

# Shadow-Free Projection with Blur Mitigation on Dynamic, Deformable Surfaces

Andre Muehlenbrock  
muehlenb@uni-bremen.de  
University of Bremen  
Bremen, Germany

Yaroslav Purgin  
purgin@uni-bremen.de  
University of Bremen  
Bremen, Germany

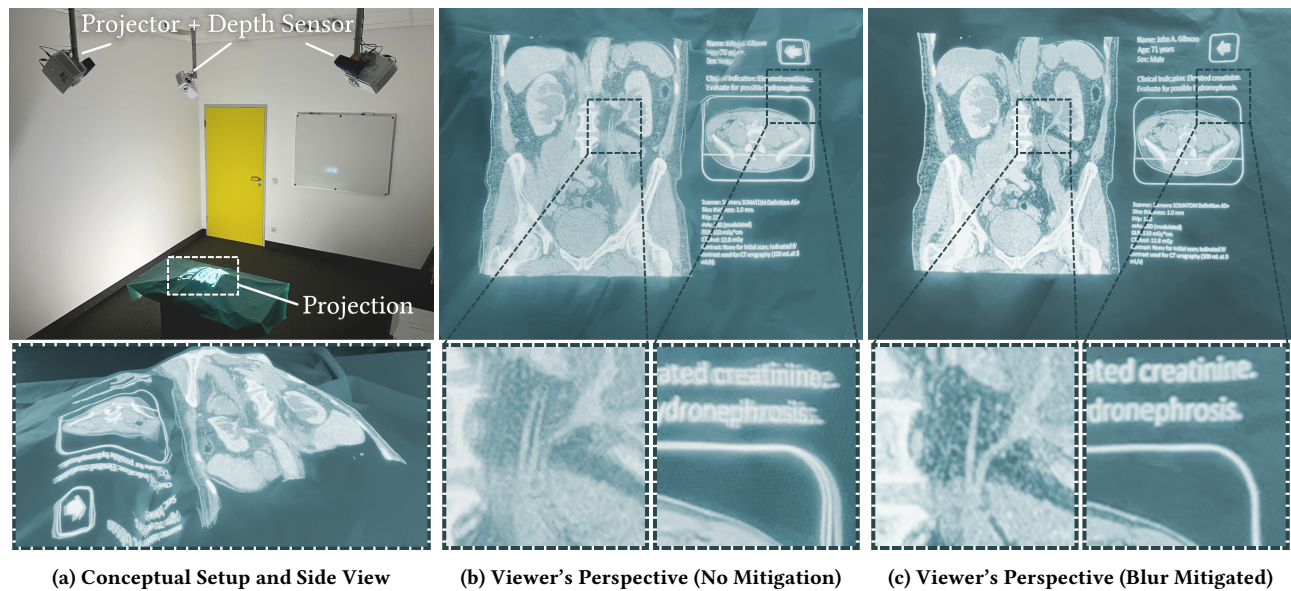
Nicole Steinke  
nicstein@uni-bremen.de  
University of Bremen  
Bremen, Germany

Verena Uslar  
verena.uslar@uol.de  
University of Oldenburg  
Oldenburg, Germany

Dirk Weyhe  
dirk.alfons.veyhe@uol.de  
University of Oldenburg  
Oldenburg, Germany

Rene Weller  
weller@cs.uni-bremen.de  
University of Bremen  
Bremen, Germany

Gabriel Zachmann  
zach@cs.uni-bremen.de  
University of Bremen  
Bremen, Germany



**Figure 1: Real-time shadow-free projection on non-planar surfaces using a multi-projector system and depth-based geometric correction. (a) Conceptual setup with three projectors, depth sensors, and a non-flat surface. (b) Without blur mitigation, but with geometric correction; artifacts from overlapping projections are still clearly visible. (c) With our blur mitigation technique: sharp and consistent projection at ~2 m projection distance.**

## Abstract

We present a real-time projection mapping system for visualizing information and displaying user interfaces on uneven, deformable surfaces in dynamic environments, where the surface gets partially and dynamically occluded. An important application area is the operating room, where this technology would allow for projection onto the surgical drapes. To achieve precise and adaptive geometric correction in setups with multiple projectors and overlapping, partially occluded projection regions, we adapt a point cloud rendering technique that accurately and efficiently reconstructs surface

geometry in the projectors' image space. This enables an overlap precision of 1.6 mm at a projection distance of 2 m, even on uneven surfaces. In addition, we propose two novel GPU-based blur mitigation methods that address blur caused by inevitable inaccuracies of the depth sensors and the overlapping projector images. A user study ( $n = 23$ ) shows that our blur mitigation strategies significantly enhance perceived readability, reduce visual artifacts, and lower user workload compared to conventional multi-projector blending. Our system supports projection on arbitrary surfaces, without requiring explicit segmentation and is well-suited to meet the demands of sensitive environments, including those with sterility constraints or limited display access.



This work is licensed under a Creative Commons Attribution 4.0 International License.  
VRST '25, Montreal, QC, Canada

© 2025 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-2118-2/2025/11  
<https://doi.org/10.1145/3756884.3766018>

## CCS Concepts

• **Computing methodologies** → **Mixed / augmented reality**;  
**Visibility**; Point-based models.

## Keywords

Dynamic projection mapping, virtual displays, occlusion, shadow-free projection, blur mitigation

### ACM Reference Format:

Andre Muehlenbrock, Yaroslav Purgin, Nicole Steinke, Verena Uslar, Dirk Weyhe, Rene Weller, and Gabriel Zachmann. 2025. Shadow-Free Projection with Blur Mitigation on Dynamic, Deformable Surfaces. In *31st ACM Symposium on Virtual Reality Software and Technology (VRST '25), November 12–14, 2025, Montreal, QC, Canada*. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3756884.3766018>

## 1 Introduction

In many professional environments, rapid and unobtrusive access to information is essential to maintain workflow and avoid distractions. This is particularly true in technical fields, high-risk environments, or critical settings such as operating rooms. During surgery, for instance, teams must continuously observe essential data such as vital signs, neuro-monitoring information, or CT images. However, spatial constraints and sterility requirements often prevent the placement of conventional monitors near the surgical site. As a result, surgeons must frequently shift their attention away from the operative field to consult distant displays. These distractions have been linked to decreased performance and may contribute to surgical errors, including visceral injuries or bleeding [28, 30].

While augmented reality (AR) glasses offer in-sight data access, they introduce limitations such as a narrow field of view, potential color distortions that hinder tissue differentiation, and added ergonomic strain. To address these challenges, we propose a projection-based system that displays information directly within the user's field of view by projecting onto irregular, dynamic surfaces, such as surgical drapes, which may shift due to patient repositioning, drape movement, or change of table height. Although prior work has demonstrated projection mapping on deformable surfaces [17–20], these systems typically use multiple projectors only to extend the projection area and do not consider shadowing effects caused by people moving within the environment.

Projection onto dynamic, non-flat surfaces using multiple, overlapping projectors introduces specific technical challenges: accurate surface reconstruction is required to ensure seamless overlay between projectors, and even minor reconstruction errors can lead to visible blurring artifacts, see Figure 1. Also, dynamic shadowing (e.g., caused by moving people in the projectors' pathways) has to be resolved in real-time in order to present a consistent, homogeneous display on the surface.

In this paper, we address these issues by the following main scientific contributions:

- A novel, general-purpose, and low-latency pipeline for shadow and blur mitigation in projection mapping, introducing the concept of *Shadow Distance Maps*.
- An efficient and marker-free geometric correction technique that uses multi-camera point-cloud rendering for precise correction and pre-warping. It operates directly in projector image space and is robust against partial occlusions.
- The *Projection Overlap Divergence (POD)*, an error metric reflecting the visible divergence (or blurring) of overlapping projections, expressed in physical units.

- A user study ( $n = 23$ ) validating the effectiveness and perceptual advantages of the proposed blur mitigation techniques.

