

Using a Robotic Arm for Measuring BRDFs

Rasmus Ahrenkiel Lyngby, Jannik Boll Matthiassen, Jeppe Revall Frisvad,
Anders Bjorholm Dahl, Henrik Aanæs

Technical University of Denmark
<http://eco3d.compute.dtu.dk/>

Abstract. The measurement of a bidirectional reflectance distribution function (BRDF) requires a purpose-built instrument. We ease this requirement by presenting a relative method for measuring a BRDF using a multipurpose robotic arm. Our focus is on the alignment of the system to perform accurate camera positioning and orientation. We use a six degrees of freedom robotic arm to move a camera on a hemisphere surrounding a flat material sample. Point-like light sources, fixed on a quarter circle arc, sequentially illuminate the sample from different directions. The resulting images are used to reconstruct the material BRDF. We limit ourselves to tristimulus (RGB) isotropic BRDF acquisition.

Keywords: Robot arm · BRDF · camera pose estimation.

1 Introduction

Material appearance is important when using image analysis for quality assurance [12], and one way to represent the appearance of materials is by means of their reflectance properties. To understand the reflectance properties of materials, we need devices for measuring them. A gonioreflectometer is a device for measuring reflectance distribution functions [3,6,8]. Such an instrument is usually an expensive, purpose-built machine that does only one job. While gonioreflectometers have important applications in computer graphics and computer vision, purchasing one is a substantial investment (around €150k). To pursue scientific discoveries within these areas while limiting the amount of cash tied up in specialized equipment, we propose to use a multipurpose robot arm (around €40k) as a gonioreflectometer. Our point is not so much the lower cost of the robot arm, as other equipment needs to be added to make it useful as a gonioreflectometer. Our point is to describe a setup where the robot arm can be used for various other research purposes as well.

Many six-axis robots, including the one we use, have a low absolute position precision. They do, however, have very high repeatability. We exploit this feature in our BRDF acquisition by measuring the angles between camera and light directions relative to a 3D artifact using pose estimation. We avoid measuring the light source radiant exitance and the absolute camera intensity response, as well as exact distances between camera and surface and between surface and light source. This is done by normalizing observed pixel intensities using Spectralon, which is an almost perfectly diffuse (Lambertian) material.

2 Previous Work

Several precise, robot-based gonioreflectometers exist today [5,1,16,4,15,8]. At Physikalisch-Technische Bundesanstalt (PTB), for example, they have developed a robot-based gonioreflectometer with a five-axis robot holding the sample and a homogeneous sphere radiator mounted on a ring-shaped rotation stage for illumination [5,4]. This system relies on calibrated incident radiance and is generally comprised of more purpose-built components than our setup. Baribeau et al. [1] also use a five-axis robot for holding the sample and others [16,15] use a six-axis robot for the same task. These systems employ a ring-shaped rotation stage for holding either a light source [1,15] or a spectroradiometer [16]. Instead of holding the sample, we let the robot hold a camera, as we can then use the robot in a wider range of applications [18]. The facility at Fraunhofer [8] is similar to ours but more elaborate. They have the robotic arm hanging upside down holding a spectrometer, the sample on a rotation table, and a light source moving on a quarter circle. This enables full 5D (wavelength and two directions) measurement of BRDFs. This is however a one-purpose instrument, and the authors do not discuss how they accurately position and orient the detector.

Techniques for positioning and orientation of a camera mounted on a robot’s end effector are available when the camera observes an object with known features [19]. However, feature extraction from the acquired images is oftentimes the main source of error [10]. Feature extraction is typically based on 3D features from a CAD model [19] or features in a 2D calibration object [10,2]. Our samples vary in shape and a 2D calibration grid is not observed well from all the viewing angles that we need. We therefore introduce a 3D artifact and rely on the high repeatability of the robot. We consider our method for ensuring that the robot accurately positions and orients a camera our main contribution. To enable the same robot to be used for multiple purposes, we built a simple light arch that can be moved away from the system when the robot with camera mount is needed for other purposes. Moreover, as stated above, our method is relative and does not rely on calibration of light emission or camera response.

