

1 La Coupole: A SVBRDF measurement device for 2 large and non-planar objects

3 ANTOINE LUCAT,^{†1,3} PIERRE MÉZIÈRES,^{†2} FRANÇOIS MARGALL,²
4 LOUIS DE OLIVEIRA^{2,3}, MARJORIE PAILLET², ARNAUD TIZON⁴,
5 PIERRE BÉNARD^{2,3} AND ROMAIN PACANOWSKI^{2,*}

6 [†] Joint first authors

7 ¹ Institut d'Optique Graduate School,² Inria, France

8 ³ Université de Bordeaux, France,⁴ CNRS, France.

9 *romain.pacanowski@inria.fr

10 **Abstract:** We present *La Coupole*, a device for measuring SVBRDF of large and non-planar
11 objects whose ultimate goal is to reproduce the appearance of the object as it would be seen
12 under any lighting conditions and as perceived by the human eye at a distance of approximately
13 30 cm minimum; in particular, to provide virtual tools for visualization, inspection, and diagnosis
14 of ancient and fragile objects. With this in mind, we compare two BRDF representations
15 without prior assumptions (Spherical Harmonics and Radial Basis Functions) to reconstruct any
16 BRDF and a third one, a GGX microfacet-based model. The paper also provides details of the
17 geometric and radiometric calibration process, along with an in-depth study of their respective
18 uncertainties.

19 1. Motivation and Related Work

20 Digitization of real-world objects has many applications ranging from Cultural Heritage (preser-
21 vation, virtual exhibits) to the manufacturing industry where a virtual mock-up needs to be
22 validated before the mass production launches. The result of the digitization is a digital twin that
23 should look as realistic as the original object to a human observer. To produce digital twins, the
24 appearance of the original object must be acquired. Appearance is the combination of distinct,
25 but intertwined, intrinsic quantities: the shape and the optical properties of the artifact.

26 In this paper, we introduce a new measurement device, named *La Coupole*, that has been
27 conceived to acquire the appearance of large objects with the following two objectives: (i) be as
28 accurate as the human visual system, (ii) maintain metrological traceability. *La Coupole* combines
29 measurements from a 3D laser scanner, which acquires the shape of the object, and a DSLR cam-
30 era, which takes pictures in a controlled lighting environment to provide reflectance measurements.

31
32 Over the past twenty years, other acquisition systems have been developed to acquire the
33 appearance of objects. Since our acquisition device aims at measuring spatial variations of
34 reflectance, we purposely focus on the Spatially-Varying Bidirectional Reflectance Distribution
35 Function (SVBRDF). BRDF measurement devices (more or less automated) have been developed
36 since the 1960s. With the development of digital photography, measurement devices for BRDFs
37 have become imaging systems (e.g., [1, 2]). Recent and comprehensive overviews of BRDF
38 measuring instruments are presented by Weinmann and Kleiny [3] as well as Guarnera *et al.* [4].
39 It is worth noticing that, without imaging technologies, it would be too time-consuming to
40 measure an SVBRDF densely, even for small-sized objects. Mc Allister [5] and Sattler *et al.* [6]
41 were among the first to measure a nSVBRDF **on a planar surface** using a camera. Their systems
42 are quite slow due to the many mechanical movements that are needed to orient the material
43 sample. To overcome this limitation, recent systems (e.g., [7]) are doing acquisition in parallel
44 (a.k.a. multiplexing) by adding several cameras and light sources.

45 However, a major limitation remains for appearance acquisition. Objects (e.g., museum
46 artifacts, automotive parts) are rarely planar. Consequently, more recent systems (e.g., [8–11])

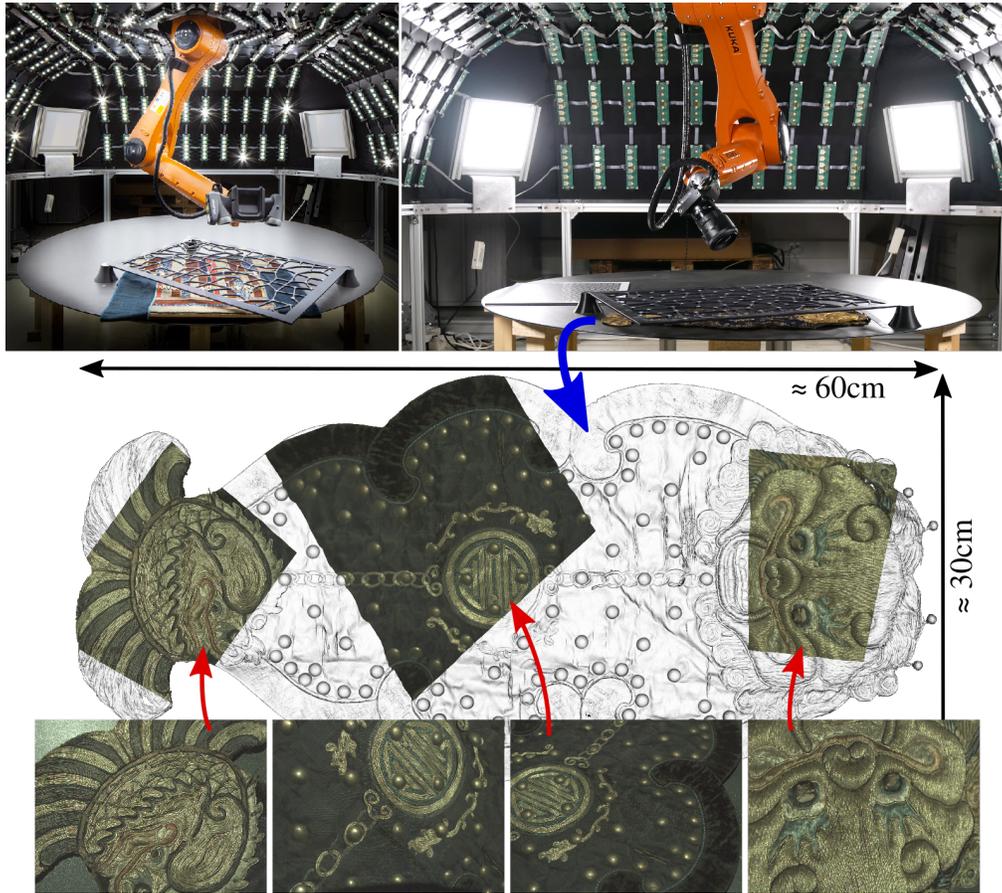


Fig. 1. La Coupole and its two measurement configurations. **Top left:** The 3D scanner is mounted on the robotic arm to acquire the shape (geometry) of the object. **Top right:** after the acquisition of the shape, a high-speed camera is mounted on the robotic arm to measure the SVBRDF with 1080 LEDs. The calibration process determines the acquisition parameters (e.g., the mesh position and camera position and rotation), which allows projecting each photo onto the mesh (**Bottom**).

47 combine structured lighting to acquire the shape of the object with multiple light sources and/or
 48 cameras. In particular, the solutions introduced by Schwartz [9, 10] and by Köhler [11] are
 49 related to ours since they rely on a dome to place multiple light sources and cameras. However,
 50 the resolution of the view direction they obtain is inferior to ours since the number of cameras is
 51 fixed (e.g., in [10]). In summary, previous dome-based systems are limited in terms of camera
 52 resolution and object size.

53 The object size limitation is mainly related to the choice of the camera field of view. It
 54 is a compromise between the spatial resolution of the measurement and the object size. In
 55 practical terms, for *La Coupole* imaging 1 m^2 at a resolution of $100 \mu\text{m}$ would require a sensor
 56 of 100 million pixels per view direction and an adequate lens to have sufficient depth-of-field,
 57 resulting in an overly expensive camera. The system proposed by Holroyd *et al.* [8] is quite
 58 original: two robotic arms are moving around the object, each arm carrying a camera and a video
 59 projector, which is used either to project a structured pattern or to illuminate the object. The
 60 main disadvantage of the system is its relative slowness, and the maximum size of the object is

61 also limited by the camera's field of view.