To ensure reproducibility and facilitate future research, we release our software as open source.<sup>1</sup>

## 2 Related Work

### 2.1 Static Projection Mapping

Projection systems utilizing multiple projectors have been developed for a wide range of applications. Multi-projector displays are a common use case, where a large surface is covered by dividing it into smaller regions for individual projectors, with overlapping areas blended to create a seamless image [38]. The primary focus of such systems is achieving accurate calibration of the projectors with respect to the static surface being projected onto, typically facilitated by additional cameras. Numerous publications have proposed improvements to these calibration techniques [4–6].

A current research question, particularly in static projection mapping, is radiometric compensation to accurately reproduce colors and materials on arbitrary surfaces. Approaches range from physically-based simulation using differentiable path tracing (DPCS) [26], to efficient scene representation with Gaussian Splatting (GS-ProCams) [7], to joint optimization of projection and scene in a neural reflectance field without explicit calibration [9]. CompenNet++ [16] addresses this challenge by learning geometric correction and radiometric compensation jointly in an end-to-end manner.

### 2.2 Dynamic Projection Mapping

Another area of research is Dynamic Projection Mapping (DPM), which involves projecting onto moving and transforming surfaces. Tracking dynamic, deformable projection surfaces poses significant challenges.

An early method encoded visible patterns that were used for 3D reconstruction and distortion correction, rendering them imperceptible by alternating the pattern with its inverse across successive frames [36]. Other systems primarily focused on tracking rigid objects, such as those demonstrated by Lincoln et al. [27] and Zhou et al. [47]. The system proposed by Siegl et al. [40] also addresses self-shadowing of tracked rigid objects. Later efforts extended to specialized deformable targets, supported by advanced high-speed detection and projection hardware. These systems enable dynamic texture projection onto hands [10], arms [37], and faces [41], with applications such as aging simulation. For such tasks, ultra-high-speed systems equipped with custom projection and tracking hardware have been developed to minimize latency when capturing continuously moving targets [15]. Other approaches aim to mitigate delays by predicting the pose of rigid objects [29].

Tracking non-rigid, deformable surfaces presents additional complexities, particularly when the surfaces lack identifiable features. These challenges are often addressed using marker-based techniques, such as point markers recognized by consumer-grade RGB-D cameras, which are subsequently used to reconstruct the surface geometry [32]. For instance, Taya and Takahashi combines marker-based tracking with skeletal data to map projections onto clothing for performances and shows [43]. Similarly, some systems employ

<sup>1</sup>Available at: <https://github.com/muehlenb/DeformableProjection>

markers to undistort projection mappings, approximating surfaces using B-splines [20][19]. These methods enable projections dynamically shifting surfaces, such as tent walls. Expanding upon this, Ibrahim et al. [18] integrate a multi-projector system to cover larger areas by blending the small transition regions between overlapping projectors. While the system uses multiple projectors and cameras to expand the projection area, it does not address occlusions caused by occluders between the projectors and the projection surface.

Because visible markers on projection surfaces can detract from the intended visual content, various efforts have sought to develop less obtrusive markers. For example, active LED markers have been used [44], or hidden markers have been embedded directly into rigid 3D-printed objects [3]. Additionally, Ahmed et al. [1] highlights the challenges of sparse marker distributions, which often result in inaccuracies during surface reconstruction and proposes its own marker strategy to address these issues.

A marker-free projection method for arbitrary surfaces is presented in [31], which can project textures and materials using normals. However, in addition to being marker-free, it is also model-free, meaning that no 3D model is created, and absolute world coordinates of the objects are not known, making it impossible to precisely undistort projections on the projection surface.

### 2.3 Shadow Mitigation

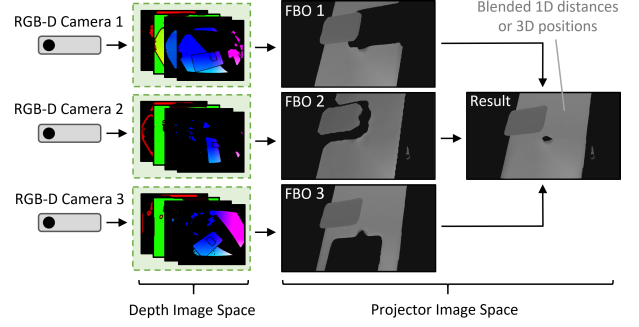
Ibrahim et al. [17] address the occlusion of markers used to define projection targets in a single-projector setup but do not implement any shadow removal. In contrast, some approaches explicitly aim to reduce or avoid shadows on projection surfaces by utilizing multiple projectors. Some techniques detect occluders, estimates the projectors that can illuminate specific target area and balancing which projector displays which part of the whole display [21–23, 42]. These techniques almost all operate only for flat front projection targets. Nomoto et al. [35] presents a shadow removal technique for a high-speed multi-projector setup, where rigid, non-deformable occluders and targets are tracked using infrared markers, with their geometry known a priori. Despite the high-speed multi-projector setup, minor artifacts are visible at the edges of objects due to delayed shadow adjustment. In the context of static projection mapping, there are approaches that reduce cast shadows by employing very large lenses to achieve a broad light distribution, similar to techniques used in surgical lighting to minimize shadows [13, 14, 25]. However, these methods are currently unsuitable for dynamic projection surfaces and themselves suffer from issues such as blurring.

To the best of our knowledge, no existing approach uses depth sensors to geometrically correct a dynamic, non-planar surface with multiple projectors, while simultaneously avoiding shadows and preventing blur artifacts caused by multi-projector overlays.

## 3 Our Method

In this section, we present our novel methods for (1) multi depth camera surface reconstruction for geometric correction (Section 3.1), and (2) blur mitigation (Section 3.2) in multi-projector systems.

Our system targets scenarios with frequent occlusions, which necessitates overlapping projection areas. As a consequence, the display surface must be reconstructed as accurately as possible, in



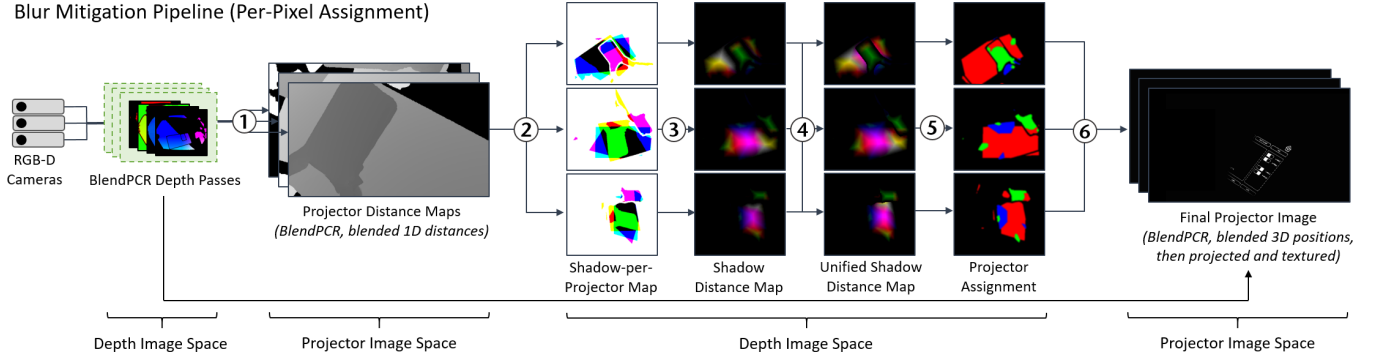
**Figure 2: Adapted point cloud rendering for geometric correction of a single projector using multiple depth cameras. Passes in *Depth Image Space* and *Projector Image Space* are described in Sections 3.1.1 and 3.1.2, respectively.**

order to achieve consistent projections, and dynamic shadowing must be handled such that users cannot see those occlusions, ideally. In the following, we assume a multi-camera, multi-projector setup, in which all projectors and cameras have been precisely calibrated using standard methods; details of our calibration procedure based on checkerboards [2, 8] are provided in the Supplemental Material.

