

Web-based Multi-layered Exploration of Annotated Image-based Shape and Material Models

A. Jaspe¹, R. Pintus¹, A. Giachetti², and E. Gobbetti¹

¹Visual Computing, CRS4, Italy

²Department of Computer Science, University of Verona, Italy

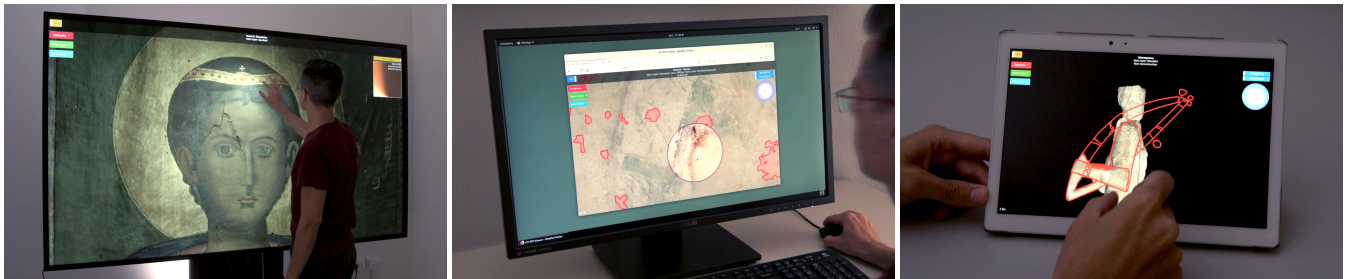


Figure 1: Web-based interactive inspection. Our system allows users to explore annotated heterogeneous datasets coming from a variety of pipelines using a web platform capable to run on several devices. Left: exploration of a multispectral acquisition of a painting on a large scale display; Center: analysis of a fragmented statue with annotations in a desktop browser; Right: analysis of an annotated fresco on a mobile platform.

Abstract

We introduce a novel versatile approach for letting users explore detailed image-based shape and material models integrated with structured, spatially-associated descriptive information. We represent the objects of interest as a series of registered layers of image-based shape and material information. These layers are represented at multiple scales, and can come out of a variety of pipelines and include both RTI representations and spatially-varying normal and BRDF fields, eventually as a result of fusing multi-spectral data. An overlay image pyramid associates visual annotations to the various scales. The overlay pyramid of each layer can be easily authored at data preparation time using widely available image editing tools. At run-time, an annotated multi-layered dataset is made available to clients by a standard web server. Users can explore these datasets on a variety of devices, from mobile phones to large scale displays in museum installations, using JavaScript/WebGL2 clients capable to perform layer selection, interactive relighting and enhanced visualization, annotation display, and focus-and-context multiple-layer exploration using a lens metaphor. The capabilities of our approach are demonstrated on a variety of cultural heritage use cases involving different kinds of annotated surface and material models.

CCS Concepts

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

1 Introduction

The virtual inspection of digital representations of objects is of fundamental importance in many application fields [PDC*19]. In particular, it is an integral part of the daily activity of Cultural Heritage (CH) scholars, as it supports the study and definition of conservation and preservation strategies by revealing geometric cues or information on the materials used to create an object [PPY*16]. Moreover, digital representations are also increasingly used to present cultural objects to experts or to a wider public, in order to replace, augment, or complement the inspection of real objects. While early approaches mainly focused on passive visual presentations, e.g. from

the browsing of image collections done by experts to the presentation of authored videos in museums, interest has now shifted to letting users directly drive exploration, since interactive methods are known to better support experts in inspection tasks, as well as to improve engagement of casual users in museums [FD00, KSZ*11].

In parallel to generic interactive viewers displaying fully-3D virtual replicas [PCD*15, Ske19], relighting interfaces, popularized by Reflectance Transformation Imaging (RTI) viewers [CHI19], have emerged as one of the most successful exploration modes. This is due to the fact that relightable models are easy to acquire and naturally support a type of visualization very appropriate to inspect fine

surface details and resembling the classical physical inspection raking light sources to reveal surface detail of actual objects under study. Moreover, the restriction of camera motion to panning and zooming is very appropriate to a variety of cultural objects and, at the same time, removes one of the main difficulties of 3D exploration applications, reducing learning curves [JH13]. For this reason, interactive inspection of relightable images has been applied to a wide range of items [AIK13]. Such viewers, however, have so far been restricted to inspecting only few types of models, mainly in the form of raw multi-light image collections or low-frequency analytical relighting representations, such as PTM or HSH [PDC*19]. This is due to the fact that, so far, more complex representations, such as spatially-varying normals and BRDFs were too difficult to reliably extract from captured data. The progress on methods for extraction of geometry [AG15] and surface appearance [WK15, DRS10, GGG*16] is, however, removing these limitations, introducing the need for tools also supporting the inspection of image-based representations of decoupled geometry and appearance parameters. Moreover, few relighting viewers support annotations on top of the relightable model, and, in most cases this is restricted to adding separate information through hyperlinks associated to point or area hotspots.

Our work aims at letting viewers explore detailed image-based shape and material models integrated with structured spatially associated visual information. This is achieved through a novel web-based framework for exploring complex image-based 3D object representations on the web. Off-line content preparation organizes data in a series of registered multiresolution layers of shape and material information. These layers are all image-based, and can come out of a variety of pipelines which produce parametric information, including both RTI representations and spatially-varying normal and BRDF fields, eventually obtained by fusing multi-spectral data. At authoring time, each layer is associated to an overlay image pyramid, which contains visual annotations to the various scales. Such a pyramid can be the outcome of other processing steps (e.g., crack detection), or hand-drawn by users. Annotations are different at the various resolution levels to provide various levels of abstractions (e.g., with a set of text and drawings associated to the global shape of the object, and a different set to the small-scale feature details). At run-time, such annotated multi-layered datasets are made available to clients using web servers. Users explore the datasets on a variety of devices, from mobile phones to large-scale displays in museum installations, using JavaScript/WebGL2 clients capable to perform layer selection, interactive relighting and enhanced visualization, annotation display, and focus-and-context exploration of multiple layers using a lens metaphor to provide alternative visual representations for selected regions of interest. Our main specific contributions are the generalization of fixed-view relighting interfaces to several representations outside of pure RTI, the support for unrestricted multi-scale visual annotations, and the application of advanced focus-and-context visual exploration means to the field of shape and material analysis in CH. A source code release of a reference implementation is available from www.crs4.it/vic/download/.

2 Related work

Image-based relighting techniques are a vast and well-known subject, and a complete coverage of the literature is out of the scope of this paper. We discuss here only the approaches most closely

related to ours, and we refer the reader to established surveys for a wider coverage [PDC*19, MMC*08, VVP*18].

