

Efficient Geometric Calibration for Mirror-Based Light-Field Projection

Kouki Fujimoto^{1,*}, Toshiyuki Amano¹

¹Wakayama University, 930, Sakaedani, Wakayama-shi, Wakayama, Japan

Abstract

In the field of Spatial Augmented Reality (SAR), methods have been studied for altering the color, material appearance, and lighting environment of real object surfaces through projection. In particular, light-field projection, which controls light rays from multiple directions, is effective for presenting material appearances and lighting effects that vary depending on the viewing direction. A multiple-mirror-based light-field projection system, which combines multiple projectors with a mirror array, can generate multidirectional light rays using a limited number of projectors. However, the geometric correspondence among the projectors, mirrors, and stage becomes complicated. In previous approaches, these correspondences were obtained manually, which required a considerable amount of calibration time and depended on the operator's skill.

In this study, we propose an efficient geometric calibration method for a multiple-mirror-based light-field projection system that uses a camera to acquire the correspondence between projector pixels and world coordinates on the stage. In the proposed method, Gray-code patterns are projected while dividing the projection area for each mirror, thereby avoiding the superposition of reflected light from multiple mirrors and obtaining the correspondence between projector pixels and camera pixels. Furthermore, by using a coordinate-defining projector to acquire the correspondence between world coordinates and camera pixels and integrating it with the pixel correspondence for each mirror, the proposed method enables projection based on world coordinates while reducing the influence of geometric distortion included in camera images. As an application example, we generated target projection images based on BRDF data and demonstrated viewpoint-dependent material appearance projection using the acquired geometric correspondences.

Keywords

Spatial Augmented Reality, Light Field Projection, Geometric Calibration, BRDF

1. Introduction

Spatial Augmented Reality (SAR) is a technology that presents visual information by directly projecting images onto the surfaces of real-world objects using projectors, without requiring users to wear special devices. In conventional SAR, many studies have focused on manipulating the appearance of objects by assuming the target surface to be diffuse and altering its color or texture. For example, Okazaki et al. [1] reproduced high-quality metallic gloss on diffuse surfaces by estimating surface normals using structured light and photometric stereo. Aliaga et al. [2] achieved high-resolution appearance editing for objects with arbitrary shapes and colors using superimposed projections from multiple projectors. Amano [3] also proposed real-time material appearance manipulation using a projector-camera feedback system. These methods are effective for appearance manipulation similar to texture mapping onto object surfaces; however, because they are mainly based on optical texture mapping for diffuse surfaces, they have limitations in presenting material appearances that vary depending on the viewing direction.

For materials such as metals, glossy surfaces, and structural colors, brightness and color change according to the viewing direction due to the surface reflectance properties of the object. To reproduce such viewpoint-dependent material appearances, it is necessary not only to present colors according to positions on the target surface, but also to control the distribution of light rays traveling from each surface point toward different viewing directions. In other words, for BRDF-based material appearance

APMAR'26: The 18th Asia-Pacific Workshop on Mixed and Augmented Reality, Aug. 3-5, 2026, Taipei, Taiwan

*Corresponding author.

✉ s276221@wakayama-u.ac.jp (K. Fujimoto); amano@wakayama-u.ac.jp (T. Amano)

id 0009-0001-8912-9382 (K. Fujimoto); 0000-0003-4146-1375 (T. Amano)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

presentation, it is important to consider not only the projection position but also the incident direction of light rays and the viewing direction.

A straightforward approach to realizing viewpoint-dependent presentation is to estimate the observer's position using eye or head tracking and switch the projected image according to the viewpoint. However, this approach is basically limited to presentation for a single viewpoint and is therefore difficult to apply to situations in which multiple observers view the projection simultaneously. Another possible approach is to combine SAR projection with an optical see-through HMD, as in HySAR [4], or to present viewpoint-dependent textures on glossy surfaces using motion capture and shutter glasses [5]. These approaches can present information according to the user's viewpoint, but they require wearable devices or additional calibration for the viewer. This undermines one of the advantages of SAR, namely that it does not require wearable devices.

In contrast, light-field projection controls light rays according to their positions and directions in space, enabling the presentation of material appearance changes caused by viewpoint movement without requiring observers to wear special devices. Amano et al. [6] proposed viewpoint-dependent appearance manipulation using multiple projector-camera feedback systems. Murakami et al. [7] achieved material appearance manipulation based on reflectance analysis, and Ohsumi et al. [8] extended this framework to anisotropic reflection. Kimura et al. [9] demonstrated BRDF-based material appearance presentation using a multiple-mirror-based light-field projection system. These studies show that controlling directional light distributions is effective for reproducing viewpoint-dependent appearances such as gloss and structural colors.

On the other hand, in a multiple-mirror-based light-field projection system, geometric calibration is essential for obtaining the correspondence among projector pixels, projection positions, and light-ray directions. In particular, it is necessary to determine not only where each projector pixel reaches on the stage, but also through which mirror the corresponding light ray travels and from which direction it enters the stage. However, in a light-field projection system, light rays reflected by multiple mirrors are superimposed on the stage. Therefore, structured-light methods commonly used in general projector-camera systems cannot be directly applied. For this reason, conventional systems have relied on manual acquisition of the correspondences among projector pixels, mirrors, and the stage. This manual calibration requires a considerable amount of time and depends on the operator's skill and visual alignment accuracy.

In this study, we propose a geometric calibration method for a multiple-mirror-based light-field projection system by combining Gray-code projection with separated projection regions for each mirror and world-coordinate estimation using a coordinate-defining projector. In the proposed method, Gray-code patterns are first projected only onto the target mirror region. This avoids the superposition of reflected light from multiple mirrors and enables the acquisition of correspondences between projector pixels and camera pixels. Next, the correspondence between world coordinates on the stage and camera pixels is acquired using a coordinate-defining projector, and this correspondence is integrated with the pixel correspondences obtained for each mirror. This makes it possible to efficiently acquire the correspondence between projector pixels and world coordinates, which has conventionally been obtained manually.

The contributions of this study are summarized as follows:

1. We propose a Gray-code-based geometric calibration method for a multiple-mirror-based light-field projection system by using projection regions separated for each mirror.
2. We introduce a coordinate-defining projector and present a method for uniformly acquiring correspondences among projector pixels, camera images, and world coordinates.
3. We apply the proposed method to a light-field projection system and confirm its effectiveness in a real system by performing BRDF-based presentation using the acquired correspondences.