### 3.1 Real-Time Geometric Correction

To enable precise alignment of multiple projectors and geometric correction from a specific point of view, a common target surface for projection must be defined. Instead of relying on physical markers [18, 19], we propose an approach based on point cloud rendering from the perspective of each projector. In our method, we define a virtual display plane  $D$ , oriented towards the viewer, and encode its UV coordinates into the point cloud. By rendering the point cloud from a projector’s perspective and using the user interface as the texture, the interface is inherently pre-warped to account for surface deformations. This results in distortion-free projection from the viewer’s perspective and fully exploits the spatial accuracy of the reconstructed surface without requiring explicit geometric correction. Our approach builds upon *BlendPCR* [34], a method for efficient, artifact-free rendering of continuous surfaces from dynamic multi-camera point clouds. It avoids complex and costly unified mesh reconstruction (e.g., via TSDF and Marching Cubes) by blending separate meshes while preserving the full resolution of each depth camera. In the following, we present this adapted rendering technique and perspective-correct texture mapping.

**3.1.1 Point Cloud Processing from Depth Cameras.** Each depth camera produces a structured point cloud aligned with its depth image, which we treat as a regular 2D grid, implicitly defining a mesh topology. Following a strategy similar to *BlendPCR*, we first apply a binary rejection mask to discard vertices exhibiting depth discontinuities beyond a threshold, thereby avoiding invalid triangles between foreground and background geometry. We then detect edges and smooth them using a weighted Moving Least Squares filter [24]. Surface normals are estimated via PCA of local neighborhoods. Finally, a reconstruction quality score is assigned to each vertex based on its distance to the camera, surface orientation,



**Figure 3: Visualization of our GPU-based blur mitigation pipeline. Steps (1)–(6) correspond to the individual stages described in Section 3.2. The illustration assumes three depth sensors and three projectors.**

and proximity to edges [34]. All computations are performed once per frame and per depth sensor on the GPU.

**3.1.2 Projector-Space Mesh Rendering and Blending.** For each of the  $P$  projectors, point clouds from  $C$  depth cameras are rendered as meshes into  $P \times C$  separate framebuffers, with geometry shaders discarding invalid triangles via the rejection mask. Instead of rendering colors, we encode either 3D positions (for geometric correction) or 1D distances (for *Projector Distance Maps*), see Section 3.2.

Unlike the original BlendPCR approach, we do not perform a fragment-wise dominant-camera selection. Instead, a soft blending strategy is applied, where we use the quality score directly for blending the separate FBOs. This results in a dense, blended surface representation from the projector’s perspective, robust to occlusions and fully leveraging the per-vertex quality information from all depth sensors, see Figure 2.

**3.1.3 Perspective-Correct Texture Mapping.** We define a virtual display plane  $D$ , oriented toward the viewer’s head position  $x$ , and compute per-fragment UV coordinates for UI texture projection:

- (1) Transform the blended 3D position  $p_C$  into the display coordinate system by applying the camera transformation  $T_C$  and the inverse display transformation  $T_D^{-1}$ :

$$p_D = T_D^{-1} \cdot T_C \cdot p_C$$

- (2) Apply a perspective projection with viewer distance  $d$ :

$$s_{\text{Perspective}} = \frac{p_D}{1 - \frac{p_D(z)}{d}} + 0.5$$

Fragments with  $s_{\text{Perspective}} \in [0, 1]^2$  are sampled from the UI texture; others are discarded. This ensures correct perspective alignment from the viewer’s viewpoint. Combined with point cloud rendering as described in Section 3.1.1 and 3.1.2, this yields an efficient and precise geometric correction that is markerless, relocatable, and robust to partial occlusions of individual depth cameras.

## 3.2 Real-Time Blur Mitigation

In dynamic projection, surface geometry and occluders are often unknown. Depth sensors can capture this geometry but inevitably introduce small measurement errors. With multiple projectors used

to avoid shadows, these errors lead to inaccuracies in surface reconstruction, causing visible blur in overlapping regions. We mitigate this by assigning each surface region to a primary projector while allowing limited overlaps, thereby reducing blurring overall. We propose two strategies:

- *Per-Pixel Assignment*: Each pixel is assigned to the projector with the greatest distance to shadows, based on per-projector visibility analysis. Blending is applied in transition zones to avoid visual seams.
- *Tile-Based Assignment*: This strategy is especially useful in cases where the content can be partitioned into discrete, meaningful tiles, which is usually the case for GUIs. Then, each tile can be assigned to the projector with the maximum shadow distance across the entire tile. During an assignment switch, a temporal crossfade between projectors is performed. If no projector can render a tile shadow-free, the fallback is *Per-Pixel Assignment* for that tile.

To implement both strategies, we designed a unified processing pipeline, illustrated in Figure 3. Note that step (1) is performed in the projector’s image space, while steps (2–5) operate in the depth image space, whose pixels correspond one-to-one to the 3D point clouds of the depth cameras:

- (1) *Projector Distance Maps*: From each projector’s perspective, we render the scene’s depth using our adapted BlendPCR technique, storing only 1D distances. We chose a larger field of view than that of the real projectors to ensure that occlusions can be detected early.
- (2) *Shadow-per-Projector Map*: Each point from the structured point cloud is compared to the *Projector Distance Map*. A point is marked as shadowed if it is occluded from the projector’s view (i.e., its depth is larger or lies outside the frustum). These results are encoded into channel-wise boolean maps with the resolution of the structured point cloud, where each channel corresponds to a projector.
- (3) *Shadow Distance Map*: Using the Jump-Flooding algorithm [39], we convert these binary maps into continuous Shadow Distance Maps, where each channel encodes the 2D image distance to the nearest shadow region for one projector.
- (4) *Multi-Camera Fusion and Smoothing*: To address inconsistencies across depth cameras, we project and blend each



camera’s Shadow Distance Map into the others. Additionally, we apply temporal smoothing to reduce flicker due to occasional holes or noise in the depth data.

- (5) *Projector Assignment Map*: Each point is assigned to the projector with the maximum shadow distance, enabling predictive occlusion handling while compensating for minor calibration inaccuracies and sensor latency. A small offset ( $ID \cdot \epsilon$ ) ensures stability when distances are nearly equal. A  $3 \times 3$  blur kernel and bilinear sampling produce soft transitions between projector regions.
- (6) *Final UI Rendering*: During a final BlendPCR screen-space pass, we blend the 3D positions of separately reconstructed meshes. These blended positions are projected according to Section 3.1.3, and the resulting  $s_{\text{perspective}}$  is used as UV coordinates to render the screen texture from the projector’s perspective, enabling geometric correction.

For *Tile-Based Assignment*, we additionally render tile IDs onto the structured point cloud, slightly downsampled for efficiency. This texture is read back to CPU memory to compute the minimum shadow distance per tile from the temporally smoothed maps in step (4). Each tile is then assigned to the projector with the highest minimum shadow clearance. If no projector can fully cover the tile, *Per-Pixel Assignment* is used as fallback. The final Projector Assignment Map is then updated accordingly. Note that both the surface reconstruction and the blur mitigation conceptually allow independent placement of projectors and depth sensors, as well as arbitrary 3D surfaces.

## 4 Measurements of the Projection Quality

To evaluate our methods, we conducted measurements of the accuracy of the geometric correction (see Section 4.2), as well as performance measurements of our pipeline (see Section 4.3).