3 Instrumentation

We use a six-axis ABB IRB 1600 10/1.45 industrial robot capable of carrying a payload of 10 kg. It has a repeatability of 50 μm , which means that the sensors are guaranteed to arrive at the same poses within 50 μm . Fig. 1 (left) shows the robot in the experimental setup. The camera we use is a Point Grey Grasshopper3 9.1MP RGB camera.

As seen in Fig. 1, the robot is painted black and our setup is completely surrounded by an enclosure also painted black. This shields our setup from any illumination from the surroundings and prevents most internal reflections from disrupting our experiments (a co-bot surrounded by blackout curtains is another option). For lighting, we built an arc-shaped light source covering 90° vertical angle with an array of light sources placed 7.5° apart, see Fig. 1. This setup provides us with control over sensor pose and lighting conditions.



Fig. 1. Left: robotic arm holding a camera in a lightproof enclosure. Right: light arch and camera system measuring the BRDF of an aluminum laptop.

4 BRDF Acquisition

The BRDF characterizes material appearance by describing the change in surface reflectance for varying view and light directions. In terms of geometry, the BRDF is a four-dimensional function. For many materials, this can be reduced to three dimensions by assuming isotropic reflectance. This means that the reflectance does not change with the orientation of the material. We make this assumption and thus exclude materials such as brushed metals and some fabrics.

The BRDF is defined by the ratio of an element of reflected radiance dL_r to an element of irradiance dE [11]:

$$f_r(\theta_i, \phi_i; \theta_o, \phi_o) = \frac{dL_r(\theta_o, \phi_o)}{dE(\theta_i, \phi_i)},$$

where θ_i, ϕ_i denote the direction of incoming light in spherical coordinates, θ_o, ϕ_o denote the direction of outgoing light in spherical coordinates, L_r is the reflected radiance, and E is the irradiance of the sample. The ratio is taken in the limit where only one direction of incidence is considered. The BRDF obeys Helmholtz reciprocity, which means that $f_r(\theta_i, \phi_i; \theta_o, \phi_o) = f_r(\theta_o, \phi_o; \theta_i, \phi_i)$.

In principle, it is easy to measure the BRDF using a gonioreflectometer, by observing a flat homogeneous material sample under all possible view and illumination combinations. Unfortunately, dense sampling of BRDFs requires many samples to accurately capture the appearance of a material, which is a comprehensive task to perform, even with an automated robot.

4.1 Isotropic BRDF Capture

By equipping a robotic arm with a camera and combining this setup with our arc-shaped point-light array, we can semi-densely sample light reflections off of a flat material sample for various configurations of incoming and outgoing light directions. See Fig. 2 for an illustration of the measurement setup.




Fig. 2. Our surface reflectance measuring system consisting of robot, light arch, table, sample stand, and sample. The latter is illustrated by the small gray box.

We move the camera around a sample placed under the light arch while imaging the reflectance. As the BRDF obeys Helmholtz reciprocity, one image provides two reciprocal configurations of light and camera. In addition, since we limit ourselves to isotropic BRDFs, consecutive rotation around the surface-normal of both incoming and outgoing rays does not change the BRDF. Thus, we need only move the camera during acquisition.

The spacing between lights in the light array and the spatial resolution of the robotic arm limits how densely we can sample the surface reflectance. The robot has a non-uniform grid of reachable positions, but, in general, the spacing between two reachable locations is sub-millimeter. The light arch has a radius of 1000 mm and has a fixed set of bulbs with a spacing of 7.5° . This configuration provides a compromise between sample rate and measurement time for semi-dense sampling within a reasonable time frame.

Ideally, lights would be point sources and the camera’s field of view would be only a fraction of a degree. However, in practice, we have to make compromises. To make the lights more point-like, we keep the sample-to-light distance at around three times the camera-to-sample distance. For the camera, we use the mean value of received light within a broader field of view to estimate reflected radiance as this eases the camera tracking.

