

Article

An Appearance Optimisation Method for Projection-Based Spatial Augmented Reality

Lunan Wu , Federico Morosi *  and Giandomenico Caruso * 

Mechanical Engineering Department, Politecnico di Milano, Via La Masa, 1, 20156 Milano, Italy; lunan.wu@polimi.it

* Correspondence: federico.morosi@polimi.it (F.M.); giandomenico.caruso@polimi.it (G.C.)

Abstract

Projection-based spatial augmented reality (P-SAR) supports appearance-oriented design evaluation by projecting digital materials onto physical mock-ups, but the projected result may deviate from the intended screen-rendered appearance in both colour distribution and normal-induced shading. This paper proposes a screen-referenced, measurement-driven appearance optimisation framework in which a calibrated monitor serves as the visual reference and the projected mock-up as the optimisation target. The workflow combines controlled Unity rendering, colourimetric measurement, D65 CIE Lab analysis, and MATLAB (R2024b)-based iterative update, and separates the problem into albedo appearance optimisation and shading appearance optimisation. The albedo branch uses dominant-colour grouping and CIEDE2000-guided group-wise correction, while the shading branch uses a lightness-contrast descriptor derived from flat and normal-modulated renderings to update the normal-map-driven shading response. Experiments on ten material textures showed that all 23 identified colour groups converged below the adopted ΔE_{00} threshold of 2.3; the mean texture-level colour difference decreased from 6.24 to 1.36, corresponding to an average reduction of 77.43%. Comparative evaluation showed that the proposed group-wise optimisation outperformed global neutral-grey and global Lab-offset correction baselines. For shading, the mean residual $r_D = |D_{\text{proj}} - D_{\text{ref}}|$ decreased from 1.164 to 0.264 L^* units, and all ten normal-map cases satisfied the 1.0 L^* tolerance. A comparison with a global luminance-contrast baseline further supported the benefit of material-level normal-map update over image-domain contrast adjustment. Additional analyses examined the sensitivity to the number of dominant colour groups and clarified the rationale and scope of the adopted thresholds. Integrated photographic examples provided qualitative illustrations of the overall appearance tendency after the complete workflow, while the quantitative assessment was based on colourimetric and lightness-domain measurements. The full workflow required approximately 26 min per material case, indicating practical feasibility for controlled or semi-controlled P-SAR material appearance preparation and design-review scenarios.



Academic Editor: Stefanos Kollias

Received: 3 April 2026

Revised: 12 May 2026

Accepted: 15 May 2026

Published: 18 May 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\) license](https://creativecommons.org/licenses/by/4.0/).

Keywords: projection-based spatial augmented reality; colourimetric measurement; appearance optimisation; material textures

1. Introduction

Projection-based spatial augmented reality (P-SAR) offers a distinctive way of supporting design visualisation by projecting digital content directly onto physical artefacts.

In such settings, the physical mock-up retains its tangible and spatial qualities, while its visible surface appearance can be modified digitally. This makes P-SAR particularly relevant to mixed-prototyping workflows, where designers need to assess visual alternatives quickly on real objects rather than on screens alone. The importance of such visual assessment is well established in product design research: product appearance strongly influences how artefacts are interpreted, evaluated and chosen, and these effects extend beyond purely decorative considerations to broader cognitive and behavioural responses [1–4].

This importance is especially clear when material-related qualities are involved. In product development, materials and surface appearance contribute not only to technical performance but also to product character, perceived quality and user experience [5,6]. Recent work on materials experience has further shown that experiential qualities of materials are a meaningful subject of design evaluation in their own right, rather than a secondary outcome of engineering selection [6]. For design-oriented P-SAR applications, this means that the projected result should ideally preserve not only the presence of a graphic overlay, but also the intended appearance of digital materials, including their colour distribution and shading behaviour.

At the same time, the literature on prototyping suggests that fidelity matters when prototypes are used as decision-support tools. Prototype fidelity and aesthetic treatment can affect how prototypes are evaluated and interpreted during design activities [7]. Similar concerns have also been raised in mixed-reality prototyping, where the balance between physical and virtual fidelity influences how a prototype is understood and how useful it is for assessment [8,9]. These observations are directly relevant to P-SAR. Even if projection is correctly aligned geometrically, a mixed prototype cannot serve as a reliable visual reference when the reproduced appearance differs substantially from the intended digital one. In other words, appearance fidelity is not an optional refinement, but a prerequisite for credible appearance-oriented evaluation.

However, achieving such fidelity remains difficult in practice. The same digital material rendered on a calibrated screen and projected onto a physical white surface may show noticeable differences in colour and shading. These discrepancies arise from the combined effects of projector response, projection surface properties, and the interaction between projected light and physically rendered appearance. As a result, conventional projector calibration alone is often insufficient when the objective is not merely to display an image, but to reproduce a screen-defined material appearance in a form suitable for design judgement. This challenge becomes even more pronounced when the appearance to be reproduced includes both albedo-related colour distribution and normal-related shading variation, since these two components do not necessarily respond to the projection process in the same way.

Since closed-loop projector compensation, radiometric correction, and projector-camera feedback are already well established in P-SAR and projector compensation literature, this work focuses on a specific problem in appearance-oriented P-SAR: how to reduce the discrepancy between a screen-rendered material reference and the corresponding projected material appearance on a physical mock-up. In this context, the calibrated screen represents the digital material appearance that would normally be inspected during design evaluation before being transferred to a projected physical prototype.

To address this problem, this paper proposes a screen-referenced, measurement-driven appearance optimisation framework for P-SAR. In the proposed workflow, the screen-rendered result is treated as the visual reference, while the projected appearance on a physical white mock-up is treated as the optimisation target. A closed loop is established between controlled rendering, colourimetric measurement, D65 CIE Lab analysis and iterative update. Within this framework, the optimisation problem is divided into two com-

plementary parts: albedo appearance optimisation, which corrects base-colour distribution errors through dominant-colour group adjustment, and shading appearance optimisation, which compensates for discrepancies in normal-induced lightness contrast. This decomposition follows the structure of the rendered material representation, where the albedo map and normal map contribute differently to the final projected appearance.

The main contributions of this work are fourfold. First, it formulates projected material appearance optimisation as a screen-referenced P-SAR design-evaluation problem, rather than as generic projector calibration alone. Secondly, it proposes a material-component-aware optimisation workflow that separately addresses albedo-related colour-distribution mismatch and normal-induced shading mismatch using different measurement descriptors and update variables. Thirdly, it evaluates the proposed workflow through objective optimisation results, comparative baseline analysis, sensitivity analysis of the dominant colour-group number, and justification of the adopted colour and shading thresholds. Fourthly, it assesses the integrated appearance tendency through qualitative photographic examples and analyses the practical time burden of the workflow under controlled or semi-controlled P-SAR material appearance preparation conditions.

The remainder of the paper is organised as follows. Section 2 reviews previous work on P-SAR design visualisation, projector compensation, and objective evaluation of appearance consistency. Section 3 describes the experimental platform and measurement workflow. Section 4 presents the proposed albedo and shading appearance optimisation methods. Section 5 reports the optimisation results, baseline comparisons, sensitivity analyses, qualitative integrated appearance illustrations and practical assessment. Section 6 discusses the implications, limitations and future directions of the work.

2. Related Works

P-SAR has been widely explored as a means of visualising digital content directly on physical artefacts, especially in design and prototyping contexts. Its value lies in combining the flexibility of digital modification with the tangibility, scale awareness and shared inspectability of real objects. However, for appearance-oriented design review, the usefulness of P-SAR depends not only on registration, interaction or display capability, but also on whether the reproduced surface appearance remains sufficiently consistent with the intended visual reference. This issue becomes particularly important when the projected content represents a digital material, because designers may need to judge colour, texture and shading qualities rather than only the presence of an overlay. Against this background, the following review focuses on three closely related areas: the role of P-SAR in design visualisation, the problem of appearance reproduction and projector compensation, and the objective and perceptual evaluation of appearance consistency.