### 4.1 Setup

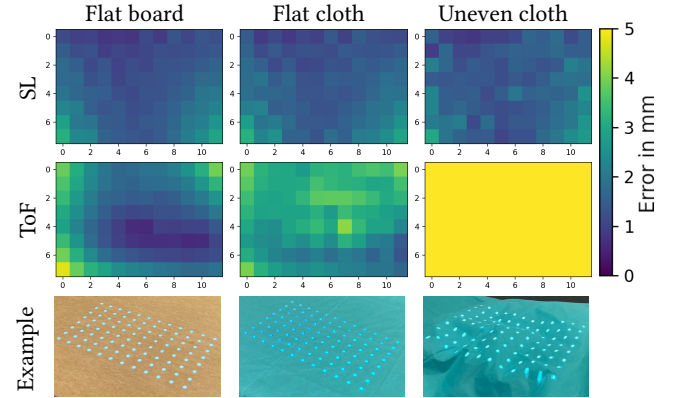
In our evaluation setup, we used three Optoma UHZ66 UHD projectors, each operating at 60 Hz with a resolution of  $3840 \times 2160$ . For depth sensing, we employed either (a) three Orbbec Femto Bolt devices at  $640 \times 576$  mounted directly on the projectors, which feature the same time-of-flight (ToF) sensor as the Microsoft Azure Kinect, or (b) two industrial Ensenso N36 structured light cameras at  $968 \times 608$ . All sensors and projectors are connected to a single computer equipped with an NVIDIA GeForce RTX 4090 GPU and an AMD Ryzen 9 3900X CPU. All projection and camera devices are placed at distances ranging from 1.9 m to 2.2 m relative to the center of the projection area (see Figure 1).

### 4.2 Precision of Geometric Correction

*Projection Overlap Divergence (POD)*. In professional contexts, precise and fast access to the information presented to the staff is essential. To evaluate the quality of overlapping projections and therefore the accessibility, prior work — such as Ibrahim et al. [18] — employed SSIM [46] scores as a perceptual similarity measure. While SSIM is valuable for quantifying visual fidelity, it is sensitive to camera distance, projection content, and projection scale, and it does not provide absolute geometric error values.

**Table 1: Projection Overlap Divergence (POD) in millimeters, obtained using two Ensenso N36 (SL) and three Orbbec Femto Bolt (ToF) cameras on different surface types.**

Condition	Ensenso N36 (SL)			Orbbec Femto (ToF)		
	N	Mean	SD	N	Mean	SD
Flat Board	4	1.51	0.87	6	1.91	1.18
Flat Cloth	4	1.60	0.87	6	2.86	1.53
Uneven Cloth	7	1.58	0.91	6	15.28	6.49



**Figure 4: Color-coded maps of average precision obtained using two Ensenso N36 (SL) and three Orbbec Femto Bolt (ToF) cameras across various surface types. The surface covered approximately  $35 \text{ cm} \times 20 \text{ cm}$ . Values exceeding a 5 mm error were clamped.**

To directly quantify the geometric misalignment of overlapping projections, we introduce the *Projection Overlap Divergence (POD)*. This metric captures the average 2D displacement between corresponding projection points as observed from an external camera  $C$ , and reports the geometric error in millimeters.

Let  $P_1 = \{p_{11}, \dots, p_{1N}\}$  denote the set of  $N$  image points of the first projection, and  $P_2 = \{p_{21}, \dots, p_{2N}\}$  the corresponding image points of the second projection, both expressed in the pixel coordinate system of camera  $C$ . Each point  $p_{ki} \in \mathbb{R}^2$  represents a 2D location in the image plane. The metric is defined as

$$POD(P_1, P_2) = \frac{1}{N} \sum_{i=1}^N \|p_{1i} - p_{2i}\| \cdot \kappa, \quad (1)$$

where

- $\|p_{1i} - p_{2i}\|$  is the Euclidean distance between a pair of corresponding points in pixels,
- $N$  is the total number of correspondences, and
- $\kappa \in \mathbb{R}^+$  is a scaling factor that converts pixel distances into millimeters, determined for a calibration plane orthogonal to the camera’s viewing direction within the working area.

In practice, we project a  $12 \times 8$  dot grid (96 points) from each projector, capture it separately using a system camera mounted approximately 2 m above the projection surface and detecting these

image points using OpenCV. For  $M$  projectors with corresponding point sets  $P_1, \dots, P_M$ , the combined divergence with respect to camera  $C$  is defined as the average over all pairwise divergences:

$$POD(P_1, \dots, P_M) = \frac{2}{M(M-1)} \sum_{1 \leq i < j \leq M} POD(P_i, P_j). \quad (2)$$

The  $POD$  metric is robust to variations in projection size, camera distance, and projected image content, and yields a reproducible geometric error value that reflects the perceivable divergence (or blurring) from a viewing direction in physical units (millimeter).

**Results.** Using two Ensenso N36 structured-light cameras and assuming a top-down viewing direction, our system achieves an average  $POD$  of only 1.5–1.6 mm across the full surface, including highly non-planar regions, see Figure 4 and Table 1. Note that this includes the accumulated errors of the calibration, depth sensors, and surface reconstruction. Considering the projection distance of approximately 2 m, this amounts to an overall error of 0.075 percent, which underscores the high accuracy of our correction pipeline. We further evaluated performance using three Orbbec Femto Bolt time-of-flight (ToF) cameras. To account for known temporal drift in ToF sensors [45], the system was re-calibrated prior to each test series. On flat surfaces, we observed alignment errors of 1.91 mm (wooden board) and 2.86 mm (surgical drape). On non-planar fabric, the average error increased to 15.3 mm, which we attribute to known multipath effects that are typical for ToF cameras [11, 33, 45]. Further ToF sensor measurements and comparative visual examples of the uneven surface are provided in the Supplemental Material.


Overall, our measurements show that we achieve highly precise results using the Ensenso N36 and our pipeline for surface reconstruction with geometric correction.

### 4.3 Runtime Measurements

To evaluate runtime performance, we measured the execution times of pipeline modules in a setup with three UHD projectors and up to three depth cameras. The BlendPCR depth passes take approximately 4.75 ms with three  $640 \times 576$  depth cameras while generating normals, smoothing edges, computing quality scores, applying temporal filtering, and filling holes. By generating point clouds directly in the shader, we avoid memory upload bottlenecks. In the blur mitigation pipeline, *Per-Pixel Assignment* (3.45 ms) proved faster than *Per-Tile Assignment* (5.5 ms) due to the additional tile processing. Geometric correction across three  $3840 \times 2160$  projectors requires on average 6.84 ms. Overall, the OpenGL rendering thread scales with the number of cameras, requiring 22.2 ms with three cameras, 16.9 ms with two, and 8.7 ms with one. A detailed breakdown of the runtime measurement is provided in the Supplemental Material.

## 5 User Study on Blur Mitigation

We conducted a user study to investigate perceived differences in workload, readability, and user preference, along with objective measures such as error rates and task completion times across different blur mitigation approaches. Section 5.1 details the study design; Section 5.2 presents the results.



**Figure 5: Study setup.** The image shows the controlled lab environment and a participant standing in front of the projected user interface.

### 5.1 Study Design

The study was conducted as a within-subjects experiment with 23 participants (7 female, 16 male), aged between 18 and 65 years. Its objective was to investigate the effects of two blur mitigation techniques in shadow-free projection — *Per-Pixel Assignment (A)* and *Per-Tile Assignment (B)* — in comparison to a baseline condition without mitigation (*No Mitigation*), where all projector images were overlaid. The three conditions are illustrated in Figure 6. The study focused on differences in perceived readability, perceived workload, user preference, as well as task completion time and error rates. To control for learning effects and ensure balanced exposure, the order of conditions was counterbalanced across all six permutations.