Representations suitable for relighting frameworks and visualization tools are computed from a variety of sources, from colored 3D models to multi-light image collections (MLICs). In any case, image-based relighting relies on a pixel-wise representation, which stores for each image coordinate a specific type of data that is used at real-time to produce the rendering under a novel, virtual illumination. Depending on the rendering strategy, this representation might be a set of resampled/compressed/rearranged input samples [Mac15, PCS18], one or more image layers [VHW*18], coefficients of a reflectance parametric model [CHI19], or shape (e.g., normal map) and material (e.g., (SV)BRDF) data [GTHD03]. In this work, we aim to provide a general multi-layered representation inspection framework which is not limited to pure reflectance fields, but also includes separate shape and material representations.

In the last decade, a wide variety of tools for performing visual inspection of image-based data have been presented and, sometimes, deployed for open, public use. In all cases, the tools target either static original data exploration (e.g., multispectral or stratigraphic data [MAD*18] or multi-light image collections [VHW*18, Mac15]) or dynamic exploration through relighting, exploiting specific RTI formats [PDC*19]. Our solution extends this approach to a larger variety of cases, including image-based representations that explicitly separate shape and appearance contributions.

While early software tools were designed for desktop use and locally resident data [CHI19], recently a lot of web-based solutions have been made available for remote access to relightable models, and for use also on mobile devices [PCS18, PBFS14]. Among desktop tools, the most popular in CH applications remains *RTI viewer* [CHI19], also due to its coupling with popular acquisition and processing software (*RTI Builder* [CHI19]). It provides various ways to visualize/inspect relightable image files both locally or remotely. It employs the classic light direction selector to allow the users to render the studied artwork under novel illuminations. In addition to a photorealistic relighting mode by using Polynomial Texture Mapping (PTM) or Hemi-Spherical Harmonics (HSH) reflectance field models, it also features several non-photorealistic or illustrative enhancements, e.g., diffuse gain, specular enhancement, unsharp masking applied to the normals, the image, the luminance or the model coefficients, static or dynamic multi-lighting, or normal map visualization. While mostly devoted to processing MLIC data, the *APTTool* [G*19] desktop tool enables a continuous visualization of original data through Radial Basis Functions (RBFs) interpolation. It also provides classic PTM and HSH modes, the display of some feature maps, and the visualization of the per-pixel interpolated reflectance maps. In the field of multi-spectral relighting, *PLDviewer* [KUL19], strongly linked to a specific processing framework [VHW*18] and a proprietary compressed file, supports photorealistic and illustrative visualizations of various 2D and 3D outputs, e.g., curvature coloring, exaggerated shading, sketching, single pixel reflection maps, and local histogram. Web-based tools for image-based relighting, i.e., *WebRTIViewer* [P*19a], *PLDWebviewer*, the *Digital Materiality Viewer* [DHL19, FBKR17], and *Relight* [P*19b, PCS18], typically support only specific parametric representations of relightable images (in particular, PTM and HSH), and provide interactive relighting and some enhancement capabilities.

R1	Inspection of geometric features at multiple scales. CH items have a high heterogeneity and present important geometric details at various scales. It should be possible to analyze shape both in conjunction and separately from material, at the micro/meso/macro level.
R2	Multi-spectral appearance. The understanding of the material behaviour is of fundamental importance to grasp the highest possible amount of insights for CH item interpretation and conservation. The visualization tool must enable the inspection not only of the visible signal, but also of invisible (e.g., from ultraviolet or infrared light) surface response. Similarly, it is important for CH scholars to visualize and analyze more complex phenomena such as object induced fluorescence.
R3	Interactive relighting. It should be possible to select a virtual light direction and to relight the captured object under that novel illumination, in order to understand how the object locally interacts with the light. This should be interactive to get motion parallax.
R4	Information spatially connected with models. Descriptive information, textual or visual, should be associated to specific parts of the presented images, since it is generally spatially connected to a region of the cultural object.
R5	Information abstraction. Different macro-structural and micro-structural views should be associated with different kinds of information (e.g., descriptions of carvings and decorations, reconstruction hypotheses, comparisons with other objects, which all refer to different scales of the objects).
R6	Information authoring. Adding textual and visual information (drawings, images) should be easily supported and should be possible for museum curators and archaeologists without particular training, without requiring intervention of specialized personnel.
R7	Information layers. Each object can be acquired and seen from different perspectives in terms of measured data or combination of different acquired signals. For each CH item, a visualization tool must provide the user with the possibility to choose a series of different layers/representations of the data (e.g., visual data, ultraviolet, infrared, normal map, fluorescence, or combinations of these modes).
R8	Stratigraphy, focus and context. It is often necessary to analyze multiple layers coming from different modalities to support interpretation. Switching among layers to show, e.g., a stratigraphy of the object is a required feature. However, toggling a single-layer views makes it difficult to understand correlations between separate layers. It is therefore necessary to support the flexible combination of at least two layers at a time.
R9	Focus on cultural object. For both the user and the experts, the important information is the visualized object itself, which should therefore not be obstructed by interaction widgets.
R10	Fast learning curve. CH domain experts, which do not have a scientific background, have to use visualization and analysis tools with the minimum effort. Moreover, software already exist that are widespread among the CH community. New solutions must take this aspect into account, in order to be compliant with standard interfaces in CH daily work or to add new features that require the smallest and easiest learning process.
R11	User interface and platform flexibility. The possible use cases of cultural object analysis are various, ranging from conservation analysis by domain experts to museum exhibits. The tool should be configurable, not constrained a particular hardware, and usable with the largest variety of setups (local/remote) and devices, ranging from common desktops to mobile devices (tablets, smartphones), and large, possibly touch, display for museum or general collaborative applications.
R12	Seamless interactive exploration. All the visualization and analysis tasks must be interactive and smooth. The relighting must be continuous and at interactive rate, and all the other inspection operations must give the end-user a smooth and real-time experience.

Table 1: Requirements. List of requirements derived from experts, analysis of related work, and past experience.

ties. An important aspect of these tools is their compression strategy used to decrease the data memory footprint for remote access; some rely on classic image compression techniques, while others explicitly exploit the nature of some input data (e.g., MLICs) to perform resampling and compression [PCS18].