2.1. P-SAR for Design Visualisation

P-SAR belongs to the broader family of Augmented Reality (AR) technologies, but differs from head-mounted and handheld approaches in one important respect: digital content is registered directly onto real surfaces through projection, so the physical artefact itself becomes the display medium. In the wider AR literature, Azuma's survey remains a standard reference for the definition and scope of AR [10]. Within this broader context, early projector-based systems such as The Office of the Future and Shader Lamps established important technical foundations for spatially aligned projection and for altering the perceived appearance of physical objects by projected imagery [11,12]. These early systems were not yet focused on industrial design workflows, but they demonstrated the core principle that later enabled P-SAR to be used as a design visualisation medium.

From a design perspective, the interest in P-SAR lies in its ability to combine the tangibility and scale awareness of a physical mock-up with the flexibility of digital modification. This directly addresses a familiar limitation of conventional prototyping: physical prototypes are intuitive to inspect and discuss, but they are costly and slow to remake, whereas fully digital representations are easy to modify but may be harder for non-expert stakeholders to interpret. Research on augmented prototyping therefore argued early on that mixed physical–digital representations should be assessed in relation to actual design practice rather than technical novelty alone [13]. More recent review work on mixed reality in design prototyping has reinforced this view, showing that such systems can create value not only for the prototype itself, but also for communication, iteration and decision-making within the wider design process [14].

This rationale is reflected in a series of application-oriented studies. Nam and Lee explored AR-based prototyping for digital products, emphasising the early integration of hardware and software representations [15]. Akaoka et al. then showed that projected interfaces could be applied to inexpensive physical models made from Styrofoam, paper or cardboard, making functional prototyping quicker and more accessible [16]. In a more appearance-oriented direction, Park et al. proposed the use of SAR specifically for product appearance design evaluation, demonstrating that projected visualisation on a physical mock-up can support more intuitive judgement of surface appearance [17]. Later studies extended this logic to co-creative and industrial settings. For example, Cascini et al. investigated projection-based AR for co-creative product and packaging design, while Poulin et al. examined how SAR affects user participation in co-design sessions [18,19]. Recent scoping work has likewise emphasised that SAR remains valuable in application settings where shared, non-wearable viewing and direct projection onto physical surfaces are important [20].

Taken together, these studies show that P-SAR has evolved from a technical display concept into a credible design representation medium, especially in tasks where surface look, graphics, colour, texture and shared inspection of a real object are important. At the same time, the existing design-oriented literature also reveals a clear limitation. Most studies evaluate P-SAR in terms of collaboration, usability, process integration, ideation support or general representational value [13–15,19]. Comparatively less attention has been paid to a narrower but critical question for appearance-oriented design review: whether the projected appearance is sufficiently faithful to a digital visual reference to support reliable judgement of materials and surface qualities. This issue is particularly relevant when the projected surface appearance is generated from digital material maps, because designers are not only evaluating whether projection is engaging or usable, but whether the reproduced material alternative remains visually comparable with the screen-rendered reference. For this reason, the present work moves from the general design-use perspective of P-SAR towards the more specific problem of screen-referenced appearance reproduction and compensation, which is reviewed in the next section.

2.2. Appearance Reproduction and Projector Compensation

Work on appearance reproduction in projection systems initially developed from a colour-characterisation perspective. Early studies treated the projector primarily as a colour output device whose response had to be modelled before reliable reproduction could be attempted. Liaw et al., for example, characterised the forward colour behaviour of a liquid crystal projector using regression- and lookup-table-based approaches [21]. As projector–camera systems became more widely studied, this device-level view expanded into a broader research area concerned with displaying images on surfaces that are not ideal projection screens. Reviews by Bimber et al. and Grundhöfer and Iwai show that

projector compensation has since evolved into a rich field covering geometric warping, radiometric correction, calibration, surface reflectance effects and projection mapping on coloured, textured or geometrically complex surfaces [22,23]. These studies established the device-level and system-level basis for later projector compensation methods.

A major strand of this literature focuses on radiometric compensation, that is, modifying the projector input so that the physical projection more closely matches a desired appearance despite surface imperfections or environmental disturbances. Nayar et al. addressed this problem for non-ideal screens through offline calibration and compensation [24]. Grossberg et al. then framed the problem more directly as one of controlling appearance, aiming to make one object look like another by projecting a carefully computed compensation image [25]. Later work pushed this idea towards greater practicality. In particular, Grundhöfer and Bimber proposed real-time adaptive radiometric compensation, showing that compensation can be adjusted dynamically according to both the projection surface and the projected content [26]. Together, these studies established the central logic of projector compensation: rather than projecting the desired image directly, the input image must be altered so that the resulting appearance on the physical surface better approximates the intended target. The present work adopts this general principle of measurement-guided correction, but applies it to a different level of representation, namely rendered material components rather than only the final projected image.

As compensation research moved closer to SAR, the problem also became more clearly one of appearance fidelity rather than simple projector correction. Wetzstein and Bimber generalised radiometric compensation through inverse light transport, allowing the compensation process to account for more complex optical effects such as inter-reflection [27]. Sheng et al. further addressed indirect illumination in SAR, showing that global illumination can substantially disrupt the final appearance if it is not explicitly compensated [28]. Menk and Koch then moved even closer to design-relevant SAR by focusing on truthful colour reproduction, explicitly addressing whether the reproduced colour appearance on physical objects matches the intended visual result [29]. A related but slightly different direction was explored by Jones et al. in *Projectibles*, where the optimisation problem involved both projected light and the colour of the display surface itself in order to improve the final visual result [30]. These works are particularly relevant because they move beyond simple device calibration and address the perceived appearance of projected physical objects. They also show that, once projection is applied to real objects and real surfaces, appearance fidelity depends not only on projector response, but also on light transport, surface properties and the interaction between projected imagery and the physical display medium.

More recent work has introduced data-driven and full-compensation formulations. Huang and Ling proposed *CompenNet*, which models photometric compensation as an end-to-end learning problem [31]. This line was then extended to joint geometric and photometric compensation in *CompenNet++* and later in the journal version on end-to-end full projector compensation [32,33]. Li et al. recently revisited full compensation from a physics-based perspective and proposed an efficient framework that estimates projector response, geometry and surface reflectance using only natural images [34]. These advances have substantially improved projector compensation on challenging surfaces. Related work in *Electronics* has also addressed the geometric side of projection correction, for example through binocular-vision-based multi-projection correction on special-shaped screens, further highlighting the importance of accurate projector–surface mapping in non-planar display conditions [35]. These learning-based, physics-based and geometry-aware compensation frameworks are highly relevant to the broader projector compensation problem. However, they are usually formulated as camera-based image-level or full

projector compensation problems, in which a projector input image is optimised so that the camera-observed projection matches a desired target image. Their acquisition and evaluation protocols are therefore different from the colourimeter-based, screen-referenced material-component evaluation adopted in the present work.

This distinction is central to the positioning of the present study. Existing compensation methods have established powerful ways to correct projection errors at the projector-input, camera-observed image, surface-reflectance or light-transport level. However, they do not usually formulate the optimisation in terms of rendered material components such as albedo-related colour distribution and normal-induced shading response. In design-oriented P-SAR, the screen-rendered material appearance is not merely an arbitrary target image. It is the visual reference that designers normally inspect before applying the same material alternative to a projected physical mock-up. The projected result is therefore expected to remain comparable with this screen-side material reference, not only in overall image appearance but also in the separate contributions of base colour and shading.

For this reason, the present work differs from general image-level projector compensation in three related respects. First, the reference condition is a calibrated screen-rendered material appearance used for design judgement, rather than only a camera-observed target image. Secondly, the representation level is material-based: the albedo map and normal map are treated as distinct contributors to the final projected appearance. Thirdly, the correction variables and evaluation criteria are component-specific. The albedo branch operates on dominant colour groups and uses CIE Lab/CIEDE2000-based convergence, whereas the shading branch uses a uniform-albedo baseline and a lightness-domain contrast descriptor to adjust normal-induced shading response. The proposed workflow should therefore be understood not as a replacement for general projector compensation, but as a material-component-aware extension of appearance compensation principles to screen-referenced P-SAR design-review scenarios.

