based saliency reveal a more geometric nature of the object. A large number of papers completely focus their attention to geometry and strive to compute a visualization that contains a pure geometric-based signal, by removing as much as possible the information about shading, color, and general material BRDFs. Normal computation from MLICs (a.k.a. Photometric Stereo) is a well-known topic with an overwhelming amount of literature [AG15]; besides the computational

aspect, a big bunch of works uses this signal for visualization purposes among a really heterogeneous set of applications (Sec. 6.1). In some cases the original [BCDG13, DGL*14, EBB*11, CDG*18] or the enhanced [GCD*17] normal map display can be an additional option for object inspection. In other cases, this kind of visualization is necessary due to extreme conditions, such as visual clutter due to other appearance components. For instance, highly spatially-varying and noisy albedo over a smoothed surface geometry (e.g., due to material aging) makes the aforementioned sharpening algorithms almost useless. Only by separating color and normals one can produce a visualization that is capable of revealing the underlying scene content [MCPC17]. Compared to the literature on Photometric Stereo, which aims at exactly computing normal maps or shape geometry from MLICs, normal map based visualization can relax the requirement of an accurate normal vector field, given that the visualization is still suitable to understand the geometric content of the surface. When the image signal has a strong bias due to particular acquisition noise and environmental conditions (e.g., underwater acquisition with high turbidity), normal maps, although not precisely reconstructed, can provide better and cleaner visualization of the finest surface details [SSM*17]. Similarly, in a multi-spectral MLIC framework, different wavelengths produce slightly different normal maps of the same object, due to the fact that each light frequency reacts differently at different sub-surface layers of the object [TML*13]. Forgetting the fact that, from a Bayesian approach, the definition of a “true” normal per pixel is a meaningless concept, and it makes sense only from a probabilistic point of view, those sets of “pseudo-normals” (per pixel and per wavelength) can be employed to visualize and render a lot of information about both surface behaviour at a mesoscale (nor globally, neither locally), and local, per-pixel scale, tiny details across the object. For instance, since ultra violet (UV) light does not penetrate the surface, it can be used to visualize very sharp normal maps from the outer object layer; conversely, infrared MLIC produces a normal-based visual signal which renders the low-frequency component of the object shape [VdPHB*16, VHW*18]. Rather than providing normal map visualization from different wavelengths, some papers extract and visualize normals from a single optical component of the surface; more specifically, some do that from the diffuse nature of the object [BCDG13, DGL*14, EBB*11, CDG*18], while others only from the specular component of the reflectance [MHP*07]. Whatever is the source of the normal vector field, several visualizations can be done by enhancing and sharpening this signal. Willems et al. [WVM*05] apply a slope exaggeration operator to the normal vector field, by increasing the angle between the z-axis (i.e., the view direction) and each per-pixel normal. Since this operation is very fast, it can also be employed in real-time surface visualization and inspection [MWGA06]. The local standard deviation of the normal orientation has been also considered a valuable way to visualize morphological anomalies in the case of expected low-frequency geometry [MBW*14]. Normal map processing can also produce several types of derived visualizations, ranging from surface roughness to tangent and bi-tangent vectors [GCP*09]; by integrating normals one can also display 2D images of depth maps (or height fields) [DGL*14, TLQ06]. If the MLICs are acquired by using polarization principles, it is possible to visualize different normals from, for instance, circularly polarized and unpolarized illumination, as



(a) Visible light, infrared, and ultraviolet photos of a painting



(b) Lambertian Outlier map (top) and Outlier Direction map (bottom)

Figure 10: Derived maps. Given a multi-spectral acquisition, consisting in three (visible light, infrared, and ultraviolet) MLICs (a), it is possible to extract derived maps that highlight different responses of the object surfaces (b). Here we present the Lambertian Outlier maps of the visible and infrared signal, and the Outlier Direction maps for the same channels. Courtesy of Giachetti et al. [GCD*17].

well as to produce the visualization of other feature maps such as circularly polarized and unpolarized stokes parameters [GPDG12].

Different signal and interesting visual cues emerge as a side product of classic fitting or MLIC processing. During the computation of matte component from MLICs a lot of computation is devoted to spotting deviations from the imposed model or simply data outliers. These are not binary measures (e.g., outlier/not outlier), but, obviously, they represent a probabilistic estimation of labeling confidence. Besides helping in the computation framework, those values used as visualization matter proved to be a good way to reveal the spatially varying behaviour of the surface, e.g., in terms of diffuse and specular component distribution [CDG*18]. Giachetti et al. [GCD*17] present two kinds of maps derived from those numerical pipelines. The first map, called the *Lambertian Outlier* map, can be obtained by summing the squared differences between the measured reflectance and the result of matte fitting (e.g., by using Photometric Stereo or PTM); this difference is a visual representation of the high-frequency component of the surface optical material response. After the matte component fitting, another map can be computed by counting for each pixel the number of reflectance values that deviate from the fitted curve; this multiplicity map of deviations from the model has been called *Outlier Direction* map. Fig. 10 shows those derived maps computed from the visible and infrared signals of a multi-spectral five channel MLIC (the five channels are the infrared, the three visible, and the ultraviolet signal). It is clear how they highlight different types of surface behaviours that are poorly visible or invisible at a naked eye. Look, for instance, at the *Outlier Direction* map of the infrared channel, which makes a goldish region of the knight's hat pop up; this particular behaviour is really hard to spot in common photographs of that painting. Finally, other approaches visualize several maps obtained by combining

relightings, e.g., the difference between the same relighting of two acquisitions of the same object made at different times [PSM05], or computing Di Zenzo gradients [Pan16, GCD*17]. Fig. 11 presents a comparative view of different signals; one is the *Static Multi-Light* enhancement presented in Sec. 5.5, and the other are the Di Zenzo gradient map [Pan16, GCD*17], a contrast-enhanced normal map, and the *Outlier Direction* map. Arrows in the latter prove how those types of derived data are capable of make relief patterns across the object surface more visible and readable.

5.6.2 Visualization and analysis

Many sorts of feature maps have been devised to provide insights on the shape and appearance of the imaged scene. Most methods are designed to be robust to errors, non-idealities in the acquisition pipeline, and inaccurate calibration of light, camera and radiometric parameters. Unlike most single-image methods (Sec. 5.2) or illustrative techniques (Sec. 5.5), many feature maps bear little relation with the original photorealistic view, but directly identify and highlight the regions of interest and their information content.

For shape understanding, feature maps are used in support of visualization and analysis at all scales. Approaches include both micro- and meso-details (sub-pixel geometric structures [TLQ06], normal/gradient maps or discontinuity maps [Pan16, BCDG13]), and more global shape information (e.g., visualization and analysis of structures such as leaf venations [ZHS*18] or inscription [Ham15] supported by depth maps obtained from the integration of normal and/or shadow maps. When dealing with appearance, the main use is in support of material characterization and segmentation. For instance, the *Lambertian Outlier* or the *Outlier Direction* maps can visually reveal the distribution of optical behaviour across the surface (e.g., material discontinuities [RTF*04]).

The capability of pre-computed or interactively computed feature maps to improve the understanding of scenes makes them a component of many tools, which use them either as direct display, or as a component of illustrative techniques for a variety of specialized applications (Sec. 6.1). Moreover, the extracted feature maps are not solely used for direct visualization, but are also components of more elaborate processing and visualization pipelines. For instance, supervised or unsupervised machine learning systems for shape-based or material classification [BCDG13, WGS09] often take as input feature maps in replacement or in addition to the original image data.

6 Applications and tools

MLIC techniques have been mixed and matched to solve several problems in a variety of application domains, and they have been also integrated into user-friendly software tools aimed at supporting both visualization and analysis.

6.1 Applications

The visual analysis of MLICs is at the basis of several applications and tasks, ranging from pure visualization to geometric analysis [PPY*16]. We provide here a structured overview of the areas where MLICs have been successfully used for daily visualization and analysis work. Table 2 presents a schematic view of the distribution of the techniques presented in Sec. 5 with respect to the main application domains they are employed in.




Figure 11: Combination of Features and Relighting. The extraction of features can also enhance interactive relighting by superimposing those two components in one dynamic visualization. Here we show how those techniques are used to better display brush strokes in a detail of a painting. We compare (a) Static Multi-Light enhancement (Sec. 5.5) with (b) a di Zenzo gradient map, (c) contrast-enhanced normal map, and (d) an Outlier Direction map. Arrows in the latter show how this type of derived map is capable of make relief patterns more visible. Courtesy of Giachetti et al. [GCD*17].

6.1.1 Cultural Heritage

CH appears, by far, to be the application domain where MLICs are more popular. In particular, this is due to both the versatility of the capturing approach (capable of working at a large number of scales), and the large quantity of information that can be gathered about object surfaces by using relatively low-cost equipment. Moreover, using moving lights fits well with traditional CH analysis means, and it is thus well received by CH experts. Nowadays MLIC based acquisition is considered a well-assessed digitization method in CH, and specific surveys and tutorials have been dedicated to them [MMC*08, VVP*18].

Since their introduction, relightable images based on MLICs (Sec. 5.3) have been used for the analysis of CH objects like tablets, statues [MGW01, GTHD03, MMSL06, MVSL05, SOSI03, DHOMH12], paintings [PSM05, PCS18], coins [PBFS14, PCS18], and manuscripts [WVP15]. The interactive relighting is particularly appreciated by CH experts as it also enables the simulation of the typical raking light surface inspection [PRA15]. Another common practice in the CH community is related to the use of illustrative enhancements (Sec. 5.5), such as *Diffuse gain*, *Specular enhancement*, and *Unsharp masking*, to intensify the detail perception on tablets [WVM*05, PCC*10], petroglyphs or written stones [UW13, MCPC17], mural graffiti [CSS15], coins [KK13], or, generally, metallic alloys [Mac18]. In this context, a major research need is the enhancement of legibility and readability of written texts or engraved symbols [EBB*11], and the processing of MLICs can lead to the extraction of meaningful lines and inscriptions on tablets [BCDG13, Pan16]. Further, an advanced application to automatically recover line drawings of inscriptions has been demonstrated on a collection of Mesopotamian clay cones [Ham15].

The combination of the surface penetration given by the infrared light together with the non-photorealistic exaggeration facilitated by the multi-directional light images has proven effective for recovering invisible texts of the Dead Sea Scrolls [CM11], where the deterioration of the parchment caused its blending with the black iron-gall ink. Similarly, the engravings on a very small golden lamina [CDG*18] benefited from the single visualization of the albedo (Sec. 5.2) corresponding to the infrared and ultraviolet signals, where the contrast between the engraved symbols and the background is higher than in the visible case. Furthermore, by using RBF interpolation (Sec. 5.1), the raking light analysis is enabled, and certain engravings around the margins of the lamina gain more visibility.