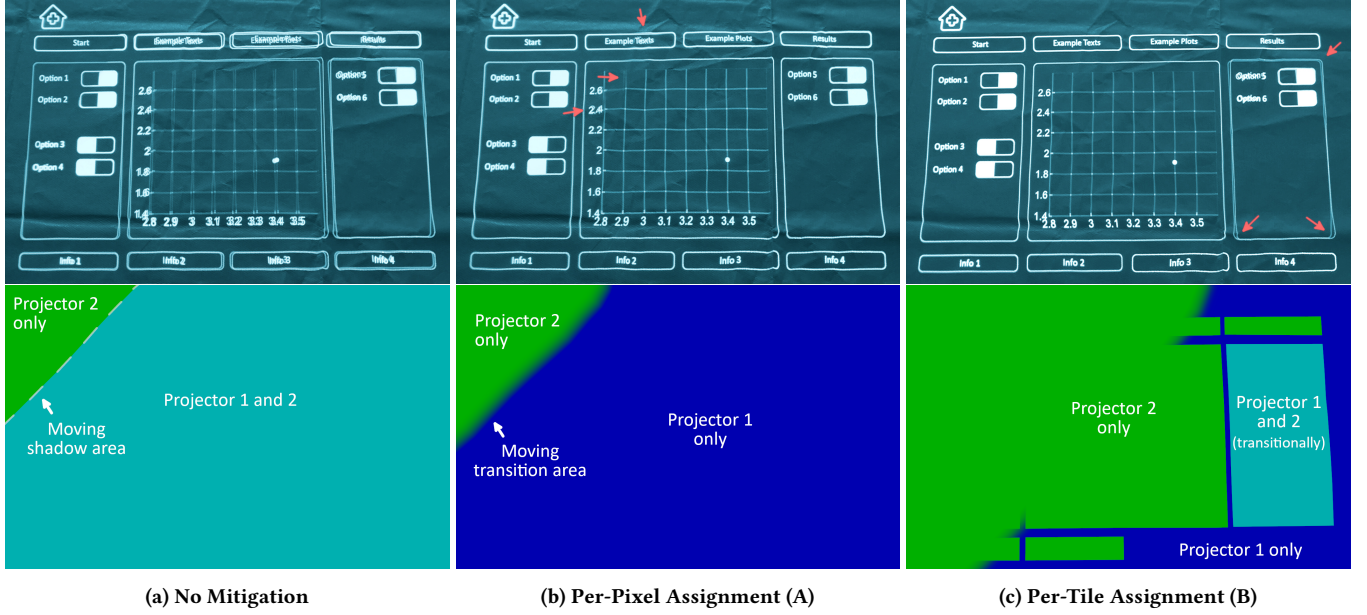
**5.1.1 Tasks and Measurements.** In each of the three conditions, participants were asked to read aloud ten 12-character passwords and the x/y values of ten 2D plots. The passwords and the plots were presented in alternating order. After each condition, the NASA-TLX (Raw) test was administered, along with the following three questions (which we evaluated separately from the NASA-TLX):

- *Character Readability*: “How difficult was it to read characters?”
- *Plot Readability*: “How difficult was it to read plot values?”
- *Artifacts*: “How disturbing were visual artifacts (e.g., blur, motion)?”

Task time was automatically recorded during transitions between password and plot reading. Errors were logged via audio recordings and manually reviewed to ensure accurate timing and error identification. The projector-camera calibration was repeated prior to each participant to account for potential drifts or misalignments over time. Every time, we also checked the projection errors during the calibration: across all study trials, the error was on average of 1.1 mm over a  $26 \times 24$  cm area, and 0.5 mm within the focus region of  $9 \times 8$  cm on a flat wooden calibration board.

**5.1.2 Questionnaires.** After each condition, participants completed a questionnaire containing the six NASA-TLX (Raw) items, followed by the three readability-related items listed above. After completing all three conditions, participants ranked them from 1 (best) to 3 (worst) based on overall preference. After that, the participants were allowed to directly compare the two blur mitigation methods again without time constraints and switch freely between them. Finally, qualitative feedback on their preferred method was collected.

**5.1.3 Simulation of Reproducible Occlusions.** To ensure reproducible occlusions and eliminate variability due to participant position



**Figure 6: Comparison of conditions in the experimental setup. Top: Viewer’s perspective, showing the rendered user interface as presented during the user study. Bottom: Projector assignment, with each surface area color-coded according to the projector responsible for its illumination. In all conditions, a shadow was simulated by a moving occluder. While in (a) and (b) the dynamic transitions are clearly visible, they are not as noticeable in (c) since the tiles of the user interface switch projectors as entire units. In (b) and (c), the arrows highlight remaining blurring areas.**

— especially since small posture changes can have large effects on shadowing — we deactivated the projector positioned behind the participant and used only the two front-facing projectors. A virtual occluder with the shape of a mannequin was simulated opposite the participant, moving on a reproducible path. We captured the point cloud of the virtual occluder using simulated depth sensors and fused it with the real point cloud of the projection surface. To ensure consistency across all three conditions, including the *No Mitigation* baseline, the occluder’s shadow was also simulated in the rendered projector images, resulting in partial occlusion even in the unmitigated condition.

**5.1.4 Mitigation of Multipath Interference.** While performing the user study, the structured light cameras were not available to us; therefore, we had to use the ToF cameras. This technology is prone to cause multipath artifacts; in order to eliminate this potentially confounding factor, we froze the real depth image of the projection surface after calibration and before participants stood in front of the display. This ensures that our findings hold without loss of generality and apply to interference-free sensors such as structured light systems.

## 5.2 Results

Our results show significant differences in the NASA-TLX (Raw) scores, the readability-related questionnaire items, and user preferences regarding blur mitigation approaches. *Per-Pixel Assignment (A)* performed significantly better than the baseline across all three

**Table 2: P-values for the individual items of the NASA-TLX questionnaire, together with our three additional questions. Significant differences are highlighted. Overall comparisons across conditions were performed using the Friedman test, and pairwise differences were examined with Wilcoxon tests applying Bonferroni correction.**

	Friedman	Pairwise Wilcoxon		
	No, A, B	No vs. A	No vs. B	A vs. B
Mental Demand	0.009 (**)	0.03 (*)	0.066	1.0
Physical Demand	0.6	0.47	0.81	1.0
Temporal Demand	0.76	1.0	1.0	1.0
Performance	0.008 (**)	0.094	0.16	1.0
Effort	0.012 (*)	0.023 (*)	0.75	0.68
Frustration	0.093	0.094	0.29	1.0
Char Readability	2.4e-06 (***)	0.0002 (***)	0.0004 (***)	0.34
Plot Readability	1.8e-05 (***)	0.0004 (***)	0.021 (*)	0.18
Artifacts	0.0001 (***)	0.0002 (***)	0.068	0.59
NASA-TLX (Raw)	0.009 (**)	0.041 (*)	0.184	0.964

measures, while *Per-Tile Assignment (B)* also outperformed the baseline in most cases, with statistically significant differences in several metrics. Although a clear trend favoring *Per-Pixel Assignment (A)* over *Per-Tile Assignment (B)* was observed, the measured differences between the two blur mitigation techniques were not statistically significant. No significant differences were found in terms of error rates or task completion times across conditions.

Section 5.2.1 presents the detailed results of the NASA-TLX scores. Section 5.2.2 reports on the readability-related questionnaire items. Section 5.2.3 summarizes participants' preferences, and Section 5.2.4 details the measured completion times and error rates.

**5.2.1 Nasa-TLX (Raw) Results.** A Friedman test revealed significant differences in perceived workload across the three conditions. Post-hoc pairwise comparisons were conducted using Wilcoxon signed-rank tests with Bonferroni correction to control for Type I errors.

The results show a significant difference in workload between *Per-Pixel Assignment (A)* and the baseline ( $p = 0.041$ ), where *Per-Pixel Assignment (A)* was perceived as less demanding. Additionally, there was a trend indicating that *Per-Tile Assignment (B)* was perceived as less demanding than the baseline, although this difference did not reach statistical significance ( $p = 0.184$ ). An overview of the results is shown in Figure 8.