2.3. Objective Evaluation of Appearance Consistency

The evaluation of appearance consistency is closely related to human visual perception, because the usefulness of a visual representation ultimately depends on whether it can support reliable interpretation and design judgement. This is especially important in appearance-oriented design review, where the relevant question is not only whether two outputs are numerically close, but whether the reproduced appearance remains sufficiently comparable with the intended reference for material and surface-quality evaluation. At the same time, an iterative optimisation workflow requires stable, repeatable and device-independent quantities that can be measured after each update. For this reason, the image-quality literature has long treated objective assessment as a practical basis for repeated comparison, parameter tuning and algorithmic optimisation, while recognising that such metrics are approximations of visual appearance rather than complete descriptions of perception [36,37].

In the context of the present work, this distinction is particularly important. The aim is to reduce the appearance discrepancy between a screen-rendered material reference and the corresponding projected physical appearance, but the optimisation loop itself must be driven by measurable quantities rather than by direct design judgement at every iteration. Objective metrics are therefore used here as operational criteria for guiding the optimisation process. These metrics are selected because they are physically measurable, repeatable across iterations, and perceptually motivated in the sense that they are defined in colour and lightness spaces related to visual appearance. However, they should not be interpreted as complete substitutes for all possible forms of human judgement. In the present study, the quantitative evidence is therefore based on objective colourimetric and

lightness-domain measurements, while photographic examples are used only as qualitative integrated illustrations of the final appearance tendency.

For colour-related appearance evaluation, the objective basis is typically grounded in colourimetry. The CIE 1976 $L^*a^*b^*$ colour space was standardised as an approximately uniform space for representing colour in terms of lightness and chromatic components [38]. It also provides a more perceptually meaningful framework for expressing colour differences than raw RGB or XYZ coordinates alone. Nevertheless, later work showed that simple Euclidean distances in CIELAB do not always correspond equally well to perceived colour differences across different regions of colour space. This led to the development of CIEDE2000 (ΔE_{00}), which introduces corrections for lightness, chroma, and hue, and their interactions, in order to improve agreement with perceived colour difference [39]. The formulation and practical implementation of ΔE_{00} have since become well established through both the original derivation and later implementation notes [39,40]. These developments are directly relevant to the albedo branch of the present study, where the optimisation target concerns colour consistency between the screen reference and the projected result.

At the same time, appearance consistency is not determined by colour difference alone. The broader image-quality literature has shown that perceived similarity also depends on luminance, contrast and spatial structure. A representative example is the structural similarity index (SSIM), which evaluates similarity by considering structural relationships rather than relying only on pointwise error [41]. Related work on colour image quality has further attempted to combine structural information with colour-difference measures such as ΔE_{00} in order to improve agreement with subjective perception [42,43]. A comparable use of CIEDE2000 has also been reported in immersive graphics research to quantify the fidelity between measured physical samples and their virtual reproductions under controlled lighting conditions [44]. These studies are useful because they show that appearance evaluation cannot be reduced to a single scalar colour error without considering the broader visual context.

For the present study, this observation motivates the use of different objective descriptors for different material components. Albedo-related mismatch is evaluated using CIEDE2000 because it concerns the base-colour appearance of the material. By contrast, normal-induced shading mismatch is evaluated under a uniform-albedo condition using a lightness-domain contrast descriptor. This choice reflects the role of the normal map in the controlled rendering setup: once chromatic texture variation is suppressed, the remaining difference of interest is mainly the strength of lightness modulation caused by normal-dependent shading. The proposed lightness-based descriptor is therefore used as an operational measure of shading consistency, rather than as a general image-quality metric for all aspects of material appearance.

For the problem addressed in this paper, image-based similarity metrics alone are also not sufficient. The task is not simply to compare two digital images generated in the same medium, but to compare a screen-rendered reference with a physically projected appearance produced by a different device on a real surface. Under these conditions, a physically grounded cross-device comparison is required. The present work therefore combines colourimeter-based measurements in CIE XYZ and D65 CIE Lab space with component-specific appearance descriptors. This measurement-driven evaluation provides a repeatable basis for optimisation and allows the albedo and shading branches to be assessed under a common screen-to-projection reference–target framework.

3. Experimental Platform and Workflow

This section describes the experimental platform developed for the screen-referenced appearance optimisation framework. The platform integrates controlled rendering, pro-

jection display, objective colour measurement, and iterative optimisation within a unified workflow. The rendering environment is implemented in Unity, while colour measurements are obtained using a calibrated monitor, a projector, and a tristimulus colourimeter under controlled laboratory conditions. The setup was designed as a controlled planar validation environment for evaluating the proposed material appearance optimisation method, rather than as a complete deployment solution for arbitrary real-world SAR conditions. The following sections present the system overview and data flow, the hardware and measurement configuration, the Unity-based rendering environment, and the reference–target definition and measurement data handling used throughout the optimisation process.

3.1. System Overview and Data Flow

The experimental platform developed in this study was designed as a closed-loop system for screen-referenced material appearance optimisation in P-SAR. The central idea was to treat the screen-rendered material appearance as the reference condition and the projected appearance on a physical white mock-up as the optimisation target. Rather than performing projector calibration alone, the system aimed to reduce the appearance discrepancy between these two outputs through repeated measurement and material-level update. In this context, the calibrated screen served as the design-side visual reference, while the projector and the physical projection surface formed the reproduction branch to be optimised.

At a system level, the workflow consisted of six main stages: controlled rendering in Unity, display on the reference screen and projector, colour measurement using a colourimeter, conversion of the measured data from CIE XYZ to D65 CIE Lab, optimisation in MATLAB, and re-rendering with the updated texture or parameter set. The loop was repeated until the projected appearance approached the screen reference under the selected branch-specific evaluation criterion. The overall closed-loop structure of the experimental platform is illustrated in Figure 1.

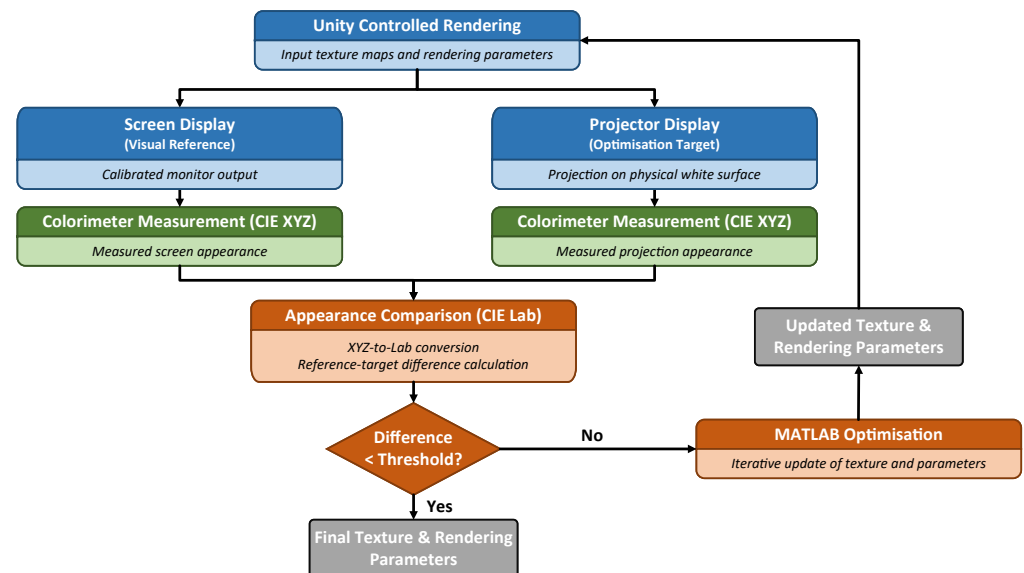


Figure 1. Closed-loop workflow of the proposed screen-referenced appearance optimisation system.