Most of the CH data presentation tools also support the integration of interconnected text and model information with bidirectional navigation [GVH*07, JD12, CLDS13, PCD*15], and several tools are emerging that also offer interfaces for multi-user annotation creation (e.g., Aioli [Aio19]). In this work, we focus on image-based representations, and, therefore, do not place 3D annotations on objects but on the image plane. This leads to some limitations, since the annotations carried out on a representation are transferable from one layer to another of the same multi-layered dataset, but are not automatically transferable to another view of the same object. On the other hand, we also achieve more flexibility, as we expand over present approaches by supporting the authoring and integration of free-form overlays at multiple scales. Freeform annotations are not forced to be just 3D ones, but can be, e.g., sketches of possible reconstructions of missing parts. In typical image viewers, multiple layer presentation is generally supported by the controlled blending of data coming from multiple modalities. We provide a more general solution by adopting an interactive lens approach [TGK*14], whose main purpose is to support multi-faceted data exploration, by providing, on top of a defined visualization mode, an alternative visual representation of a local area of the data. This method has proved successful for multi-field visualization, but has not been used, so far, in the CH field for relightable image exploration.

3 Overview

Even though the proposed framework is of general use, its creation has been motivated and inspired by a large series of use cases and collaborative efforts in the CH field. From interactions with domain experts, in particular conservators and museum curators, we derived a list of requirements to guide our development. Additional requirements stem from our own past experience developing interactive systems for cultural heritage [BGM*09, MAG*12, BGMT13, BAB*16]. The final list of requirements, summarized in Table 1, as well as our analysis of related work presented in Sec. 2, were taken as guidelines for our development process, which resulted in the definition of a client-server approach based on a multi-layered representation of data. The overall architecture is depicted in Fig. 2. In order to handle a large variety of use-cases (R11), we designed our system on a web-based platform, with data preparation done off-line and resulting in a repository made available by standard web server to a web client running in a browser on top of WebGL2, a JavaScript API for rendering interactive 3D and 2D graphics within any compatible web browser without the use of plug-ins. WebGL2 does so by introducing an API that closely conforms to OpenGL ES 3.0 that can be used in HTML5 <canvas> elements. Since the API enjoys a wide support and exploits hardware acceleration on several platforms, we can support at interactive rates both remote exploration on mobile clients and local exploration on museum setups, using a full-screen interface (R12).

Off-line content preparation organizes data in a series of registered multiresolution layers of shape and material information




Figure 2: Architecture overview. Off-line data preparation phases assemble different data representations into an image-based multi-layered dataset composed of shape and material layers coupled with per-scale annotation overlays and additional metadata. The dataset is made available on the net using a standard web server. A JavaScript/WebGL2 client loads data and interactively supports layer selection, interactive relighting and enhanced visualization, annotation display, and focus-and-context multiple-layer exploration using a lens metaphor.

combined with manually added overlay information. For maximum flexibility and to support a simple interface (R10), we assume that all data layers are all image-based, and can come out of a variety of pipelines producing relightable shape and material information, either from processing capture data, or by off-line data-fusion operation that merge layers together (R7). Shape and material data are, at this level, represented in several forms, either explicitly separating shape and material (e.g., normal+brdf), or fusing them in a reflectance field (e.g., PTM or HSH) (R1, R2). Annotations are generated in the form of per-layer per-level-of-detail overlay images corresponding to levels of abstraction (R5) of spatially registered information (R4). Standard drawing tools can be exploited for the annotation operation (R6), and a processing script takes care of transforming data into publishable format. The final multi-layered dataset is, thus, a list of layers, each composed of a combination of several multiresolution data channels and an optional multiresolution annotation layer, and enriched with other additional metadata (such as textual descriptions). Details on formats and data preparation are further provided in Sec. 4.

At run-time, each annotated multi-layered dataset is made available to clients using a standard web server. The JavaScript/WebGL2 client queries the server to receive metadata as well as image information for each of the layers. The data images are interpreted as parametric fields, depending on the particular data representation associated to the layer (e.g., normal+brdf or PTM). The user interface is oriented towards panning, zooming, and relighting operations (R3, R9, R10), and makes it possible to select the primary layer of interest, using a lens metaphor to provide a different visual representations coming out of a secondary layer in an interactively defined region of interest. The display of overlay information is synchronized with the zoom level of the current active layers (R8). Further details on viewer design and behavior are provided in Sec. 5.

4 Data representation and preparation

Our datasets are represented as a set of *layers* associated to general metadata. Each layer is defined as standalone *image-based renderable representation* (Sec. 4.1) coupled to an *annotation overlay* (Sec. 4.2), and a hypertext description.

4.1 Image-based renderable representations

We currently support two main types of *image-based renderable representation*: *shape and material representations*, and *relightable*

images. The kind of representation is specified by a *tag*, and each representation has different sets of images defining it, with each image stored in a mipmapped pyramid format.

Shape and material representations explicitly decouple geometry and appearance, represented in images as regularly sampled spatially-varying fields. For geometry, we support *depth images* and *normal maps*, while for materials we currently support a Lambertian + isotropic Ward BRDF representation [NDM05], with coefficients stored in separate image maps: two RGB images for the K_d and K_s components and a 1-component image for the α specular lobe parameter. Both shape and material component are optional in the definition of a layer. When material is missing, the shape component is implicitly associated to a default constant BRDF in the viewer, making it possible to define a pure geometric layer. When shape is missing, the viewer only displays the diffuse color without shading, allowing for the definition of pure image layers. This sort of general representation can come out of a variety of acquisition and processing pipelines. These include single-image capture [LSC18], multi-light image capture [HS17], and combinations of scanning and reflectance field capture [PDCS13].

Relightable images representations, instead, do not separate shape from material. These representations naturally come out of common RTI processing pipelines [CHI19, GCD*18]. In our first implementation, we support the common LRGB PTM format, in which a single biquadratic polynomial is used to determine a luminance value that is then used to modulate a companion RGB texture.

Note that, in all cases, we decided to only support RGB layers. Multispectral support is provided by defining multiple layers (e.g., a monochrome one for each spectral band), as well as by performing data fusion at data preparation time to produce high-quality monochrome or false-color representations using advanced global methods rather than the simple local techniques supportable within the real-time constraints of the image viewer [ZBL18].

4.2 Annotation overlays

Annotation overlays provide additional information to layers in the form of drawings and hypertexts associated to specific hierarchical levels of the image-based renderable representations. The images, which can be either rasterized or scalable vector drawings, are either the results of processing steps (e.g., crack identification) or, more frequently, created by users with standard image editors. In the authoring process, each annotation overlay is created in association with a specific zoom level of the corresponding renderable layer, and is supposed to be rendered at the same resolution. This possibility to link a single layer to multiple annotation overlay allows for several level of abstraction in conveying CH content related to a particular object.