We move the camera on the surface of a hemisphere with the sample surface at the center. The robot’s arm-length resulted in an optimal hemisphere radius of 350 mm. Due to reciprocity, an isotropic BRDF is mirror symmetric around the vertical plane of the light arc. In other words, the sample’s appearance at the right-hand side is identical to its appearance on the left-hand side. Thus, slicing the hemisphere by this plane, we need only sample half of it. The camera path is defined in spherical coordinates with a resolution of 7.5° in both azimuth and

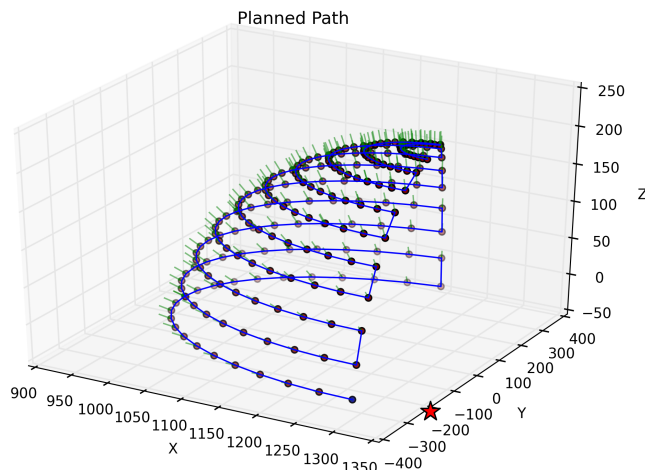


Fig. 3. The camera’s trajectory in the robot’s coordinate frame. The red star indicates the sphere center, the purple dots indicate sample positions, and the green lines indicate the camera orientation. Sampling starts at an elevation of 7.5° . To avoid kinematic singularities, the trajectory’s center is not at the origin.

elevation. The Cartesian equivalent of a given spherical camera position is then

$$\mathbf{x}_{ij} = \mathbf{x}_0 + r (\sin \theta_i \cos \phi_j, \sin \theta_i \sin \phi_j, \cos \theta_i),$$

where θ_i is the inclination angle of sample row i and ϕ_j is the azimuthal angle of sample column j . The position \mathbf{x}_0 is the center of the sphere, and r is the radius. The orientation of the camera is then

$$(\beta_{ij}, \gamma_{ij}) = \left(\frac{\pi}{2} + \theta_i, \phi_j \right),$$

where β_{ij} is the camera pitch and γ_{ij} is the camera yaw. For each sample point, we compute the Cartesian coordinates and orientations and store them in a matrix. We feed these to the robot one by one at run time. The camera trajectory with sample positions is illustrated in Fig. 3.

The orientation of the robots coordinate system must be taken into account. In practice, we let points on the robot’s x -axis correspond to an azimuth of 0° . In order to operate the robot within its working area, we had to position the light arc along the robot’s y -axis. This is illustrated in Fig. 4 (left).

To avoid absolute measurements of radiance emitted from light sources and radiance received in camera pixels, we use Spectralon (patented and manufactured by Labsphere). This is a very white material exhibiting $> 99\%$ diffuse reflectance, and the reflection is fairly direction-independent (except at very grazing angles). Thus, with very little loss, Spectralon uniformly spreads incident flux across the hemisphere surrounding the point of incidence. This enables us to measure the BRDF for a given material and pair of incoming and outgoing

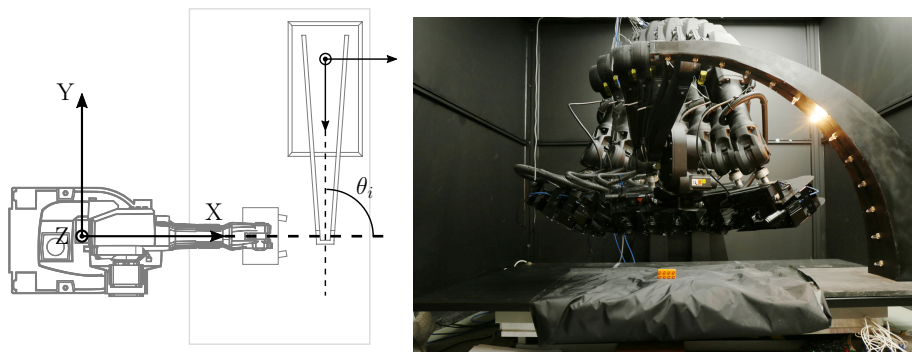


Fig. 4. The position of the light arc in the robot’s coordinate system (left). The sample is at the intersection of the two dashed lines, which indicates the center of the spherical robot path. The azimuth of the incoming light is here $\theta_i = 90^\circ$. We merged a sequence of images as the robot sweeps across one row of the path (right) to illustrate the measurement process for a single light source.