62 Recently, Garcès et al. [12] have introduced a dome system, with a wingspan of 30 cm,
63 specialized for a small textile patch of 11x11cm. The acquisition is done under 127 LED and is
64 constrained to 3 camera positions. As done in VarIS [13], they assume a planar geometry and a
65 SVBRDF model with a normal map, and optimize the geometry and the SVBRDF at the same time.
66 In other words, as usually accomplished by photogrammetric-based approaches(*e.g.*, [14, 15]),
67 shape and reflectance measurements are intertwined and tracking the uncertainties of the shape
68 or the SVBRDF is never established.

69 Our solution, to combine measurements of different equipment while keeping track of the
70 uncertainties, is presented in more detail in Section 4 where we give more insights into the
71 calibration procedure and in Section 6.1, where we provide a comprehensive analysis of the
72 radiometric and geometric uncertainties of the measured SVBRDF. Furthermore, instead of
73 assuming a BRDF model, we evaluate two mathematical representations (Spherical Harmonics
74 and linear Radial Basis Functions, Sec.5) to approximate the measurements and reconstruct
75 the BRDF as a continuous function, and compare (Section 6) them to a classical BRDF model
76 (GGX) based on the micro-facet theory [16]. Finally, we conclude by discussing (Section 7) the
77 different limitations and sketch future work directions in terms of calibration and uncertainty
78 reduction, as well as the data processing pipeline.

79 2. Acquisition Setup

80 The aforementioned constraints led us to conceive and build *La Coupole* (*cf.* Fig. 1) that consists
81 mainly in the following parts:

82 **The support** for the material sample is a disk of 1.48 m diameter and 0.3 mm thickness made of
83 black anodized Dibond. The Dibond has the advantage of being lightweight and the anodization
84 process makes the support more absorbing, thus reducing parasite lighting.

85 **The 6-axis robot,** KUKA Model Agilus kr10r1100sixx C, is used to move either the camera or
86 the 3D scanner. With a maximum reach of 1.1 m, the robot can reach all positions of the support.
87 According to the manufacturer, its pose repeatability is 30 μm .

88 **The 3D laser scanner,** HandyScan 700 (Creaform), is used to acquire the shape (geometry)
89 of the sample (see Fig. 1 left). The manufacturer announces a resolution of 50 μm (distance
90 between two points) with an accuracy of 30 μm (*i.e.*, the position of a point measured from the
91 scanner is known up to 30 μm). The speed of the scanner (480 000 points per second) makes
92 it very competitive compared to photogrammetric approaches: a 1 m² surface can be acquired
93 in 15 minutes. From a metrological point of view, the HandyScan precision is traced back to
94 the meter standard delivered by the French *National Institute of Metrology*. The scanner needs
95 to see retro-reflective stickers to orient itself. These stickers are pasted on a calibration target
96 positioned around or above the material sample.

97 **The lighting system** consists of 1080 LEDs and 4 flat white LED panels. For cultural heritage
98 applications, LEDs have the key advantage of not emitting any UV radiation, which would
99 degrade textile artifacts, for example. The four LED Panels (manufactured by Airis) generate
100 an almost perfect uniform lighting and serve during the calibration step to calibrate the spatial
101 sensitivity of the camera sensor (see below). The 1080 LEDs are distributed across 180 electronic
102 Printed Circuit Boards (PCBs) mounted on an aluminum structure that surrounds the support.
103 Each LED is 7 mm square and emits around 1800 lumens. An LED can be turned on for at
104 least 1 ms (electrical current rise constraint) and up to 2 seconds (heat dissipation constraint).

105 An electronic control box allows us to control each LED independently and also triggers the
 106 camera via a physical cable. This hardware triggering mechanism ensures good synchronization
 107 performance and high-speed acquisition.

108 **The camera** (Ximea Model CB120CG-CM) is used to acquire the optical properties of the
 109 material sample by taking one picture every time an LED is turned on. The camera can capture
 110 100 frames per second. Its RGB-CMOS sensor size is 4096×3072 pixels, where each pixel size
 111 is $5.5 \mu\text{m}$, with a good signal-to-noise ratio of 42 dB. The **lens** in front of the camera is a 35 mm
 112 Sigma HSM Art. The minimal focus distance is stated to be 30 cm. With a lens aperture set
 113 to at least $f/8$, we can achieve the one-degree angular resolution required to mimic the human
 114 eye. The combination of the Sigma lens and the Ximea sensor leads to an imaging system that
 115 acquires a $14.5 \times 19.3 \text{ cm}^2$ rectangular surface where each image pixel, at normal incidence,
 116 represents a square of $47 \mu\text{m}$ on a side in object space.

117
 118 In summary, *La Coupole* can acquire the appearance of an object whose projected wingspan
 119 does not exceed 1.4 m and whose height is at most 70 cm. The acquisition is a two-step process
 120 where, during the first step, the 3D scanner acquires the object shape and outputs a triangle mesh,
 121 whereas during the second step, for each position of the camera, an image is acquired each time
 122 an LED is switched on.

123 3. Measurement Equation

124 Radiometrically, the BRDF at position \mathbf{x} is defined [17] for a pair of infinitesimal outgoing \mathbf{o}
 125 and incoming \mathbf{i} directions, as the ratio of the reflected radiance and irradiance. Applied to our
 126 measurement context that uses a DSLR camera, which images the object, and an LED light
 127 source, which subtends a finite solid angle, our measurement equation of the spatially varying
 128 BRDF (SVBRDF) becomes:

$$BRDF(\mathbf{x}, \mathbf{o}, \mathbf{i}_k) = \frac{L(\mathbf{x}, \mathbf{o})}{E_k(\mathbf{x}, \mathbf{n}_x)} \quad \text{with} \quad L(\mathbf{x}, \mathbf{o}) \approx \frac{4(m+1)^2 \cdot f_{\#}^2 \cdot V_p}{\pi \cdot A_p \cdot t_e \cdot g(p)} \quad (1)$$

129 where: $L(\mathbf{x}, \mathbf{o})$ is the reflected radiance at point \mathbf{x} in direction \mathbf{o} ; $E_k(\mathbf{x}, \mathbf{n}_x)$ is the irradiance at
 130 point \mathbf{x} , with normal \mathbf{n}_x , due to the k -th LED light source, which main direction is \mathbf{i}_k (cf. Fig. 2).

131 Since the BRDF is a relative quantity, we measure the incident irradiance and the reflected
 132 radiance in the same radiometric space using the same camera. In the equation above, the pixel
 133 value is denoted V_p , A_p is the area of one pixel of the camera sensor, t_e is the exposure time, $f_{\#}$
 134 is the selected aperture, m is the optical magnification of the lens and $g(p) \in [0, 1]$ represents a
 135 flat-field correction that takes into account vignetting and other radiometric effects.

136 Without absolute calibration, $L(\mathbf{x}, \mathbf{o})$ unit is $\text{DN s}^{-1} \text{ m}^{-2} \text{ sr}^{-1}$ and $E_{i_k}(\mathbf{x}, \mathbf{n}_x)$ unit is necessarily
 137 $\text{DN s}^{-1} \text{ m}^{-2}$. We assume that an LED is a rectangular light source that emits uniformly with a
 138 cosine angular distribution, which is reasonable according to the LED manufacturer specifications.
 139 Under these assumptions, we model the irradiance $E_k(\mathbf{x}, \mathbf{n}_x)$ using Lambert's irradiance [18]:

$$E_k(\mathbf{x}, \mathbf{n}_x) = \frac{\Phi_k}{2 A_l} \langle \mathbf{G}_f(\mathbf{x}, \mathbf{l}_k), \mathbf{n}_x \rangle \langle \mathbf{n}_k, -\mathbf{i}_k \rangle \quad (2)$$

140 where: $\langle \cdot, \cdot \rangle$ represents the dot product between two vectors; \mathbf{l}_k is the position of the center of
 141 the k -th LED; A_l is the area of an LED; Φ_k is the power of the k -th LED; \mathbf{n}_k is the k -th LED
 142 surface normal (i.e., its orientation and $\langle \mathbf{n}_k, -\mathbf{i}_k \rangle$ corresponds to the cosine emission of the LED);
 143 $\mathbf{G}_f(\mathbf{x}, \mathbf{l}_k)$ is a geometrical and vectorial quantity defined in the [supplemental document](#) [19].