A detailed description of the authoring process is orthogonal to the topic of this paper. We will just mention, in general, any editor that can produce a drawing or an image can be used with our method. In this paper, we have adopted the following process, fully based on open source solutions. We choose one map/overlay from the layer under study, e.g., the diffuse or specular color, the normal map, etc. We open it into an editor, we use GIMP [GIM19], and we draw the annotation as a transparent layer. After we finish that or more annotation overlays, we save them as PNG images with transparency, and we output a dataset description with the levels associated to each

annotation (see Sec. 4.3). The same dataset description also links each level to an optional hypertext description. Note that it is not necessary, and sometimes not recommended either, to have all the levels assigned to an annotation; on the contrary, if it is needed that an annotation remains for different levels of zoom, it is advisable to map that to more than one level. In Fig. 3 we show three different levels of an original layer map (in this case the diffuse color of a Ward model of a statue), the two annotations linked to the first and third level, and, finally, the superimposed visualization of them at three different levels of zoom. Note that for the middle zoom, since no annotation has been inserted, only the original map is displayed.

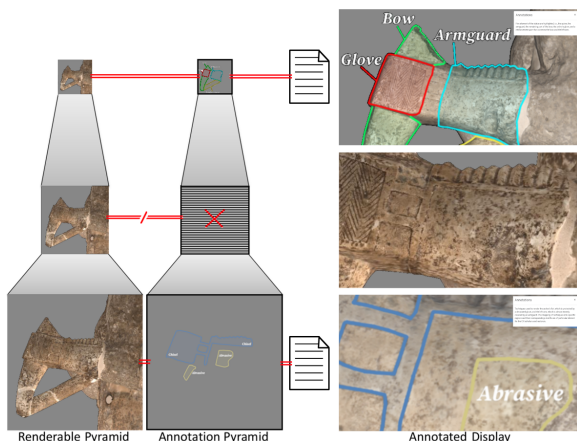


Figure 3: Multi-level Annotations. Each annotation overlay is related to a mipmap level of the Renderable model, making it possible to define scale-related levels of abstraction. Not all mipmap levels are required to have an annotation. At run-time, the zoom level selected for display activates the associated annotation. The image is displayed over the renderable layer, while the text dynamically updates an hypertext display widget.

4.3 Multi-layered data assembly

The full multi-layered data representation is assembled in a single description necessary to define a scene to initialize the viewer. All the elements, composed of scene-level data layers, are declared as JavaScript JSON structures, with properties and values, and assembled into a comprehensive scene structure, which is parsed by the viewer at initialization time to setup the viewer's interface and internal data and download the required information. Similarly to other popular solutions [PCD* 15], we chose to use JSON over XML because it is a human-readable format simple to understand, write and parse. An example of multi-layered dataset is the following:

```
{
  "name": "Sample dataset",
  "info": "<h1> ... </h1> ... ",
  "width": 2969, "height": 5469,
  "layers": [
    {
      "id": "VIS", "name": "Visible",
      "obj": {
        "kind": "nward", "resources": {
          "kd": "vis/kdMap.jpg",
          "ks": "vis/ksMap.jpg",
          "alpha": "vis/alphaMap.jpg",
          "N": "shared/N.jpg"
        }
      },
      "overlay": {
        img: "vis/overlay.png", range: [2,5]
      }
    },
    {
      id: "UV1", "name": "UV (385nm) + Fluorescence",
```

```
      "obj": {
        "kind": "nward", "resources": {
          "kd": "uvp/kdMap.jpg",
          "ks": "uvp/ksMap.jpg",
          "alpha": "uvp/alphaMap.jpg",
          "N": "shared/N.jpg"
        }
      },
      "overlay": {
        img: "uv_plus/overlay.gif",
        range: [0,3],
        shortInfo: [
          "Short annotation level 0",
          "Short annotation level 1", ...
        ],
        longInfo: [
          "Long annotation level 0",
          "Long annotation level 1", ...
        ]
      }
    }
  ]
}
```

An important feature of the scene description is that layers are allowed to share image resources. Rather than explicitly declare the sharing as in scene graphs, we decided to implement the sharing implicitly by resource URI matching (see Sec. 5). In the above example, for instance, the normal map is shared between the two layers. Moreover, for the renderable representations we assume that the full pyramid is available, while, for the overlays, only some levels of detail, within a specified range, may be available on the server. Note that the specified image filenames are templates, and the actual file names stored on the server and queried by the client are a combination of the base name of the image, a two digit level of detail, and the extension.

5 Exploring layered models locally and on the web

Our client-server viewer has been designed to provide a flexible exploration experience of a multi-layered dataset. At the core of the viewer is a basic visualization engine, which supports layer display and layer combination (Sec. 5.1). The engine is used as a basis in customized viewers, which, by configuring and selecting user interface components, target specific application use cases (Sec. 5.1). All viewers support both single-touch and multi-touch interaction so as to support a large variety of devices (Sec. 5.3), and can work both on standalone machines or connected to a remote server (Sec. 5.4).

5.1 Visualization engine

The visualization engine is developed as a JavaScript library that manages memory and display. It is initialized with the layer structure defined in Sec. 4.3. Using a hashed container, it first determines the unique set of image resources required for a given multi-layered object. A cache manager, then, is responsible to move the resources to RAM by fetching data from the server, and then to GPU texture memory, as required. The RAM cache is larger than the GPU cache, which is limited to maintaining in core at most all the textures required by all currently active layers, at the minimum the textures required by a single layer. At data fetching time, all the available mipmap pyramid is fetched. When loading to texture memory, all the needed mipmap levels are transferred if available or created using the hardware supported *generateMipmap* primitive. Automatic mipmap creation is only done for renderable layers, and not for overlay layers, which implement levels of abstraction.

The rendering process maintains a stack of active layers, which are rendered back to front and alpha-blended using OpenGL to implement layer compositing as well as interactive lenses. For each

layer, the maintained status includes, in addition to layer data and image caching status, the viewing parameters and the rendering parameters. All layers are rendered by activating a specific vertex and fragment shader pair, which depends on the layer kind.

For each currently active layer, first the GPU texture cache is updated. Then, vertex and fragment shaders are activated and configured, and rendering is triggered by drawing a geometry corresponding to the area of the viewport that should be updated (the full viewport for regular layers, and the lens geometry for lens layers). The vertex shader just computes the data required to compute the 3D and texture coordinates based on the current viewing parameters, while the fragment shader performs the actual rendering, exploiting samplers bound to each of the required texture resources, and rendering parameters associated to uniform variables. In order to enable relighting, all shaders take as input a light direction and intensity, and in order to support transparency, they also take as input a layer opacity value. In addition, per-mode specific parameters are associated to each shader (see Sec. 5.2). The end result of the shader is a RGBA pixel value, which is composited by alpha blending with the currently stored frame buffer value to implement transparency. After the renderable representation of a layer is generated, the associated annotation layered is also rendered (if activated) and composited with alpha blending, using the same geometry as the renderable layer, and a vertex and fragment shader pair which just selects the correct level of detail. By iterating the rendering of all active layers and associated annotations, the full image is produced.