Another application of MLICs coveted by the CH end-users is the visualization of the dynamic interplay between light and the complex surfaces of artworks characterized by heterogeneous optical behaviours. In a recent work [FBKR17], the authors manage to trace back the original visual impression of mosaic tesserae and early prints by unmixing the glossy details from the diffuse component as a result of fitting the light transport function formulated as a polynomial (Sec. 5.3). In this manner, they obtain a layered representation of the artwork's reflectance. Moreover, in order to allow the end-users to customize the digital surrogate according to their own artistic view, the gloss layer is enhanced in a non-photorealistic fashion (Sec. 5.5) and generated synthetically according to a Phong reflection model. Through the user input, parameters such as the degree of shininess, diffusiveness, and specularities can be interactively modified in order to render different degrees of glossiness. The separation of the specular component has been proposed as well by Macdonald [Mac18] as a way to improve the quality of CH object relighting. The importance of the specular reconstruction for the visualization of relighted paintings' surface is demonstrated also by Thanikachalam et al. [TBF*17], where the method based on the reconstruction of light transport matrix with a compressive sensing approach (Sec. 5.3) preserves the visual richness of paintings, and allows for the visualization of how the specular behaviour migrates over the painting's surface. Pitard et al. [PLGM*17a] show that a better interactive relighting model could not only improve the quality of the specular component on a jewel visualization but also allows for a better understanding of texture relief on a wall painting.

The analysis of the rich data in MLICs has given promising results in revealing the manufacturing techniques of a craftsman or the painting style of an artist. This application is especially relevant for reclaiming information of unknown or unstudied heritage, as well as very old or lost-and-found historical objects. Newman [New15], based on the specular enhancement visualizations (Sec. 5.5), distinguishes between three types of negative profiles (incisions and striations) in the creation and finishing of archaeological bone artifacts from El Zotz assemblage. At the same time, Artal and Klausmeyer [AIK13] propose a method that uses non-photorealistic enhancements and feature maps in the case of Greek red-figure vases (Sec. 5.5 and Sec. 5.6) to guide the end-user in distinguishing between various types of decoration lines with positive profiles, as well as their order of application at the fabrication moment. Similarly, others spot different orientation patterns in the clay mold-

	Cultural Heritage	Natural Sciences	Industry	Underwater	Medical Imaging
Direct exploration (Sec. 5.1)	Interpretation/Monitoring [GCD*17]; Dissemination/Inspection [GCD*17, PCS18, SOSI03].				
Single-image data fusion (Sec. 5.2)	Enhanced perception [CSS15, CDG*18]; Showing the invisible [CDG*18, VdPHB*16, VHW*18, WVP15]; Pigments [CSS15, VdPHB*16, VHW*18, WVP15].	Enhanced perception [FAR07, RTF*04, ZLR*10]; Automatic illustration [FAR07, ZLR*10, RTF*04]; Dissemination/Inspection [FAR07].	Enhanced perception [RTF*04].		Endoscopy [RTF*04]; Shape and reflectance analysis [SSS*08, SLDS14].
Relightable images (Sec. 5.3)	Surface defects [Mac15]; Material characterization [Mac18]; Interpretation/Monitoring [GLSWE18, New15, MMC*08]; Dissemination/Inspection [GLSWE18, PBFS14, Mac15, FBKR17, TBF*17, DHOMH12, MGW01, MMC*08, PLGM*17a].	Enhanced perception [MMC*08, MGW01]; Species identification/Interpretation [HP15]; Dissemination/Inspection [MMC*08, MGW01].	Surface Quality Inspection [SS00, PLGF*15, PLGM*17a].		Shape and reflectance analysis [SSS*08, SLDS14].
(SV)BRDF fitting (Sec. 5.4)	Material characterization [GTHD03]; Dissemination/Inspection [GTHD03, MHP*07].		3D Printing - Real-world visualization [TU16]; Surface design and inspection [RMS*08].		
Non-photorealistic enhancements (Sec. 5.5)	Enhanced perception [CM11, EBB*11, CSS15, UW13, MCPC17, MGW01, WVM*05, PCC*10, PSM05, MMSL06]; Showing the invisible [CM11, MCPC17]; Pigments [GCD*17, CSS15]; Surface defects [PSM05]; Interpretation/Monitoring [VdPHB*16, VHW*18, KK13, EBB*11, New15, AIK13, CM11]; Dissemination/Inspection [VHW*18, MGW01, AIK13, PCC*10, KK13, PSM05, MMSL06, UW13, EBB*11, MVSL05].	Enhanced perception [MGW01, HBMG02, MWGA06]; Automatic illustration [HBMG02]; Species identification/Interpretation [HP15]; Dissemination/Inspection [MGW01, HBMG02, MWGA06].		Enhanced perception [SSM*17].	
Feature and derived maps (Sec. 5.6)	Enhanced perception [Pan16, VVP*18, Ham15, BCDG13, CDG*18, MCPC17, EBB*11, WVM*05, PSM05]; Showing the invisible [MCPC17]; Surface defects [PSM05, MBW*14]; Material characterization [GCD*17]; Interpretation/Monitoring [VdPHB*16, New15, MBW*14, EBB*11, BCDG13]; Dissemination/Inspection [VHW*18, PSM05, EBB*11, MHP*07, Ham15].	Enhanced perception [Pan16, RTF*04, MWGA06]; Automatic illustration [RTF*04, Pan16, ZHS*18]; Species identification/Interpretation [ZHS*18, WLW17]; Dissemination/Inspection [MWGA06, Pan16].	Enhanced perception [RTF*04]; Scanning Electron Microscope Data Visualization [PPV08, PS05]; Surface Quality Inspection [PLGM*17c, PLGM*17b].	Enhanced perception [SSM*17, FQJ*17].	Enhanced perception - Endoscopy [DGL*14, PLG*13, RTF*04] - Laparoscopy [CB12]; Shape analysis [LCG12].

Table 2: Techniques vs Major Application Domains. For each class of MLIC analysis and visualization methods, the major use cases for each application domain are presented.

ing process of Execration Egyptian statuettes made of unbaked clay [VdPHB*16]. Analogously, in painted items, it is possible to render the relief and orientations of the brushstrokes related to a style or an artist. In addition, many methods aim at detecting particular pigments through direct local MLIC interpolation in the light space (Sec.5.1), by computing feature maps (e.g., *Outlier Direction* map, Sec. 5.6) [GCD*17], by mixing the multispectral data into a single false color visualization [VHW*18, VVP*18] (Sec. 5.2), or by directly analyzing histograms of the per-pixel reflection maps (Sec. 5.1). In this context, MLICs have also been used to recover lost pigments and to revive faded colorings by using non-photorealistic renderings [VdPHB*16]. Finally, fibers in historical textiles with different materials and patterns of decay can be visually detected by using magnified multi-light images under UV radiation [GLSWE18] (Sec. 5.2).

MLICs can be maneuvered as a measurement tool in assessing the conservation condition of CH materials. For instance, single image visualization and non-photorealistic enhancements based on MLICs of coins [KK13] indicate the need for material cleaning, and show whether a conservation treatment was applied successfully or not. Similarly, the incorporation of profile measurements on an exaggerated shaded view of manuscripts has shown to provide useful visual annotation with respect to the analyzed documents and their medium [VHW*18]. In addition to single image visualizations and dramatic enhancements, by deriving feature maps (Sec. 5.6) that target crack formation, Manfredi et al. [MBW*14] propose a quantitative approach to monitor damages in paintings. With a resembling underlying procedure, MacDonald [Mac15] measured the extent of cracks in a fresco at the National Gallery of London.

6.1.2 Natural Sciences

MLICs are a well-recognized imaging technique for natural sciences as well [MMC*08]. In biology, MLIC-based relightable image models (Sec. 5.3) are used to support animal species identification [HP15], while MLIC-based image fusion (Sec. 5.2) is used for detail visualization of leaves or flowers [FAR07, RTF*04, ZLR*10]. Non-photorealistic visualization methods combined with feature extraction (Sec. 5.6) have also been successfully used for the visual analysis of biological surfaces [WLW17] and leaf venations [ZHS*18]. In paleontology, MLIC visualization methods have mainly two functionality: detail visualization and whitening. The low visibility of images of fossils is hampering visual analysis. Mostly, this factor comes due to low color contrast and a corroded relief. To solve this problem, the classic RTI framework is often used to illustrate subtle features from fossils, assisting in discovering hidden details [HBMG02]. MLIC data are used not only to create interactively relightable images of the fossils [MGW01, MWGA06] (Sec. 5.3), but also to simulate fossil whitening, a task performed to get information regarding their morphology. Normally, this task is done manually using various techniques such as airbrush or chemical agent (ammonium chloride). Hammer and Spocova [HBMG02] show that this can be done digitally using MLIC-based photorealistic visualization, with a competitive performance with respect to traditional more invasive techniques. Enhancements like edge detection have been used as well to increase the perception of fossil structures [Pan16].

6.1.3 Industry

In industrial manufacturing, the inspection of an object surface is often performed to control whether the final product meets predefined tolerances. In recent years, major efforts have been made to replace, wherever possible, manual or physical testing with virtual visual inspection. MLIC visualization is emerging as a promising solution since it allows for a reliable acquisition of materials, their high-quality visualization, and an exchange of that information for industrial monitoring and design [XR19]. Enhanced single image visualization based on MLICs can be used to better understand the structure of mechanical parts [RTF*04]. Pitard et al. [PLGM*17c, PLGM*17b] have used a MLIC photorealistic visualization technique for the detection and the analysis of visual anomalies on challenging metallic surfaces. This technique is derived from DMD [PLGF*15] (Sec. 5.3), and automates the visual inspection by introducing a rotation-invariant representation of the reflectance characteristics of the inspected surface. Smith and Stamp [SS00] also use MLIC based photorealistic visualization technique for quality inspection of textured ceramic tiles, while Rump et al. [RMS*08] employed MLICs visualization for metallic car paint inspection and design. In the context of quality control and investigation of electronic reliability, MLICs captured with the Scanning Electron Microscope (SEM) [PS05] are used to produce meaningful visualization of micro- and nano-details on the surface of electronic circuits; Pintus et al. [PPV08] demonstrate an application to the visualization of damaged solid-state electron devices.