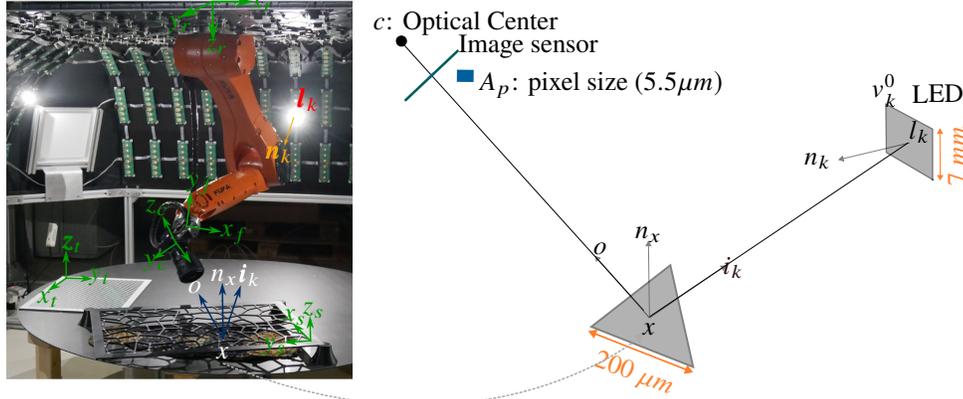


Fig. 2. **Left.** The different coordinate systems: (x_s, y_s, z_s) for the 3D scanner, (x_t, y_t, z_t) for the ChArUco board, (x_c, y_c, z_c) for the camera, (x_f, y_f, z_f) for the robot flange, and (x_r, y_r, z_r) for the robot base. **Right.** Illustration of the different mathematical notations used in the paper. The optical center c is the center of the camera coordinate system. The triangle, whose size is much smaller than an LED, belongs to the mesh acquired thanks to the 3D scanner.

144 Therefore, for a given observation direction and a given LED, our measurement equation becomes:

$$BRDF(\mathbf{x}, \mathbf{o}, \mathbf{i}_k) = \frac{8(m+1)^2 \cdot f_{\#}^2 \cdot V_p}{\pi \cdot A_p \cdot t_e \cdot g(p)} \cdot \frac{A_l}{\Phi_k \langle \mathbf{G}_f(\mathbf{x}, \mathbf{l}_k), \mathbf{n}_x \rangle \cdot \langle \mathbf{n}_l, -\mathbf{i}_k \rangle} \quad (3)$$

145 As shown by Eq. 3, one needs to calibrate for radiometric quantities (the power of an LED Φ_k ,
 146 the flat field function g) as well as geometrical quantities (e.g., $\mathbf{o}, \mathbf{x}, \mathbf{i}_k = \frac{\mathbf{l}_k - \mathbf{x}}{\|\mathbf{l}_k - \mathbf{x}\|}, \mathbf{l}_k, \mathbf{n}_k$).

147 4. Calibration Procedure

148 4.1. Geometrical Calibration

149 The definition of the SVBRDF requires the geometrical quantities $\mathbf{x}, \mathbf{o}, \mathbf{l}_k$ to be expressed in the
 150 same coordinate frame. In our setup, only the coordinate system (x_r, y_r, z_r) , attached to the
 151 base of the robot, is temporally invariant. To express all quantities in the robot base frame, we
 152 calibrate the geometrical quantities following three main steps:

- 153 1. Calibration of the optical center of the camera with respect to the robot base frame, which
 154 requires choosing a camera model and calibrating for intrinsic (effective focal length,
 155 distortions, focus distance, etc.) and extrinsic parameters of the camera.
- 156 2. Calibration of the position of each LED with respect to the robot base frame.
- 157 3. Calibration of the rigid transformation (called “mesh-to-robot” in the following) between
 158 the frame of the mesh, obtained from the 3D scanning process, and the robot base frame.

159 Once the camera, the LEDs, and the mesh are expressed in the same frame, for each pixel of
 160 an acquired image, we trace a ray to find the intersection with a triangle of the mesh, and then
 161 deduce the viewing \mathbf{o} and LED direction \mathbf{i}_k . The local surface normal, \mathbf{n}_x , is computed from the
 162 mesh vertices.

163 **Camera Intrinsic and Extrinsic Parameters calibration.** For each focusing distance, we
 164 calibrate the intrinsic and extrinsic properties by taking between forty and sixty images of

165 the ChArUco calibration board. With these images, we apply the method from Zhan [20],
166 implemented in OpenCV [21], for estimating intrinsic and extrinsic parameters of a thin-lens
167 camera model. For each image, we obtain the camera position in the ChArUco frame ($(\mathbf{x}_t, \mathbf{y}_t, \mathbf{z}_t)$)
168 in Fig. 2).

169 **Camera Position calibration.** To estimate the camera position in the robot frame, we apply a
170 simplified hand-eye (e.g., [22]) calibration procedure, yielding a positioning uncertainty of 300 μm .
171 The hand-eye calibration procedure takes advantage of the fact that the rigid transformation
172 between the robot flange center and the center of the camera frame (*i.e.*, the optical center) is
173 constant no matter the position of the robot.

174 **LED Position calibration.** Once the camera position is calibrated *w.r.t.* the robot frame, we
175 take 6 images of each switched-on LED and perform an optical triangulation by tracing camera
176 rays for each image, and finding the camera rays minimal intersection distance. As detailed in the
177 [supplemental document](#) [19], the positioning accuracy of each LED center is 700 μm whereas
178 the uncertainty associated with its normal is at worst 0.21° .

179 **Mesh-to-robot rigid transformation calibration.** To establish the transformation between
180 the scanner frame ($(\mathbf{x}_s, \mathbf{y}_s, \mathbf{z}_s)$) and the robot base, we introduce a new procedure that seeks to
181 establish the rigid transformation between the positions of the 3D scanner stickers defined in the
182 scanner frame, and the position of the stickers photographed by the Ximea camera. Our new
183 procedure is based on two algorithms: (i) an algorithm that operates in image space and detects
184 the gray stickers of the ChArUco board. This algorithm uses both the shape and color properties
185 of the stickers; (ii) an algorithm that operates in 3D space and establishes the best match between
186 the triangles formed by the stickers and the associated points. The uncertainty of the estimated
187 rigid transformation between frames ($(\mathbf{x}_s, \mathbf{y}_s, \mathbf{z}_s)$) and ($(\mathbf{x}_r, \mathbf{y}_r, \mathbf{z}_r)$) is at most 100 μm . The entire
188 geometric calibration procedure is detailed in the [supplemental document](#) [19].

189 4.2. Radiometric Calibration

190 As done classically, the camera is calibrated to account for vignetting by estimating the flat field
191 g . This is done by imaging the white LED panel and by fitting a mathematical model to filter
192 out the noise. The response curve of the sensor is also calibrated by taking multiple images for
193 different exposure times of a static scene. Our additional calibration results [19] show that our
194 camera exhibits a linear response.

195 The power of the LEDs (Φ_k in Eq. 3) is obtained from several photos of a spectralon target, by
196 averaging the pixels seen by the spectralon and illuminated by different LEDs. By using the same
197 camera between the spectralon calibration photos and the photos used to extract the SVBRDF
198 of the measured objects, Eq. 3 can be simplified. The equations used in practice for the data
199 processing pipeline (Sec. 5) and their derivations are given in the [supplemental document](#) [19].