2. Related Works

2.1. Appearance Manipulation using Light-Field Projection

In the field of Spatial Augmented Reality (SAR), many studies have investigated methods for manipulating the appearance of real objects by projecting images onto them using projectors. Early studies mainly focused on projection onto diffuse surfaces, whereas recent studies have advanced toward material appearance presentation that considers the reflectance properties of target objects, such as the bidirectional reflectance distribution function (BRDF). Amano et al. [6] proposed light-field feedback using multiple projector-camera feedback systems, enabling independent appearance manipulation for each viewpoint. Murakami et al. [7] achieved viewpoint-dependent material appearance manipulation through light-field projection based on reflectance analysis. Ohsumi et al. [8] reproduced anisotropic material appearances, such as the hairline finish of metals, by projection based on an anisotropic reflection model. Furthermore, Nagase et al. [10] realized high-quality projection onto moving objects by dynamically controlling the focal planes of multiple projectors.

These studies demonstrate that advanced material appearance representation can be achieved by precisely controlling light rays emitted from projectors. On the other hand, to realize light-field projection, geometric calibration is necessary to accurately determine the positions and directions of light rays.

2.2. Mirror-Based Light-Field Optical Systems

Zhou et al. [11] and Takeuchi et al. [12] realized high-density light-field generation and lighting control using lens arrays and lenticular lenses. To present viewpoint-dependent material appearances and reproduce complex lighting environments, it is necessary to control light rays from multiple directions; that is, light-field projection is required. To achieve this using a limited number of devices, various optical systems employing mirrors and lens arrays have been proposed.

Han et al. [13] proposed a method for acquiring the bidirectional texture function (BTF) of object surfaces from multiple directions using a single camera based on the principle of a kaleidoscope. Mukaigawa et al. [14] realized confocal imaging with a hemispherical field of view by using a polyhedral mirror called a Turtleback Reflector.

Kubo et al. [15] generated virtual lighting environments using a multiple-mirror-based light-field projection system, and Kimura et al. [9] realized viewpoint-dependent BRDF presentation using the same type of system.

The multiple-mirror-based light-field projection system targeted in this study is positioned within this line of research.

2.3. Structured-Light-Based Calibration

In general projector-camera systems, structured-light methods are widely used to acquire correspondences between projector pixels and observed images. Salvi et al. [16, 17] systematically reviewed three-dimensional measurement and calibration methods using structured-light patterns. In particular, the Gray-code method is robust against illumination changes and noise, and is widely used for acquiring correspondences between projector pixels and camera pixels. Moreno et al. [18] proposed a highly accurate projector-camera calibration method that treats the projector as an inverse camera. These methods assume that each observed pixel corresponds to a single projected light ray, allowing correspondences to be acquired efficiently and accurately.

In systems using multiple mirrors, light rays reach the target object via mirrors. Therefore, to determine where each projector pixel reaches in space, strict correspondence mapping that considers the position, orientation, and reflection properties of each mirror is required. When this correspondence is acquired manually, the amount of work increases as the number of projectors and mirrors increases, making simplification of the calibration procedure an important issue. Therefore, in this study, we

introduce a camera-based method for acquiring correspondences and aim to improve the efficiency of the calibration process in a light-field projection system.

3. Overview of the Light-Field Projection System and Conventional Calibration

The multiple-mirror-based light-field projection system targeted in this study was developed in previous studies [9, 15] and has been used for lighting-environment reproduction and BRDF-based appearance presentation. In this study, we propose an efficient geometric calibration method for acquiring correspondences between projector pixels and world coordinates on the stage for this existing light-field projection system. Therefore, this section describes the existing system on which the proposed method is based, including its configuration, previous applications using the system, the conventional procedure for generating projection images, and the conventional geometric calibration procedure. The contents described in this section are based on previous studies, whereas the newly proposed calibration method is presented in Section 5.

3.1. System Configuration

To present color changes caused by wide viewing-direction variations, as in BRDF-based appearance reproduction, it is necessary to superimpose multiple projections from different directions onto the same point on the stage. Kimura et al. [9] employed the light-field projection system shown in Fig. 1. This system consists of four ultra-short-throw projectors (RICOH PJWX4125, 1280×800 , 3,000 lm) and a 3×3 planar mirror array installed above the projection surface, consisting of nine square mirrors with a side length of 300 mm. The mirror array is placed approximately 850 mm above the projection surface, and the light emitted from the projectors is reflected by each mirror and then incident on a 300×300 mm stage. In this way, the projection light is divided by the mirrors and the reflected rays are superimposed on the stage, enabling the projection of light rays from multiple directions at relatively low cost.

The light traveling from a projector to the display surface via the mirror array is treated as a directional light ray. Even when the rays arrive at the same point on the display surface, they have different incident angles if they are reflected by different mirrors. Furthermore, within a zenith angle range of approximately $-40^\circ \leq \theta \leq 40^\circ$, 3 to 13 light rays can be independently controlled and projected onto each 1 mm^2 region on the stage.




Figure 1: Appearance of the light-field projection system [15]

3.2. Studies Using the Light-Field Projection System

The light-field projection system described above can also be used to reproduce lighting environments. Kubo et al. [15] used the developed light-field projection system to reproduce a lighting environment in which multiple point light sources were present above the stage. In their experiment, a plaster object with Lambertian reflectance was used as the projection target, and a dynamic lighting environment was presented in which three colored point light sources moved independently. They confirmed that the luminance gradients of attached shadows, corresponding to variations in the surface normal of the plaster object, and the directions of cast shadows, corresponding to the positions of the light sources, were appropriately reproduced.

In addition, Kimura et al. [9] applied this system to manipulate the surface appearance of objects by reproducing materials based on BRDFs. In their experiment, materials with characteristic reflectance properties included in the RGL Material Database [19] were used, and material appearance reproduction by light-field projection was investigated. As a result of projecting onto curved objects, BRDF data with sharp gloss, such as Golden Yellow, Silk Blue, and Aurora White, enhanced specular reflection regions. They also confirmed that complex material appearances, such as structural colors, could be reproduced, where the specular region and color changed according to the observer's viewpoint. On the other hand, for BRDF data with matte appearance, such as Pink Acrylic Felt, the color became relatively uniform, and a matte appearance with small luminance changes under viewpoint movement was reproduced.