6.1.4 Underwater inspection

An image-based surface analysis in the underwater environment is a difficult task, due to factors such as water clarity/turbidity,

excessive algae, and other participating media. This highly affects the light phenomena, and usually results in dark, low contrast, and bluish tone-image. A specialized approach of MLICs applied to that type of environment has been proposed by Selmo et al. [SSM*17]. Authors use RTI to study underwater CH objects, e.g. details from historic shipwrecks. They also perform a study on the effects of turbidity on relighting quality. Another issue in underwater MLIC processing is related to light refraction effects. Fan et al. [FQJ*17] propose a flat refractive model to cope with this phenomenon by computing the mapping of virtual and real points; they are capable of removing distortions in underwater acquisitions both for 3D reconstruction and MLIC visualization.

6.1.5 Medical imaging

The application of multi-light visualization in the area of medical science strives to increase readability of medical images. It is mostly used in the type of examination where the internal organ such as the esophagus, stomach, and upper part of the small intestine is visually inspected, often in real-time during examinations [PLG*13, CB12]. This is typically achieved by modifying existing instruments (e.g., endoscopes [DGL*14]) to perform multi-light acquisitions, and by fusing several images to help in visualization. The typical approach is to reconstruct normals and remove highlights in the visualization [PLG*13]. By using a simple MLIC capture with few light sources and single image enhanced visualization it is possible to assist doctors in the understanding of the 3D structure of anatomical parts in endoscopic views [RTF*04]. MLIC analysis is also used for 3D reconstruction of different organs. For instance, Lv et al. [LCG12] used Photometric Stereo to retrieve the shape of a tongue, which is used for accurate diagnosis. To support visualization, the reconstructed organ is mapped with the surface albedo. MLIC-based visualization is also used in dermatology, which often uses interactive visual inspection to support diagnostic tasks. In this area, a photorealistic visualization with fitted relightable models (Sec. 5.3) has been proposed for pigment lesion analysis [SSS*08], or to find biomarkers for early skin cancer detection [SLDS14]. Sun et al. [SSS*08] have shown that MLIC-based non-photorealistic enhancement (Sec. 5.5) is able to generate an efficiently encoded and reasonably complete representation of skin appearance, which assists in analysis and diagnosis of pigmented skin lesions.

6.2 Interactive visual inspection tools

Several interactive MLIC-based visualization tools have been described in the literature, and a few of them have been made publicly available and are actively maintained. Some of them are generically applied to datasets captured with different cameras and lighting setups, while others are used to visualize only data acquired by a specific acquisition device (e.g., a light dome) and/or pre-computed versions of the relighted image [VHW*18, Mac15]. While some of them allow for visualization on the web or on portable devices [PCS18, PBFS14], others are developed considering only desktop applications with fully resident data. A summary of the major available tools is shown in Table 3.

The most popular application is *RTI viewer* (Fig. 12a). It has been primarily developed by the ISTI/CNR Visual Computing Laboratory, and financed and distributed now by Cultural Heritage Imaging [CHI19]. Complete credits and executable codes can be found at

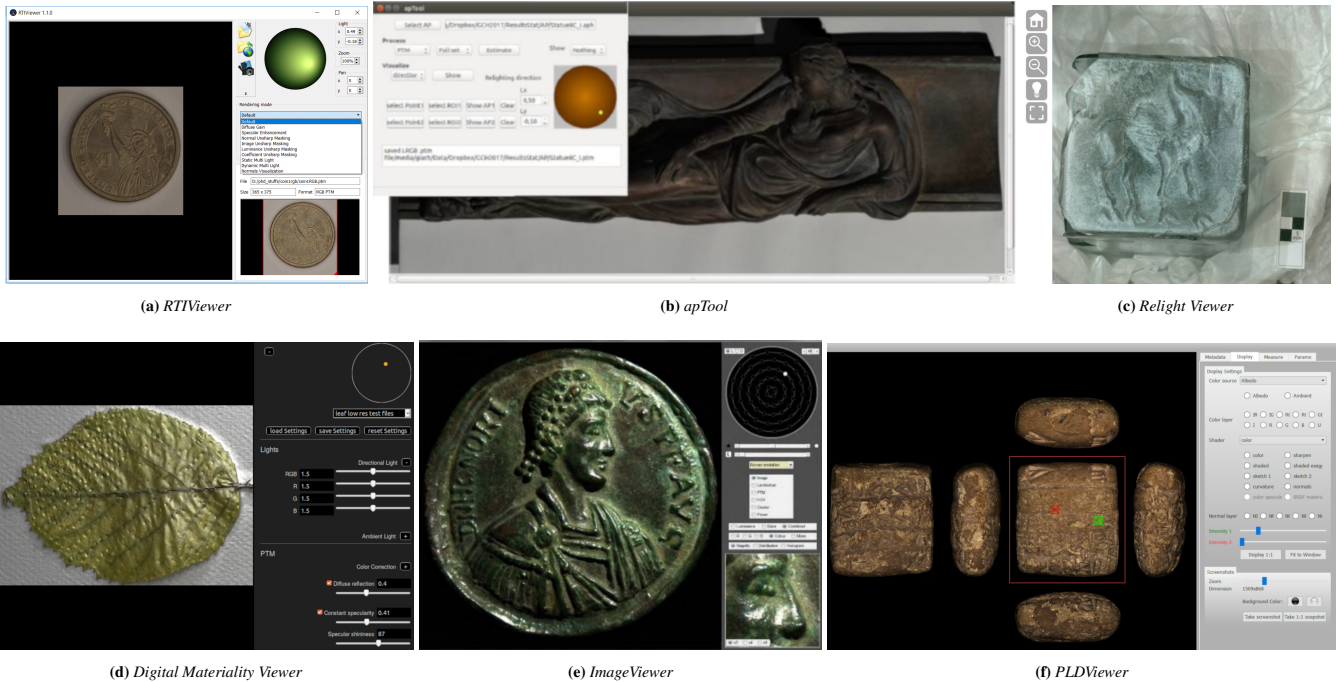


Figure 12: Screenshots of visualization tools: (a) RTIViewer [CHI19], the most popular solution for relightable image visualization; (b) apTool, a research tool developed by University of Verona [G*19]; (c) Relight Viewer, a tool part of a powerful library developed at ISTI/CNR Pisa [P*19b]; (d) Digital Materiality Viewer application developed at the Digital Humanity Labs (University of Basel) [DHL19]; (e) GUI of desktop visualization tools described by MacDonald [Mac15] (this tool does not allow interactive relighting); (f) PLDViewer GUI, which handles dense acquisition realized with the Portable Light dome [KUL19].

the CHI web site. This tool allows for the visualization of relightable image files stored on the local file system and remote files through HTTP. It features a visible virtual torch, which allows the user to perform photorealistic image relighting based on PTM or HSH (Sec. 5.3), and includes several non-photorealistic enhancements (Sec. 5.5), including specular enhancement and image unsharp masking. Due to its free availability and ease of use, together with the CHI [CHI19] capture software, it is widely utilized in different areas of study presented in Sec. 6.1. Another desktop application tool, but mainly used in image processing research works, is *apTool* [G*19] (Fig. 12b). This tool allows for interactive exploration of raw MLIC data using RBF interpolation; it also creates relightable images by using PTM or HSH, and extracts and visualizes feature maps. It can also interactively show per-pixel interpolated reflectance maps. This tool has been used in CH for the study of various objects such as coins, rock art, statues, and paintings [GCD*17, CDG*18]. A multi-functionality application tool called *PLDviewer* (Portable Light Dome viewer) has been developed at KU Leuven [KUL19], based on a full processing framework [VHW*18]. This tool accepts as input proprietary compressed file generated with a specific dome device and an undisclosed processing pipeline. It provides interactive relighting (photorealistic visualization), 3D model construction, and non-photorealistic visualization. The first can be performed with the help of two virtual light sources with tunable intensity (Fig. 12f). The viewer supports several rendering modes, e.g. albedo and ambient, and allows for multi-spectral visualizations in five spectral bands (i.e., Infrared, Red, Green, Blue, and UltraViolet). Six views of an object are shown in the interface for full 3D objects creation, and the user can interactively visualize one at a time by switching between

different views using arrow-keys. In addition, this tool allows for performing non-photorealistic enhancement (Sec. 5.5) such as curvature coloring, exaggerated shading, and sketching. Furthermore, it allows for displaying the map of the response to the set of lights on single pixels (reflection maps) and local histograms. This tool is mainly developed for CH applications but can be extended to other application domains. The desktop visualization tool presented by MacDonald [Mac15] (*ImageViewer*, Fig. 12e) is capable of displaying original photographic images and different kinds of modified renderings with the same lighting. Unlike RTIViewer, *apTool*, and *PLDViewer* this tool does not support interactive relighting. Several tools have been proposed also for web-based visualization of relightable images. *WebRTIViewer* [P*19a] is a web-based multi-light visualization tool that allows for the visualization of the relightable images created using PTM or HSH in a standard web page, without the installation of any plugin and with a recent web-browser. *PLDWebviewer* is the web-based counterpart of the *PLDViewer* with similar functionality. Other web-based tools have been developed by the University of Basel [DHL19]. In particular, the *Digital Materiality Viewer* (Fig. 12d) allows for interactive relighting and glossy enhancement, as described by Fornaro et al. [FBKR17]. *Relight* [P*19b] is a library to create and view (on the web) relightable images (RTI), which supports the encodings described by Ponchio et al. [PCS18]. The library can be used to create relightable images based on PTM, HSH, and RBF, and exploits different compression and color coding options. The visualization methods feature both an OpenGL Javascript library for RTI rendering on the web, which can be used to create web pages with single or multiple relightable

Tool	Access	Visualization	Platform
RTIViewer [CHI19]	Public	Photorealistic relighting (PTM, HSH); Specular enhancement; Diffuse gain; Normal, Image & coeff unsharp mask	Desktop
APTool [G*19]	Private	Photorealistic relighting (PTM, HSH, RBF); Normal/Albedo Maps; Appearance profile vis.	Desktop
PLDViewer [KUL19]	Private	Photorealistic relighting; (Exaggerated)Shading; Curvature; Sketch	Desktop
ImageViewer [Mac15]	Private	Original images; Outline, albedo and normal map; Specular components; Low-frequency fitting (PTM, HSH)	Desktop
WebRTIViewer [P*19a]	Public	Photorealistic relighting (PTM, HSH);	Web
Relight-viewer [P*19b]	Public	Photorealistic relighting (PTM, HSH, RBF, Bilinear)	Web
PLDWebViewer [KUL19]	Public	Photorealistic relighting; (Exaggerated) Shading; Curvature; Sketch	Web
DigitalMateriality Viewer [DHL19]	Public	Photorealistic relighting; Specular enhancements	Web

Table 3: Tools. Major MLIC-based visualization software tools described in the literature.

images, and the Relight Viewer, a multiresolution viewer with a toolbar controlling a few basic options (Fig. 12c).