200 5. Data Processing Pipeline

201 In this section, we focus on single-view and multi-light acquisition. In other words, we
202 only perform BRDF slice acquisition as the camera direction is fixed for every point on
203 the object. Section 5.1 describes the construction of our SVBRDF and the processing of
204 images to extract SVBRDF samples. Since the extracted data are angularly sparse, we need
205 a continuous representation to synthesize a new image for a non-acquired light direction.
206 Therefore, using different material samples, we analyze three types of representations: Spherical
207 Harmonics (Sec. 5.2), Radial Basis function (Sec. 5.3), and a microfacet-based BRDF model
208 using the GGX [23] distribution (Sec. 5.4).

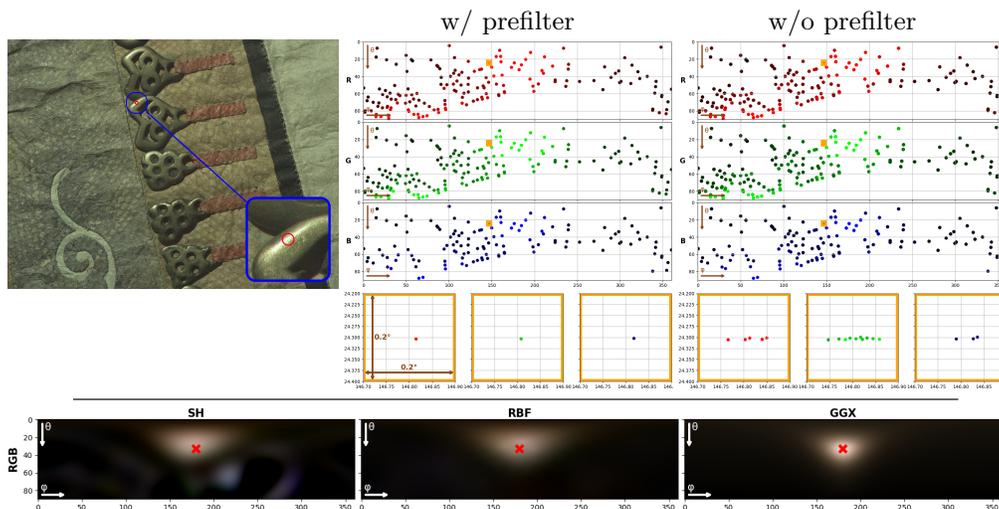


Fig. 3. Visualization of samples (latitude-longitude parameterization) from a triangle located within the red circle (top). Pre-filtering avoids redundancy of samples that are too close together angularly (around 0.1°). The sample numbers for the RGB channels are, respectively, 168, 170, 129 with the prefilter and 770, 1505, 325 without the prefilter. The RGB reconstruction using the SH and RBF shows that the specular peak in the measurements is slightly offset from the mirror direction (red cross) (bottom).

209 5.1. From Acquired Images to BRDF data

210 To extract BRDF samples from the images, each pixel (using RAW images encoded with GBRG
 211 Bayer pattern) is processed and converted into a BRDF value (Eq. 3). Hence, each pixel
 212 corresponds to a single RGB color channel, yielding three separate sets of samples (one per
 213 channel). Since this paper is limited to the study of fixed viewpoint, storing a triplet is sufficient:
 214 $\theta_i, \theta_o, \Delta\phi = |\phi_i - \phi_o|$, where \mathbf{i} corresponds to the light input direction and \mathbf{o} is the view
 215 output direction. The reflectance value f is also stored, and all these parameters are stored in a
 216 single-precision floating-point format. We remove too grazing samples ($\theta_i > 89^\circ$), this choice
 217 discussed in Section 6.4.

218
 219 We apply a 2% exposure threshold to reject pixels that are underexposed or overexposed.
 220 Thanks to the robot 3D model, a shadow mask for the robot arm is computed during the
 221 acquisition; hence, pixels belonging to these shadow areas are removed. However, shadows cast
 222 by the cables remain hard to predict. The pixel value threshold alleviates this problem and is
 223 discussed further in Section 6.4. In this paper, each dataset corresponds to N photos taken, at
 224 fixed exposure time, from the same viewpoint and only the LED position varies between each
 225 picture.

226 BRDF Samples are obtained by casting rays from photo pixels to identify the triangle that is
 227 "seen" by the pixel. Considering a single photo (a single LED lit) and since the triangles are
 228 small compared to the distance between the triangle and the camera, or between the triangle and
 229 the LED, this creates samples that are angularly close together when several pixels are linked to
 230 the same triangle (see the theoretical analysis in the [supplemental document](#) [19]). This angular
 231 redundancy is avoided by adding a sample pre-filtering step (Fig. 3).

232 **5.2. From BRDF data to Spherical Harmonics (SH) representation**

233 Any spherical function f can be represented by a weighted sum of spherical harmonics Y .

$$f(\theta, \phi) = \sum_{l=0}^{+\infty} \sum_{m=-l}^l \mathbf{F}_l^m Y_l^m(\theta, \phi) \approx \sum_{l=0}^n \sum_{m=-l}^l \mathbf{F}_l^m Y_l^m(\theta, \phi) . \quad (4)$$

234 where \mathbf{F} is a vector of coefficients. The index l plays the role of an angular frequency: higher
 235 values of l correspond to higher-frequency components on the sphere. This representation is
 236 exact only when an infinite number of basis functions is used. When using a finite number of
 237 SH, the reconstruction is bandlimited, with a band limit equal to n . In practice, this means
 238 that all SH components with $l > n$ are abruptly discarded. Such a sharp truncation in the SH
 239 frequency domain introduces a discontinuity in the spectrum, which produces oscillations in
 240 the reconstructed signal, commonly called *ringing*, analogous to the *Gibbs phenomenon*. The
 241 oscillations are stronger when the reconstruction band limit is small compared to the band limit
 242 of the projected signal [24]. From equation 4, the coefficient vector \mathbf{F} can be obtained by writing
 243 a minimization problem solved by least squares: $\|\mathbf{Y}\mathbf{F} - F\|^2 + \lambda \|\Gamma\mathbf{F}\|^2$, where Γ is the Thikonov
 244 matrix and λ control the amount of regularization ($\lambda = 0.1$ by default). Thus, the least squares
 245 system is solved for each triangle in the mesh.

246 **5.3. From BRDF data to Radial Basis Function (RBF) representation**

247 The RBF representation of our measurements is a linear combination of radial-basis functions,
 248 such as

$$BRDF_{RBF}(\mathbf{x}, \mathbf{o}, \mathbf{i}) = \sum_{n=0}^{N_{\mathbf{x}}} \omega_{\mathbf{x},n} K(\|\mathbf{o}, \mathbf{i}\|; (\mathbf{o}, \mathbf{i})_n) \quad (5)$$

249 where:

- 250 - $\|\cdot\|; \cdot\| : \mathbb{S}_+^2 \times \mathbb{S}_+^2 \rightarrow \mathbb{R}_+$ is the distance function between two bidirections, one (\mathbf{o}, \mathbf{i}) being the
 251 argument passed to $BRDF_{RBF}$, the other $(\mathbf{o}, \mathbf{i})_n$ being one bidirection of the interpolant;
- 252 - $K : \mathbb{R}_+ \rightarrow \mathbb{R}$ is the kernel, a radial basis (linear in this paper) function centered on a given
 253 bidirection $(\mathbf{o}, \mathbf{i})_n$;
- 254 - $\{(\mathbf{o}, \mathbf{i})_n\}_{n=0}^{N_{\mathbf{x}}}$ is the set of all bidirections used in the interpolant associated with the \mathbf{x}
 255 position;
- 256 - $\omega_{\mathbf{x}}$ is the vector of size $N_{\mathbf{x}}$ containing all the weights associated with each radial-basis
 257 function of the interpolant, each one centered on a given bidirection $(\mathbf{o}, \mathbf{i})_n$.