3.3. Procedure for Generating Target Projection Images

3.3.1. Definition of Coordinate Systems

To consistently describe the geometric relationship of the system, we define the world coordinate system, the projector image coordinate system, and the normalized coordinate systems on the mirrors and the stage as follows, as shown in Fig. 1.

- World coordinate system: A device-fixed orthogonal coordinate system (X, Y, Z) is defined with its origin at the center of the stage.
- Projector image coordinates: The pixel coordinates on the projector image are denoted by (x_p, y_p) .
- Mirror normalized coordinates: Normalized coordinates (x_m, y_m) are defined on each mirror surface, where $(0, 0) \leq (x_m, y_m) \leq (1, 1)$.
- Stage normalized coordinates: Normalized coordinates (x_s, y_s) are defined on the stage, where $(0, 0) \leq (x_s, y_s) \leq (1, 1)$.

3.3.2. Calculation of the Correspondence between Projector Pixels and Mirror/Stage Coordinates

To project light onto an intended position, it is necessary to identify which point (x_m, y_m) on the mirror and which point (x_s, y_s) on the stage are reached by a projector pixel (x_p, y_p) . In previous studies [9, 15], this pixel correspondence was obtained using homographies. Since both the mirror surface and the stage surface are assumed to be planar, their correspondence with the projector image plane can be represented by projective transformations, namely homographies.

The relationship between the projector pixel coordinates (x_p, y_p) and the mirror normalized coordinates (x_m, y_m) is described using the homography matrix H_{pm} as follows:

$$(x_m, y_m, 1)^T \propto H_{pm}(x_p, y_p, 1)^T. \quad (1)$$

Similarly, the relationship between the projector pixel coordinates and the stage normalized coordinates (x_s, y_s) is represented using H_{ps} as follows:

$$(x_s, y_s, 1)^T \propto H_{ps}(x_p, y_p, 1)^T. \quad (2)$$

The procedure for estimating the homography matrices H_{pm} and H_{ps} is described later in Section 3.4.

3.3.3. Conversion to World Coordinates and Calculation of the Incident Ray Direction

To present color changes according to viewpoint movement, the incident direction of each light ray must be calculated. First, the mirror normalized coordinates (x_m, y_m) and the stage normalized coordinates (x_s, y_s) are converted into the world coordinate system. For the mirror, the world coordinates of the four mirror corners are measured, and a homography matrix H_{mw} obtained from the corresponding points is used for the coordinate transformation. The transformation from the normalized coordinates (x_m, y_m) to the point (X_m, Y_m, Z_m) on the mirror surface in the world coordinate system is expressed by Eq. (3):

$$(X_m, Y_m, Z_m, 1)^T \propto H_{mw}(x_m, y_m, 0, 1)^T. \quad (3)$$

For the stage, the coordinate transformation is performed using the scale factors (S_x, S_y) based on the physical size of the stage and the offsets (v_x, v_y) with respect to the origin of the world coordinate system. The transformation from the stage normalized coordinates (x_s, y_s) to the point (X_s, Y_s, Z_s) on the stage in the world coordinate system is expressed as follows:

$$\begin{pmatrix} X_s \\ Y_s \\ 1 \end{pmatrix} = \begin{pmatrix} S_x & 0 & v_x \\ 0 & S_y & v_y \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x_s \\ y_s \\ 1 \end{pmatrix}. \quad (4)$$

Since the stage surface is defined on the $Z = 0$ plane in the world coordinate system, $Z_s = 0$.

From the above, the incident direction of the light ray $I(x_p, y_p)$ corresponding to each projector pixel (x_p, y_p) is calculated as

$$\begin{cases} \theta_r = \cos^{-1} \left(Z_r / \sqrt{X_r^2 + Y_r^2 + Z_r^2} \right) \\ \phi_r = \tan^{-1}(Y_r/X_r) \quad (= \text{atan2}(X_r, Y_r)) \end{cases} \quad (5)$$

where

$$(X_r, Y_r, Z_r) = (X_m - X_s, Y_m - Y_s, Z_m) \quad (6)$$

is the vector from the stage point to the mirror point. Here, θ_r denotes the zenith angle from the Z axis, and ϕ_r denotes the azimuth angle measured counterclockwise from the X axis.

Through this procedure, the arrival position and direction of the light ray generated by each projector pixel are determined. Consequently, projection images containing the angular components required for viewpoint-dependent presentation can be generated.

3.4. Conventional Geometric Calibration Procedure

To calculate the homography matrices H_{pm} and H_{ps} , four points on the mirror (x_m, y_m) or four points on the stage (x_s, y_s) , together with their corresponding projector image coordinates (x_p, y_p) , are required. These correspondences must be obtained individually for every combination of projectors and mirrors. This calibration procedure is not automated and requires manual adjustment as described below.

For H_{pm} , a rectangular marker with diagonal and center guide lines is first generated on the projector image, as shown in the upper-left panel of Fig. 2. In the initial state, the projected rectangle does not coincide with the outline of the mirror, as shown in the upper-right panel of Fig. 2. Therefore, the rectangle is translated and deformed using mouse and keyboard operations so that the four corners of the projected rectangle accurately coincide with the four corners of the actual mirror, as shown in the lower-right panel of Fig. 2. At this time, the rectangle in the projector image takes a distorted shape, as shown in the lower-left panel of Fig. 2, in order to maintain geometric consistency. After the adjustment is completed, the projector pixel coordinates (x_p, y_p) at that point are recorded, thereby obtaining the four corner points (x_m, y_m) on the mirror surface and their corresponding projector image coordinates (x_p, y_p) . Depending on the actual hardware arrangement, the light from the projector may not reach all four corners of the mirror. In such cases, instead of using the entire mirror area, a rectangle is drawn within the projectable region, and the corresponding points are obtained by visually adjusting the rectangle so that its edges become parallel to the outline of the mirror.