Within this framework, Unity was used to provide a controlled rendering environment. It generated the material appearance shown on both the screen and the projection surface under predefined rendering conditions. The same scene structure, shader configuration and material representation were used for both branches, so that the differences measured during the experiment were mainly reflected by the screen-to-projection transfer rather

than changes in the digital material definition itself. The screen output served as a stable visual reference, while the projector reproduced the same material appearance on a physical white surface. The resulting appearances were then measured with a colourimeter at selected locations, producing device-independent XYZ tristimulus values. These measurements were subsequently converted into D65 CIE Lab space, where perceptual colour and lightness differences could be analysed more appropriately and used to guide the optimisation process.

For each optimisation iteration, the data flow followed a fixed sequence. First, Unity loaded the current material input, including the albedo map, normal map, and relevant rendering parameters. Secondly, the required colour or shading patches were generated and displayed on the calibrated monitor and on the projection surface. Thirdly, the screen-side and projection-side appearances were measured separately using the colourimeter, with the raw values recorded as CIE XYZ tristimulus data. Each measurement was repeated three times and averaged before being used in the optimisation. Fourthly, the averaged XYZ values were converted to D65 CIE Lab in MATLAB, where the branch-specific residual was calculated. For the albedo branch, the residual was evaluated using the CIEDE2000 colour difference for each dominant colour group. For the shading branch, the residual was evaluated using a lightness-based contrast descriptor derived from flat and normal-modulated patches. Finally, MATLAB generated the updated material input or rendering-related parameter state, which was reloaded into Unity for the next iteration.

MATLAB was employed as the optimisation backend. Based on the measured discrepancy between the reference and projected appearances, it generated updated texture data or rendering-related parameters, which were then fed back into the Unity scene. In this way, the platform established a measurement-driven implementation loop linking digital rendering, physical projection, objective colour acquisition, and iterative appearance compensation. The purpose of this design was to provide a reproducible basis for improving projected material appearance in controlled P-SAR scenarios, while keeping the rendering and measurement stages clearly separated from the optimisation algorithms themselves.

For reproducibility, the iteration process recorded the material identifier, optimisation branch, iteration number, measured XYZ values, converted Lab values, residual metric, update parameter, output file or parameter state, and convergence status. These records made it possible to trace how the measured screen–projection discrepancy was transformed into each subsequent material update.

For clarity, the following sections describe the platform from four complementary perspectives. Section 3.2 introduces the hardware configuration and measurement conditions. Section 3.3 presents the Unity-based rendering environment. Section 3.4 defines the reference–target relationship and the colourimetric basis used throughout this study. The optimisation methods built upon this platform are then described in Section 4.

3.2. Hardware and Measurement Setup

The experimental platform combined a reference display, a projector, a tristimulus colourimeter, a matte projection surface, and a workstation running the rendering and optimisation software. The reference display was a Philips 288P6L 4K liquid crystal display (LCD) monitor with a white-LED (W-LED) backlight. It was used as the screen-side visual reference throughout the study. The projector was an Epson EB-PQ2008B based on 3LCD technology with a laser phosphor light source and a native resolution of 3840 × 2160. Both devices were connected to the same workstation, which was equipped with an Intel Core i7-6800K processor, dual NVIDIA GTX 1080 graphics cards, 64 GB of memory, and Windows 11. Figure 2 illustrates the overall hardware configuration used in the experiments.




Figure 2. Hardware and measurement setup used in the experiments. The calibrated monitor provides the screen-side reference, while the projector displays the corresponding material appearance on a matte white projection surface.

The monitor was configured to its maximum brightness, corresponding to approximately 300 candela per square metre (cd/m^2), and was calibrated using DisplayCAL with ArgyllCMS as the measurement backend. The target conditions were a white point of 6500 Kelvin (K) and gamma 2.2. After calibration, the resulting ICC profile was loaded into the system. The average calibration error was approximately $\Delta E = 0.2$ and the maximum error was $\Delta E = 1.6$, indicating that the display provided a sufficiently stable and colour-accurate reference for the subsequent screen-to-projection comparison tasks.

The projector was operated in economy mode. Under this mode, its maximum output was lower than its nominal full-power specification and therefore more suitable for matching the luminance level of the reference monitor. To further reduce the luminance gap between the two devices, the projector light output was set to its minimum value of 8%. By contrast, the image brightness control was not adjusted, so that the video signal level and projected contrast would not be altered unnecessarily. The projector was configured in Presentation colour mode with a colour temperature of 6500 K and gamma 2.2. Dynamic Contrast and Image Enhancement were disabled in order to avoid automatic image modifications during the experiments. In addition, the projector's built-in camera-assisted colour calibration procedure was used as an initial practical adjustment step to reduce obvious colour bias before the formal measurements.

The measuring instrument was a Calibrite Display Pro HL, which is a tristimulus colourimeter. Unlike a spectrophotometer, a colourimeter estimates colour by using filtered sensors to approximate human visual response. Because the spectral power distributions of different display technologies can differ substantially, direct colourimeter readings may be biased if no spectral correction is applied. For this reason, colour measurements were performed with display-specific Colourimeter Calibration Spectral Sample (CCSS) correction files. Separate CCSS corrections were loaded for the reference monitor and the projector before measurement, so that the colourimeter response could better match

the spectral characteristics of each device. These corrections were used to reduce device-dependent measurement bias within the available setup, although they do not replace full spectrophotometric characterisation of the display and projector.

The projection target was a white matte cardboard surface measuring 1.2 m × 0.6 m. This surface was chosen to approximate the neutral white physical mock-ups commonly used in mixed prototyping applications of P-SAR, while minimising substrate-induced colour variation during this first-stage validation of the optimisation method. The board was mounted vertically on a desk rather than being aligned orthogonally to the projector optical axis. This arrangement was adopted mainly for practical reasons, including easier access for colour measurement and a stable central measurement region. The projector was mounted on a tower stand with a downward tilt of approximately 45°, and the zoom setting was adjusted so that the central measurement region on the cardboard remained sharply focused.

All experiments were conducted in a darkened indoor environment. Ambient lighting was removed by switching off all room lights other than the monitor and projector and by closing the curtains. Before each measurement session, both the monitor and the projector were warmed up for more than 30 min in order to reduce short-term instability in luminance and chromaticity. The same display and projector settings were then maintained throughout the corresponding measurement session. This procedure was used consistently across the albedo and shading optimisation experiments.

Formal colour measurements were performed with the ArgyllCMS command-line tool `spotread`. The `-X` option was used to load the appropriate CCSS correction file before measurement. For the monitor, the colourimeter was used in contact mode, with the sensor placed directly against the display surface. For the projector, non-contact measurement was adopted: the sensor was mounted on a tripod and positioned approximately 30 cm in front of the projection surface, facing the centre of the measured region. During projection measurements, the same central region of the projection surface was used to reduce the influence of spatial variation in focus, luminance and projection geometry.

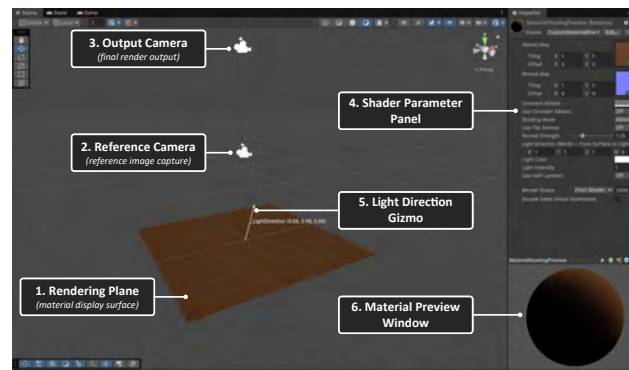
The screen-side and projection-side measurements followed the same acquisition protocol. For each patch or rendered measurement condition, the screen-side value was measured as the reference and the projection-side value was measured as the current target state. Each measurement was repeated three times, and the average value was used in the subsequent analysis. This averaging procedure was adopted to reduce the effect of short-term measurement fluctuations while keeping the measurement time manageable.