258 The distance function $\|\cdot\|; \cdot\|$, depending on its definition, allows us, on the one han,d to represent
 259 the bispherical topology $\mathbb{S}_+^2 \times \mathbb{S}_+^2 \rightarrow \mathbb{R}_+$ of the BRDF, but also some of its properties, such as
 260 bilateral symmetry (*i.e.*, with respect to the plane of incidence) or Helmholtz reciprocity.

261 To use the model, the weights need to be computed, which is done by solving the linear system
 262 $\mathbf{K}\boldsymbol{\omega} = \mathbf{f}$, where:

- 263 - \mathbf{K} is the kernel matrix of size $(N_{\mathbf{x}}, N_{\mathbf{x},m})$, each element defined as $K_{i,j} = K(\|\mathbf{o}, \mathbf{i}\|; (\mathbf{o}, \mathbf{i})_i)$;
- 264 - $\boldsymbol{\omega}$ is the vector of size $N_{\mathbf{x}}$ containing all the weights associated to every RBF bidirection;
- 265 - \mathbf{f} is the vector containing the input measured data associated with the position \mathbf{x} , of size
 266 $N_{\mathbf{x},m}$.

267 The [supplemental document](#) [19] details how non-negativity and extrapolation to large angles
 268 are handled.

269 **5.4. From BRDF data to Microfacet GGX BRDF Model**

270 Cook and Torrance [25] model the optical response as a set of Fresnel mirrors, called microfacets,
 271 and thus define specular reflections as the product of three functions: D , the distribution of

272 microfacet normals, F , the Fresnel term, and G , the shadowing term between microfacets. The
 273 GGX model by Walter *et al.* [23] introduces an analytical formulation to compute the terms D
 274 and G . The complete BRDF model used in this paper is the sum of a Lambertian reflection and
 275 the specular GGX model:

$$f_{\text{GGX}}(\mathbf{i}, \mathbf{o}) = \underbrace{\frac{k_d}{\pi}}_{\text{Diffuse}} + \underbrace{\frac{D_\alpha(\mathbf{h})F(\mathbf{o}, \mathbf{h}, \eta)G(\mathbf{i}, \mathbf{o}, \mathbf{h})}{4(\mathbf{n} \cdot \mathbf{i})(\mathbf{n} \cdot \mathbf{o})}}_{\text{Specular}} \quad (6)$$

276 where \mathbf{n} is the surface normal, $\mathbf{h} = \frac{\mathbf{i} + \mathbf{o}}{\|\mathbf{i} + \mathbf{o}\|}$ is the half vector, η the index of refraction (RGB
 277 value), α the roughness of the surface, and k_d is an RGB value, yielding a BRDF model with 7
 278 coefficients. The fitting is performed with the Levenberg-Marquardt algorithm implemented in
 279 the Eigen library [26] using an error metric similar to one introduced by Löw *et al.* [27].

280 6. Results

281 This section presents the results obtained for the acquisition and reconstruction of the SVBRDF.
 282 We first summarize the main sources of uncertainty affecting the BRDF measurements (Sec. 6.1).
 283 We then introduce the results obtained with single-view and multi-light datasets, where the
 284 reflectance function is estimated at the triangle level and represented using SH, RBF, and GGX
 285 models. The sampling characteristics of the reflectance measurements are analyzed in section
 286 6.2. Next, the three continuous representations are compared (Sec. 6.3). Finally, section 6.4
 287 examines the influence of the various parameters that guide the entire processing pipeline.

288 Each dataset used here corresponds to approximately 300-500 photos (where each corresponds
 289 to a single LED being lit) and approximately 500,000 to 1,000,000 triangles are visible in
 290 each viewpoint of the datasets. Complementary information on the datasets is provided in the
 291 [supplemental document](#) [19]. All tests conducted in this paper are performed on a computer with
 292 an Intel® Xeon(R) W-2155 CPU @ 3.30GHz × 20 cores, and data are stored on NVME.

293 6.1. Uncertainty Summary

294 We conducted a comprehensive uncertainty analysis (see also [supplemental document](#) [19]) of
 295 our BRDF acquisition system by propagating all identified uncertainty sources through the full
 296 measurement model using Monte-Carlo sampling [28]. This evaluation encompasses both the
 297 angular uncertainties in the incidence and observation directions, as well as the radiometric
 298 uncertainty in the BRDF values.

299 For angular uncertainty, we compute the variance of the directional spread induced by the
 300 finite size of the LEDs, triangles, and pixels. Across our complete dataset, both incidence and
 301 observation directions are found to be stable within approximately $\pm 0.5^\circ$, owing to the large
 302 LED-triangle and pixel-triangle distances compared to the small triangle size.

303 For radiometric uncertainty, we evaluate the full probability distribution of each BRDF
 304 measurement. We further analyse the global distributions of the asymmetric uncertainty bounds
 305 across the entire dataset. These distributions reveal that:

- 306 - in the best configurations, radiometric uncertainty can be as low as $[-1.06\%; +1.23\%]$;
- 307 - 50% of all measurements lie below $[-10.66\%; +14.00\%]$ at 68% confidence;
- 308 - 95% of all measurements lie below $[-15.23\%; +22.42\%]$ at 68% confidence.

309 A source-wise sensitivity analysis shows that most uncertainty sources are negligible. Two
 310 contributions dominate:

- 311 1. geometric uncertainty on triangle vertices (caused by scanner 3D accuracy limit of 30 μm),
 312 which alone accounts for a reduction of 57.22% to 74.01% of the total uncertainty when
 313 removed;

314 2. camera pixel-value uncertainty, contributing to a global reduction of 5.6 % to 12.5 % when
315 removed.

316 6.2. BRDF Samples Analysis

317 **Angular prefiltering** avoids sample redundancy (Sec. 5.1) (Fig. 3). Depending on the case,
318 not applying pre-filtering multiplies the number of samples by approximately 2 to 9. This has
319 an important cost in terms of storage and computation time, specifically for the continuous
320 representations computations (SH, RBF, or GGX) for all triangles, without improving the quality
321 of the results. We typically observe a greater spread of the samples in ϕ than in θ . This is
322 because the spread in θ is calculated in the space of the plane defined by the normal and the
323 sample direction, so in this plane, the distance between the LED and the triangle is very large
324 compared to the size of the triangle, which implies a small spread. For ϕ , the spread is calculated
325 in tangent space, so the projection of the center of the LED on the tangent plane can be very close
326 to the triangle, causing a large spread of samples. An illustration is provided in the [supplemental](#)
327 [document](#) [19], along with a more in-depth theoretical analysis of the angular spread of the
328 samples.

329 **Angular uncertainty.** The reconstruction of the specular lobe is slightly shifted (with SH and
330 RBF) from the expected main reflection direction (mirror direction) (Fig. 3). This is mainly due
331 to the fact that only an approximation of the surface is available, so the uncertainty about the
332 local surface coordinates is reflected in the direction of the samples, thus creating directional
333 uncertainty. Note that this discrepancy does not appear with GGX, because the model enforces
334 the main reflection lobe to be reconstructed in the mirror direction, which is not the case with SH
335 or RBF.

336 **Memory cost and timings.** In terms of memory usage, we typically observed sample sets
337 of around 5-15 GB with angular prefiltering (30-60 GB without) for meshes natively acquired
338 by 3D scanner. The meshes we process range from 1 million to 40 million triangles in total,
339 corresponding to file sizes of approximately 25 MB to 1 GB. For each individual viewpoint, the
340 portion of mesh involved typically contains 500,000 to 1 million triangles, amounting to about
341 9 MB to 17 MB. In terms of computational cost, the samples are exported in approximately
342 100-200 seconds using approximately 300 photos.

343 6.3. Comparisons of the three continuous representations

344 This section compares the results obtained from the three representations defined in Section 5.
345 As shown in Fig. 4, SH and RBF behave similarly during image reconstruction, whereas the GGX
346 model tends to produce more error. However, note that in our experiments, we fitted the three
347 models (SH, RBF, and GGX) to spherical slices of the BRDF. SH and RBF approximations can
348 only reproduce the fitted slice, since they are general-purpose function approximations, whereas
349 the parametric GGX model retains the physical formulation of the BRDF quantity. Therefore,
350 fitting the GGX model to a single slice permits generating an image from a new viewpoint,
351 whereas SH and RBF are not suited for this application case.