For H_{ps} , the same procedure is performed by manipulating the projected rectangular marker and identifying the projector pixel coordinates that reach the four corners of the stage after being reflected by the mirror. Specifically, starting from the initial state shown in Fig. 3, the rectangle is adjusted so that the projected rectangle coincides with the four corners of the stage, as shown in Fig. 4, and the corresponding coordinates (x_p, y_p) are recorded. If the projector light reflected by the mirror cannot cover the entire stage because of hardware constraints, a rectangle is set within the projectable region, and the corresponding points are obtained by adjusting the rectangle so that its edges become parallel to the stage outline, as shown in Fig. 4.

The calibration described above must be performed manually using the keyboard and mouse while visually checking the projected results, which can introduce human error. In addition, this procedure must be repeated for all combinations of the four projectors and nine mirrors, and the entire calibration process requires several hours to complete. As a result, it is difficult to frequently change the system configuration according to the presentation range or experimental purpose. To address these problems, this study aims to automate the calibration of the light-field projection system.




Figure 2: Manual calibration procedure for aligning a projected rectangular marker with the mirror surface.




Figure 3: Procedure for projecting a rectangular marker onto the stage via a mirror.




Figure 4: Manual adjustment of the rectangular marker for estimating H_{ps} .

4. Limitations of Existing Calibration Methods for Multiple-Mirror Systems

In general projector-camera systems, geometric calibration is often performed by acquiring pixel correspondences using structured light. A straightforward approach is to sequentially project each projector pixel and capture the projected result with a camera. However, scanning all pixels at a resolution of 1280×800 requires a considerable amount of time. Moreover, when only a single pixel is illuminated, the observed light intensity becomes extremely low because of mirror reflection and scattering. Therefore, this approach is not suitable for calibrating the target system.

Line projection is a more efficient alternative. This method acquires correspondences by projecting slit-shaped light while scanning its position. However, in the proposed system, even a single projected line is reflected by multiple mirrors and appears as multiple lines on the stage, as shown in Fig. 5. This produces many candidate correspondences and makes it difficult to uniquely determine the intended projector coordinates. Therefore, line projection is also unsuitable for this system.


Gray-code projection is widely used to efficiently acquire correspondences between projector and camera pixels. In this method, binary patterns are projected, and projector coordinates are decoded from the observed bit sequence. The number of required patterns is on the order of $O(\log_2 N)$ with respect to the projector resolution N , which is much smaller than that required for pixel-by-pixel scanning.

However, a standard Gray-code method cannot be directly applied to the target light-field projection system. As shown in Fig. 6a, reflected light from multiple mirrors is superimposed on the stage. Therefore, when Gray-code patterns are projected over the entire projector image, as shown in Fig. 6b, light rays passing through different mirrors reach the same stage position and the projected patterns are mixed. For example, one mirror may project a bright region while another mirror projects a dark region onto the same point. The camera then observes their superposition as an intermediate luminance value, which makes accurate binarization and decoding difficult. Thus, the camera image alone cannot determine through which mirror and from which projector pixel the observed light has arrived.

For this reason, accurate geometric calibration requires a projection method that prevents interference among light rays reflected by different mirrors.



Figure 5: Captured image showing that a single projected line appears as multiple lines on the stage after reflection by multiple mirrors.



(a) White projection

(b) Gray-code projection

Figure 6: Captured images when projection is performed using a single projector. (a) shows that light from multiple mirrors overlaps on the stage, and (b) shows that Gray-code patterns are also mixed by the overlapping reflected light.

5. Proposed Method

5.1. Overview of the Proposed Method

As described in Section 4, when pattern projection is uniformly applied over the entire projector image in a light-field projection system, light reflected through multiple mirrors spatially overlaps on the stage. As a result, the projected patterns corresponding to different mirrors are mixed, making decoding difficult. Therefore, in this study, we use "Per-mirror Gray-code Projection," in which the projection region is divided for each mirror and Gray-code projection is performed individually.

In the proposed method, a mask is applied so that only the region corresponding to each mirror is transmitted. This allows only the light passing through one mirror to illuminate the stage at a time. As a result, the superposition of projected light from different mirrors is avoided, and pixel correspondences can be obtained for each mirror.

However, the correspondences obtained through camera capture include geometric distortion caused by lens aberration and other factors. To solve this problem, we propose introducing an additional "coordinate-defining projector" as a reference for evaluation and coordinate definition. This makes it possible to compensate for distortion in the camera image and realize geometric calibration based on accurate world coordinates.

5.2. Experimental Environment

The experimental environment is shown in Fig. 7. To implement the proposed method, the following two calibration devices were newly added to the light-field projection system described in Section 3. First, a camera (Ximea MQ013CG-E2, 1280×1024) was installed at a position where the entire stage could be observed. The pixel coordinates on the captured camera image are denoted by (x_c, y_c) . In addition, a coordinate-defining projector (EPSON EB-W12, 1280×800 , 2,800 lm) was installed to directly project images onto the entire stage without using the mirrors and to define the world coordinates.




Figure 7: Experimental environment

5.3. Per-mirror Gray-code Projection

5.3.1. Generation of Mask Images

To perform Gray-code projection individually for each mirror, a mask image is required to allow only the light passing through the target mirror to be projected while blocking the light passing through the other mirrors. In this study, a mask image corresponding to each mirror is interactively generated using an image consisting of a white quadrilateral on a black background.


Fig. 8 shows the scene in which the white quadrilateral is projected onto the target mirror and the resulting mask image. First, the white quadrilateral is projected onto the mirror array from the target projector. Next, while visually observing the projected quadrilateral, the four vertices are moved using keyboard operations so that the white region approximately overlaps the outline of the target mirror. The adjusted projection image is then saved and used as the mask image corresponding to the target mirror. During Gray-code projection, the Gray-code pattern is displayed only within the white region of the mask image, while no light is projected in the black region.

The edges of the quadrilateral do not need to be aligned precisely with the outline of the mirror. The purpose of this procedure is to isolate the projected light that reaches the stage through the target mirror. Therefore, even if part of the white region overlaps an adjacent mirror, the mask image can still be used as long as the light reflected from that region does not reach the stage and no light passing through mirrors other than the target mirror appears on the stage. This procedure is repeated for each combination of projector and mirror used in the system.