The raw measurement data were recorded as CIE XYZ tristimulus values. XYZ was selected as the acquisition format because it is device-independent and directly provided by the measuring instrument. However, since XYZ is not perceptually uniform, the measured values were subsequently converted into D65 CIE Lab for appearance comparison and optimisation. The conversion and branch-specific residual calculations were performed in MATLAB as part of the optimisation workflow described in Section 4.


3.3. Rendering Environment in Unity

The rendering component of the experimental platform was implemented in Unity (version 6000.0.58f2) using the Universal Render Pipeline (URP). Unity was employed to generate the visual appearance of digital materials under controlled conditions and to display the rendered results on both the reference monitor and the projector. The same Unity scene, shader configuration and material input structure were used for the screen and projection branches, so that the digital material representation remained identical while the differences introduced by the display and projection processes could be measured and analysed.


A dedicated virtual scene was constructed to provide a controlled rendering environment. At the centre of the scene, a planar geometry was used to display the material textures. This surface represented the digital material under investigation and served as the rendering target for both the screen-side reference and the projection-side output. The use of a planar rendering surface also ensured a direct correspondence between the rendered material region and the physical projection region used in the controlled validation setup. The geometry was rendered using a custom shader designed specifically for controlled appearance analysis. Figure 3a shows the Unity scene constructed for controlled material rendering.




(a)



(b)



(c)



(d)

Figure 3. Unity-based controlled rendering environment. (a) Annotated view of the Unity scene used for controlled material appearance rendering. (b) AlbedoOnly mode, in which only the albedo texture is rendered without shading influence. (c) NormalOnly mode, in which only the shading generated by the normal map is visualised. (d) AlbedoAndNormal mode, in which the albedo texture and the normal-based shading are rendered together.

The custom shader enabled configurable rendering of material textures. The shader accepted an albedo map and a normal map as its primary inputs and supported three rendering modes: AlbedoOnly, NormalOnly, and AlbedoAndNormal, as illustrated in Figure 3b–d. The first mode displayed only the albedo texture without lighting influence. The second mode visualised the shading produced solely by the normal map under the controlled lighting model. The third mode combined both components to produce the final shaded material appearance. This separation of rendering modes was important for the proposed optimisation framework, because it allowed albedo-related colour appearance and normal-induced shading behaviour to be generated, measured and analysed separately before being recombined at the integrated appearance level.

The Unity project was operated using a consistent linear colour-space workflow throughout the experiments. Albedo maps were treated as colour texture inputs, whereas normal maps were treated as normal-map inputs and decoded in the shader before shading calculation. The rendered RGB values were used as commanded display inputs for the screen and projector branches, while the quantitative comparison was based on the mea-

sured CIE XYZ values and their conversion to D65 CIE Lab. In this way, the optimisation did not rely on Unity RGB values as direct colour evidence, but used them as controlled input states whose physical outputs were measured with the colourimeter.

Lighting within the scene was defined by a single directional light model implemented directly inside the shader in world space. The direction and intensity of this light source could be adjusted during runtime, enabling controlled variation of shading conditions. During each experimental comparison, the lighting configuration was kept fixed so that changes in the measured appearance could be attributed to the material input state and the screen-to-projection transfer rather than to uncontrolled changes in scene illumination. In addition, the strength of the normal map could be modified to control the amplitude of surface detail and therefore the contrast of the resulting shading.

The shader also provided an optional constant albedo condition in which the spatially varying albedo texture was replaced by a uniform colour. This condition was used in the shading-oriented part of the workflow to suppress texture-dependent colour variation and isolate the effect of lighting and normal-map modulation. In particular, it supported the uniform-albedo baseline calibration and the construction of flat and mean patches described in Section 4.2. This terminology is used here to describe a rendering condition within the proposed workflow, rather than a comparative baseline method.

The scene provided runtime control over material textures and rendering parameters. Multiple material texture sets could be loaded and switched during execution, allowing different materials to be evaluated under identical rendering conditions. Parameters such as rendering mode, constant albedo state, normal strength, and light intensity could be adjusted interactively. Once a measurement sequence began, these parameters were kept fixed unless they were explicitly updated by the corresponding optimisation step. This ensured that the measured differences between iterations reflected the intended material update rather than accidental changes in the rendering environment.

Two cameras were defined for different purposes. A preview camera provided a real-time view of the rendered material for the operator. A second camera was used to generate standardised render outputs for analysis and patch construction. This camera employed an orthographic projection positioned above the rendering surface in order to eliminate perspective distortion and ensure that the rendered surface consistently filled the image frame. The rendered output could then be exported as PNG images for further processing, patch generation and comparison in MATLAB. The same export procedure was used across the tested materials to maintain consistency between iterations.

To ensure that the lighting conditions were fully controlled, all default Unity light sources were removed from the scene. Environmental illumination was also disabled by turning off global ambient lighting and removing the skybox. As a result, the shading observed in the rendered images originated solely from the directional lighting model defined in the custom shader. This configuration ensured that the rendering conditions remained deterministic and reproducible throughout the experiments, providing a stable basis for the screen-referenced measurement and optimisation workflow.

3.4. Reference–Target Definition and Data Handling

The optimisation framework used a consistent reference–target definition throughout the study. For each measurement condition, the screen-side output was treated as the reference, and the corresponding projection-side output was treated as the current target state. The term “reference” therefore refers to the material appearance rendered on the calibrated monitor, whereas the term “target state” refers to the appearance reproduced by the projector on the physical white surface under the same controlled rendering condition.

This definition was used independently of the specific optimisation branch, so that all later comparisons followed the same screen-to-projection logic.

The comparison was performed on measured colourimetric data rather than on rendered RGB values or camera-recorded photographs. For a given measurement condition i and iteration t , the screen-side measurement was denoted as

$$\mathbf{X}_i^{\text{ref}} = [X_i^{\text{ref}}, Y_i^{\text{ref}}, Z_i^{\text{ref}}]^T. \quad (1)$$

and the corresponding projection-side measurement was denoted as

$$\mathbf{X}_i^{\text{proj}}(t) = [X_i^{\text{proj}}(t), Y_i^{\text{proj}}(t), Z_i^{\text{proj}}(t)]^T. \quad (2)$$

Here, $\mathbf{X}_i^{\text{ref}}$ represents the measured screen reference for the selected condition, while $\mathbf{X}_i^{\text{proj}}(t)$ represents the measured projected result at the current iteration. When no update had yet been applied, $t = 0$ denoted the initial projection state.

As described in Section 3.2, each XYZ measurement was repeated three times. The average XYZ value was used for all subsequent calculations:

$$\bar{\mathbf{X}} = \frac{1}{3} \sum_{r=1}^3 \mathbf{X}_r, \quad (3)$$

where r denotes the repeated measurement index. This averaging step was applied before colour-space conversion, so that the optimisation process used a single stabilised XYZ triplet for each screen-side or projection-side measurement condition.

The averaged XYZ values were then converted into D65 CIE Lab space:

$$\mathbf{C}_i^{\text{ref}} = [L_i^{*\text{ref}}, a_i^{*\text{ref}}, b_i^{*\text{ref}}]^T, \quad \mathbf{C}_i^{\text{proj}}(t) = [L_i^{*\text{proj}}(t), a_i^{*\text{proj}}(t), b_i^{*\text{proj}}(t)]^T. \quad (4)$$

The notation \mathbf{C} is used here to distinguish the full Lab colour vector from the lightness component L^* . The D65 reference white was used consistently because both the monitor calibration and the projector colour-temperature setting were defined around 6500 K. This provided a common colourimetric basis for comparing the screen and projection branches.