352 **Image generation process and comparison method** Once the SH, RBF, or GGX coefficients
353 have been obtained for each triangle in the dataset, a synthetic image, under measured or
354 novel lighting conditions, can be generated since we have an estimation of the BRDF for each
355 triangle of the mesh (Fig. 5) [29]. To fairly compare between the acquired photos and the
356 renderings generated with the SH, RBF, and GGX models, the images are generated at the same
357 initial resolution (4096×3072) as the photos and following the same Bayer pattern. Then the

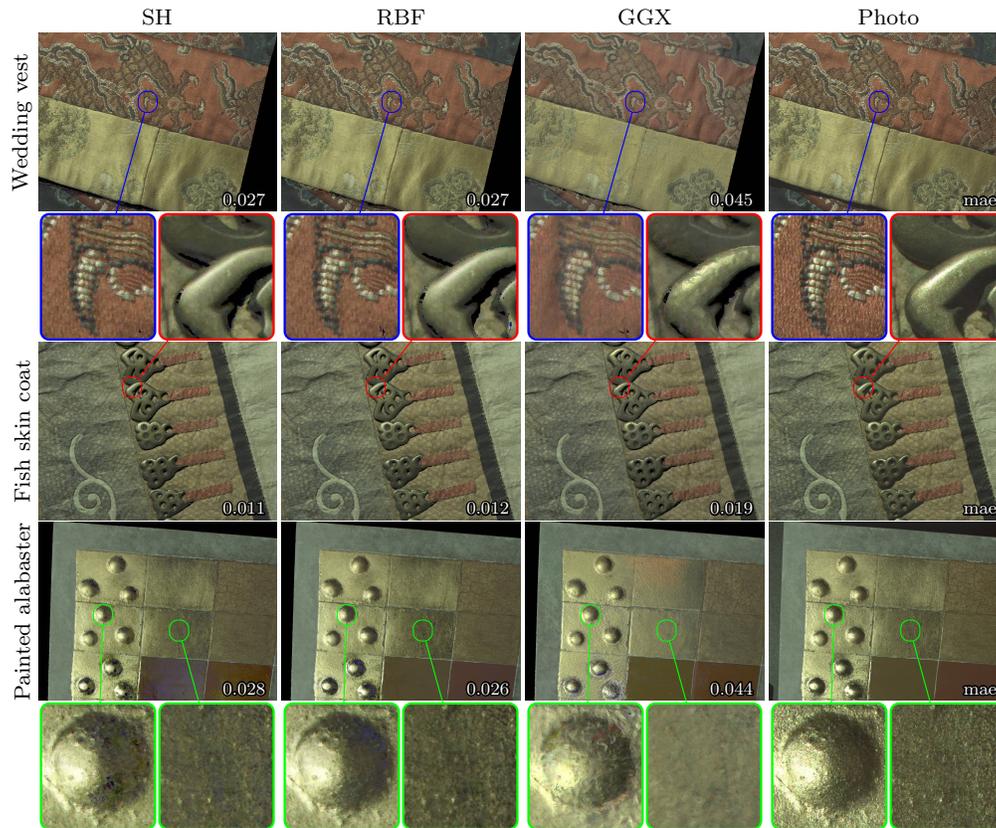


Fig. 4. Reconstruction of a photo using the different continuous representations: SH, RBF, and GGX. SH and RBF each use 100 coefficients per RGB color channel and produce very similar results, whereas GGX tends to produce more errors, which is particularly noticeable in the Mean Absolute Error (MAE)

358 photographs and renderings are demosaiced for each block of 4 pixels (with the GRBG Bayer
 359 pattern), a single RGB pixel is obtained by copying the red and blue and taking a single green
 360 value. In other words, the number of pixels is divided by 4. Therefore, the RGB images obtained
 361 can be directly compared to generate difference images or error metrics (Mean Absolute Error in
 362 this paper). In this way, each RGB channel is represented uniformly by the error metric, which
 363 would not be the case if the RAW images were compared (due to the presence of twice as much
 364 green in the Bayer pattern). Finally, to display the images, and only for display purposes, a
 365 gamma correction ($\gamma = 2.2$) is applied to the photos and renderings.

366 **Spherical Harmonics and Radial Basis Function** are used to approximate the measured
 367 BRDF values with N coefficients. Empirically, we found that $N = 100$ is a good compromise
 368 between computation time and accuracy (Figs. 3 and 4). To avoid *Gibbs phenomenon* induced
 369 by SH representation, we apply a window filter that removes strong visual artifacts at the price
 370 of a small increase of the MAE error. In contrast, approximation with RBF is more stable but
 371 has difficulty reproducing high frequencies accurately. This is illustrated in further detail in the
 372 [supplemental document](#) [19].

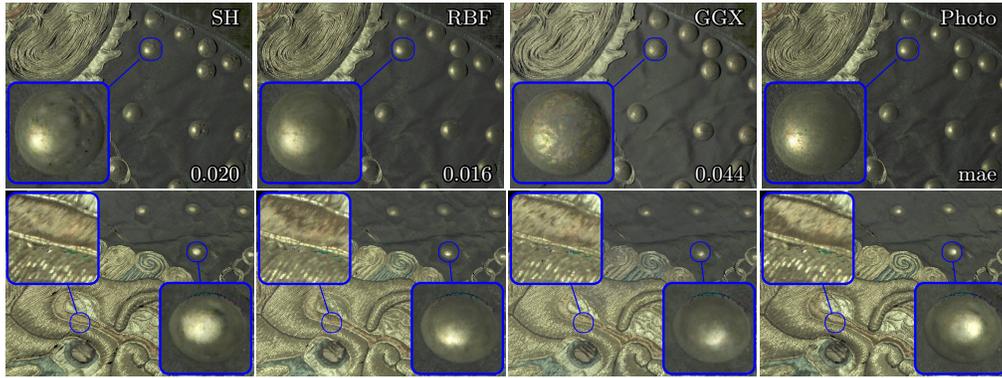


Fig. 5. **Top:** Ablation study where only 50% of the photos are used to approximate the BRDF. In that case, RBF give the lowest MAE and, as illustrated below, they are very stable, in terms of visual appearance, when computing new synthetic images under novel LED positions and orientations. **Bottom:** Four new synthetic images computed using the RBF representation under new (never measured) lighting conditions. [A video showcasing other synthetic results is available \[29\].](#)

373 **GGX BRDF Model** tends to generate more errors than the SH and RBF models. This is mainly
 374 due to angular uncertainties affecting the directions \mathbf{o} , \mathbf{i} and most importantly the surface normal
 375 \mathbf{n}_x . Unlike SH and RBF, which reconstruct the signal without relying on an explicit normal
 376 direction the GGX model is highly sensitive to the accuracy of \mathbf{n}_x . Even small uncertainties
 377 in the estimated normal propagate when transforming the measured samples into the local
 378 coordinate system defined by \mathbf{n}_x . This perturbs the relative angular positions of the samples,
 379 whereas the GGX model assumes precise, uncertainty-free orientations to perform optimally.
 380 As a consequence, GGX tends to make fabrics appear faded (*cf.* the wedding vest in Fig. 4) and
 381 metallic parts appear grainy (*cf.* the Fish skin coat).