5.3.2. Generation and Verification of P2C Maps

First, as shown in Fig. 9, Gray-code images are projected within the region of each mirror represented by the generated mask image in the projection image of one projector, and this projection is captured by the camera. By decoding the Gray codes from the captured image sequence, the correspondence between the camera image coordinates (x_c, y_c) and the projector image coordinates (x_p, y_p) is obtained, and a pixel map describing the pixel correspondence for each mirror is generated. Fig. 10 shows the integrated pixel map obtained from these individual pixel maps.

To verify the geometric consistency between the camera image and the projection image in the obtained pixel map, a stripe pattern with a 15-pixel period along the x_c direction in the camera image coordinate system was projected by referring to the camera coordinates (x_c, y_c) described in the pixel



(a) Scene during mask adjustment for the central mirror.



(b) Generated mask image.

Figure 8: Interactive generation of a mask image for Per-mirror Gray-code Projection. The mask region is adjusted while observing the projection on the target mirror and is then saved as a mask image.




Figure 9: Example of a camera image captured during Per-mirror Gray-code Projection





Figure 10: Per-mirror and integrated P2C maps.

map. However, as shown in Fig. 11a, which was captured by the installed camera, and Fig. 11b, which was captured by another camera located above the center of the stage, the stripe patterns are projected in parallel in the camera image, but appear radial with respect to the stage coordinate system. This indicates that the pixel map obtained in this manner represents the correspondence between the projector image coordinates and the camera image coordinates, but not the correspondence with the stage coordinates. In addition, because the shape of the stage region captured in Fig. 11a is curved, if the camera coordinates (x_c, y_c) described in this P2C map are transformed into the stage coordinates (X, Y) using a homography, distortion in the stage coordinates is caused by the lens distortion of the camera.



(a) Image captured by the system camera



(b) Image captured from above the stage

Figure 11: Projection result using the P2C map obtained by Per-mirror Gray-code Projection. Although the stripe patterns are parallel in the camera image, they appear distorted in the stage coordinate system.

5.4. Definition of World Coordinates Using a Coordinate-Defining Projector

5.4.1. Generation of a World-Coordinate Pixel Map

To solve the problem of camera distortion described in the previous section, coordinate correction using a coordinate-defining projector is performed. Specifically, geometric adjustment was performed using the keystone correction function so that the four corners of the projection image from the coordinate-defining projector accurately coincided with the physical four corners of the stage. Here, the world coordinate system on the stage is defined as (X_s, Y_s) . Since the physical size of the stage is $300 \text{ mm} \times 300 \text{ mm}$, the coordinate system was defined so that one pixel in the projection image corresponds to 1 mm in physical space, resulting in a coordinate range of $-150 \leq X_s \leq 150$ and $-150 \leq Y_s \leq 150$. Under this setting, Gray-code patterns are projected using the coordinate-defining projector and captured by the camera, thereby acquiring the correspondence between the world coordinates (X_s, Y_s) and the camera image coordinates (x_c, y_c) , as shown in Fig. 12.




Figure 12: Example of a camera image captured during Gray-code projection using the newly installed projector

Owing to the optical property of the projector, which projects images onto the screen with the correct shape while compensating for lens distortion, the Gray-code patterns projected from the coordinate-defining projector onto the stage are not distorted. However, the correspondence between the world coordinates and the camera coordinates obtained by capturing this projection with the system camera includes coordinate distortion caused by the camera lens, as in the pixel map obtained in Section 5.3. Nevertheless, when the coordinates of each projector are obtained from the stage coordinates through the inverse transformation of this pixel map, the coordinate distortion caused by the camera lens and

its inverse distortion cancel each other out. As a result, a world-coordinate transformation in which the influence of the common camera lens distortion is reduced can be obtained. This principle can be explained by the following composition of mappings.

Let $\mathbf{X}_s$ denote the world coordinate system on the stage, $\mathbf{x}_c$ denote the camera image coordinate system, and $\mathbf{x} * p$ denote the pixel coordinate system of the system projector. The transformation function $f * s \rightarrow c$ from world coordinates to camera image coordinates using the coordinate-defining projector is defined as the composition of the ideal perspective projection transformation H_s and the nonlinear distortion function D caused by the camera lens:

$$\mathbf{x} * c = f * s \rightarrow c(\mathbf{X}_s) = D(H_s(\mathbf{X}_s * s)) \quad (7)$$

Similarly, the transformation function $f * p \rightarrow c$ from projector pixels to camera image coordinates for the system projector is also described using the perspective projection transformation H_p from the projector to the screen and the same camera lens distortion D :

$$\mathbf{x} * c = f * p \rightarrow c(\mathbf{x}_p) = D(H_p(\mathbf{x}_p)) \quad (8)$$

The objective here is to obtain the projector pixel $\mathbf{x} * p$ corresponding to the specified world coordinate $\mathbf{X}_s * s$. That is, we derive the transformation function F from world coordinates to projector coordinates. This corresponds to transforming the camera coordinates obtained by Eq. (7) back to projector coordinates using the inverse transformation of Eq. (8), and is expressed by the following composition:

$$F = f * p \rightarrow c^{-1} \circ f * s \rightarrow c \quad (9)$$

By substituting Eqs. (7) and (8) into the above equation, we obtain


$$\begin{aligned} F(\mathbf{X}_s * s) &= f * p \rightarrow c^{-1}(D(H_s(\mathbf{X}_s * s))) \\ &= H_p^{-1}(D^{-1}(D(H_s(\mathbf{X}_s * s)))) \\ &= H_p^{-1}(H_s(\mathbf{X}_s * s)). \end{aligned} \quad (10)$$

In the intermediate term of Eq. (10), the camera lens distortion D and its inverse transformation D^{-1} are composed adjacently, and therefore cancel each other out. As a result, the final coordinate transformation does not include the distortion function D . This enables world-coordinate transformation while reducing the influence of camera lens distortion, without requiring explicit estimation of the lens-distortion parameters.


5.4.2. Integration of Correspondences and Final Projection Result

As described in Section 5.4.1, the final geometric calibration is achieved by sequentially combining the two correspondences obtained above. Specifically, the target world coordinates (X_s, Y_s) are first transformed into distorted camera coordinates (x_c, y_c) using the correspondence obtained by the coordinate-defining projector. Then, using the pixel correspondence obtained for each mirror, the derived coordinates (x_c, y_c) are transformed into the pixel coordinates (x_p, y_p) of the projector used for actual projection. Through this procedure, pixels can be accurately projected onto the intended positions in the world coordinate system while the lens distortion is canceled.