The resulting paired Lab values,

$$(\mathbf{C}_i^{\text{ref}}, \mathbf{C}_i^{\text{proj}}(t)), \quad (5)$$

formed the input data for the branch-specific residual calculations described in Section 4. The present subsection only defines the common measurement representation and reference–target convention; the specific residual metrics, update rules, and convergence criteria are introduced separately in the corresponding optimisation methods. This separation ensured that the experimental platform provided a unified data basis, while the albedo and shading procedures could use different descriptors appropriate to their respective appearance components.

All measurement records were associated with the material identifier, measurement condition, display branch, iteration index, averaged XYZ values, converted Lab values, and output state. In this way, the data handling process preserved a direct trace from the physical measurement to the subsequent optimisation input, while avoiding reliance on device-dependent RGB values or photographic documentation as quantitative evidence.

4. Appearance Optimisation Methods

This section presents the methods used to optimise the appearance of projected materials within the experimental framework described in Section 3. The optimisation process addresses two complementary aspects of material appearance: the colour distribution represented by the albedo texture and the shading behaviour associated with surface normals. For this purpose, two proposed optimisation procedures are introduced. The first focuses on correcting albedo appearance through a screen-referenced, group-wise adjustment of dominant colour groups. The second focuses on normal-map-based shading appearance optimisation, in which the normal-induced lightness contrast is adjusted according to screen-to-projection measurements.

The detailed implementation parameters and branch-specific pseudo-code of the two proposed procedures are provided in Appendix A. The main text focuses on the methodological rationale, measurement descriptors, update equations, and convergence criteria. In addition, Section 4.3 defines the comparative baseline methods used in the experimental evaluation. These baselines are not proposed as alternative contributions of this paper, but are included to test whether the proposed material-component-aware formulation provides advantages over simpler global correction strategies under the same screen-referenced, colourimeter-based measurement protocol.

4.1. Albedo Appearance Optimisation

This section describes the proposed optimisation method used to correct albedo-related colour appearance differences between the screen reference and the projected result. Instead of applying pixel-level corrections, the method represents the source albedo texture as a set of dominant colour groups and performs measurement-driven compensation at the group level. Each group is measured and updated independently according to its own screen-to-projection colour residual. The final group-wise corrections are then propagated back to the full-resolution texture through soft spatial masks, producing an optimised albedo map for projection rendering. The overall workflow of the proposed albedo optimisation method is illustrated in Figure 4, and the detailed implementation sequence is summarised in Algorithm A1 in Appendix A.2.

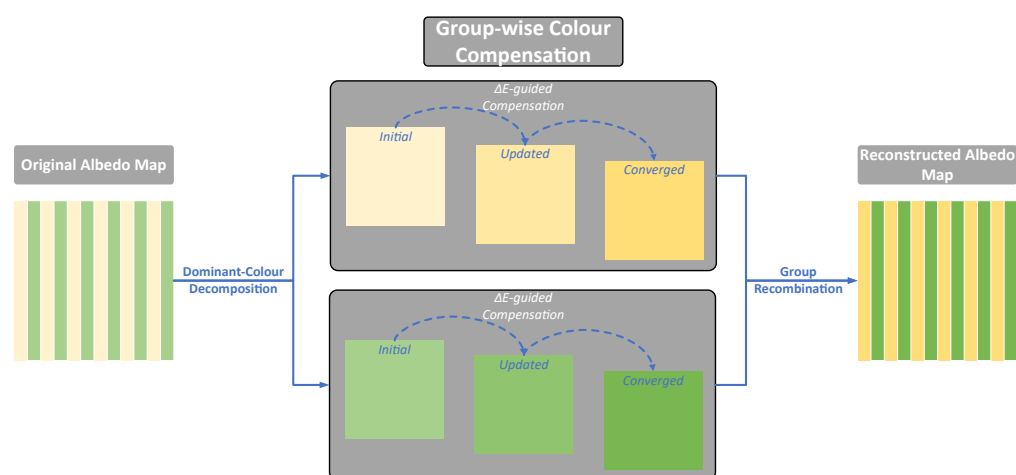


Figure 4. Overview of the albedo map optimisation workflow. The source albedo map is first decomposed into a small number of dominant colour groups. Each group is then optimised independently through a screen-referenced iterative compensation process driven by colour measurements. Finally, the measured group-wise corrections are propagated through soft spatial masks to reconstruct the optimised albedo map for projection rendering.

4.1.1. Dominant Colour Decomposition of the Source Albedo Map

Direct pixel-wise optimisation was not adopted for albedo correction, because such an approach would require an impractically large number of measurements and would be highly sensitive to local noise, minor spatial inconsistencies, and measurement-position variations. Instead, the source albedo texture was first converted into a compact spatial-chromatic representation composed of a small number of dominant colour groups. This representation preserved the main colour structure of the texture while reducing the optimisation problem to a manageable set of group-level variables.

The original albedo map was first transformed from RGB into D65 CIE Lab space, so that the subsequent grouping process was performed in a colour space that better reflects perceptual differences. In accordance with the notation introduced in Section 3.4, the Lab colour vector at pixel location (x, y) is denoted as

$$\mathbf{C}_{\text{src}}(x, y) = [L_{\text{src}}^*(x, y), a_{\text{src}}^*(x, y), b_{\text{src}}^*(x, y)]^T. \quad (6)$$

The texture pixels were then clustered in Lab space using k -means, producing K initial colour groups. In the present workflow, K was kept small, typically between one and three, in order to maintain a balance between representational capacity and measurement efficiency. In the current implementation, K was selected by visually inspecting the clustering result and the number of perceptually distinct dominant colour regions. Therefore, K should be understood as a representational parameter controlling how much of the texture's chromatic structure is retained, rather than only as a numerical clustering parameter. Its influence is further examined in the sensitivity analysis reported in Section 5.

Each cluster represented a dominant colour component of the texture, rather than a fine-grained semantic region. After clustering, the groups were sorted by area in descending order to provide stable group identities across subsequent optimisation and reconstruction stages. The area proportion of group k is denoted as a_k , and is later used in the comparative global-offset baseline defined in Section 4.3.

The initial clustering result yielded hard masks, in which each pixel belonged to only one group. However, a strictly hard partition was not suitable for later reconstruction, because abrupt mask boundaries would produce discontinuities when group-level corrections were mapped back to the full-resolution texture. To avoid this problem, the hard masks were converted into soft spatial masks by Gaussian smoothing. The smoothed masks were then normalised so that, at every pixel location, the weights of all groups summed to one:

$$\sum_{k=1}^K w_k(x, y) = 1. \quad (7)$$

In this way, each pixel was represented not by a single discrete label, but by a set of continuous group weights, allowing gradual transitions between neighbouring colour regions.

Based on these normalised soft masks, a representative source Lab centre was computed for each group as the weighted average of the original texture values:

$$\mathbf{C}_k^{\text{src}} = \frac{\sum_{x,y} w_k(x, y) \mathbf{C}_{\text{src}}(x, y)}{\sum_{x,y} w_k(x, y)}. \quad (8)$$

These weighted centres served as the initial colour descriptors of the source texture and formed the starting point for the screen-referenced compensation described in the next section. At the same time, the soft masks preserved the spatial contribution of each group over the entire image, thereby establishing the link between later group-level optimisation and final texture reconstruction. The decomposition stage therefore produced three outputs:

a set of representative group colours in Lab space, the corresponding normalised spatial weights, and the area proportion of each group. Together, these outputs transformed the source texture into a structured representation suitable for iterative optimisation and subsequent full-resolution correction. The dominant-colour decomposition process is illustrated in Figure 5.




Figure 5. Dominant-colour decomposition of the source albedo texture. The texture is first clustered in D65 CIE Lab space using *k*-means to identify a small number of dominant colour groups. The resulting hard masks are then smoothed and normalised to form soft spatial masks, from which representative group colour centres are computed as weighted averages of the original texture values.

4.1.2. Group-Wise Colour Compensation

Once the source texture had been decomposed into dominant colour groups, the optimisation was performed at the group level through a measurement-driven iterative process. The objective of this stage was to determine, for each group, a compensated Lab colour value that would reduce the perceptual difference between the projected appearance and the corresponding screen reference.