5.2 Configuration and interaction components

The visualization engine is exploited as a low-level library to implement several viewer kinds. We currently support two specifically targeted implementations. The first one is an advanced viewer which exposes, for each displayed layer, all the shading parameters (See Fig. 4 top). The second, discussed in this paper, is a simplified viewer which only enables a base and a single lens layer fully configured using presets (see Fig. 4 top). The idea behind supporting only one layer of interactive lenses is to provide on demand an alternative visual representation of the data underlying a local area of the screen. By selecting which layer and presentation mode is used as a base to fill the frame and which layer and presentation mode is, instead, rendered within the possibly semitransparent lens, many common inspection use cases can be covered, from focus and context visualization of multiple acquisition modalities, to selective shape or material representation, to attribute enhancement.

Presets and viewer configurations are described in JSON files loaded at viewer startup. A preset is simply a named set of values of shader-specific uniform variables for a particular shader class.

The viewer configuration, then, just points to a multilayered dataset description and to a preset library, and provides two list of configurations, one for the base layer and one for the lens layer. Each entry in the configuration list consists in a label associated to a rendering configuration, consisting in a layer id in the multilayered dataset, a preset id in the preset library, and, for the lens layer, an opacity value.

The simple viewer, at start-up time, extracts from the configuration the list of needed layers to pass to the visualization engine, and creates two context menus pointing to the associated configurations, one for the base layer and one for the lens layer.

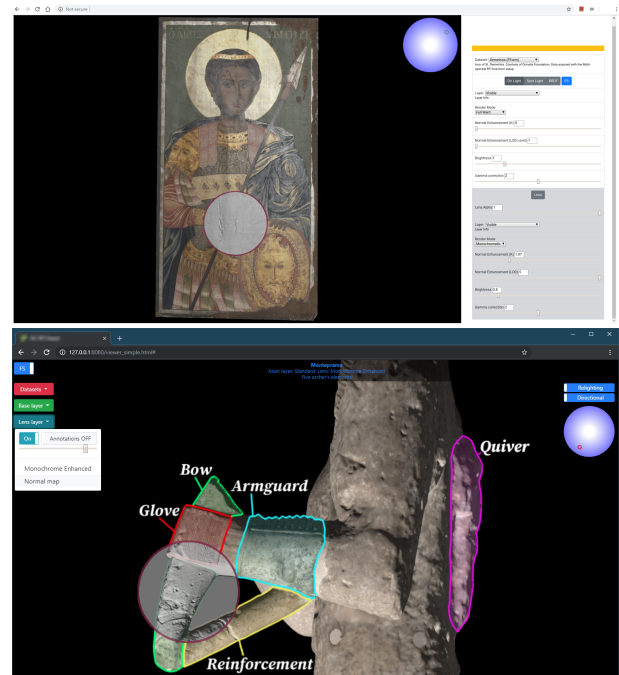


Figure 4: User interfaces. Top: advanced configuration. Bottom: simplified configuration.



Figure 5: BRDF slice. Visualization of BRDF slice under the cursor.

5.3 Pointer-based and multi-touch interaction

The minimalistic interface, using only the base and context menu, is complemented by only two other interaction widgets: one for enabling/disabling the lens and one for switching from relighting mode to measurement mode. All other actions are performed by single or multi-touch interaction. The single touch interface is meant to be operated with a 2D mouse, while multi-touch input is the default operation on most interactive surfaces. To reduce training, we use analogs of common 2D RST gestures to operate the interface, tuning the behavior depending on the current mode and the location of gesture start. Two-finger pan and pinch are mapped to translation and scaling of the lens when the gesture starts inside the lens, and to pan and zoom of the image otherwise. One-finger dragging is mapped to light direction control using the standard trackball interface when in relighting mode, and to appearance display when in measurement mode. The appearance display is only active when the current layer

contains a BRDF, and displays the $\Phi_d = 90$ slice of the BRDF under the cursor, which is a good indicator of the main reflectance features of the underlying material (see Fig. 5).

5.4 Online and offline deployment

The framework has been designed to easily support both online and local deployment. On the server side, there is currently no need for dedicated servers with specific server-side computation. All communication can be handled by a standard web server (we use Apache in our implementation) accessing a storage location with image data, metadata, and configurations. The exact same setup can be achieved on a single mini-PC without internet connection, using a local client transparently communicating with a local server. This makes deployment easier also for typical small institutions without any IT infrastructure (like most museums) and limits problems related to the need to export possibly protected models outside the institutions (as needed, e.g., for cloud-based systems such as SketchFab [Ske19]). Moreover, the simplified touch-based interface is well adapted for walk-up-and-use interaction modalities, making this work suitable for kiosks or other installations, where the system can be simply run within a browser in full-screen mode.

6 Results

An experimental software library and viewer applications have been implemented both using JavaScript and the WebGL2 environment. We have extensively tested our system with a number of complex heterogeneous datasets. In this paper, we discuss the results obtained on three different use cases: the multispectral exploration of a painting (Sec. 6.1), the analysis of a fragmented statue coming out of a 3D scanning pipeline (Sec. 6.2), and the analysis of a damaged fresco coming out of an RTI pipeline (Sec. 6.3). The accompanying video shows the system in action on all the datasets on several platforms, which were selected to demonstrate both the portability of the code and the adaptability to several different use cases. In particular, we use a museum setup based on a 98 inch UHD 4K touchscreen display driven by Chrome 75 running on Windows 10, a desktop setup using Firefox 66 on the Arch Linux distribution, a Tablet setup using a Teclast T10 Master device with Chrome 74 on Android 7.0, as well as a smartphone setup using a V Mobile J7 device running Chrome mobile 75 on android 7.0

6.1 Multi-spectral data exploration

In this use case we show how the proposed exploration tool is capable of efficiently visualize and inspect a spatially varying multi-spectral signal of a painting. The artwork is the Icon St. Demetrios (17th - 18th century). The icon is interesting from the material analysis point of view; it is an egg tempera painting on a wood support, and it contains regions where gold leaf has been employed for artistic reasons.