We also analyzed the individual items of the NASA-TLX questionnaire. The p-values are reported in Table 2, while additional plots and effect sizes are provided in the Supplemental Material. The results show that the overall difference in workload can be attributed primarily to the *Mental Demand* and *Effort* items, both of which were rated substantially higher in the *No Mitigation (Blur)* condition compared to either blur mitigation approach. While the differences between *Per-Pixel Assignment (A)* and the baseline were also statistically significant according to Wilcoxon tests, the differences between *Per-Tile Assignment (B)* and the baseline did not reach significant levels after the Bonferroni correction.

**5.2.2 Readability.** The results of the three additional questions regarding readability and visual artifacts are shown in Figure 7 and Table 2. Compared to the NASA-TLX items, the differences observed here are considerably stronger and, in most cases, highly significant. The only exception is the third question concerning visual artifacts, where the difference between *Per-Tile Assignment (B)* and the baseline condition did not remain statistically significant after Bonferroni correction.

Overall, both blur mitigation approaches — particularly *Per-Pixel Assignment (A)*, but also *Per-Tile Assignment (B)* — were perceived as substantially more readable and less affected by visual artifacts than the baseline method.

**5.2.3 Ranking.** The quantitative ranking of the blur mitigation methods revealed a clear overall preference for *Per-Pixel Assignment (A)*, which was ranked first by 70% of participants. It was followed by *Per-Tile Assignment (B)* with 17%, and the baseline condition (*No Mitigation*) with 13%, see Figure 9. As the second preference, *Per-Tile Assignment (B)* was most frequently selected (61%), followed by the other conditions. The baseline method was ranked last by 70% of participants, indicating a strong overall rejection.

In terms of statistical significance, both *Per-Pixel Assignment (A)* and *Per-Tile Assignment (B)* showed clear and statistically significant differences compared to the baseline method, with  $p = 0.005$  and  $p = 0.033$  respectively (Wilcoxon test, Bonferroni-corrected). Although a trend was observed between the two blur mitigation methods, the difference did not reach significance ( $p = 0.158$ ).

**5.2.4 Timing and Errors.** Task completion times and error counts showed no statistically significant differences between conditions; further details are provided in the Supplemental Material.

**5.2.5 Qualitative Results.** After completing the quantitative questionnaires, participants were invited to compare *Per-Pixel Assignment (A)* and *Per-Tile Assignment (B)* once more, without time constraints. They were allowed to freely switch between both conditions to directly assess differences.

In the qualitative feedback, participants particularly noted that the transition areas between projectors in the *Per-Pixel Assignment (A)* condition appeared more continuous and predictable, and were therefore perceived as less distracting. In contrast, the transitions between tiles in the *Per-Tile Assignment (B)* condition were sometimes described as abrupt or unexpected, which some users found distracting during reading tasks. For example, one participant noted, "Mode A is better because the transitions are smoother and easier to anticipate, whereas Mode B makes my eyes jump around due to unexpected changes at the edges." Another participant commented, "I find A slightly better and easier to read. In B, it looks like the projection shifts a bit more — but I'm not entirely sure, as the difference is very subtle." The subtlety of the difference was mentioned more frequently; a difference was not always noticeable to participants while performing the tasks.

### 5.3 Discussion

In our lab experiment, simulating a uniformly and reproducibly moving virtual shadow, both blur mitigation strategies substantially improved the perceived quality of shadow-free projection compared to a baseline condition without mitigation. *Per-Pixel Assignment (A)* consistently outperformed the other methods in terms of workload, readability, and user preference, and was rated as significantly less demanding than the baseline condition. *Per-Tile Assignment (B)* also showed significant improvements over the baseline, although to a slightly lesser extent.

While subjective ratings showed clear advantages for both mitigation approaches, especially in terms of readability and perceived artifacts, objective measures such as task completion time and error rate did not show statistically significant differences. In the case of simple reading tasks, blur mitigation appears to primarily affect user experience and cognitive load, rather than task performance. We hypothesize that performance-related effects could be observed with more complex or extended tasks, for instance, the diagnostic interpretation of a CT scan.


While participants in our study generally preferred the *Per-Pixel Assignment (A)* results, some noted that this preference was partly due to the predictability of the subtle transition-line artifacts. They speculated that for different shadow types, their preference might shift towards *Per-Tile Assignment (B)*, suggesting that the perceived visual continuity could depend on the specific motion patterns involved, pointing to the need for further investigation into how blur mitigation strategies interact with various shadow types and movement patterns.

Overall, our findings suggest that blur mitigation can meaningfully enhance perceived visual quality and reduce subjective workload in shadow-free multi-projector systems in cases with projector misalignments of small misalignments. Future work should investigate these effects in more demanding and dynamic scenarios that impose higher cognitive and temporal demands.






**Figure 7: Comparison of questions regarding the *Readability* of characters, plots, and the disturbance caused by *Artifacts*. All three questions show highly significant differences among the conditions. For detailed p-values, see Table 2.**



**Figure 8: The results of the NASA-TLX in its raw version. The distributions across conditions differ significantly according to the Friedman test ( $p = 0.009$ ). Pairwise Wilcoxon tests, adjusted with the Bonferroni correction, reveal that only the comparison between No Mitigation (Blur) and Per-Pixel (A) is significant ( $p = 0.041$ ), while No Mitigation (Blur) vs. Per-Tile (B) ( $p = 0.184$ ) and Per-Pixel (A) vs. Per-Tile (B) ( $p = 0.964$ ) are not.**



**Figure 9: Rankings of the conditions based on participant preferences. It is evident that the Per-Pixel (A) condition is the most preferred method, Per-Tile (B) is the second most preferred, and the absence of blur mitigation is clearly the least preferred. The Friedman test indicates that the differences between the groups are highly significant ( $p = 0.0004$ ). In pairwise comparisons, there are significant differences between Per-Pixel (A) and the baseline ( $p = 0.005$ ), as well as between Per-Tile (B) and the baseline ( $p = 0.033$ ), but not between Per-Pixel (A) and Per-Tile (B) ( $p = 0.158$ ).**

## 6 Limitations

While our system performs well in mitigating blur and shadows on deformable surfaces, some practical limitations remain. On the sensor side, Time-of-Flight (ToF) cameras provide high frame rates (30 fps) for responsive occlusion handling, but may introduce depth

artifacts on uneven surfaces due to multipath effects. The structured-light (SL) sensor we used offers more accurate geometry but runs at lower frame rates ( $< 15$  fps), which can limit responsiveness in dynamic scenes. A combined ToF–SL approach, or the use of high-speed SL sensors such as those presented by Heist et al. [12], might offer a practical combination of speed and accuracy. Additionally, the system does not yet include photometric compensation. Minor brightness differences between projectors can occur, particularly with reflective surfaces or at steep angles. However, in our top-down clinical setup with mostly matte materials, these effects remained minimal. Finally, without the use of polarization or other separation techniques, the system can only optimize the projection for a single viewer. How to select and dynamically adapt this viewer position remains an open research problem.

## 7 Conclusion

We present a real-time projection system for dynamic, deformable surfaces that achieves precise alignment of overlapping projectors (1.6 mm error at 2 m distance) and effective shadow avoidance through a multi-projector setup. To address remaining blur from minor misalignments in overlapping projections, we introduce a novel blur mitigation pipeline that dynamically assigns surface regions to the most suitable projector – either per-pixel or per-tile – based on real-time shadow distance maps. A user study ( $n = 23$ ) showed both strategies significantly enhance readability, reduce cognitive load, and improve user preference over conventional blending.

Designed for high-stakes, constrained environments such as operating rooms, our system maintains real-time performance ( $\leq 25$  ms per frame) in a full setup with three UHD projectors and three depth sensors, running entirely on a single workstation. It requires no markers, generalizes to arbitrary 3D content, and operates robustly in dynamic scenes. In future work, we aim to explore gesture-based interaction, integrate observer-aware rendering, and conduct context-specific evaluations through focus groups and expert interviews, for example in operating room scenarios, including comparisons with conventional displays.