directions by calculating the ratio of intensities received by the camera when observing material $I^{(M)}$ and Spectralon $I^{(S)}$ from the same direction:

$$f_r(\theta_i, \theta_o, \phi_o) = \frac{I^{(M)}(\theta_i, \theta_o, \phi_o)}{\pi I^{(S)}(\theta_i, \theta_o, \phi_o)}. \quad (1)$$

As this is a relative measure, the unit used to measure the intensity is irrelevant. It could be radiant exposure, but it might as well be pixel intensity in $[0, 1]$ or in $[0, 255]$. This concludes the description of the different parts of our BRDF measurement system. The work flow is defined in Table 1 and the process is illustrated in Fig. 4 (right).

The robot is guided in its own coordinate system (Fig. 4, left) by specifying positions of its end effector. However, we would like to directly specify the camera’s position and orientation, which requires that we find the spatial transform from end effector to camera. This transformation has to be known rather precisely, as we want the camera to keep pointing as closely as possible toward the same surface point. We deal with this issue in the following section.

5 Hand-Eye Calibration

The position of the robot’s end effector is defined for a standard tool. For BRDF measurement, we instead have a camera that we need to position relative to a sample. We thus need to accurately determine the six degrees of freedom (DoF) transformations between the robot’s standard tool and the camera. Having these transformations, we can determine the position of the camera within the robot’s coordinate system by $\mathbf{P}_c = \mathbf{P}_{rt} \mathbf{T}_{ct}$, where \mathbf{T}_{ct} is the robot tool to camera transformation, \mathbf{P}_{rt} is the six DoF standard tool position, and \mathbf{P}_c is the resulting six DoF camera position.

Table 1. Work flow to measure a BRDF.

Step	Time	Notes
1 Tool transform	1[h]	Ideally, we would only have to estimate the tool transform once. However, as the tool changes with temperature, humidity, and wear and tear, it is good practice to re-estimate it regularly.
2 Aligning Spectralon to arch	5[min]	The surface of the Spectralon sample has to be carefully aligned with the center of the semicircle formed by the light arch.
3 Spectralon measurement	3[h]	As with the tool transform, this process can be omitted. However, it should be conducted at least once for every new measurement day.
4 Aligning material to arch	5[min]	Replace the spectralon sample with the material sample. Make sure the surface of the material align exactly with the surface of the Spectralon.
5 Material measurement	3[h]	
6 BRDF calculation	5[min]	Use Equation 1 to calculate the BRDF of the material.
Total	7[h]:15[min]	
Another sample	3[h]:10[min]	

We estimate the transformation \mathbf{T}_{ct} using hand-eye calibration, which is based on a set of relative camera and tool motions. Specifically, the solution \mathbf{X} to the system $\mathbf{AX} = \mathbf{XB}$ is found, where \mathbf{A} represents the relative motions of the camera, \mathbf{B} represents the corresponding relative motions of the tool, and \mathbf{X} is the transformation between the camera and the tool. Several algorithms exist to solve this problem. We use the one by Liang and Mao [7]. Using a Kronecker product and singular value decomposition (SVD), we find a closed-form solution to the rotation part of the transformation. This becomes an initial guess in a least-squares optimization used to find the optimal rotation. We then find the translation using a least-squares optimization based on the optimal rotation. This procedure combines the performance of the closed-form solutions with the accuracy and noise invariance of the least-squares solutions, thus providing a good estimate of the transformation in a short amount of time.

To practically solve this transformation, a checkerboard is positioned in the working area. Then, 20 images of the checkerboard are captured from different locations with the camera. For each location, we record two positions: (1) the position of the robot tool with respect to the robot’s base and (2) the position of the camera with respect to the checkerboard. The latter is obtained through camera calibration [20]. We manually select the 20 locations by jogging the robot tool so that the locations span the hemisphere over the checkerboard. We vary all six degrees of freedom of the robot and ensure that the entire checkerboard is within the field of view of the camera. The saddle points of the checkerboard must be easily identifiable in all images. Thus, the viewing angle of the checkerboard cannot be too steep. The process is illustrated in Fig. 5.