The painting data has been acquired by a custom free-form Multi-spectral RTI setup, i.e., a camera and a hand-held light source. We have chosen a standard high-resolution Nikon D810 DSLR camera without the hot mirror (the IR filter), and a Nikkor Lens 50mm $f/1.8D$, which ensures a proper transmittance characteristics across the spectrum from near-UV to near-IR. We capture seven monochromatic signals, by exploiting a set of five different LEDs



Figure 6: Illustrative lens. Different uses of the interactive lens and the context+focus metaphor. (Top) a crack visualization is more meaningful by adding, to the pure visible signal, an almost opaque lens showing an enhanced achromatic layer emphasizing shape variations. (Bottom) Areas of different materials are better understood by adding, to the visible signal, an interactive lens showing UV-induced fluorescence, which clearly differentiates the right side from the left side material.

light sources, and the three Bayer filters in the camera color filter arrays (CFAs); the light sources are two ultra-violet (UV), a white (VIS) and two infra-red (IR) emitters. The UV LEDs are centered at 385nm and 405nm with a full width at half maximum (FWHM) of about 20nm; the IR emitters are centered at 740nm and 850nm with approximately 40nm FWHM; on the other hand the VIS LED is a white light (5500K) that covers the entire visible spectrum. The entire captured scene contains both the painting and a series of calibration objects, i.e., a white planar target and four dark glossy spheres. For each wavelength, we have acquired a total number of about 60 images, with a light direction domain that covers the hemisphere above the painting. The calibration has been performed to compute a per-pixel light direction and intensity [GCD*18], and to produce a dataset suitable for physically based computation.

The final multi-layered representation of the Icon St. Demetrios dataset consists, first of all, in a single shared normal map computed with a robust photometric stereo method [PGPG17], and several material layers computed by applying a BRDF fitting procedure to the acquired data, i.e., a visible light layer, two pure monochromatic infrared layers, and the two ultraviolet layers that contain both the invisible reflectance and some induced fluorescence. The model used for BRDF fitting was Lambertian + Ward. In addition, through data fusion, we included three false color layers by combining infrared and visible, visible and UV, and all the seven signals merged together.

Due to its environmental conditions, this is a clear example of an artwork that needs to be constantly monitored through periodic

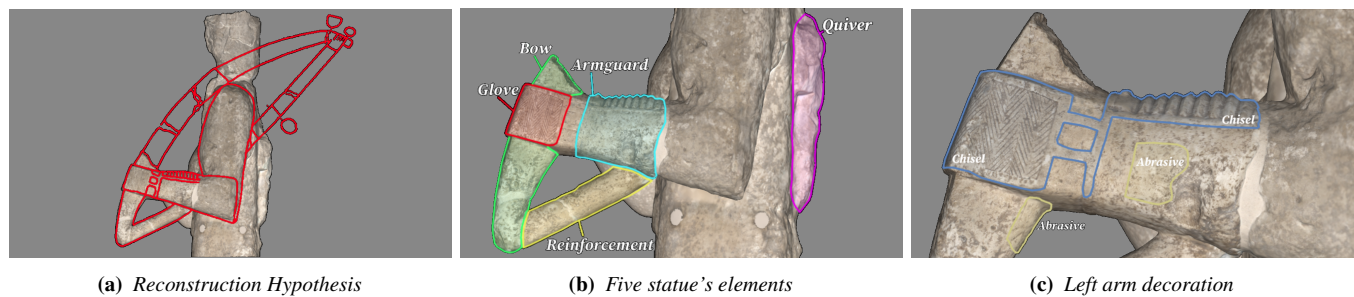


Figure 7: Annotated 3D scan model. Three different levels of annotation have been presented for a statue of an archer: (a) a hypothesis of the reconstruction of the missing archer's bow; (b) five archer's parts highlighted with contours and text annotation, i.e., the quiver, the armguard, the remaining part of the bow, the archer's glove, and a reinforcement part; (c) motifs of the decorated glove and armguard.

acquisitions, visualizations, and analysis. In fact, the artwork is affected by a series of damages and conservation issues. The painting contains cracks, layer separations/detachments, fungi, retouching areas from previous restorations; the protective layer undergoes oxidation, soot or dust deposits, fine cracks, and discoloration. All of those defects involve the appearance of the object and its geometric behavior, so it is a good example of how the proposed visualization tool can assist the scholars in examining the painting, and taking a decision on future interventions. In Figure 6 (Top) we show how the proposed interactive lens tool is capable of revealing the information about a crack on the surface in an extremely clearer way compared to a plain visualization. Here, the context layer is the full material rendering, while within the lens a monochromatic enhanced signal has been activated. The presence of dust, fungi, or other superficial contaminant could be often analyzed through ultraviolet imaging. In this case (Fig. 6, Bottom), by only inspecting the visible signal (left), we could not spot any difference in material behavior across the shield zone of the painting. Conversely, giving, as before, the visible signal as a context, the superimposed lens (right) displays the surface fluorescence induced by the ultraviolet light. With this tool it is possible to see how the left part of the face has a completely different chromatic value (more gray) than the right part (more green). In this illustrative representation of the induced ultraviolet fluorescence, the more gray is the surface, the less fluorescence we have.

6.2 Annotated 3D scan model

The exploration of an annotated 3D model has been conducted with one of the statues from the Mont'e Prama collection [BJM*15], which includes 25 life-size human figures, depicting archers, boxers and warriors, and 13 building models representing typical Nuragic towers. It has been acquired using a combination of laser scanning and flash photography to recover shape and albedo of complete 3D models. The example discussed here is Archer *n.5* (*Prexiau*). The fixed view used to produce the image data has been chosen to display a set of important features from the CH point of view, and has been resampled at the resolution of approximately 0.25mm/pixel. In this view, it is possible to see the quiver (with the resin box and the small sword), and the left arm, which holds the bow. Three different annotations have been provided. Since the bow has not been entirely preserved (only the lowest part is visible in the statue), at the highest level the annotation aims at depicting a hypothesis of reconstruction of the missing bow (Fig. 7a). It should have lied on the archer

shoulder, and the presence of a tight bowstring remains in some details in the back of the statue. At the middle level, five element of the statue are highlighted (Fig. 7b), i.e., the quiver, the armguard, the remaining part of the bow, the archer's glove, and a reinforcement part that connects the bow and the left arm. By zooming further, the annotation presents in detail the archer's fist, which is protected by a decorated glove, and the left arm almost entirely covered by an armguard (Fig. 7c). As shown in the video, the annotations naturally adapt to the viewed area, and level-specific hypertext description can be accessed on demand to gain further information. A particularly interesting usage of the tools has been, moreover, the activation of annotation overlays inside the lens, connected with their deactivation in the main view. This permits to have a clean context presentation, without any clutter due to annotation, combined with an enhanced view just in the small lens area, which acts as a focus dynamically moved by the user.