382 **Memory cost and timings.** Considering a mesh of approximately 500,000 triangles, we
 383 generally observe in our datasets that computing the parameters of the SH representation was the
 384 fastest, whereas the GGX model was approximately 1.4 times slower and RBFs were twice as
 385 slow, but this depends greatly on the number of BRDF measures per triangle. On the other hand,
 386 for image reconstruction the GGX model is the fastest, while SH is twice as slow and RBF is
 387 six times slower. In terms of memory, the GGX model (with only 7 parameters) is obviously
 388 the least expensive, whereas the SH model is $3N/7$ times more expensive. The cost of RBF is
 389 similar to SH, with one exception: RBF also requires storing the directions of the radial-based
 390 functions. However, as done in this paper, if the directions are the same for all triangles, then this
 391 additional cost is negligible. To give an idea, using half a million triangles, GGX represents a
 392 cost of 14MB, while SH and RBF represent a cost of 600 MB for $N = 100$.

393 6.4. Fitting parameters sensitivity

394 The data processing is guided by several parameters that are more or less tunable by the user.
 395 First, a threshold on the theta angle is used to remove samples located at too grazing angles, which
 396 tend to generate outliers in the data. However, a threshold of 89° was sufficient in our experiment
 397 to remove most of the artifacts. The threshold on the pixel value removes noise and shadows from
 398 poorly exposed pixels. We noticed that the RBF and GGX models were particularly robust to this
 399 threshold, whereas this is not the case for SH model. We noticed the same trend for the number
 400 of photos given as input to the data processing (Fig. 5). Finally, the size of the triangles plays an

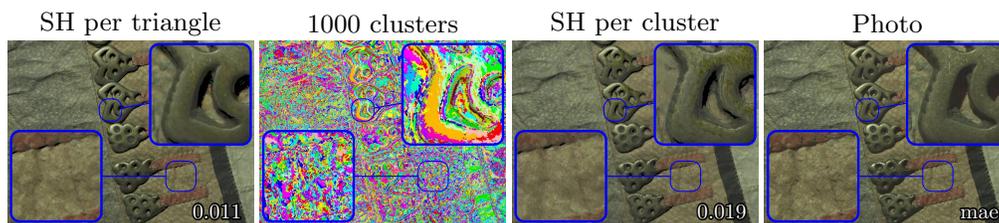


Fig. 6. Segmentation test applying k-means with 1,000 clusters on the SH coefficients per triangle. The samples are then concatenated by cluster to be projected onto the SHs in order to render with only 1,000 SH sets instead of 416,859 in this case. Despite this huge/bug compression factor induced by the clustering, the visual error, reported as MAE, remains very low (0.019 vs. 0.011).

401 important role in the accuracy of the renderings. Indeed, the smaller the triangles, the greater the
 402 spatial accuracy. On the other hand, triangles will tend to collect fewer BRDF samples. While
 403 this does not seem to be a particular problem for SH or GGX, it is for RBF, which does not give
 404 consistent results when the number of BRDF directions becomes extremely sparse. All these
 405 parameters are discussed and illustrated in more detail in the [supplemental document](#) [19].

406 7. Discussion and Future Work

407 **Uncertainties of BRDF radiometric values.** As explained in section 6, the positions of the
 408 vertices acquired by the 3D scanner are known with an accuracy of $30\ \mu\text{m}$, which is responsible
 409 for 57% of the total BRDF radiometric uncertainty. Using a more recent 3D scanner ($20\ \mu\text{m}$
 410 accuracy) would only reduce the overall BRDF median uncertainty from $[-10.66\%; +14.00\%]$
 411 to $[-8.11\%; +9.51\%]$. By combining this with a very low-noise camera, such as a Hamamatsu
 412 ORCA-Quest 2, our system could achieve a median measurement uncertainty of less than 3.6%.

413 **Per triangle BRDF data densification.** The BRDFs we obtain per triangle are particularly
 414 sparse (Fig. 3). Therefore, we would like to densify the BRDFs without necessarily enlarging
 415 the initial dataset (*i.e.*, without taking additional photos). Segmentation in BRDF space can be
 416 a way to increase data density per triangle. If two triangles are identified as having the same
 417 BRDF, the direction samples can be merged to create a more highly sampled BRDF. This would
 418 be particularly interesting for the *La Coupole* device, as it has a high spatial density of data
 419 (around half a million triangles seen by each photo). However, the uncertainty of the normal
 420 \mathbf{n}_x poses a major problem for merging BRDF directions from different triangles. Furthermore,
 421 automatic segmentation of materials based on sparse measurement data remains a challenge
 422 without using an a priori BRDF model. Nevertheless, preliminary tests show encouraging results,
 423 with acceptable visual quality achieved by drastically reducing the number of BRDFs (Fig. 6).

424 **Multi-view reconstruction** approaches are typically based on a single sensor, which greatly
 425 facilitates data calibration and data processing by combining geometry and materials [30, 31]
 426 at the price of sacrificing the tracking of the uncertainties of the SVBRDF. In our device,
 427 measurement data come from two different sensors: the camera ($50\ \mu\text{m}$ resolution) and the 3D
 428 scanner ($200\ \mu\text{m}$ resolution). Our calibration allows us to realign the photos from two different
 429 viewpoints and reproject them on the mesh with maximum precision of one to two millimeters,
 430 which corresponds to an offset of around 5-10 triangles (Fig. 7). As future work, we would like
 431 to add image-based information to approximate simultaneously the geometry and the SVBRDF
 432 while keeping track of the uncertainties. All recent literature on auto-differentiable tools for

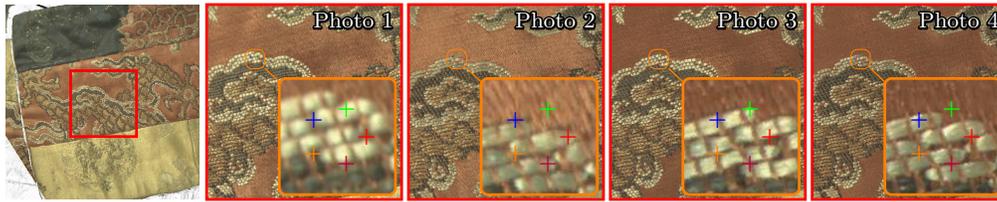


Fig. 7. Illustration of the problem of misalignment by projecting photos from different points of view on the mesh. In practice, for the test to be successful, the colored crosses must point exactly to the same features in the photos. Additional results (available in the [supplemental document](#) [19]) show that the offset is of the order of ten triangles (approximately 1-2 mm).

433 rendering could thus be leveraged in this application context [32] but within the limits of the
 434 assessed uncertainties above. The optimization process could be made more robust by taking
 435 into account other parameters, such as extrinsic camera parameters, at the cost of complexity.

436 Multi-view reconstruction raises other issues in terms of memory management and continuous
 437 representations. For large 3D objects, the cost to export the BRDF samples is around 10 GB of
 438 memory per viewpoint. This memory cost becomes redemptive for multiple viewpoints, especially
 439 when using RBF or SH representations that need *all* data to provide the best approximation
 440 of the SVBRDF. One way to circumvent this memory problem is to compute the RBF or SH
 441 representations only for a subset of triangles; however, the processing time will drastically
 442 increase since photos need to be read multiple times. In terms of mathematical representations,
 443 only SH need to be extended to 4D (*cf.* [33]) to handle multi-view reconstructions. However,
 444 the number of required coefficients will increase quadratically for SH. An interesting venue to
 445 explore would be to develop a processing pipeline that fetches the data from an image, treats it,
 446 and discards it immediately afterward. This will also require a novel and compatible SVBRDF
 447 representation.

448 **Progressive acquisition.** Currently, data acquisition and reconstruction are two separate
 449 processing pipelines that require a significant amount of memory. Furthermore, acquired data
 450 may exhibit redundancy. Therefore, merging the two processing pipelines to enable progressive
 451 acquisition would allow for better measurement guidance through on-the-fly data processing [34].

452 8. Conclusion

453 We have introduced *La Coupole*, a device that combines shape measurements from a 3D
 454 Scanner, with optical measurements, from a DLSR camera, to reproduce the appearance of
 455 real-world 3D objects. In conjunction with the heterogeneous equipment comprising *La Coupole*,
 456 we have also established a comprehensive calibration method and a statistical analysis of the
 457 geometric and radiometric uncertainties. Furthermore, we have also compared 3 different
 458 continuous representations (SH, RBF, GGX) for SVBRDF approximations and shown that they
 459 can reconstruct highly detailed photographs with excellent visual accuracy even in the case of
 460 novel lighting configurations.