Fig. 5. The hand-eye calibration process for two robot/camera positions. The coordinate systems of robot (R) and checkerboard (C) are marked. The red dots represent the position of the robot tool relative to its coordinate system, and the blue dots represent the camera position relative to the checkerboard. The pair of relative transformations are calculated based on these positions. We use multiple pairs to do the calibration.

5.1 Photometric Optimization

Although the hand-eye calibration provides a good estimate of the tool transformation of the camera, we wish to further refine it through optimization. We do this to ensure that the uncertainty in camera position is as close as possible to be within the size of a pixel.

From the hand-eye calibration, we get an estimate of the location of the camera's optical axis relative to the end effector. We now capture two images of a checkerboard: before and after rotating the camera 90° around its estimated optical axis. The second image is digitally counter-rotated by 90° . The difference between these two images assesses the correctness of the estimated camera axis. If the difference image is dark with no apparent edges, the optical axis is estimated to within a pixel. Suppose $\mathbf{M}_{0^\circ} \in \mathbb{N}^{n \times n}$ is the square-cropped first image and $\mathbf{M}_{90^\circ} \in \mathbb{N}^{n \times n}$ is the second image, then the transformation error is $E_{\mathbf{T}_{ct}} = \sum |\mathbf{M}_{0^\circ} - \mathbf{M}_{90^\circ}^T|$. This error depends on \mathbf{T}_{ct} . Fig. 6 shows examples of difference images. We use the BFGS algorithm [14] to minimize $E_{\mathbf{T}_{ct}}$ by optimizing \mathbf{T}_{ct} , which has six dimensions in total.

5.2 Pose Estimation

While the robot has a repeatability of $50 \mu\text{m}$, the systematic error of its absolute position is much larger. As we need a more precise measurement of the camera pose in order to assign the measured BRDF values to the correct angles, we constructed the 3D calibration object shown in Fig. 7 for camera pose estimation.




Fig. 6. Difference of orthogonal images acquired by the camera before and after a roll of 90° . When taking the difference, the second image is digitally counter-rotated by 90° . We show the result before (left) and after refinement to minimize the error (right). This ensures an accurate tool offset.



Fig. 7. 3D artifact used for camera pose estimation.

We manually position our artifact underneath the light arch. The robot then traverses the camera trajectory while capturing an image of the artifact at each sample position. Precision of the artifact positioning is unimportant as long as the artifact is visible in all camera views. After pose estimation, all positions and directions will be in the coordinate system of the artifact. The sphere centers of the artifact are annotated manually in the first image. A least-squares optimization then minimizes the error between the back-projection of the spheres onto the image plane and the segmented spheres by adjusting the estimated camera pose. As the relative transformation between two positions on the planned trajectory can be calculated roughly, the estimated sphere centers can be projected into subsequent images and used as an initial guess. Finally, the absolute camera poses for all positions, relative to the pose at the first position, are computed from the poses estimated relative to the artifact.

In a second step, the stereo camera system observes the light bulbs through a mirror at positions from the sample trajectory with now known poses. The pose of each light bulb is then estimated using triangulation.

Pose estimation was found to increase the quality of the BRDF measurements significantly. Better highlights were observed, due to the increased angle precision. Note that the sample itself was not pose estimated. The vertical position of the material surface was mechanically calibrated to precisely align with the base of the calibration artifact.

6 Representation and Sample Interpolation

A common format for storing BRDFs is the MERL binary format [9]. Data is stored in a non-linear voxel-grid using the 3D Rusinkiewicz parametrization $(\theta_h, \theta_d, \phi_d)$ [17], allowing a fine data-resolution around specular highlights. Most BRDF tools and physically based renderers support this format, why it is a convenient way of representing the measured BRDFs.



Fig. 8. Photo of a yellow toy brick made from ABS polymer. A point light source is positioned directly opposite the camera. Two diffuse light sources, positioned above and below the camera, illuminate the brick from the front. The specular highlight due to the point light is clearly visible on the top surface. The image is slightly blurred due to use of a wide aperture. Black tape was added to remove glares from the edges.

Although our BRDFs consists of hundreds to thousands of samples, they are still somewhat sparse compared to the MERL cubes’ $90 \times 90 \times 180 \approx 1.5 \cdot 10^6$ values. To convert the observations to this much finer resolution format, we thus impose two different strategies: an unbiased and a biased approach.