6.3 Annotated RTI model

The last example of usage of our tool concerns the analysis of a damaged fresco. At the third-floor of *Casa dei Canonici* in Credarelo (Italy), in 1947 the painter Ettore d'Iseppi stripped from the wall a huge fresco-paint (1,40 x 11 meters) with the technique of the *strappo*, dividing the surface in five smaller two-meters-paints. The paint is unique in the iconography of the late XVI-century Verona for his size and his secular subject. Unfortunately, we have no photographs of the paint before the *strappo*. After that operation, the paint was probably so damaged to request the restoration by d'Iseppi. He added chalk-plaster where paint was lost during the *strappo*, and he heavily re-painted the whole surface to hide the damage. Now the painting is hard to read: the shapes are vague with no depth, the color is flat and faint, the paint layer is fragile and easy to detach. A preliminary survey is necessary to study the actual state of the paint and design a careful conservation job. The first task is to map the distribution of re-paintings and plasters before to plan the removal of the plaster and the stabilization of the paint layer.

The fresco was thus measured with the multispectral RTI technique, using the same setup as used in Sec. 6.1, using only the VIS and IR lighting. The relighting model was a pure Lambertian one.

The visualization technique has been a powerful tool in surveying. The possibility to relight the surface while enhancing normal variations through unsharp masking clearly permits to observe the variation of the surface shape due by presence of plasters and over-painted layers. This presence is also aided by the analysis of the




Figure 8: Annotated RTI model. The area outside the lens shows the re-lighted and annotated version of a damaged fresco detail. The annotation refer to regions where new materials, due to re-painting or plaster deposition, have been added. Inside the lens, a sharpened FCIR image makes the difference between pigments more evident.

appearance (in this case, pure albedo). In particular, the analysis was performed by producing a false-color-infrared image by merging the infrared channel with the visible channel (FCIR). An annotation overlay is added to identify the regions where new material (re-paint/plaster) has been added (see Fig. 8). As shown in the video, using a lens layer displaying enhanced FCIR on top of the visible layer with annotations makes it possible to analyze the areas where the restorer placed the annotation, perceiving through lens movement the subtle changes in shading and color due to plaster addition and changed materials.

7 Conclusions

We introduced a novel versatile approach for letting users explore detailed shape and material models integrated with structured spatially-associated descriptive information. Our approach extends the classic relighting viewers by combining and extending state-of-the-art results in several areas, in particular by generalizing fixed-view relighting interfaces to several representations, by supporting multi-scale visual annotations at several levels of abstractions, and by applying advanced focus-and-context visual exploration means to the field of shape and material analysis in cultural heritage, namely through the use of interactive lenses. Our results demonstrate how our approach can be applied to a variety of exploration use cases, ranging from casual inspection on mobile phones to exhibit-scale exploration on large-scale displays, a variety of objects, including statues and paintings, and a variety of representations, from analytical relighting models to decoupled shape and material data.

The presented approach is very versatile and can be applied to a variety of representations, including plain images, relightable images storing parametric representations of the reflectance fields, and explicit shape and material models in terms of spatially-varying normal and BRDF fields. Such data can come out of MLIC pipelines, as well as 3D scanning and photogrammetry pipelines. The multi-layered representation is easy to author. In particular, authors can use standard image editing tools for content preparation, leading to a simple but effective procedure to create spatially relevant rich visual information in forms of overlays at different levels of abstraction. The user interface, despite the additional capabilities, is very similar

to the relighting interfaces already popular in the CH field, thus reducing the learning curve.

The image-based data representation, and the related user interface, avoiding the complexity of full 3D exploration are simple and effective, but, at the same time, limits by design the application of the method to single-view exploration. The approach is thus meant to complement, rather than replace, fully 3D solutions. This image-based approach covers, however, a large portion of CH use-cases, where preferential view directions of cultural objects are, by design, very common.

Moreover, this work only targets the problem of model exploration with image/text overlays. Using other associated multimedia information (e.g., hyperlinks or video) and/or supporting very complex narratives are orthogonal problems not treated in this work, and can be integrated in the system with standard means (e.g., a canvas overlay). Similarly, the current proof-of-concept implementation uses single-image streaming, and is thus limited in terms of image size. A more scalable implementation would involve integrating a tiled pyramid image service, such as IIF [III19]. Finally, the current evaluation focuses mostly on demonstrating the capabilities of our approach in several uses cases. More work is required to objectively assess the effectiveness of the user interface in a variety of settings. Addressing this would require cognitive measures that are beyond the scope of the paper, and are an important avenue for future work.

Despite the limitations of our current implementation, including the loading of full-scale images instead of adaptive streaming, our results show that our solution is of immediate practical interest for many use cases. Our web-based implementation, moreover, simplifies a widespread adoption for both local and remote-viewing settings. A source code release of our reference implementation is available from the download web site at www.crs4.it/vic/download/. We believe that this could be a helpful instrument to support the CH community in the creation and dissemination of advanced contents on the web. We are currently working on expanding its scalability by using tiled images, and its applicability beyond CH, in particular for the visualization of simulation data at a urban scale.

Acknowledgments. The authors thank CRBC Sassari, Accademia delle Belle Arti di Verona, and Ormylia Foundation for the access to the artworks for the purpose of digitization. Stefano Marziali, from the Accademia delle Belle Arti di Verona, is also warmly thanked for his collaboration to the analysis and for helpful comments and suggestions. The project received funding from Sardinian Regional Authorities under projects VIGELAB and TDM (POR FESR 2014-2020 Action 1.2.2).

References

- [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. 2
- [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
- [Aio19] AIOLI: Aioli - a reality-based 3d annotation platform for the collaborative documentation of heritage artefacts, 2019. [Online; accessed 28-Aug-2019]. URL: <http://www.aioli.cloud/>. 3
- [BAB*16] BALSÀ RODRIGUEZ M., AGUS M., BETTIO F., MARTON F., GOBBETTI E.: Digital Mont'è Prama: Exploring large collections of detailed 3D models of sculptures. *ACM Journal on Computing and Cultural Heritage* 9, 4 (September 2016), 18:1–18:23. 3
- [BGM*09] BETTIO F., GOBBETTI E., MARTON F., TINTI A., MERELLA