7 Discussion

Our survey demonstrates that a vast amount of research results can be exploited for visually displaying and analyzing the information encoded in multiple images of a surface captured from a single point of view and changing the illumination. These solutions have been used standalone, or mixed and matched in various ways, to the benefit of scholars and practitioners in a wide range of practical applications. Moreover, due to the availability of low-cost protocols or automatic systems to capture MLICs at many scales, the applicability of these methods is constantly growing. The analysis of the literature outlined several interesting points for discussion and revealed relevant challenges that need to be faced in future research in order to improve the effectiveness of MLIC analysis and visualization.

7.1 Analysis of current state-of-the-art

In the following assessment of the current state-of-the-art, we focus on the main issues that we find important for using MLICs in visualization and analysis settings.

7.1.1 Light space sampling, data storage, and compression

MLIC data are practically multidimensional vectors storing per-pixel information sampling the space of a light transmission encoding. The amount of information captured on the analyzed surface by a MLIC depends clearly on the number of images stored, e.g. on

the sampling density in the light parameters' space. With a dense sampling, it is possible to have a good understanding of the surface properties just by looking at the original images, or to have a simple interactive relighting by interpolating pixel values in the light parameter space (Sec. 5.1). However, this requires to store all the original information in memory to enable interactive surface inspection, with limitations on interactivity as well as on usability in non-local situations (e.g., web viewers). The use of data compression (e.g., PCA [PCS18]) can reduce the memory requirements for data storage and exchange, at the expense, however, of some loss in quality. Relightable images consisting of interpolation coefficients can be efficiently stored and rendered (Sec. 5.3), but, since they rely on low-frequency models, they are not able to reproduce complex material behaviors. While several solutions have been presented to extend them to higher-frequency models, these are far from perfect, and incur a considerable storage overhead. BRDF models, coupled with mesostructure information, could provide the best trade-off between encoding size and relight quality, but are difficult to be fitted from a sparse sampling of complex materials (Sec. 5.4).

The sampling of the data is critical. While it is possible to obtain edge information [FAR07], or even a photorealistic relighting [XSHR18], with very few images, the recovery of reliable BRDF parameters or non-hallucinated details requires a quite dense sampling. Other applications, like normal map reconstruction, may be realized with a limited number of optimally sampled directions, and the addition of further samples does not necessarily improve the quality of results in a relevant manner [BZS18].

Some methods to derive relighting or material descriptors rely on the assumption of predefined and fixed light directions or light bases. This may guarantee optimal results for both material segmentation and relighting. Algorithms can learn sets of optimal predefined lights to maximize the extracted information. However, in many cases, the lighting cannot be accurately controlled, but only measured a posteriori with calibration procedures. More flexible algorithms can deal with different sampling, and also with the inaccuracy of simplified light models (like directional and point lights), to produce relightable images independent of original light direction sampling and light beam shape [GDR*15]. One of the most interesting research directions in MLIC processing techniques is the investigation of optimal resampling of the data in light space. The idea of resampling the lights in a regular pre-defined grid has been recently proposed for compression purposes [PCS18], but the proposed approach, based on RBF interpolation, is effective only if the original sampling is sufficiently dense.

7.1.2 Handling color channels and multispectral data

MLIC capture can result in a set of monochromatic images, color images, or even multispectral images, depending on the light sources and the camera employed. Most of the techniques for visualization and analysis (Sec. 5) are described for single-channel data.

While the most common approach is to simply apply them to color or multispectral data by handling each channel separately, a popular approach is to separate luminance and chromaticity in order to reduce memory and increase speed. In the RTI framework, relighting is often performed in Luminance-RGB (LRGB) mode, by multiplying the relighted luminance by an average/median per-pixel

chromaticity (Sec. 6.2). This simplification, reasonable for matte surfaces, is typically wrong for glossy objects, where the hue typically changes with lighting direction. The LRGB approximation should thus be used with particular care. When using BRDF extraction, a common simplification is to use only a Lambertian contribution plus one specular lobe with a shape identical across the color channels, so that only two chromaticities need to be specified [NDM05]. This dichromatic approach is closer to represent a larger variety of reflectance behaviors.

The separation of chromaticity from luminance is also used in compression methods, e.g., performing PCA for dimensionality reduction [PCS18], even if this approach generates artifacts for objects with varying reflectance properties. For compression, the separation is often done in the $YCbCr$ color space, which reduces cross-correlations so as to increase compression while improving chroma fidelity [SZV*14]. As an alternative, other approaches have also proposed to perform computation in perceptual color spaces to improve compression or provide fittings that most closely match human perception [DPF01].

Moreover, some MLIC acquisition frameworks capture images in different spectral bands [RM07, GCD*17, WHV*16]. A useful application of multispectral imaging is in photometric stereo setups, since the extra channels lead to improved separation of the shape and albedo contribution for polychromatic surfaces [TML*13, NK14]. In terms of visualization and analysis, several approaches have been presented to merge channels into false-color images, and to analyze individual spectral response maps [WHV*16]. However, so far multispectral visualizations do not appear to have been fully exploited together with relighting.

7.1.3 Relighting accuracy and evaluation

The quality of novel images of the captured surfaces relighted with the methods described in Sec. 5.3 can be evaluated by measuring the similarity with corresponding ground truth images (captured with the same similar light and not included in the data used to create the relighting). The similarity can be assessed with standard image quality measuring metrics such as Peak Signal to Noise Ratio image (PSNR), and/or Structural Similarity Index (SSIM). These approaches are used in many papers dealing with relightable images [DHOMH12, ZD14, JA11, GCD*17, PCS18].

According to those metrics, on most images local interpolation (e.g., RBF) provides better results than HSH and PTM [PDJ*18]. Looking at several works [ZD14, PLGM*17a], it seems that, using a similar number of coefficients, PTM/HSH and DMD provide similar results. The results reported by Xu et al. [XSHR18] show that the CNN based relighting presented may provide higher PSNR than PTM and BRDF based methods [HS17].

However, simple comparison metrics cannot account for the many different aspects of the relighting quality. Algorithms can be evaluated for their ability in reproducing the low-frequency reflectance behavior, in replicating the high-frequency specular effects, and in correctly representing non local effects (e.g., shadows and inter-reflections). For the matte behavior, low order parametric functions (e.g. PTM) seem sufficient to obtain good image quality, but, in the case of glossy or metallic surfaces, it is evident that the global fitting

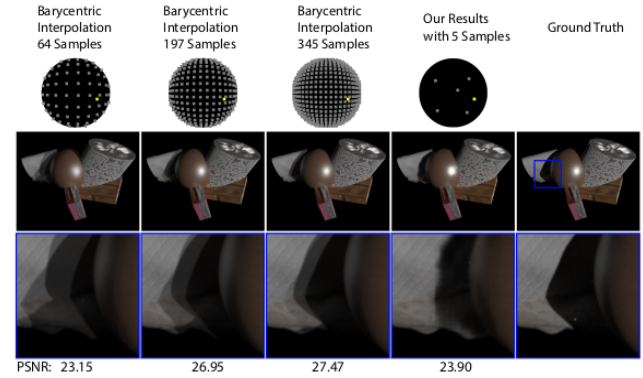


Figure 13: Global CNN-based relighting. The blended shadows artifacts obtained in interpolation-based relighting (left). The method proposed by Xu et al. [XSHR18] avoids it by using a global CNN based relighting, but the shadow shape is clearly different from the real one. Courtesy of Xu et al. [XSHR18].

with low-frequency functions cannot provide the correct appearance. For this reason, if the sampling is sufficiently dense, a local interpolation provides results that are closer to the target image, even if the location of the specular reflections is not accurate. All the algorithms interpolating in the light direction space are "local" and cannot reproduce correctly shadows and interreflections. This results in the typical "blended" shadows, less evident in the local interpolation of dense sampling, but still clearly visible. Global CNN-based relighting methods [XSHR18] solve this problem, but create sharp shadows with hallucinated shape, not corresponding to the real one (Fig. 13).

Global illumination effects, correct glossiness and even relighting from different viewpoints could be handled by reliable methods that recover both the shape and the BRDF from MLIC data only. However, at the moment, such methods seem still not available for complex objects.

7.1.4 Photorealism, accuracy, and task specific evaluation

Given these considerations, the quality of the "photorealistic" relighting should be evaluated in a very task-specific way. First, it should be stated if the goal of the relighting is to create just a photorealistic image or an accurate simulation of a real image of the same object under novel light (that is the case of many applications shown in Sec. 6.1). In the first case, we could rely also on methods that potentially create hallucinated details, otherwise, they should be avoided. Furthermore, we should consider the object materials and application task. In Pintus et al. [PDJ*18] subjective tests are performed in addition to the objective ones, by engaging users of different background to evaluate both the similarity to target images and the perceived quality. It turned out that the evaluation criteria may depend on users, tasks, and materials. For example, even if CH experts consider the RBF based relighting to better simulate target image of glossy objects, when the relighted images are presented without a ground truth reference, a matte-like PTM rendering is considered of higher quality. For less specular materials this was no longer true. Non-expert people, however, gave an opposite evaluation.

	Arbitrary light	Sampling	Relight efficiency	Compact enc.	Photorealism	Accuracy	Arbitrary view
Original images	no	fixed, dense	-	no	total	perfect	no
PTM [MGW01]	yes	free, sparse	good	yes	+	+	no
HSH [MMSL06]	yes	free, sparse	good	yes	+	++	no
DTD [PLGF*15]	yes	free, sparse	good	yes	+	++	no
RBF [GCD*18]	yes	free, sparse	bad	no	++	+++	no
RBF+PCA [PCS18]	yes	free, sparse	good	yes	++	++	no
Resampling+Bilinear [PCS18]	yes	free, sparse	good	yes	++	++	no
RelightNET [XSHR18]	yes	fixed, 5 images	good	yes	+++	++	no
BRDF Dictionary [HS17]	yes	free, sparse	good	yes	+	+	yes

Table 4: Comparison of photorealistic visualization approaches working on MLIC data.