Fig. 13 shows the result of projecting a stripe pattern with a 15-mm period in the world-coordinate system using the proposed method. The curvature observed in Fig. 11 is eliminated, and the projected lines are drawn parallel to the outline of the stage at equal intervals. This result demonstrates that the proposed coordinate definition using a coordinate-defining projector is effective for reducing geometric distortion.



(a) Projection image of stripe patterns



(b) Camera image captured when (a) is projected

Figure 13: Projection result after applying coordinate transformation using the coordinate-defining projector.

6. Projection Example Using Calibration Results

In this section, we describe the results of generating BRDF-based target projection images and performing light-field projection using the geometric correspondences obtained by the proposed method. To reproduce natural material appearances using light-field projection, it is essential to determine not only where each projector pixel reaches on the stage, but also from which direction the corresponding light ray enters the stage. Therefore, in this section, we confirm that projection images considering both position and direction can be generated using the correspondences between projector pixels and world coordinates obtained by the proposed method.

In the experiment, a two-dimensional object with specular reflection properties was used as the projection target. The following sections describe the detailed algorithm for generating the target projection images and the actual projection results.

6.1. BRDF Reproduction by Controlling Light-Ray Distribution

In this study, as a method for reproducing complex color changes caused by viewpoint movement using specular reflection, light rays are projected under the assumption that the surface reflection of the target object is specular. Specifically, the color obtained from the BRDF model, the light source direction θ_i , and the viewing direction θ_o is projected from the specular reflection direction calculated from the surface normal $\mathbf{n}$ of the target surface and θ_o . However, accurate BRDF reproduction requires detailed information about the shape and position of the target object, and it is difficult in practice to dynamically measure such information and reflect it in the presentation. Therefore, considering versatility, this study determines the projected color by approximating the projection surface as a horizontal plane, that is, by assuming that the normal $\mathbf{n}$ is uniformly directed vertically upward as $(0, 0, 1)$. This assumption is referred to as the plane approximation assumption.

6.2. Generation of Target Projection Images

This section describes the procedure for generating BRDF-based target projection images using the acquired geometric correspondences.

First, using the geometric correspondences constructed by the proposed method, the world coordinates on the stage corresponding to each projector pixel (x_p, y_p) are obtained as $(X_s, Y_s, 0)$. Here, based on the assumption described in the previous section, the target surface is assumed to lie on the stage plane, namely the $Z = 0$ plane.

Next, the position on the mirror through which the projector pixel passes is identified using the region occupied by the target mirror in the projector image and the world coordinates of the corresponding four corners of the mirror. Specifically, the projector image coordinates are transformed into normalized

coordinates on the mirror by a homography. Then, the world coordinates on the mirror, (X_m, Y_m, Z_m) , are calculated by bilinear interpolation based on the obtained normalized coordinates and the world coordinates of the four mirror corners.

Based on the resulting three-dimensional direction vector, the zenith angle θ_r and azimuth angle ϕ_r are given by

$$\begin{cases} \theta_r = \cos^{-1} \left(Z_r / \sqrt{X_r^2 + Y_r^2 + Z_r^2} \right) \\ \phi_r = \text{atan2}(X_r, Y_r) \end{cases} \quad (11)$$

where

$$(X_r, Y_r, Z_r) = (X_m - X_s, Y_m - Y_s, Z_m). \quad (12)$$

Here, θ_r denotes the zenith angle from the Z axis, and ϕ_r denotes the azimuth angle measured counter-clockwise from the X axis.

For the obtained light-ray direction (θ_r, ϕ_r) , the corresponding specular reflection direction is defined as $(\theta_o, \phi_o) = (\theta_r, \phi_r + \pi)$ under the plane approximation assumption, in which the surface normal is regarded as vertically upward. The target color assigned to each projector pixel is then determined by referring to the BRDF value corresponding to the virtual lighting direction (θ_i, ϕ_i) and the viewing direction $(\theta_r, \phi_r + \pi)$ obtained from the plane approximation.

By sequentially applying this process to all combinations of mirrors and projectors, target projection images are generated for all pixels of each projector. Fig. 14 shows an example of a target projection image generated using the RGL Material Database [19]. Even when pixels correspond to the same position on the stage, the referenced BRDF values differ if they pass through different mirrors, because the incident and reflected directions of the light rays are different.




Figure 14: Example of a target projection image generated based on BRDF data

6.3. Projection Results

Projection experiments were conducted using the generated images. Fig. 15 shows the projection targets. All targets had specular reflection components. Targets A and D were flat surfaces, target B was a curved surface, and target C had a complex shape. To enhance immersion, mask processing was applied according to the shape of each projection target so that light was not projected onto regions other than the target. In addition, flocked paper was placed on the stage.

Fig. 16 shows the results of projection onto target A, observed from $V = (-500, 0, 500)$ mm, $V = (0, 500, 500)$ mm, and $V = (500, 0, 500)$ mm, from left to right. The results show that the brightness and color distributions changed depending on the viewing direction. In particular, glossy and structural-color materials exhibited noticeable viewpoint-dependent color changes, whereas the matte material showed relatively small changes.

Fig. 17 shows the results of projection onto targets B, C, and D, observed from $V = (-500, 0, 500)$ mm. The results show that BRDF-based color distributions were projected not only onto the flat targets but also onto the curved and complex-shaped targets. Although the projection images were generated under the plane approximation assumption, material-dependent color and specular changes were observed for these targets.




Figure 15: Projection targets used in the experiment

6.4. Discussion

6.4.1. BRDF-Based Projection Results and Limitations

The projection results shown in Fig. 16 confirm that viewpoint-dependent appearances corresponding to the BRDF characteristics of the materials were presented. This is because the target projection images were generated using the geometric correspondences obtained by the proposed method, taking into account not only the projection position of each projector pixel but also the direction of the light ray incident on the stage via the mirror. In other words, the directional distribution of projected light rays was controlled by light-field projection, and the variation in BRDF values according to the viewing direction was represented in the projection results.