For each colour group *k*, a uniform colour patch was generated from the current commanded Lab value and displayed through the Unity rendering pipeline. The same patch condition was presented on both the calibrated monitor and the projector. The screen patch represented the reference appearance, while the projected patch represented the current state of that group under projection. Both patches were measured using the colourimeter described in Section 3.2. The measurements were obtained as XYZ tristimulus values, averaged over three repeated readings, and converted to D65 CIE Lab for comparison.

Let C_k^{ref} denote the Lab value measured from the screen reference patch of group *k*, and let $C_k^{proj}(t)$ denote the Lab value measured from the corresponding projected patch at iteration *t*. The colour residual was defined as

$$\Delta C_k(t) = C_k^{ref} - C_k^{proj}(t). \tag{9}$$

This residual indicates the direction in Lab space required to move the measured projected appearance towards the screen reference.

The commanded colour of the patch was then updated according to

$$C_k^{out}(t + 1) = C_k^{out}(t) + \alpha \Delta C_k(t), \tag{10}$$

where $C_k^{out}(t)$ is the Lab colour used to generate the projection-side patch at iteration *t*, and α is a relaxation factor controlling the update magnitude. In the experiments, $\alpha = 0.5$, as reported in Appendix A. This relaxation factor was used to reduce the risk of unstable colour jumps caused by measurement noise or local nonlinearity in the screen-to-projection

response. The updated Lab value was also constrained to the valid colour range before being converted back to RGB for rendering.

After each update, a new projection-side patch was rendered and measured again, producing a new projected Lab value. The perceptual colour difference between the projected patch and the reference patch was evaluated using the CIEDE2000 colour-difference metric:

$$\Delta E_{00,k}(t) = \Delta E_{00}(\mathbf{C}_k^{\text{ref}}, \mathbf{C}_k^{\text{proj}}(t)). \quad (11)$$

The optimisation continued iteratively until

$$\Delta E_{00,k}(t) < \tau_A, \quad (12)$$

where $\tau_A = 2.3$ was used as an acceptability-oriented perceptual tolerance for albedo appearance matching. This value was not tuned from the present experimental results. It was adopted from closely related P-SAR colour reproduction work, where $\Delta E_{00} = 2.3$ has been used to judge whether projected physical colour references are sufficiently close to the intended reference in design-oriented SAR conditions [45]. The same order of magnitude is also consistent with broader CIEDE2000 colour-difference literature, in which acceptability thresholds are commonly reported in the vicinity of this range, although exact values depend on stimulus type, viewing condition and observer population [39,40]. Therefore, $\tau_A = 2.3$ is interpreted here as a practical acceptability-level criterion rather than as a universal just-noticeable-difference threshold.

This iterative procedure was applied independently to each dominant colour group. Unlike a global-offset correction, each group therefore followed its own measured residual and update trajectory. A texture was considered fully converged only when all of its dominant colour groups satisfied $\Delta E_{00} < \tau_A$. This group-wise criterion avoided the situation in which a low texture-level average colour difference could hide a poorly corrected colour group. Because the number of groups was small, the optimisation required only a limited number of measurements while still capturing the main colour structure of the texture.

4.1.3. Reconstruction of the Optimised Albedo Texture

After the group-wise compensation described in Section 4.1.2, each colour group was associated with a final compensated Lab value $\mathbf{C}_k^{\text{out}}$. However, these values corresponded only to the compact group-level representation introduced in Section 4.1.1. To obtain a usable full-resolution texture map for rendering, the group-level corrections had to be propagated back to the original image domain through the soft spatial masks.

To preserve the spatial detail of the source texture, the reconstruction was implemented as a smooth colour-correction field applied to the original albedo map, rather than as a direct synthesis of the texture from the optimised group colours alone. For each group k , the final group-wise Lab correction was first defined as

$$\Delta \mathbf{C}_k^{\text{out}} = \mathbf{C}_k^{\text{out}} - \mathbf{C}_k^{\text{src}}. \quad (13)$$

This correction represents the measured compensation required for the representative colour of group k .

The full-resolution Lab correction field was then obtained by blending the group-wise corrections with the normalised soft masks:

$$\Delta \mathbf{C}_{\text{field}}(x, y) = \sum_{k=1}^K w_k(x, y) \Delta \mathbf{C}_k^{\text{out}}. \quad (14)$$

Because the masks were normalised so that their weights summed to one at every pixel, this formulation produced smooth transitions between neighbouring colour regions and avoided discontinuities at mask boundaries.

The optimised Lab texture was finally reconstructed as

$$\mathbf{C}_{\text{opt}}(x, y) = \mathbf{C}_{\text{src}}(x, y) + \Delta \mathbf{C}_{\text{field}}(x, y). \quad (15)$$

In this way, the original spatial texture detail was preserved, while the colour distribution was adjusted according to the measured screen-to-projection residuals of the dominant colour groups.

The reconstructed Lab texture was converted back to RGB for rendering in the Unity environment. The resulting image constituted the optimised albedo map and replaced the original albedo texture used by the Unity material during projection. Because the reconstruction was derived from measured discrepancies between the screen reference and the projected patches, the optimised texture implicitly compensated for the colour shifts introduced by the projector and projection surface under the tested setup.

4.2. Normal-Map-Based Shading Appearance Optimisation

While the albedo optimisation described above corrects colour-distribution discrepancies in the projected texture, it does not directly address differences in shading behaviour caused by surface normals and lighting interactions. This section therefore introduces a complementary optimisation procedure targeting normal-induced shading appearance. The method first establishes a calibrated uniform-albedo condition to suppress spatial colour variation, and then compares screen-side and projector-side shading contrast using a lightness-domain descriptor. When the projected shading contrast differs from the screen reference, the tangential components of the normal map are updated iteratively so that the projected shading response approaches the reference. An overview of the normal-map-based shading optimisation workflow is shown in Figure 6, and the detailed implementation sequence is summarised in Algorithm A2 in Appendix A.3.

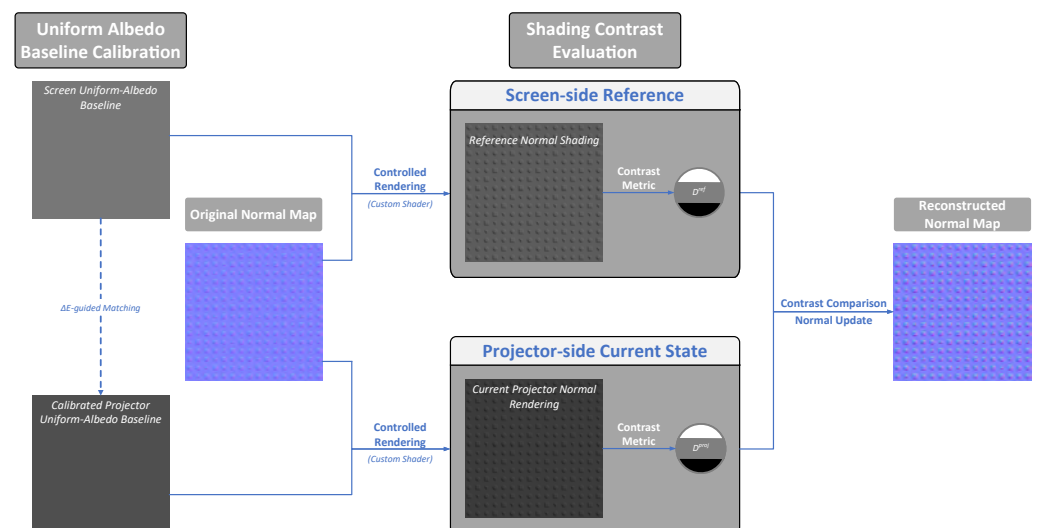


Figure 6. Overview of the normal-map-based shading appearance optimisation workflow. A uniform-albedo condition is first calibrated to reduce the influence of texture-dependent colour variation and provide a stable reference state. Based on this condition, the screen-side and projector-side shading contrasts are compared, and the normal map is iteratively updated until the projected normal-induced shading response approaches the screen reference.