7.1.5 Comparative evaluation of relighting methods

Considering the critical properties described above, we summarize in Table 4 the main features of the relighting methods that have been described in Sec. 5. The evaluation of the different properties is just indicative: photorealism and accuracy may vary by changing the order of PTM or HSH, or by using different color/luminance compression. Photorealism here is referred to the natural aspect of the image, looking correct for both matte aspect, gloss and shadow reproduction quality. Accuracy is related to the ability to reproduce original image details. We penalize the global CNN based approach due to detail hallucination. According to the evaluation of Xu *et al.* [XSHR18], we considered not very good, at the current state of work, both photorealism and accuracy for the BRDF based methods, mainly because of the difficulty of their practical application to common settings (multi-material object with complex shapes and limited numbers of available views). We expect that, at least for selected object and material classes, BRDF based method providing photorealistic and accurate results will emerge in the near future.

7.2 Challenges for future research

Our comparative survey of methods, applications, and tools, and our analysis of the current state-of-the-art makes it possible to highlight promising and challenging research areas for future work. We summarize the main identified ones in the following sections.

7.2.1 Standardization of protocols, data exchange formats, and visualization approaches

One relevant challenge for the MLIC data visualization side is the development of standardized acquisition, processing, encoding, and visualization methods for this kind of data. Most image capture setups share a light-direction based parameterization, even if with different density and constraints. However, there are not well-defined standards to encode the calibrated data stack. Moreover, in all the intermediate steps of the MLIC processing pipeline (e.g. light calibration, data storage, reflectance models, and parameters encoding) there is a clear lack of guidelines and widely accepted exchange formats. Also, for the relightable image visualization, even if the requirements for most practical applications seem to be quite similar (users need to be able to relight the surfaces with novel illumination, using false colors to represent attributes over the surface, visualize material signatures clicking on specific points), there are not standard container formats for the data that allow for interoperability of different solutions. The RTI framework [CHI19] can be considered a de facto standard for some applications in CH due to the fact that it is based on free and known software, but many more recent and

performing solutions for relightable images and enhancements are not integrated and not compatible. Moreover, creators of light dome applications typically develop proprietary solutions for encoding and rendering data. It would be therefore useful in the near future to spend some efforts for the development of standardized protocols and data formats. This would also make it possible to experiment with novel visualization and user interface techniques.

7.2.2 Improving direct visualization

The improvement in capture techniques is making dense acquisition more practical. Direct visualization techniques (Sec. 5.1) have shown their applicability in several domains, but are currently underdeveloped. The need to access at visualization time large amounts of data has, in particular, restricted their applicability, especially in remote settings, where bandwidth is limited. Relighting based on low-frequency fitting (Sec. 5.3) is by far the most commonly employed visualization, but significant inaccuracies have been shown to exist [PDJ*18]. Early attempts at using resampling and compression methods to support direct interpolation are promising [PCS18], but this research is in early stage and only supports aggressive lossy compression. It is reasonable to expect that the area of direct visualization can be significantly improved, tackling at least two different directions. On one hand, interpolation of nearby images can be improved by exploiting more prior knowledge instead of using plain smoothness constraints. Moreover, multiresolution structures and compression techniques can be exploited, as done, e.g., in massive volumetric rendering [BGI*14], to allow for the real-time browsing of full MLICs on bandwidth- and storage-limited settings.

7.2.3 Smart use of machine learning techniques for relighting and BRDF estimation

Deep learning is a powerful learning method inspired by how the brain works. (Convolutional) Neural Network based methods have recently replaced traditional methods in a large part of Computer Vision applications. For MLIC processing aimed at recovering shape (Photometric Stereo), CNN based methods have recently demonstrated the possibility of recovering normals better than traditional methods [Ike18]. It is expected that these methods could be used also to recover other intrinsic properties of the material for visualization and also for interactive relighting, as shown in Sec. 5.3, Sec. 5.3 and Sec. 5.4. The use of deep learning approaches could solve also current issues encountered by methods estimating BRDF from sparse MLIC acquisitions [HS17]. Recent approaches proposed for SVBRDF acquisition use training data and models designed to incorporate physical insights for material estimation, using specialized loss functions [LSC18]. If normals and BRDF coefficients could

produce a synthetic rendering comparable with the interpolation based one, this would be clearly preferred due to its compact encoding and flexibility; in this case, multiview rendering could be supported as well, and standard 3D visualizers could be used instead of the specific 2.5D viewer. A key factor for the development of novel machine learning solutions will consist in the creation of large datasets of real and simulated MLIC dataset to train the algorithms. Accurate and fast rendering techniques will be required to quickly generate images, but also, for example, to compute loss functions in CNN training based on a similarity between rendered and reference images.

7.2.4 Improved interaction and data fusion for illustrative and perceptually-motivated presentation

While a variety of data fusion and illustrative techniques have been presented, this area seems under-developed with respect to other visualization areas, such as, e.g., volumetric data exploration, where these techniques have shown to significantly improve data understanding. Several data-fusion techniques have been introduced, e.g., in multispectral/hyperspectral-only data-fusion and visualization [YGC17] that could be explored also to enhance the perception of details in MLICs. Moreover, besides few examples (Sec. 5.5 and Sec. 5.6), most of the used techniques perform just overlays of feature maps and/or contrast enhancement using manipulations of some shape and appearance parameters. User interaction, in these cases, remains very limited.

In the area of visualization, however, many promising solutions have been explored, including a vast array of perceptually-motivated illustrative rendering approaches [PBC*16], and several interactive tools connected to illustrative methods, such as interactive lenses [TGK*17]. We foresee as a promising research direction the integration and further development of these techniques into the MLIC visualization and analysis pipeline.

7.2.5 Availability of specific benchmarks and datasets

Another important goal for future research could be the validation of the visualization tools for specific user tasks. As previously shown, the usefulness of the different visualization approaches is not easily evaluated with a simple quantitative comparison of relighted images or normal maps with a ground truth reference by using common similarity metrics. From a visualization point of view, we should be able to evaluate the possibility of performing different tasks, e.g., searching for specific information, or evaluating details without distortion or artifacts. The task should be performed on different materials in accordance with the end user applications. The availability of specific benchmarks for different surface analysis tasks will be fundamental to validate methods and develop guidelines for the use of MLIC visualization software in different domains.

7.3 Guidelines for end user applications

Considering the surveyed literature, we can also derive some guidelines for people exploiting MLIC based visualization in different applied sciences. In particular, we find important to take into account the following issues:

- **Maximize data quality and perform accurate calibration.** It is true that interactive relighting can be done without the need of having a very accurate calibration of light directions and intensity; in many of the applications presented in Sec. 6.1, little or no calibration was performed. However, it has been shown that the quality of the relighting and the accuracy of all the other estimated parameters (e.g., normals) is heavily dependent on the correct light characterization. Robust methods to estimate normals should also be applied to remove outliers, increasing the quality of normals and albedo estimation [DHOMH12]. All the methods performing enhancements or relighting based on normals and BRDF critically depend on the quality of the estimation of the model parameters, on the accuracy of the calibration, and on the robustness of the fitting procedure.
- **Adapt the methods to task requirements.** The choice of the encoding and relighting techniques, as well as parameters' tuning and the customization of algorithm options, should be carefully done according to the specific task. For instance, the number of coefficients in PTM/HSR/DMD can make a difference in conveying the proper rendering in the case of highlights or other high-frequency or global effects. The choice of a compression strategy can be also controlled by finding a trade-off between data size and information preservation. Choosing the color space can affect the final result a lot; for instance, LRGB color encoding is not suitable for colored metallic objects as the hue changes in the highlight region. Moreover, in the medical domain, specific protocols could be defined for the different material under study.
- **Check the relight quality.** A good evaluation practice for people that apply a MLIC visualization pipeline in a particular context is to check the quality obtained by each available method in a leave-one-out testing setup, e.g., like that used in Pintus et al. [PDJ*18]. This could be useful to optimize the parameters for the specific materials, and does not require the availability of ground truth parameters (normals, shape).

8 Conclusion

We have provided a structured survey on the use of MLICs as a mean to gain insight into scenes and objects through visual means. The sheer amount of research and applied work surveyed demonstrates the very high interest in the topic, as well as the many connections between methodological and application areas. We expect that this survey will provide a fresh view of the subject, helping both researchers and practitioners to navigate through the ever-increasing literature. Our analysis has highlighted, moreover, that, even though the domain of MLIC capture and processing is mature and has a long history, open problems remain, especially connected to the analysis and visualization area.

Acknowledgments. We acknowledge the contribution of Sardinian Regional Authorities (project VIGELAB). This work has been partially supported by the DSURF (PRIN 2015) project funded by the Italian Ministry of University and Research and the project MIUR Excellence Departments 2018-2022.

References

- [AFG13] ACKERMANN J., FUHRMANN S., GOESELE M.: Geometric point light source calibration. In *Proc. VMV* (2013), pp. 161–168. 3
- [AG15] ACKERMANN J., GOESELE M.: A survey of photometric stereo techniques. *Foundations and Trends in Computer Graphics and Vision* 9, 3-4 (2015), 149–254. 1, 2, 4, 6, 13