461 9. Acknowledgments

462 This project was supported by the French national research agency project Material (ANR-15-
 463 CE38-0005), the region Nouvelle-Aquitaine through the project VESPAA, the European project
 464 23IND14 xDDiff, and the French Ministry of Culture through the project ANIS. Special thanks
 465 to Philippe Teulat for his help in designing and manufacturing mechanical parts.

466 **Disclosures.** The authors declare no conflicts of interest

467 See [Supplement 1](#) for supporting content.

468 References

- 469 1. S. R. Marschner, "Inverse rendering for computer graphics," Ph.D. thesis, Cornell University (1998).
- 470 2. W. Matusik, H. Pfister, M. Brand, and L. McMillan, "A data-driven reflectance model," *ACM Trans. Graph.* **22**,
- 471 759–769 (2003).
- 472 3. M. Weinmann and R. Klein, "Advances in geometry and reflectance acquisition (Course Notes)," in *SIGGRAPH*
- 473 *Asia 2015 Courses, SA 2015*, (Association for Computing Machinery, Inc, 2015).
- 474 4. D. Guarnera, G. C. Guarnera, A. Ghosh, *et al.*, "BRDF representation and acquisition," *Comput. Graph. Forum* **35**,
- 475 625–650 (2016).
- 476 5. D. K. McAllister, "A generalized surface appearance representation for computer graphics," thesis, Chapel Hill p.
- 477 129 (2002).
- 478 6. M. Sattler, R. Sarlette, and R. Klein, "Efficient and realistic visualization of cloth," in *Proceedings of the 14th*
- 479 *Eurographics Workshop on Rendering*, (Eurographics Association, 2003), EGRW '03, pp. 167–177.
- 480 7. V. Havran, J. Hošek, Š. Němcová, *et al.*, "Lightdrum—Portable light stage for accurate BTF measurement on site,"
- 481 *Sensors (Switzerland)* **17** (2017).
- 482 8. M. Holroyd, J. Lawrence, and T. Zickler, "A coaxial optical scanner for synchronous acquisition of 3D geometry and
- 483 surface reflectance," *ACM SIGGRAPH 2010 Pap. SIGGRAPH 2010* **29**, 1 (2010).
- 484 9. C. Schwartz, M. Weinmann, R. Ruiters, and R. Klein, "Integrated High-Quality Acquisition of Geometry and
- 485 Appearance for Cultural Heritage," in *The 12th International Symposium on Virtual Reality, Archeology and Cultural*
- 486 *Heritage VAST 2011*, (2011), pp. 25–32.
- 487 10. C. Schwartz, R. Sarlette, M. Weinmann, and R. Klein, "DOME II: a parallelized BTF acquisition system," *Workshop*
- 488 *on Material Appear. Model.* (2013) pp. 25–31 (2013).
- 489 11. J. Kohler, T. Noll, G. Reis, and D. Stricker, "A full-spherical device for simultaneous geometry and reflectance
- 490 acquisition," in *Proceedings of IEEE Workshop on Applications of Computer Vision*, (2013), pp. 355–362.
- 491 12. E. Garces, V. Arellano, C. Rodriguez-Pardo, *et al.*, "Towards Material Digitization with a Dual-scale Optical System,"
- 492 *ACM Trans. on Graph. (Proc. SIGGRAPH)* **4** (2023).
- 493 13. J. Baron, X. Li, P. Joshi, *et al.*, "VarIS: Variable Illumination Sphere for Facial Capture, Model Scanning, and
- 494 Spatially Varying Appearance Acquisition," in *Smart Tools and Applications in Graphics - Eurographics Italian*
- 495 *Chapter Conference*, (The Eurographics Association, 2023).
- 496 14. G. Nam, J. H. Lee, D. Gutierrez, and M. H. Kim, "Practical svbrdf acquisition of 3d objects with unstructured flash
- 497 photography," *ACM Trans. Graph.* **37** (2018).
- 498 15. M. Tetzlaff, "High-fidelity specular svbrdf acquisition from flash photographs," *IEEE Trans. on Vis. Comput. Graph.*
- 499 **30**, 1885–1896 (2024).
- 500 16. K. E. Torrance and E. M. Sparrow, "Theory for off-specular reflection from roughened surfaces*," *J. Opt. Soc. Am.*
- 501 **57**, 1105–1114 (1967).
- 502 17. F. E. Nicodemus, J. C. Richmond, J. J. Hsia, *et al.*, *Geometrical Considerations and Nomenclature for Reflectance*
- 503 (National Bureau of Standards, 1977).
- 504 18. J. Arvo, "Applications of irradiance tensors to the simulation of non-lambertian phenomena," in *Proceedings of the*
- 505 *22nd Annual Conference on Computer Graphics and Interactive Techniques*, (1995), SIGGRAPH '95, p. 335–342.
- 506 19. A. Lucat, P. Mézières, F. Margall, *et al.*, "La coupole: A svbrdf measurement device for large and non-planar objects:
- 507 Supplemental document," https://people.bordeaux.inria.fr/romain.pacanowski/lacoupole/lacoupole_supp.pdf (2025).
- 508 20. Z. Zhang, "A flexible new technique for camera calibration," *Tech. Rep. 11*, Microsoft Research (2000).
- 509 21. O. Authors, "Opencv pinhole camera model," https://docs.opencv.org/4.x/d9/d0c/group__calib3d.html (2025).
- 510 22. R. Horaud and F. Dornaika, "Hand-Eye Calibration," *Technical report* (1994).
- 511 23. B. Walter, S. R. Marschner, H. Li, and K. E. Torrance, "Microfacet models for refraction through rough surfaces." *Render. techniques* **2007**, 18th (2007).
- 512 24. P.-P. Sloan, "Stupid spherical harmonics (sh) tricks," *Game Dev. Conf.* **9**, 320–321 (2008).
- 513 25. R. L. Cook and K. E. Torrance, "A reflectance model for computer graphics," *ACM Trans. Graph.* **1**, 7–24 (1982).
- 514 26. G. Guennebaud, B. Jacob *et al.*, "Eigen v3," <http://eigen.tuxfamily.org> (2010).
- 515 27. J. Löw, J. Kronander, A. Ynnerman, and J. Unger, "Brdf models for accurate and efficient rendering of glossy
- 516 surfaces," *ACM Trans. Graph.* **31** (2012).
- 517 28. ISO/IEC, "Propagation of distributions using a Monte Carlo method," *Tech. Rep. 98-3:2008/Suppl 1:2008* (2008).
- 518 29. A. Lucat, P. Mézières, F. Margall, *et al.*, "La coupole : relighting," <https://youtu.be/sfWFZEBEWP4> (2025).
- 519 30. B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis, "3d gaussian splatting for real-time radiance field rendering." *ACM Trans. Graph.* **42**, 139–1 (2023).
- 520 31. M. Tetzlaff, "High-fidelity specular svbrdf acquisition from flash photographs," *IEEE Trans. on Vis. Comput. Graph.*
- 521 **30**, 1885–1896 (2023).
- 522 32. C. Zhang, B. Miller, K. Yan, *et al.*, "Path-space differentiable rendering," *ACM Trans. Graph.* **39**, 143:1–143:19
- 523 (2020).
- 524 33. C. Soler, M. M. Bagher, and D. Nowrouzezahrai, "Efficient and accurate spherical kernel integrals using isotropic
- 525 decomposition," *ACM Trans. Graph.* **34** (2015).
- 526 34. R. Wiersma, J. Philip, M. Hašan, *et al.*, "Uncertainty for svbrdf acquisition using frequency analysis," in *Proceedings*
- 527 *of the SIGGRAPH Conference Papers*, (2025), pp. 1–12.
- 528
- 529