For example, in Golden Yellow, different color distributions were observed as the viewing direction changed. In ordinary projection mapping, where a fixed image is projected onto a diffuse surface, such color changes basically do not occur when the viewing direction changes. This suggests that the geometric correspondences obtained by the proposed method are effective for presenting strongly




Figure 16: Projection results for target A observed from different viewpoints. The columns show the results observed from $V = (-500, 0, 500)$ mm, $V = (0, 500, 500)$ mm, and $V = (500, 0, 500)$ mm, respectively.




Figure 17: Projection results for targets B, C, and D observed from $V = (-500, 0, 500)$ mm. The columns show the results for the curved target, complex-shaped target, and flat target, respectively.

viewpoint-dependent material appearances, such as structural colors. In contrast, for a matte BRDF such as Acrylic Felt Pink, the color variation was relatively small even when the viewing direction changed, and the projected appearance was observed as almost uniform. This is because matte materials have weak angular dependence in their BRDFs and therefore exhibit small changes in reflectance with respect to viewpoint movement.

However, discontinuities in the color and brightness distributions were observed in some parts of the projection results. Such discontinuities may be caused not only by errors in the calculation of light-ray directions but also by differences in the number of projected light rays superimposed at each position on the stage. In this system, the number of projected rays reaching the stage varies depending on the position: some regions are illuminated by rays from multiple mirrors, whereas other regions are illuminated by rays from only a limited number of mirrors. Therefore, the total amount of projected light may vary depending on the position, which can appear as brightness discontinuities. To address this problem, optical correction considering the number of overlapping projected rays and the diffuse reflection component is required in addition to geometric calibration. For example, when the projection target is approximated as a plane, the number of projected rays superimposed at each stage position can be estimated in advance. Based on this number, the luminance of each projector pixel could be weighted so that the total light intensity for the diffuse reflection component becomes uniform. On the other hand, for the specular reflection component, the reflection distribution itself must be presented according to the viewing direction. Therefore, simple luminance equalization is not sufficient, and introducing an optical correction method that considers the BRDF characteristics remains future work.

As shown in Fig. 17, although the target projection images in this study were generated based on the plane approximation, it was confirmed that BRDF-based color distributions could be presented to some extent even on targets with curved or complex shapes. For target B, material-dependent color differences appeared in the projection results even on a convex surface. In addition, for target C, which has a complex uneven shape, color distributions following the shape of the target were observed. Furthermore, for target D, clear specular components were observed for BRDFs with metallic gloss, whereas such strong specular components were suppressed for Acrylic Felt Pink.

On the other hand, the limitations of target projection image generation based on the plane approximation were also confirmed for targets with curved or complex shapes. In the proposed method, the surface normal of the target is assumed to be uniformly directed vertically upward. Therefore, when the actual local surface normals deviate significantly from this assumption, discrepancies occur between the assumed BRDF values and the reflection directions actually observed. As a result, especially for targets with complex uneven shapes, locally bright regions or deviations in the color distribution may occur. To achieve more accurate BRDF presentation, it is necessary to generate target projection images that take into account the target shape and surface normal directions.

6.4.2. Practical Efficiency and Trade-offs of the Proposed Calibration

In the conventional method described in Section 3.4, corresponding points had to be acquired individually for each valid projector–mirror pair by visually aligning projected rectangular markers with the boundaries of the mirror and the stage. In contrast, in the proposed method, once a mask image has been prepared for each pair, the projection of Gray-code patterns, image capture by the camera, decoding of projector coordinates, generation of the P2C maps, and integration of the coordinate correspondences are performed automatically. Therefore, the repetitive manual acquisition of corresponding points required in the conventional method can be eliminated.

On the other hand, several manual operations remain in the current implementation. Specifically, the mask region must be adjusted for each projector–mirror pair, the projection region of the coordinate-defining projector must be aligned with the boundary of the stage, and the world coordinates of the four corners of each mirror must be acquired manually. However, the mask region does not need to be aligned precisely with the mirror boundary; an approximate adjustment is sufficient as long as the light reaching the stage through the target mirror can be separated from the light passing through the other mirrors.

Because the time required for calibration depends on the device arrangement and the operator's experience, an absolute comparison of calibration time was not conducted in this study. The efficiency improvement achieved in this work lies in replacing the repetitive manual acquisition of corresponding points in the conventional method with automatic correspondence acquisition based on Gray-code projection and image processing. The stripe projection result shown in Fig. 13 demonstrates that the spatial relationships specified in the world-coordinate system can be reflected in the actual projection using the acquired correspondences. In addition, the results shown in Figs. 16 and 17 qualitatively confirm that the acquired correspondences can be used to generate projection images that consider not only the projection positions but also the directions of the light rays passing through the mirrors. A quantitative comparison of positional errors with the conventional method remains future work.

The proposed method requires a camera for observing the stage and a coordinate-defining projector in addition to the existing light-field projection system, which increases the hardware cost, installation space, and constraints on the device arrangement. Moreover, if the position of either device changes after calibration, the coordinate correspondences must be acquired again. On the other hand, by observing the projections from both the system projectors and the coordinate-defining projector with the same camera and integrating the resulting correspondences, the proposed method reduces repetitive manual calibration and enables pixel correspondences in the world-coordinate system to be obtained while reducing the influence of distortion in the camera image, without requiring explicit lens-distortion calibration. The additional camera and coordinate-defining projector are used only during calibration and do not need to operate during the subsequent BRDF-based light-field projection once the correspondences have been obtained.

7. Conclusion

In this study, we proposed an efficient geometric calibration method for a light-field projection system consisting of multiple projectors and a mirror array, with the aim of acquiring correspondences between projector pixels and world coordinates on the stage. In a multiple-mirror-based light-field projection system, light emitted from the projectors reaches the stage after being reflected by multiple mirrors. Therefore, when ordinary structured-light projection is directly applied, projected light from different mirrors is superimposed, making it difficult to accurately acquire pixel correspondences. In addition, when correspondences are acquired based on camera images, distortion caused by the camera lens can introduce errors in the correspondence with world coordinates on the stage.