- [AIK13] ARTAL-ISBRAND P., KLAUSMEYER P.: Evaluation of the relief line and the contour line on greek red-figure vases using reflectance transformation imaging and three-dimensional laser scanning confocal microscopy. *Studies in Conservation* 58, 4 (2013), 338–359. 2, 4, 11, 15, 16
- [ASSS14] AHMAD J., SUN J., SMITH L., SMITH M.: An improved photometric stereo through distance estimation and light vector optimization from diffused maxima region. *Pattern Recognition Letters* 50 (2014), 15–22. 3
- [AWL13] AITTALA M., WEYRICH T., LEHTINEN J.: Practical svbrdf capture in the frequency domain. *ACM TOG* 32, 4 (2013), 110:1–110:12. 4, 10
- [AWL15] AITTALA M., WEYRICH T., LEHTINEN J.: Two-shot svbrdf capture for stationary materials. *ACM TOG* 34, 4 (2015), 110:1–110:13. 4, 10
- [BBP14] BELCOUR L., BARLA P., PACANOWSKI R.: ALTA: a BRDF analysis library. In *EGSR Workshop on Material Appearance Modeling* (2014), Eurographics. 11
- [BCDG13] BROGNARA C., CORSINI M., DELLEPIANE M., GIACHETTI A.: Edge detection on polynomial texture maps. In *International Conference on Image Analysis and Processing* (2013), Springer, pp. 482–491. 4, 13, 14, 15, 16
- [BGI*14] BALSÁ RODRIGUEZ M., GOBBETTI E., IGLESIAS GUITIÁN J., MAKHINYA M., MARTON F., PAJAROLA R., SUTER S.: State-of-the-art in compressed gpu-based direct volume rendering. *Computer Graphics Forum* 33, 6 (September 2014), 77–100. 21
- [BJK07] BASRI R., JACOBS D., KEMELMACHER I.: Photometric stereo with general, unknown lighting. *International Journal of Computer Vision* 72, 3 (2007), 239–257. 4, 7
- [BS12] BURLEY B., STUDIOS W. D. A.: Physically-based shading at disney. In *Proc. SIGGRAPH* (2012), vol. 2012, pp. 1–7. 11
- [Buh03] BUHMANN M. D.: *Radial basis functions: theory and implementations*, vol. 12. Cambridge university press, 2003. 4, 5
- [BZS18] BRENNER S., ZAMBANINI S., SABLATNIG R.: An investigation of optimal light source setups for photometric stereo reconstruction of historical coins. In *Proc. GCH* (2018). 19
- [CB12] COLLINS T., BARTOLI A.: 3D reconstruction in laparoscopy with close-range photometric stereo. In *Proc. International Conference on Medical Image Computing and Computer-Assisted Intervention* (2012), Springer, pp. 634–642. 4, 16, 17
- [CCC08] CORSINI M., CALLIERI M., CIGNONI P.: Stereo light probe. In *Computer Graphics Forum* (2008), vol. 27, pp. 291–300. 3
- [CDG*18] CIORTAN I., DULECHA T., GIACHETTI A., PINTUS R., JASPE-VILLANUEVA A., GOBBETTI E.: Artworks in the spotlight: characterization with a multispectral led dome. In *IOP Conference Series: Materials Science and Engineering* (2018), vol. 364, IOP Publishing, p. 012025. 4, 6, 13, 14, 15, 16, 18
- [CGS06] CHEN T., GOESELE M., SEIDEL H.-P.: Mesostructure from specularity. In *Proc. CVPR* (2006), vol. 2, pp. 1825–1832. 2
- [CHI19] CHI: Cultural heritage imaging website, 2019. [Online; accessed-March-2019]. URL: <http://culturalheritageimaging.org>. 2, 3, 17, 18, 19, 21
- [CM11] CAINE M., MAGEN M.: Pixels and parchment: The application of RTI and infrared imaging to the Dead Sea scrolls. In *EVA* (2011). 4, 11, 15, 16
- [CSS15] COSENTINO A., STOUT S., SCANDURRA C.: Innovative imaging techniques for examination and documentation of mural paintings and historical graffiti in the catacombs of San Giovanni, Syracuse. *International Journal of Conservation Science* 6, 1 (2015). 4, 6, 7, 11, 15, 16
- [DCCS06] DELLEPIANE M., CORSINI M., CALLIERI M., SCOPIGNO R.: High quality PTM acquisition: Reflection transformation imaging for large objects. In *Proc. VAST* (2006), pp. 179–186. 2
- [DGL*14] DURR N. J., GONZÁLEZ G., LIM D., TRAVERSO G., NISHIOKA N. S., VAKOC B. J., PAROT V.: System for clinical photometric stereo endoscopy. In *Advanced Biomedical and Clinical Diagnostic Systems XII* (2014), vol. 8935, International Society for Optics and Photonics, p. 89351F. 4, 13, 16, 17
- [DHL19] DHLAB: RTI tools at DHLAB Basel, 2019. [Online; accessed-March-2019]. URL: <https://github.com/dhlab-basel/rti>. js. 18, 19
- [DHOMH12] DREW M. S., HEL-OR Y., MALZBENDER T., HAJARI N.: Robust estimation of surface properties and interpolation of shadow/specularity components. *Image and Vision Computing* 30, 4–5 (2012), 317–331. 4, 7, 8, 15, 16, 20, 22
- [DPF01] DUMONT R., PELLACINI F., FERWERDA J. A.: A perceptually-based texture caching algorithm for hardware-based rendering. In *Proc. Rendering Techniques* (2001), pp. 249–256. 20
- [DRS10] DORSEY J., RUSHMEIER H., SILLION F.: *Digital modeling of material appearance*. Elsevier, 2010. 1, 2
- [EBB*11] EARL G., BASFORD P., BISCHOFF A., BOWMAN A., CROWTHER C., DAHL J., HODGSON M., ISAKSEN L., KOTOULA E., MARTINEZ K., PAGI H., PIQUETTE K. E.: Reflectance transformation imaging systems for ancient documentary artefacts. In *Proc. International Conference on Electronic Visualisation and the Arts* (Swindon, UK, 2011), BCS Learning & Development Ltd., pp. 147–154. 2, 4, 11, 13, 15, 16
- [ERF11] ELHABIAN S. Y., RARA H., FARAG A. A.: Towards accurate and efficient representation of image irradiance of convex-Lambertian objects under unknown near lighting. In *Proc. ICCV* (2011), IEEE, pp. 1732–1737. 8
- [FAR07] FATTAL R., AGRAWALA M., RUSINKIEWICZ S.: Multiscale shape and detail enhancement from multi-light image collections. *ACM TOG* 26, 3 (2007), 51:1–51:9. 1, 4, 6, 7, 16, 17, 19
- [FBKR17] FORNARO P., BIANCO A., KAISER A., ROSENTHALER L.: Enhanced RTI for gloss reproduction. *Electronic Imaging* 2017, 8 (2017), 66–72. 2, 4, 8, 9, 15, 16, 18
- [FH09] FILIP J., HAINDL M.: Bidirectional texture function modeling: A state of the art survey. *IEEE TPAMI* 31, 11 (2009), 1921–1940. 2
- [FPBP09] FORÉS A., PATTANAIK S. N., BOSCH C., PUEYO X.: BRD-FLab: A general system for designing BRDFs. *CEIG* 9 (2009), 153–160. 11
- [FQJ*17] FAN H., QI L., JU Y., DONG J., YU H.: Refractive laser triangulation and photometric stereo in underwater environment. *Optical Engineering* 56, 11 (2017), 113101. 16, 17
- [G*19] GIACHETTI A., ET AL.: apTool project, 2019. [Online; accessed-March-2019]. URL: <https://github.com/giach68/apTool>. 18, 19
- [GCD*17] GIACHETTI A., CIORTAN I. M., DAFFARA C., PINTUS R., GOBBETTI E.: Multispectral RTI analysis of heterogeneous artworks. In *Proc. GCH* (2017). 3, 4, 5, 12, 13, 14, 15, 16, 18, 20
- [GCD*18] GIACHETTI A., CIORTAN I. M., DAFFARA C., MARCHIORO G., PINTUS R., GOBBETTI E.: A novel framework for highlight reflectance transformation imaging. *Computer Vision and Image Understanding* 168 (2018), 118–131. 2, 3, 21
- [GCP*09] GHOSH A., CHEN T., PEERS P., WILSON C. A., DEBEVEC P.: Estimating specular roughness and anisotropy from second order spherical gradient illumination. *Computer Graphics Forum* 28 (2009), 1161–1170. 3, 4, 10, 13
- [GCP*10] GHOSH A., CHEN T., PEERS P., WILSON C. A., DEBEVEC P.: Circularly polarized spherical illumination reflectometry. *ACM TOG* 29, 6 (2010), 162:1–162:11. 3, 4, 10
- [GDR*15] GIACHETTI A., DAFFARA C., REGHELIN C., GOBBETTI E., PINTUS R.: Light calibration and quality assessment methods for reflectance transformation imaging applied to artworks' analysis. In *Optics for Arts, Architecture, and Archaeology V* (2015), vol. 9527, International Society for Optics and Photonics, p. 95270B. 19