In our unbiased interpolation strategy, all values in the MERL cube are obtained through 1-Nearest-Neighbour (1NN) interpolation. Euclidean distances between neighbours are however not necessarily meaningful in a 4D angular domain. Instead, we perform the 1NN search in a 6D Euclidean space based on the 3D direction vectors toward light and camera. To ensure a meaningful search, we calculate the vectors in the same hemisphere of a Cartesian coordinate system. We keep the light direction vector on the semicircle through the zenith of the hemisphere and place it in the first three dimensions of the 6D space. The camera direction vector is placed in the other three dimensions. From 1NN interpolation in this space, a MERL cube is filled with information by copying the exact measurements to their neighbouring voxels without introducing any new values. The resulting MERL cube accurately reflects the raw measurements, and can thus be used to visualize the data. The cost is that the information is very discrete and does not utilize the full potential of the MERL representation. We therefore only use this interpolation for visualization of the measurements, and not as actual BRDFs. It should also be noted that by applying nearest neighbor interpolation, the property of energy conservation is likely to be violated as $\iint f_r(\theta_i, \theta_o, \phi_o) \cos \theta_o \sin \theta_o d\theta_o d\phi_o$ may become greater than 1.

In our biased interpolation strategy, we use the principal component based method of Nielsen et al. [13] to reconstruct the missing information in the MERL cube using a prior learned from the MERL dataset [9]. This method creates a smooth and continuous interpolation of the MERL volume with high quality highlights. It is thus a plausible estimate of the missing values, but confined to the data-variation in the MERL database. More info on this biased interpolation method is available in our previous work [18].

7 Results and Discussion

We use the yellow toy in Fig. 8 as a challenging test case. The brick is made of acrylonitrile butadiene styrene (ABS) polymer with a very smooth specular surface. The measured BRDF is illustrated in Fig. 9. The unbiased interpolation (leftmost) visualizes the raw data. The left hemisphere is a reversed mirroring




Fig. 9. Visualization of the BRDF measured from the toy brick shown in Fig. 8. We rendered a sphere using the unbiased interpolated BRDF (left) and a sphere and a dragon using the biased interpolated BRDF (middle and right).

of the right. The dark area originates from an impossible camera position (the camera can not be at the same position as the light source). The measurements look consistent, with no apparent discontinuities other than the dark spot. We rendered a sphere and a dragon using the biased interpolation (middle and right-most). This reconstruction method struggles to correctly represent the specular highlight, as indicated by the small dark spot visible exactly in the middle of the specular highlight of the sphere. This might indicate that the highlight should be sampled more densely [13]. It could also stem from a bias in the MERL dataset. Reconstructions of other materials with less defined highlights were much better: BRDFs of white and gray cardboard and white cloth were measured with our system in our previous work [18]. This work did not present the calibration and unbiased interpolation techniques covered here, but it does provide a larger selection of measured BRDF data delivered by the system and validated to some extent through comparison of rendered images with photographs.

In some tests, we found that the relative intensity method (Eq. 1) can fail for materials that absorb all wavelengths inside either of the red, green or blue color channel. The measured intensity of $I^{(M)}$ in Eq. 1 then drops below the camera’s noise floor, which results in a noisy BRDF. Furthermore, Spectralon is nearly but not perfectly diffuse. When observed at grazing angles, $I^{(S)}$ tended to drop below the noise floor, especially for the blue channel. The denominator in Eq. 1 was then mostly noise tending to zero. In a case where $I^{(M)}$ was also mostly noise, the effect of noise divided by noise was seen as a purple tint in the BRDF at grazing angles of observation. The purple color was due to the blue channel being more susceptible to noise, which is most likely due to the quantum efficiency of the blue channel being slightly lower than that of the red and green channels. For most of our tests, this was not a problem. We see it as a limitation that applies to certain types and colors of materials. More powerful light sources or more sensitive cameras may be employed to address this in future work.