## Acknowledgments

We would like to thank our SISOPS project<sup>2</sup> partners (KIZMO GmbH, inomed Medizintechnik GmbH and Dr. Mach GmbH & Co. KG) for their valuable collaboration. Special thanks also go to Amelie Polster from Dr. Mach GmbH & Co. KG and to IDS Imaging Development Systems GmbH for kindly providing the Ensenso cameras. This work was partially supported by BMFTR grant 16SV9239.

<sup>2</sup>SISOPS project: <https://cgvr.cs.uni-bremen.de/research/sisops>

## References

- [1] Bilal Ahmed and Kwan H. Lee. 2019. Projection mapping onto deformable nonrigid surfaces using adaptive selection of fiducials. *Journal of Electronic Imaging* 28, 6 (2019), 063008. doi:10.1117/1.JEL.28.6.063008
- [2] K. S. Arun, T. S. Huang, and S. D. Blostein. 1987. Least-Squares Fitting of Two 3-D Point Sets. *IEEE Transactions on Pattern Analysis and Machine Intelligence* PAMI-9, 5 (1987), 698–700. doi:10.1109/TPAMI.1987.4767965
- [3] Hirotaka Asayama, Daisuke Iwai, and Kosuke Sato. 2018. Fabricating Diminishable Visual Markers for Geometric Registration in Projection Mapping. *IEEE Transactions on Visualization and Computer Graphics* 24, 2 (2018), 1091–1102. doi:10.1109/TVCG.2017.2657634
- [4] Ezekiel Bhasker, Ray Juang, and Aditi Majumder. 2007. Registration Techniques for Using Imperfect and Partially Calibrated Devices in Planar Multi-Projector Displays. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1368–1375. doi:10.1109/TVCG.2007.70586
- [5] Ezekiel S. Bhasker, Pinaki Sinha, and Aditi Majumder. 2006. Asynchronous Distributed Calibration for Scalable and Reconfigurable Multi-Projector Displays. *IEEE Transactions on Visualization and Computer Graphics* 12, 5 (2006), 1101–1108. doi:10.1109/TVCG.2006.121
- [6] Han Chen, Rahul Sukthankar, Grant Wallace, and Kai Li. 2002. Scalable alignment of large-format multi-projector displays using camera homography trees. In *Proceedings of the Conference on Visualization '02* (Boston, Massachusetts) (VIS '02). IEEE Computer Society, USA, 339–346.
- [7] Qingyue Deng, Jijiang Li, Haibin Ling, and Bingyao Huang. 2024. GS-ProCams: Gaussian Splatting-based Projector-Camera Systems. arXiv:2412.11762 [cs.CV] <https://arxiv.org/abs/2412.11762>
- [8] Alexander Duda and Udo Frese. 2018. Accurate Detection and Localization of Checkerboard Corners for Calibration. In *Proc. of the British Machine Vision Conference*.
- [9] Yotam Erel, Daisuke Iwai, and Amit H. Bermanno. 2023. Neural Projection Mapping Using Reflectance Fields. *IEEE Transactions on Visualization and Computer Graphics* 29, 11 (Nov. 2023), 4339–4349. doi:10.1109/TVCG.2023.3320256
- [10] Yotam Erel, Or Kozlovsky-Mordenfeld, Daisuke Iwai, Kosuke Sato, and Amit H. Bermanno. 2024. Casper DPM: Cascaded Perceptual Dynamic Projection Mapping onto Hands. In *SIGGRAPH Asia 2024 Conference Papers* (SA '24). Association for Computing Machinery, New York, NY, USA, Article 137, 10 pages. doi:10.1145/3680528.3687624
- [11] S.A. Guomundsson, Henrik Aanæs, and Rasmus Larsen. 2007. Environmental Effects on Measurement Uncertainties of Time-of-Flight Cameras. 1–4. doi:10.1109/ISSCS.2007.4292664
- [12] Stefan Heist, Peter Lutzke, Ingo Schmidt, Patrick Dietrich, Peter Kühmstedt, Andreas Tünnermann, and Gunther Notni. 2016. High-speed three-dimensional shape measurement using GOBO projection. *Optics and Lasers in Engineering* 87 (2016), 90–96. doi:10.1016/j.optlaseng.2016.02.017
- [13] Kosuke Hiratani, Daisuke Iwai, Yuta Kageyama, Parinya Pumpongsanon, Takefumi Hiraki, and Kosuke Sato. 2023. Shadowless Projection Mapping using Retrotransmissive Optics. *IEEE Transactions on Visualization and Computer Graphics* PP (05 2023), 1–11. doi:10.1109/TVCG.2023.3247104
- [14] Kosuke Hiratani, Daisuke Iwai, Parinya Pumpongsanon, and Kosuke Sato. 2019. Shadowless Projector: Suppressing Shadows in Projection Mapping with Micro Mirror Array Plate. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. 1309–1310. doi:10.1109/VR.2019.8798245
- [15] Sora Hisaichi, Kiwamu Sumino, Kunihiro Ueda, Hidenori Kasebe, Tohru Yamashita, Takeshi Yuasa, Uwe Lippmann, Petra Aswendt, Roland Höfling, and Yoshihiro Watanabe. 2021. Depth-Aware Dynamic Projection Mapping using High-speed RGB and IR Projectors. In *SIGGRAPH Asia 2021 Emerging Technologies* (Tokyo, Japan) (SA '21). Association for Computing Machinery, New York, NY, USA, Article 3, 2 pages. doi:10.1145/3476122.3484843
- [16] Bingyao Huang and Haibin Ling. 2019. CompenNet++: End-to-end Full Projector Compensation. In *IEEE International Conference on Computer Vision (ICCV)*.
- [17] Muhammad Twaha Ibrahim, M. Gopi, and Aditi Majumder. 2023. Self-Calibrating Dynamic Projection Mapping System for Dynamic, Deformable Surfaces with Jitter Correction and Occlusion Handling. In *2023 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. 293–302. doi:10.1109/ISMAR59233.2023.00044
- [18] Muhammad Twaha Ibrahim, M. Gopi, and Aditi Majumder. 2024. Real-Time Seamless Multi-Projector Displays on Deformable Surfaces. *IEEE Transactions on Visualization and Computer Graphics* 30, 5 (2024), 2527–2537. doi:10.1109/TVCG.2024.3372097
- [19] Muhammad Twaha Ibrahim, Aditi Majumder, and M. Gopi. 2022. Dynamic projection mapping on deformable stretchable materials using boundary tracking. *Computers & Graphics* 103 (2022), 61–74. doi:10.1016/j.cag.2022.01.004
- [20] Muhammad Twaha Ibrahim, Gopi Meenakshisundaram, and Aditi Majumder. 2020. Dynamic Projection Mapping of Deformable Stretchable Materials. In *Proceedings of the 26th ACM Symposium on Virtual Reality Software and Technology* (Virtual Event, Canada) (VRST '20). Association for Computing Machinery, New York, NY, USA, Article 35, 5 pages. doi:10.1145/3385956.3418970
- [21] Daisuke Iwai, Momoyo Nagase, and Kosuke Sato. 2014. Shadow removal of projected imagery by occluder shape measurement in a multiple overlapping projection system. *Virtual Reality* 18, 4 (01 Nov 2014), 245–254. doi:10.1007/s10055-014-0250-4
- [22] C. Jaynes, S. Webb, and R.M. Steele. 2004. Camera-based detection and removal of shadows from interactive multiprojector displays. *IEEE Transactions on Visualization and Computer Graphics* 10, 3 (2004), 290–301. doi:10.1109/TVCG.2004.1272728
- [23] Jaedong Kim, Hyunggoog Seo, Seunghoon Cha, and Junyong Noh. 2019. Real-Time Human Shadow Removal in a Front Projection System. *Computer Graphics Forum* 38, 1 (2019), 443–454. doi:10.1111/cgf.13541 arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.13541>
- [24] Jan Klein and Gabriel Zachmann. 2004. Point Cloud Surfaces using Geometric Proximity Graphs. *Computers & Graphics* 28, 6 (Dec. 2004), 839–850. <http://dx.doi.org/10.1016/j.cag.2004.08.012>
- [25] Hiroki Kusuyama, Yuta Kageyama, Daisuke Iwai, and Kosuke Sato. 2024. A Multi-aperture Coaxial Projector Balancing Shadow Suppression and Deblurring. *IEEE transactions on visualization and computer graphics* PP (09 2024). doi:10.1109/TVCG.2024.3456170
- [26] Jijiang Li, Qingyue Deng, Haibin Ling, and Bingyao Huang. 2025. DPCS: Path Tracing-Based Differentiable Projector-Camera Systems. *IEEE Transactions on Visualization and Computer Graphics* (2025), 1–11. doi:10.1109/TVCG.2025.3549890
- [27] Peter Lincoln, Greg Welch, and Henry Fuchs. 2011. Continual surface-based multi-projector blending for moving objects. In *2011 IEEE Virtual Reality Conference*. 115–118. doi:10.1109/VR.2011.5759447
- [28] Sonia Mackenzie and Paula Foran. 2020. The impact of distractions and interruptions in the operating room on patient safety and the operating room team: An integrative review. *Journal of Perioperative Nursing* 33 (09 2020). doi:10.26550/2209-1092.1098
- [29] Kosuke Maeda and Hideki Koike. 2020. MirAIProjection: Real-time Projection onto High-speed Objects by Predicting Their 3D Position and Pose using DNNs. In *Proceedings of the 2020 International Conference on Advanced Visual Interfaces* (Salerno, Italy) (AVI '20). Association for Computing Machinery, New York, NY, USA, Article 59, 5 pages. doi:10.1145/3399715.3399839
- [30] Helena M Mentis, Amine Chellali, Kelly Manser, Caroline G L Cao, and Steven D Schwaiblmair. 2015. A systematic review of the effect of distraction on surgeon performance: directions for operating room policy and surgical training. *Surg Endosc* 30, 5 (July 2015), 1713–1724.
- [31] Leo Miyashita, Yoshihiro Watanabe, and Masatoshi Ishikawa. 2018. MIDAS projection: markerless and modelless dynamic projection mapping for material representation. *ACM Trans. Graph.* 37, 6, Article 196 (Dec. 2018), 12 pages. doi:10.1145/3272127.3275045
- [32] Daiki Miyazaki and Naoki Hashimoto. 2018. Dynamic projection mapping onto non-rigid objects with dot markers. In *2018 International Workshop on Advanced Image Technology (IWAIT)*. 1–4. doi:10.1109/IWAIT.2018.8369679
- [33] Andre Muehlenbrock, Roland Fischer, Christoph Schröder-Dering, René Weller, and Gabriel Zachmann. 2022. Fast, accurate and robust registration of multiple depth sensors without need for RGB and IR images. *The Visual Computer* (17 May 2022). doi:10.1007/s00371-022-02505-2
- [34] Andre Muehlenbrock, Rene Weller, and Gabriel Zachmann. 2024. BlendPCR: Seamless and Efficient Rendering of Dynamic Point Clouds captured by Multiple RGB-D Cameras. In *ICAT-EGVE 2024 - International Conference on Artificial Reality and Telexistence and Eurographics Symposium on Virtual Environments*, Shoichi Hasegawa, Nobuchika Sakata, and Veronica Sundstedt (Eds.). The Eurographics Association. doi:10.2312/egve.20241366
- [35] Takashi Nomoto, Wanlong Li, Hao-Lun Peng, and Yoshihiro Watanabe. 2020. Dynamic Projection Mapping with Networked Multi-projectors Based on Pixel-parallel Intensity Control. In *SIGGRAPH Asia 2020 Emerging Technologies* (Virtual Event, Republic of Korea) (SA '20). Association for Computing Machinery, New York, NY, USA, Article 11, 2 pages. doi:10.1145/3415255.3422888
- [36] Hanhoo Park, Moon-Hyun Lee, Byung-Kuk Seo, and Jong-Il Park. 2006. Undistorted Projection onto Dynamic Surface. 582–590. doi:10.1007/11949534\_58
- [37] Hao-Lun Peng and Yoshihiro Watanabe. 2020. High-Speed Human Arm Projection Mapping with Skin Deformation. In *SIGGRAPH Asia 2020 Emerging Technologies* (Virtual Event, Republic of Korea) (SA '20). Association for Computing Machinery, New York, NY, USA, Article 13, 2 pages. doi:10.1145/3415255.3422887
- [38] R. Raskar, M.S. Brown, Ruigang Yang, Wei-Chao Chen, G. Welch, H. Towles, B. Scales, and H. Fuchs. 1999. Multi-projector displays using camera-based registration. In *Proceedings Visualization '99 (Cat. No.99CB37067)*. 161–522. doi:10.1109/VISUAL.1999.809883
- [39] Guodong Rong and Tiow-Seng Tan. 2006. Jump flooding in GPU with applications to Voronoi diagram and distance transform. In *Proceedings of the 2006 Symposium on Interactive 3D Graphics and Games* (Redwood City, California) (I3D '06). Association for Computing Machinery, New York, NY, USA, 109–116. doi:10.1145/1111411.1111431
- [40] Christian Siegl, Matteo Colaianni, Lucas Thies, Justus Thies, Michael Zollhöfer, Shahram Izadi, Marc Stamminger, and Frank Bauer. 2015. Real-time pixel luminance optimization for dynamic multi-projection mapping. *ACM Trans. Graph.* 34, 6, Article 237 (Nov. 2015), 11 pages. doi:10.1145/2816795.2818111