To address these issues, this study introduced a method that divides the projection region for each mirror and projects Gray-code patterns individually. By projecting only the target mirror region using mask processing, we showed that the superposition of reflected light from multiple mirrors can be avoided and that correspondences between projector pixels and camera pixels can be acquired for each mirror. Furthermore, by acquiring the correspondence between world coordinates on the stage and camera image coordinates using a coordinate-defining projector and integrating it with the pixel correspondences obtained for each mirror, we realized projection based on world coordinates while reducing the influence of geometric distortion included in camera images. In the actual projection results after applying the coordinate transformation, stripe patterns were projected onto the stage in the intended direction, confirming that the proposed method is effective for acquiring geometric correspondences in a multiple-mirror-based light-field projection system.

In addition, using the geometric correspondences obtained by the proposed method, we generated target projection images based on BRDF data and performed material appearance projection on several targets with specular reflection components. The results confirmed that projection images can be generated by considering not only where each projector pixel reaches on the stage, but also through which mirror the corresponding light ray passes and from which direction it enters the stage. The projection results showed viewpoint-dependent color and specular changes for glossy and structural-color materials, while matte BRDFs exhibited relatively small color changes. These results suggest that the proposed method can be used not only for geometric calibration of light-field projection systems,

but also for direction-dependent material appearance presentation based on BRDFs.

On the other hand, several issues remain. Some parts of the calibration procedure, such as the generation of mask images that define the projection region for each mirror and the acquisition of the world coordinates of the four corners of each mirror, are still performed manually. In addition, discontinuities in brightness were observed in some projection results, which suggests that optical correction considering the number of overlapping projected rays and the diffuse reflection component is required in addition to geometric calibration. Furthermore, the target projection images in this study were generated based on a plane approximation. To perform more accurate BRDF presentation on curved or three-dimensional objects, it is necessary to generate target projection images that consider the shape and surface normal directions of the target surface. In future work, we will further automate the calibration procedure by introducing automatic extraction of mirror projection regions and automatic estimation of mirror poses. We will also combine the proposed calibration method with optical correction and target shape acquisition to realize more accurate light-field projection and material appearance presentation.

Acknowledgments

This work was supported by JSPS KAKENHI Grant Numbers JP24K02979, JP25H01886.

Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT for grammar and spelling checks. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

- [1] T. Okazaki, T. Okatani, K. Deguchi, A projector-camera system for high-quality synthesis of virtual reflectance on real object surfaces, *IPSJ Transactions on Computer Vision and Applications* 2 (2010) 71–83.
- [2] D. G. Aliaga, Y. H. Yeung, A. Law, B. Sajadi, A. Majumder, Fast high-resolution appearance editing using superimposed projections, *ACM Transactions on Graphics* 31 (2012) 1–13.
- [3] T. Amano, Projection based real-time material appearance manipulation, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2013, pp. 918–923.
- [4] Y. Hiroi, Y. Itoh, T. Hamasaki, D. Iwai, M. Sugimoto, HySAR: Hybrid material rendering by an optical see-through head-mounted display with spatial augmented reality projection, in: *Proceedings of IEEE Virtual Reality*, 2017, pp. 211–212.
- [5] T. Kaminokado, D. Iwai, K. Sato, Augmented environment mapping for appearance editing of glossy surfaces, in: *Proceedings of IEEE International Symposium on Mixed and Augmented Reality*, 2019, pp. 55–65.
- [6] T. Amano, S. Ushida, Y. Miyabayashi, Viewpoint-dependent appearance-manipulation with multiple projector-camera systems, in: *Proc. ICAT - EGVE 2017*, 2017, pp. 101–107.
- [7] K. Murakami, T. Amano, Materiality manipulation by light-field projection from reflectance analysis, in: *International Conference on Artificial Reality and Telexistence and Eurographics Symposium on Virtual Environments (ICAT-EGVE 2018)*, 2018, pp. 99–105.
- [8] S. Ohsumi, T. Amano, Manipulation of anisotropic reflection based on optical models using light-field projection, *Transactions of the Virtual Reality Society of Japan* 29 (2024) 89–97.
- [9] K. Kimura, T. Amano, Perceptual brdf manipulation by 4-dof light field projection using multiple mirrors and projectors, *19th EuroXR International Conference - EuroXR 2022* (2022) 95–99.

- [10] M. Nagase, D. Iwai, K. Sato, Dynamic control of multiple focal-plane projections for eliminating defocus and occlusion, in: *Proceedings of IEEE Conference on Virtual Reality (IEEE, 2010)*, 2010, pp. 293–294.
- [11] Z. Zhou, T. Yu, X. Qiu, R. Yang, Q. Zhao, Light field projection for lighting reproduction, in: *2015 IEEE Virtual Reality Conference (VR)*, 2015, pp. 135–142.
- [12] Y. Takeuchi, K. Nagamine, Theory and implementation of integral illumination, *IEEE Access* 10 (2022) 939–950.
- [13] J. Y. Han, K. Perlin, Measuring bidirectional texture reflectance with a kaleidoscope, *ACM Transactions on Graphics (SIGGRAPH 2003)* 22 (2003) 741–748.
- [14] Y. Mukaigawa, S. Tagawa, J. Kim, R. Raskar, Y. Matsushita, Y. Yagi, Hemispherical confocal imaging using turtleback reflector, in: *ACCV 2010*, volume 6492, 2011, pp. 336–349.
- [15] R. Kubo, T. Amano, A multiple mirror based arbitrary lighting environment reproduction, *19th EuroXR International Conference - EuroXR 2022* (2022) 81–85.
- [16] J. Salvi, J. Pagés, J. Batlle, Pattern codification strategies in structured light systems, *Pattern Recognition* 37 (2004) 827–849. doi:10.1016/j.patcog.2003.10.002.
- [17] J. Salvi, S. Fernandez, T. Pribanić, X. Lladó, A state of the art in structured light patterns for surface profilometry, *Pattern Recognition* 43 (2010) 2666–2680. doi:10.1016/j.patcog.2010.03.004.
- [18] D. Moreno, G. Taubin, Simple, accurate, and robust projector-camera calibration, in: *2012 Second International Conference on 3D Imaging, Modeling, Processing, Visualization and Transmission (3DIMPVT)*, 2012, pp. 464–471. doi:10.1109/3DIMPVT.2012.77.
- [19] J. Dupuy, W. Jakob, An adaptive parameterization for efficient material acquisition and rendering, *Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 37 (2018) 274:1–274:18. doi:10.1145/3272127.3275059.