- [GGG*16] GUARNERA D., GUARNERA G. C., GHOSH A., DENK C., GLENCROSS M.: BRDF representation and acquisition. *Computer Graphics Forum* 35, 2 (2016), 625–650. 2
- [GLSWE18] GOLDMAN Y., LINN R., SHAMIR O., WEINSTEIN-EVRON M.: Micro-RTI as a novel technology for the investigation and documentation of archaeological textiles. *Journal of Archaeological Science-Reports* 19 (2018), 1–10. 16
- [GPDG12] GUARNERA G. C., PEERS P., DEBEVEC P., GHOSH A.: Estimating surface normals from spherical stokes reflectance fields. In *Proc. ECCV* (2012), pp. 340–349. 4, 14
- [GTHD03] GARDNER A., TCHOU C., HAWKINS T., DEBEVEC P.: Linear light source reflectometry. *ACM TOG* 22, 3 (2003), 749–758. 3, 4, 10, 15, 16
- [Ham15] HAMEEUW H.: Mesopotamian clay cones in the ancient near east collections of the royal museums of art and history. *Bulletin van de Koninklijke Musea voor Kunst en Geschiedenis* 84 (2015), 5–48. 2, 4, 13, 14, 15, 16
- [HBMG02] HAMMER Ø., BENGTSON S., MALZBENDER T., GELB D.: Imaging fossils using reflectance transformation and interactive manipulation of virtual light sources. *Palaeontologia Electronica* 5, 1 (2002), 1–9. 2, 4, 11, 12, 16, 17
- [HNI08] HARA K., NISHINO K., IKEUCHI K.: Mixture of spherical distributions for single-view relighting. *IEEE TPAMI* 30, 1 (2008), 25–35. 4, 10
- [HP15] HARRIS S., PIQUETTE K. E.: Reflectance transformation imaging (RTI) for visualising leather grain surface morphology as an aid to species identification: a pilot study. *Archaeological Leather Group Newsletter* 42, 2015. 2, 4, 16, 17
- [HS17] HUI Z., SANKARANARAYANAN A. C.: Shape and spatially-varying reflectance estimation from virtual exemplars. *IEEE TPAMI*, 10 (2017), 2060–2073. 4, 10, 20, 21
- [HWBC15] HUANG X., WALTON M., BEARMAN G., COSSAIRT O.: Near light correction for image relighting and 3D shape recovery. In *Digital Heritage* (2015), vol. 1, IEEE, pp. 215–222. 2, 3
- [Ike18] IKEHATA S.: Cnn-ps: Cnn-based photometric stereo for general non-convex surfaces. In *Proc. ECCV* (2018), pp. 3–18. 3, 21
- [JA11] JOHNSON M. K., ADELSON E. H.: Shape estimation in natural illumination. In *Proc. CVPR* (2011), IEEE, pp. 2553–2560. 20
- [KK13] KOTOULA E., KYRANOUDI M.: Study of ancient greek and roman coins using reflectance transformation imaging. *E-conservation magazine* 25 (2013), 74–88. 2, 4, 11, 15, 16
- [KUL19] KUL: PLD software KU-Leuven, 2019. [Online; accessed-March-2019]. URL: <https://portablelightdome.wordpress.com/software>. 18, 19
- [LBAD*06] LAWRENCE J., BEN-ARTZI A., DECORO C., MATUSIK W., PFISTER H., RAMAMOORTHY R., RUSINKIEWICZ S.: Inverse shade trees for non-parametric material representation and editing. *ACM TOG* 25, 3 (2006), 735–745. 10
- [LCG12] LV H., CAI Y., GUO S.: 3D reconstruction of tongue surface based on photometric stereo. In *Proc. Signal Processing (ICSP)* (2012), vol. 3, IEEE, pp. 1668–1671. 16, 17
- [LSC18] LI Z., SUNKAVALLI K., CHANDRAKER M.: Materials for masses: Svbrdf acquisition with a single mobile phone image. *arXiv preprint arXiv:1804.05790* (2018). 21
- [Mac15] MACDONALD L. W.: *Realistic visualisation of cultural heritage objects*. PhD thesis, UCL (University College London), 2015. 3, 4, 5, 7, 8, 16, 17, 18, 19
- [Mac18] MACDONALD L. W.: Visual realism in digital heritage. In *Heritage Preservation*. Springer, 2018, pp. 21–45. 4, 8, 15, 16
- [MBW*14] MANFREDI M., BEARMAN G., WILLIAMSON G., KRONKRIGHT D., DOEHNE E., JACOBS M., MARENGO E.: A new quantitative method for the non-invasive documentation of morphological damage in paintings using RTI surface normals. *Sensors* 14, 7 (2014), 12271–12284. 2, 4, 13, 16
- [MCPC17] MYTUM H., CHAPMAN K., PETERSON J., CROSS A.: Reflectance transformation imaging (RTI): Capturing gravestone detail via multiple digital images. *Association for Gravestone Studies* 42, 2 (2017). 4, 11, 12, 13, 15, 16
- [MDA02] MASSELUS V., DUTRÉ P., ANRYS F.: The free-form light stage. In *Proc. ACM SIGGRAPH Abstracts and Applications* (2002), ACM, pp. 262–262. 2, 3
- [MGW01] MALZBENDER T., GELB D., WOLTERS H.: Polynomial texture maps. In *Proc. SIGGRAPH* (2001), pp. 519–528. 1, 2, 3, 4, 7, 11, 15, 16, 17, 21
- [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 *Proc. Rendering Techniques* (2007), Eurographics Association, pp. 183–194. 3, 4, 10, 13, 16
- [MMC*08] MUDGE M., MALZBENDER T., CHALMERS A., SCOPIGNO R., DAVIS J., WANG O., GUNAWARDANE P., ASHLEY M., DOERR M., PROENCA A., BARBOSA J.: Image-based empirical information acquisition, scientific reliability, and long-term digital preservation for the natural sciences and cultural heritage. In *Eurographics (Tutorials)* (2008). 3, 4, 8, 9, 15, 16, 17
- [MMSL06] MUDGE M., MALZBENDER T., SCHROER C., LUM M.: New reflection transformation imaging methods for rock art and multiple-viewpoint display. In *Proc. VAST* (2006), vol. 6, pp. 195–202. 1, 4, 11, 15, 16, 21
- [MVSL05] MUDGE M., VOUTAZ J.-P., SCHROER C., LUM M.: Reflection transformation imaging and virtual representations of coins from the hospice of the grand st. bernard. In *Proc. VAST* (2005), vol. 6, pp. 29–40. 2, 4, 11, 15, 16
- [MWGA06] MALZBENDER T., WILBURN B., GELB D., AMBRISCO B.: Surface enhancement using real-time photometric stereo and reflectance transformation. In *Proc. Rendering techniques* (2006), pp. 245–250. 4, 11, 12, 13, 16, 17
- [NDM05] NGAN A., DURAND F., MATUSIK W.: Experimental analysis of BRDF models. *Rendering Techniques 2005*, 16 (2005), 117–126. 3, 10, 20
- [New15] NEWMAN S. E.: Applications of reflectance transformation imaging (RTI) to the study of bone surface modifications. *Journal of Archaeological Science* 53 (2015), 536–549. 4, 15, 16
- [NK14] NAM G., KIM M. H.: Multispectral photometric stereo for acquiring high-fidelity surface normals. *IEEE computer graphics and applications* 34, 6 (2014), 57–68. 20
- [P*19a] PALMA G., ET AL.: WebRTI Viewer, 2019. [Online; accessed-March-2019]. URL: <http://vcg.isti.cnr.it/rti/webviewer.php>. 18, 19
- [P*19b] PONCHIO F., ET AL.: Relight website, 2019. [Online; accessed-March-2019]. URL: <http://vcg.isti.cnr.it/relight/>. 18, 19
- [Pan16] PAN R.: Detection of edges from polynomial texture maps. *3D Research* 7, 3 (2016), 81:1–81:8. 4, 12, 14, 15, 16, 17
- [PBC*16] PREIM B., BAER A., CUNNINGHAM D., ISENBERG T., ROPINSKI T.: A survey of perceptually motivated 3D visualization of medical image data. *Computer Graphics Forum* 35, 3 (2016), 501–525. 22
- [PBFS14] PALMA G., BALDASSARRI M., FAVILLA M., SCOPIGNO R.: Storytelling of a coin collection by means of RTI images: the case of the Simoneschi collection in Palazzo Blu. In *Museums and the Web* (2014). 2, 15, 16, 17
- [PCC*10] PALMA G., CORSINI M., CIGNONI P., SCOPIGNO R., MUDGE M.: Dynamic shading enhancement for reflectance transformation imaging. *ACM JOCCH* 3, 2 (2010), 6:1–6:20. 4, 12, 15, 16