References

1. Baribeau, R., Neil, W.S., Côté, É.: Development of a robot-based gonioreflectometer for spectral BRDF measurement. *J. of Modern Optics* **56**(13), 1497–1503 (2009)

2. Bobby, R.A., Saha, S.K.: Single image based camera calibration and pose estimation of the end-effector of a robot. In: IEEE International Conference on Robotics and Automation (ICRA 2016). pp. 2435–2440 (2016)
3. Erb, W.: Computer-controlled gonioreflectometer for the measurement of spectral reflection characteristics. *Applied Optics* **19**(22), 3789–3794 (1980)
4. Höpe, A., Atamas, T., Hünerhoff, D., Teichert, S., Hauer, K.O.: ARGon³: “3D appearance robot-based gonioreflectometer” at PTB. *Review of Scientific Instruments* **83**(4), 045102 (2012)
5. Hünerhoff, D., Grusemann, U., Höpe, A.: New robot-based gonioreflectometer for measuring spectral diffuse reflection. *Metrologia* **43**(2), S11 (2006)
6. Li, H., Foo, S.C., Torrance, K.E., Westin, S.H.: Automated three-axis gonioreflectometer for computer graphics applications. *Opt. Eng.* **45**(4), 043605 (2006)
7. Liang, R.h., Mao, J.f.: Hand-eye calibration with a new linear decomposition algorithm. *Journal of Zhejiang University SCIENCE A* **9**(10), 1363–1368 (2008)
8. Martínez, M.L., Hartmann, T.: Multispectral gonioreflectometer facility for directional reflectance measurements and its use on materials and paints. In: Target and Background Signatures IV. Proceedings of SPIE, vol. 10794, p. 107940V (2018)
9. Matusik, W., Pfister, H., Brand, M., McMillan, L.: A data-driven reflectance model. *ACM Transactions on Graphics (SIGGRAPH 2003)* **22**(3), 759–769 (July 2003)
10. Meng, Y., Zhuang, H.: Autonomous robot calibration using vision technology. *Robotics and Computer-Integrated Manufacturing* **23**(4), 436–446 (2007)
11. Nicodemus, F.E., Richmond, J.C., Hsia, J.J., Ginsberg, I.W., Limperis, T.: Geometrical considerations and nomenclature for reflectance. Tech. Rep. NBS Monograph 160, National Bureau of Standards (US) (October 1977)
12. Nielsen, J.B., Eiriksson, E.R., Kristensen, R.L., Wilm, J., Frisvad, J.R., Conradsen, K., Aanæs, H.: Quality assurance based on descriptive and parsimonious appearance models. In: Material Appearance Modeling (MAM 2015). pp. 21–24 (2015)
13. Nielsen, J.B., Jensen, H., Ramamoorthi, R.: On optimal, minimal BRDF sampling for reflectance acquisition. *ACM Transactions on Graphics (SIGGRAPH Asia 2015)* **34**(6), 186:1–186:11 (October 2015)
14. Nocedal, J., Wright, S.J.: *Numerical Optimization*. Springer, second edn. (2006)
15. Obein, G., Audenaert, J., Ged, G., Leloup, F.B.: Metrological issues related to BRDF measurements around the specular direction in the particular case of glossy surfaces. In: Measuring, Modeling, and Reproducing Material Appearance 2015. Proceedings of SPIE, vol. 9398, p. 93980D (2015)
16. Rabal, A.M., Ferrero, A., Campos, J., Fontecha, J.L., Pons, A., Rubiño, A., Corróns, A.: Automatic gonio-spectrophotometer for the absolute measurement of the spectral BRDF at in- and out-of-plane and retroreflection geometries. *Metrologia* **49**(3), 213–223 (2012)
17. Rusinkiewicz, S.: A new change of variables for efficient BRDF representation. In: *Rendering Techniques '98 (EGWR 1998)*. pp. 11–22. Springer (1998)
18. Stets, J.D., Dal Corso, A., Nielsen, J.B., Lyngby, R.A., Jensen, S.H.N., Wilm, J., Doest, M.B., Gundlach, C., Eiriksson, E.R., Conradsen, K., Dahl, A.B., Bærentzen, J.A., Frisvad, J.R., Aanæs, H.: Scene reassembly after multimodal digitization and pipeline evaluation using photorealistic rendering. *Applied Optics* **56**(27), 7679–7690 (September 2017)
19. Wilson, W.J., Hulls, C.C.W., Bell, G.S.: Relative end-effector control using cartesian position based visual servoing. *IEEE Transactions on Robotics and Automation* **12**(5), 684–696 (1996)
20. Zhang, Z.: A flexible new technique for camera calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **22**(11), 1330–1334 (November 2000)