- E., COMBET R.: A point-based system for local and remote exploration of dense 3d scanned models. In *The 10th International Symposium on Virtual Reality, Archaeology and Cultural Heritage* (October 2009), pp. 25–32. 3
- [BGMT13] BALSAL RODRIGUEZ M., GOBBETTI E., MARTON F., TINTI A.: Compression-domain seamless multiresolution visualization of gigantic meshes on mobile devices. In *Proc. ACM Web3D International Symposium* (2013), pp. 99–107. 3
- [BJM*15] BETTIO F., JASPE VILLANUEVA A., MERELLA E., MARTON F., GOBBETTI E., PINTUS R.: Mont'e Scan: Effective shape and color digitization of cluttered 3d artworks. *ACM JOCCH* 8, 1 (2015), 4:1–4:23. 8
- [CHI19] CHI: Cultural heritage imaging website, 2019. [Online; accessed-March-2019]. URL: <http://culturalheritageimaging.org>. 1, 2, 4
- [CLDS13] CALLIERI M., LEONI C., DELLEPIANE M., SCOPIGNO R.: Artworks narrating a story: a modular framework for the integrated presentation of three-dimensional and textual contents. In *Proc. ACM Web3D* (2013), pp. 167–175. 3
- [DHL19] DHLAB: RTI tools at DHLAB Basel, 2019. [Online; accessed-March-2019]. URL: <https://github.com/dhlab-basel/rti.js>. 2
- [DRS10] DORSEY J., RUSHMEIER H., SILLION F.: *Digital modeling of material appearance*. Elsevier, 2010. 2
- [FBKR17] FORNARO P., BIANCO A., KAISER A., ROSENTHALER L.: Enhanced RTI for gloss reproduction. *Electronic Imaging 2017*, 8 (2017), 66–72. 2
- [FD00] FALK H. J., DIERKING L. D.: *Learning from Museums: Visitor Experience and the Making of Meaning*. Rowman & Littlefield, 2000. 1
- [G*19] GIACHETTI A., ET AL.: apTool project, 2019. [Online; accessed-March-2019]. URL: <https://github.com/giach68/apTool>. 2
- [GCD*18] GIACHETTI A., CIORTAN I., 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. 4, 7
- [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
- [GIM19] GIMP: GNU Image Manipulation Program, 2019. [Online; accessed 13-June-2019]. URL: <https://www.gimp.org/>. 4
- [GTHD03] GARDNER A., TCHOU C., HAWKINS T., DEBEVEC P.: Linear light source reflectometry. *ACM TOG* 22, 3 (2003), 749–758. 2
- [GVH*07] GÖTZELMANN T., VÁZQUEZ P.-P., HARTMANN K., NÜRNBERGER A., STROTHOTTE T.: Correlating text and images: Concept and evaluation. In *Proc. Smart Graphics* (Berlin, Heidelberg, 2007), pp. 97–109. 3
- [HS17] HUI Z., SANKARANARAYANAN A. C.: Shape and spatially-varying reflectance estimation from virtual exemplars. *IEEE TPAMI* 39, 10 (2017), 2060–2073. 4
- [III19] IIIF: International image interoperability framework, 2019. [Online; accessed-June-2019]. URL: <https://iiif.io/>. 9
- [JD12] JANKOWSKI J., DECKER S.: A dual-mode user interface for accessing 3D content on the world wide web. In *Proc. WWW* (2012), pp. 1047–1056. 3
- [JH13] JANKOWSKI J., HACHET M.: A survey of interaction techniques for interactive 3D environments. In *Eurographics STAR* (2013). 2
- [KSZ*11] KUFLIK T., STOCK O., ZANCANARO M., GORFINKEL A., JBARA S., KATS S., SHEIDIN J., KASHTAN N.: A visitor's guide in an active museum: Presentations, communications, and reflection. *JOCCH* 3, 3 (2011), 11:1–11:25. 1
- [KUL19] KUL: PLD software KU-Leuven, 2019. [Online; accessed-March-2019]. URL: <https://portablelightdome.wordpress.com/software>. 2
- [LSC18] LI Z., SUNKAVALLI K., CHANDRAKER M.: Materials for masses: Svbrdf acquisition with a single mobile phone image. In *Proc. ECCV* (2018), pp. 72–87. 4
- [Mac15] MACDONALD L. W.: *Realistic visualisation of cultural heritage objects*. PhD thesis, UCL (University College London), 2015. 2
- [MAD*18] MOUTAFIDOU A., ADAMOPOULOS G., DROSOU A., TZOVARAS D., FUDOS I.: Multiple material layer visualization for cultural heritage artifacts. In *Proc. GCH* (2018), pp. 155–159. 2
- [MAG*12] MARTON F., AGUS M., GOBBETTI E., PINTORE G., BALSAL RODRIGUEZ M.: Natural exploration of 3D massive models on large-scale light field displays using the FOX proximal navigation technique. *Computers & Graphics* 36, 8 (2012), 893–903. 3
- [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). 2
- [NDM05] NGAN A., DURAND F., MATUSIK W.: Experimental analysis of BRDF models. *Rendering Techniques 2005*, 16 (2005), 2. 4
- [P*19a] PALMA G., ET AL.: WebRTI Viewer, 2019. [Online; accessed-March-2019]. URL: <http://vcg.isti.cnr.it/rti/webviewer.php>. 2
- [P*19b] PONCHIO F., ET AL.: Relight website, 2019. [Online; accessed-March-2019]. URL: <http://vcg.isti.cnr.it/relight/>. 2
- [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
- [PCD*15] POTENZIANI M., CALLIERI M., DELLEPIANE M., CORSINI M., PONCHIO F., SCOPIGNO R.: 3DHOP: 3D heritage online presenter. *Computers & Graphics* 52 (2015), 129–141. 1, 3, 5
- [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
- [PDC*19] PINTUS R., DULACHE T., CIORTAN I., GOBBETTI E., GIACHETTI A.: State-of-the-art in multi-light image collections for surface visualization and analysis. *Computer Graphics Forum* 38, 3 (2019). To appear. 1, 2
- [PDCS13] PALMA G., DESOGUS N., CIGNONI P., SCOPIGNO R.: Surface light field from video acquired in uncontrolled settings. In *Proc. Digital Heritage* (2013), vol. 1, pp. 31–38. 4
- [PGPG17] PINTUS R., GIACHETTI A., PINTORE G., GOBBETTI E.: Guided robust matte-model fitting for accelerating multi-light reflectance processing techniques. In *Proc. BMVC* (2017). 7
- [PPY*16] PINTUS R., PAL K., YANG Y., WEYRICH T., GOBBETTI E., RUSHMEIER H.: A survey of geometric analysis in cultural heritage. *Computer Graphics Forum* 35, 1 (2016), 4–31. 1
- [Ske19] SKETCHFAB: Sketchfab - publish and find 3d models online, 2019. [Online; accessed 13-June-2019]. URL: <https://sketchfab.com/>. 1, 7
- [TGK*14] TOMINSKI C., GLADISCH S., KISTER U., DACHSELT R., SCHUMANN H.: A Survey on Interactive Lenses in Visualization. In *EuroVis - STARS* (2014), Borgo R., Maciejewski R., Viola I., (Eds.). 3
- [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
- [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
- [WK15] WEINMANN M., KLEIN R.: Advances in geometry and reflectance acquisition (course notes). In *SIGGRAPH Asia 2015 Courses* (2015), ACM. 2
- [ZBL18] ZHENG Y., BLASCH E., LIU Z.: *Multispectral Image Fusion and Colorization*. SPIE Press, 2018. 4