- [41] Christian Siegl, Vanessa Lange, Marc Stamminger, Frank Bauer, and Justus Thies. 2017. FaceForge: Markerless Non-Rigid Face Multi-Projection Mapping. *IEEE Transactions on Visualization and Computer Graphics* 23, 11 (2017), 2440–2446. doi:10.1109/TVCG.2017.2734428
- [42] R. Sukthankar, Tat-Jen Cham, and G. Sukthankar. 2001. Dynamic shadow elimination for multi-projector displays. In *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*, Vol. 2. II–II. doi:10.1109/CVPR.2001.990943
- [43] Yukine Taya and Tokiichiro Takahashi. 2023. Dynamic projection mapping onto non-rigid objects using markers and skeletal information. In *International Workshop on Advanced Imaging Technology (IWAIT) 2023*. Masayuki Nakajima, Jae-Gon Kim, Kwang deok Seo, Toshihiko Yamasaki, Jing-Ming Guo, Phooi Yee Lau, and Qian Kemao (Eds.), Vol. 12592. International Society for Optics and Photonics, SPIE, 125922H. doi:10.1117/12.2666971
- [44] Daiki Tone, Daisuke Iwai, Shinsaku Hiura, and Kosuke Sato. 2020. FibAR: Embedding Optical Fibers in 3D Printed Objects for Active Markers in Dynamic Projection Mapping. *IEEE Transactions on Visualization and Computer Graphics* 26, 5 (2020), 2030–2040. doi:10.1109/TVCG.2020.2973444
- [45] Michal Tölgyessy, Martin Dekan, Luboš Chovanec, and Peter Hubinský. 2021. Evaluation of the Azure Kinect and Its Comparison to Kinect V1 and Kinect V2. *Sensors* 21, 2 (2021). doi:10.3390/s21020413
- [46] Zhou Wang, Alan Conrad Bovik, Hamid Rahim Sheikh, and Eero P Simoncelli. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE Trans Image Process* 13, 4 (April 2004), 600–612.
- [47] Yi Zhou, Shuangjiu Xiao, Ning Tang, Zhiyong Wei, and Xu Chen. 2016. Pmomo: Projection Mapping on Movable 3D Object. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (CHI '16). Association for Computing Machinery, New York, NY, USA, 781–790. doi:10.1145/2858036.2858329