- [PCGG16] PINTUS R., CIORTAN I., GIACHETTI A., GOBBETTI E.: Practical free-form RTI acquisition with local spot lights. In *Proc. STAG* (2016). 2, 3
- [PCS18] PONCHIO F., CORSINI M., SCOPIGNO R.: A compact representation of relightable images for the web. In *Proc. ACM Web3D* (2018), pp. 1:1–1:10. 2, 3, 4, 5, 15, 16, 17, 18, 19, 20, 21
- [PDJ*18] PINTUS R., DULECHA T., JASPE A., GIACHETTI A., CIORTAN I., GOBBETTI E.: Objective and Subjective Evaluation of Virtual Relighting from Reflectance Transformation Imaging Data. In *Proc. GCH* (2018). 5, 20, 21, 22
- [PF14] PAPADHIMITRI T., FAVARO P.: Uncalibrated near-light photometric stereo. In *Proc. BMVC* (2014), BMVA Press. 3
- [PGPG17] PINTUS R., GIACHETTI A., PINTORE G., GOBBETTI E.: Guided robust matte-model fitting for accelerating multi-light reflectance processing techniques. In *Proc. BMVC* (September 2017). 4, 7
- [PLG*13] PAROT V., LIM D., GONZÁLEZ G., TRAVERSO G., NISHIOKA N. S., VAKOC B. J., DURR N. J.: Photometric stereo endoscopy. *Journal of biomedical optics* 18, 7 (2013), 076017. 16, 17
- [PLGF*15] PITARD G., LE GOÏC G., FAVRELIÈRE H., SAMPER S., DESAGE S.-F., PILLET M.: Discrete modal decomposition for surface appearance modelling and rendering. In *Optical Measurement Systems for Industrial Inspection IX* (2015), vol. 9525, International Society for Optics and Photonics, pp. 952523:1–952523:10. 3, 4, 8, 16, 17, 21
- [PLGM*17a] PITARD G., LE GOÏC G., MANSOURI A., FAVRELIÈRE H., DESAGE S.-F., SAMPER S., PILLET M.: Discrete modal decomposition: a new approach for the reflectance modeling and rendering of real surfaces. *Machine Vision and Applications* 28, 5–6 (2017), 607–621. 4, 8, 15, 16, 20
- [PLGM*17b] PITARD G., LE GOÏC G., MANSOURI A., FAVRELIÈRE H., PILLET M., GEORGE S., HARDEBERG J. Y.: Reflectance-based surface saliency. In *Proc. ICIP* (2017), IEEE, pp. 445–449. 4, 12, 16, 17
- [PLGM*17c] PITARD G., LE GOÏC G., MANSOURI A., FAVRELIÈRE H., PILLET M., GEORGE S., HARDEBERG J. Y.: Robust anomaly detection using reflectance transformation imaging for surface quality inspection. In *Scandinavian Conference on Image Analysis* (2017), Springer, pp. 550–561. 4, 12, 16, 17
- [PPV08] PINTUS R., PODDA S., VANZI M.: An automatic alignment procedure for a four-source photometric stereo technique applied to scanning electron microscopy. *IEEE Transactions on Instrumentation and Measurement* 57, 5 (2008), 989–996. 16, 17
- [PPY*16] PINTUS R., PAL K., YANG Y., WEYRICH T., GOBBETTI E., RUSHMEIER H.: A survey of geometric analysis in cultural heritage. In *Computer Graphics Forum* (2016), vol. 35, pp. 4–31. 14
- [PRA15] PIRES H., RUBIO J. M., ARANA A. E.: Techniques for revealing 3D hidden archeological features: morphological residual models as virtual-polynomial texture maps. *The International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences* 40, 5 (2015), 415–421. 15
- [PS05] PALUSZYŃSKI J., SŁOWKO W.: Surface reconstruction with the photometric method in sem. *Vacuum* 78, 2–4 (2005), 533–537. 16, 17
- [PSM05] PADFIELD J., SAUNDERS D., MALZBENDER T.: Polynomial texture mapping: a new tool for examining the surface of paintings. *ICOM Committee for Conservation 1* (2005), 504–510. 2, 4, 12, 14, 15, 16
- [RDL*15] REN P., DONG Y., LIN S., TONG X., GUO B.: Image based relighting using neural networks. *ACM TOG* 34, 4 (2015), 111:1–111:12. 3, 4, 8, 9
- [RHD*10] REINHARD E., HEIDRICH W., DEBEVEC P., PATTANAIK S., WARD G., MYZKOWSKI K.: *High dynamic range imaging: acquisition, display, and image-based lighting*. Morgan Kaufmann, 2010. 2
- [RM07] REDMAN J., MUDGE M.: The simultaneous capture of spectral and textural information. In *Archiving Conference* (2007), vol. 2007, Society for Imaging Science and Technology, pp. 2–5. 20
- [RMS*08] RUMP M., MÜLLER G., SARLETTE R., KOCH D., KLEIN R.: Photo-realistic rendering of metallic car paint from image-based measurements. In *Computer Graphics Forum* (2008), vol. 27, pp. 527–536. 4, 16, 17
- [RPG16] RIVIERE J., PEERS P., GHOSH A.: Mobile surface reflectometry. In *Computer Graphics Forum* (2016), vol. 35, pp. 191–202. 4, 10
- [RTF*04] RASKAR R., TAN K.-H., FERIS R., YU J., TURK M.: Non-photorealistic camera: depth edge detection and stylized rendering using multi-flash imaging. In *ACM TOG* (2004), vol. 23, pp. 679–688. 2, 3, 4, 6, 7, 12, 14, 16, 17
- [RWS*11] REN P., WANG J., SNYDER J., TONG X., GUO B.: Pocket reflectometry. *ACM TOG* 30, 4 (2011), 45:1–45:10. 3
- [SLDS14] SUN J., LIU Z., DING Y., SMITH M.: Recovering skin reflectance and geometry for diagnosis of melanoma. In *Computer Vision Techniques for the Diagnosis of Skin Cancer*. Springer, 2014, pp. 243–265. 16, 17
- [SOSI03] SATO I., OKABE T., SATO Y., IKEUCHI K.: Appearance sampling for obtaining a set of basis images for variable illumination. In *Proc. ICCV* (2003), pp. 800–807. 2, 5, 15, 16
- [SS00] SMITH M. L., STAMP R. J.: Automated inspection of textured ceramic tiles. *Computers in Industry* 43, 1 (2000), 73–82. 16, 17
- [SSM*17] SELMO D., STURT F., MILES J., BASFORD P., MALZBENDER T., MARTINEZ K., THOMPSON C., EARL G., BEVAN G.: Underwater reflectance transformation imaging: a technology for in situ underwater cultural heritage object-level recording. *Journal of Electronic Imaging* 26, 1 (2017), 011029. 4, 11, 13, 16, 17
- [SSS*08] SUN J., SMITH M., SMITH L., COUTTS L., DABIS R., HARRLAND C., BAMBER J.: Reflectance of human skin using colour photometric stereo: with particular application to pigmented lesion analysis. *Skin research and technology* 14, 2 (2008), 173–179. 16, 17
- [SWM*16] SHI B., WU Z., MO Z., DUAN D., YEUNG S.-K., TAN P.: A benchmark dataset and evaluation for non-lambertian and uncalibrated photometric stereo. In *Proc. CVPR* (2016), pp. 3707–3716. 1, 10
- [SZV*14] SCHUSTER C., ZHANG B., VAISH R., GOMES P., THOMAS J., DAVIS J.: RTI compression for mobile devices. In *Proc. Information Technology and Multimedia (ICIMU)* (2014), IEEE, pp. 368–373. 20
- [TBF*17] THANIKACHALAM N., BABOULAZ L., FIRMENICH D., SÜSTRUNK S., VETTERLI M.: Handheld reflectance acquisition of paintings. *IEEE Transactions on Computational Imaging* 3, ARTICLE (2017), 580–591. 4, 8, 9, 15, 16
- [TFG*13] TUNWATTANAPONG B., FYFFE G., GRAHAM P., BUSCH J., YU X., GHOSH A., DEBEVEC P.: Acquiring reflectance and shape from continuous spherical harmonic illumination. *ACM TOG* 32, 4 (2013), 109:1–109:12. 3, 4, 10
- [TGK*17] TOMINSKI C., GLADISCH S., KISTER U., DACHSELT R., SCHUMANN H.: Interactive lenses for visualization: An extended survey. *Computer Graphics Forum* 36, 6 (2017), 173–200. 22
- [TLQ06] TAN P., LIN S., QUAN L.: Resolution-enhanced photometric stereo. In *Proc. ECCV* (2006), pp. 58–71. 4, 13, 14
- [TML*13] TAKATANI T., MATSUSHITA Y., LIN S., MUKAIGAWA Y., YAGI Y.: Enhanced photometric stereo with multispectral images. In *MVA* (2013), pp. 343–346. 4, 13, 20
- [TU16] TANKSALE T. M., URBAN P.: Trichromatic reflectance capture using a tunable light source: Setup, characterization and reflectance estimation. *Electronic Imaging* 2016, 9 (2016), 1–7. 11, 16
- [USTY15] UKIDA H., SANO T., TANIMOTO Y., YAMAMOTO H.: 3D shape and color estimation using linear light sources and cameras. In *Proc. Imaging Systems and Techniques (IST)* (2015), IEEE, pp. 1–5. 3
- [UW13] URIBE M. D.-G., WHEATLEY D. W.: Rock art an digital technologies: the application of reflectance transformation imaging (RTI) and 3D laser scanning to the study of late Bronze Age Iberian Stelae. *Menga: Revista de prehistoria de Andalucía*, 4 (2013), 187–203. 4, 11, 15, 16

- [VdPHB*16] VAN DER PERRE A., HAMEEUW H., BOSCHLOOS V., DELVAUX L., PROESMANS M., VANDERMEULEN B., VAN GOOL L., WATTEEUW L.: Towards a combined use of IR, UV and 3D-imaging for the study of small inscribed and illuminated artefacts. In *Lights On... Cultural Heritage and Museums!* (2016), FLUP, University of Porto, pp. 163–192. 4, 6, 7, 11, 13, 16
- [VHW*18] VANDERMEULEN B., HAMEEUW H., WATTEEUW L., VAN GOOL L., PROESMANS M.: Bridging multi-light & multi-spectral images to study, preserve and disseminate archival documents. In *Proc. Archiving Conference* (2018), vol. 2018, Society for Imaging Science and Technology, pp. 64–69. 2, 4, 6, 7, 12, 13, 16, 17, 18
- [VVP*18] VANWEDDINGEN V., VASTENHOUD C., PROESMANS M., HAMEEUW H., VANDERMEULEN B., VAN DER PERRE A., LEMMERS F., WATTEEUW L., VAN GOOL L.: A status quaestionis and future solutions for using multi-light reflectance imaging approaches for preserving cultural heritage artifacts. In *Euro-Mediterranean Conference* (2018), Springer, pp. 204–211. 2, 4, 13, 15, 16
- [WGS09] WANG O., GUNAWARDANE P., SCHER S., DAVIS J.: Material classification using BRDF slices. In *Proc. CVPR* (06 2009), vol. 0, pp. 2805–2811. 14
- [WHV*16] WATTEEUW L., HAMEEUW H., VANDERMEULEN B., VAN DER PERRE A., BOSCHLOOS V., DELVAUX L., PROESMANS M., VAN BOS M., VAN GOOL L.: Light, shadows and surface characteristics: the multispectral portable light dome. *Applied Physics A* 122, 11 (2016), 976. 20
- [WK15] WEINMANN M., KLEIN R.: Advances in geometry and reflectance acquisition (course notes). In *SIGGRAPH Asia 2015 Courses* (2015), ACM. 1, 2
- [WLL*09] WEYRICH T., LAWRENCE J., LENSCH H. P., RUSINKIEWICZ S., ZICKLER T.: Principles of appearance acquisition and representation. *Foundations and Trends in Computer Graphics and Vision* 4, 2 (2009), 75–191. 2
- [WLW17] WAINWRIGHT D. K., LAUDER G. V., WEAVER J. C.: Imaging biological surface topography in situ and in vivo. *Methods in Ecology and Evolution* 8, 11 (2017), 1626–1638. 4, 16, 17
- [WMTG05] WINNEMÖLLER H., MOHAN A., TUMBLIN J., GOOCH B.: Light waving: Estimating light positions from photographs alone. In *Computer Graphics Forum* (2005), vol. 24, pp. 433–438. 3
- [Woo80] WOODHAM R. J.: Photometric method for determining surface orientation from multiple images. *Optical engineering* 19, 1 (1980), 513–531. 1, 4, 6, 7, 10
- [WVM*05] WILLEMS G., VERBIEST F., MOREAU W., HAMEEUW H., VAN LERBERGHE K., VAN GOOL L.: Easy and cost-effective coneiform digitizing. In *Proc. VAST* (2005), Eurographics, pp. 73–80. 2, 4, 11, 12, 13, 15, 16
- [WVP15] WATTEEUW L., VANDERMEULEN B., PROESMANS M.: On the surface and beyond. a new approach with multispectral photometric stereo to assess illuminated manuscripts and their condition. In *Science and Engineering in Arts, Heritage and Archaeology, book of abstracts* (2015), vol. 1, University College London, pp. 103–103. 4, 6, 7, 15, 16
- [XDW15] XIE W., DAI C., WANG C. C.: Photometric stereo with near point lighting: A solution by mesh deformation. In *Proc. CVPR* (2015), pp. 4585–4593. 3
- [XR19] X-RITE: Total appearance capture ecosystem, 2019. [Online; accessed-March-2019]. URL: <https://www.xrite.com/categories/appearance/tac7>. 17
- [XSHR18] XU Z., SUNKAVALLI K., HADAP S., RAMAMOORTHY R.: Deep image-based relighting from optimal sparse samples. *ACM TOG* 37, 4 (2018), 126:1–126:13. 4, 8, 19, 20, 21
- [XSJ*15] XIE L., SONG Z., JIAO G., HUANG X., JIA K.: A practical means for calibrating an led-based photometric stereo system. *Optics and Lasers in Engineering* 64 (2015), 42–50. 3
- [YGC17] YOKOYA N., GROHNFELDT C., CHANUSSOT J.: Hyperspectral and multispectral data fusion: A comparative review of the recent literature. *IEEE Geoscience and Remote Sensing Magazine* 5, 2 (2017), 29–56. 2, 22
- [ZD14] ZHANG M., DREW M. S.: Efficient robust image interpolation and surface properties using polynomial texture mapping. *EURASIP Journal on Image and Video Processing* 2014, 1 (2014), 25. 4, 7, 8, 20
- [ZHS*18] ZHANG W., HANSEN M. F., SMITH M., SMITH L., GRIEVE B.: Photometric stereo for three-dimensional leaf venation extraction. *Computers in Industry* 98 (2018), 56–67. 4, 14, 16, 17
- [ZLR*10] ZHENG J., LI Z., RAHARDJA S., YAO S., YAO W.: Collaborative image processing algorithm for detail refinement and enhancement via multi-light images. In *Proc. ICASSP* (2010), IEEE, pp. 1382–1385. 4, 6, 7, 16, 17