4.2.1. Uniform Albedo Baseline Calibration

Before shading optimisation, a preparatory calibration step was performed under a uniform albedo condition. In this step, the original spatially varying albedo texture was temporarily replaced in Unity by a constant neutral grey input, so that colour variation from the texture itself was removed while the effects of illumination and normal-map modulation could still be observed under the selected shader mode. The purpose of this step was not to optimise the final integrated material appearance directly, but to establish a stable neutral condition for the subsequent shading-oriented procedure. In this subsection, the term “baseline” therefore refers to a calibrated uniform-albedo rendering condition, not to the comparative baseline methods introduced later in Section 4.3.

If shading optimisation were applied directly to a textured material, the measured response would be influenced simultaneously by the albedo distribution, the lighting term, and the normal-induced shading contrast. These factors would be difficult to separate using point-based colourimetric measurements. By replacing the texture with a spatially uniform input, the surface colour contribution became constant across the rendered region, and the remaining variation was dominated by the controlled rendering pipeline and the reproduced shading response. This made the screen–projector difference easier to interpret and provided a consistent starting point for normal-related optimisation.

The neutral grey value of sRGB [119, 119, 119] was selected as a practical mid-level input. Under the Linear colour-space setting used in the Unity project, this value approximately corresponds to an 18% grey reference, providing a balanced neutral condition while reducing the risk of unstable low-lightness measurements or premature projector saturation. The screen-side uniform grey input is denoted as G_{ref} . For the projector branch, the same neutral condition was first calibrated in a screen-referenced manner: the screen-side appearance of the uniform patch served as the reference, the projector-side appearance served as the current state, and the projected result was iteratively adjusted in D65 CIE Lab space until the colour difference fell below the same ΔE_{00} threshold used for albedo convergence. The resulting calibrated projector-side uniform input is denoted as G_{proj} , and was subsequently used as the neutral albedo input for the projector-side shading measurements.

4.2.2. Flat and Mean Patch Construction

To enable shading-oriented optimisation with point-based colour measurement, the workflow did not operate directly on the full rendered images. Instead, it constructed two representative uniform patches for each system branch: a flat patch and a mean patch. The flat patch was derived from a rendering with a uniform albedo and a fully flat normal map, so that no local normal-induced shading variation was present. The mean patch was derived from a rendering with the same uniform albedo but with the material’s normal map enabled. Since this rendering was spatially non-uniform, a single uniform patch was generated from the perceptual mean of the valid rendered region rather than from a specific point location. Under this definition, flat denotes the absence of normal modulation, whereas mean denotes the average appearance of the normal-modulated surface.

Two corresponding patch pairs were prepared: one for the screen branch and one for the projector branch. The screen-side pair served as the reference, while the projector-side pair represented the current projected state to be optimised. Although both pairs were generated through the same controlled rendering procedure, they differed in the uniform albedo input used by Unity. The screen branch used G_{ref} , whereas the projector branch used the calibrated projector-side input G_{proj} obtained from the uniform-albedo calibration step. As a result, the subsequent shading comparison was performed after the neutral-albedo mismatch between screen and projection had already been reduced.

The flat patch was obtained directly from the Unity rendering, since the combination of uniform albedo and a flat normal map produced a spatially uniform result. By contrast, the mean patch required an additional processing step. The rendered image with the normal map enabled was exported and analysed in MATLAB. The valid material region was converted into D65 CIE Lab space, and its perceptual mean Lab value was computed. This mean Lab value was then converted back into a uniform RGB patch, displayed again, and measured with the colourimeter. This procedure was applied to both the screen and projector branches so that both L_{flat}^* and L_{mean}^* were based on physically measured patches rather than on rendered RGB values alone.

The workflow therefore produced four canonical measurements for each material and lighting condition: screen flat, screen mean, projector flat, and projector mean. The screen-side measurements remained fixed for a given material and rendering condition, whereas the projector-side measurements were updated during the iterative normal-map optimisation. Together, these patch pairs provided a compact colourimeter-measurable representation of the baseline and normal-modulated shading responses required for the contrast-guided update described in the next section.

4.2.3. Contrast-Guided Normal Map Update

Using the patch pairs defined in Section 4.2.2, shading optimisation was performed by comparing the lightness contrast between the flat and mean patches on the screen and projector sides. Under the calibrated uniform-albedo condition, the flat patch represents the reference lightness level without normal modulation, whereas the mean patch represents the average appearance produced by the same rendering configuration with the normal map enabled. The difference between these two patches therefore provides a compact descriptor of the shading strength generated by the normal map. The definition and computation of this descriptor are illustrated in Figure 7.

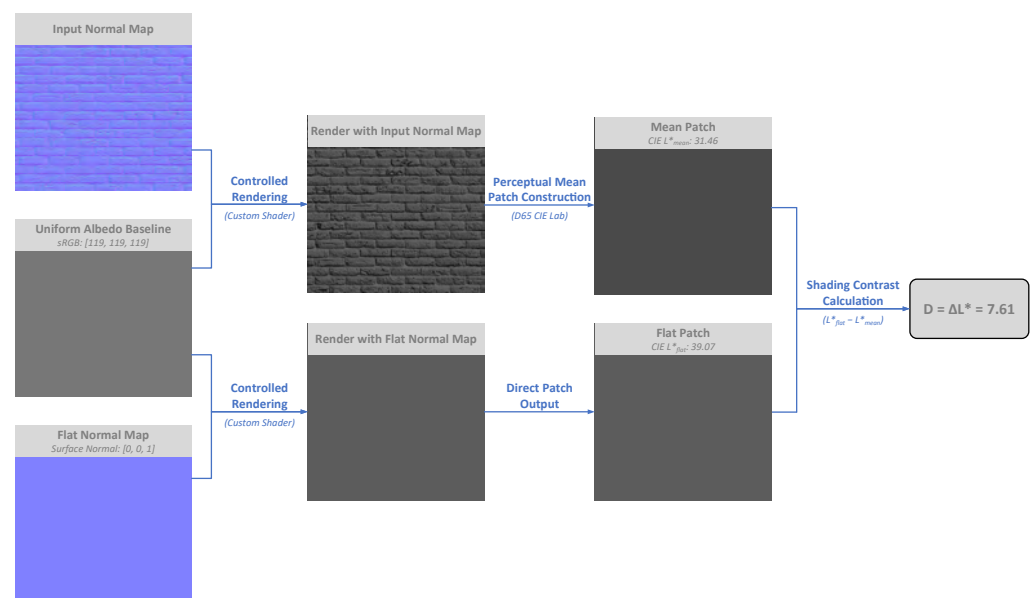


Figure 7. Definition of the lightness-domain shading contrast descriptor used for normal-map-based shading optimisation. Under a uniform-albedo condition, the scene is rendered with either the original normal map or a flat normal map to obtain the mean and flat patches in D65 CIE Lab space. The shading contrast is computed as $D = L_{\text{flat}}^* - L_{\text{mean}}^*$, which characterises the average lightness modulation produced by the normal map.

Let $L_{\text{flat}}^{*\text{ref}}$ and $L_{\text{mean}}^{*\text{ref}}$ denote the measured L^* values of the screen-side flat and mean patches, respectively. Similarly, let $L_{\text{flat}}^{*\text{proj}}(t)$ and $L_{\text{mean}}^{*\text{proj}}(t)$ denote the corresponding

projector-side measurements at iteration t . All measurements were obtained using the colourimeter and converted to D65 CIE Lab as described in Section 3.4. The screen-side reference shading contrast was defined as

$$D_{\text{ref}} = L_{\text{flat}}^{*\text{ref}} - L_{\text{mean}}^{*\text{ref}}, \quad (16)$$

and the projector-side shading contrast at iteration t was defined as

$$D_{\text{proj}}(t) = L_{\text{flat}}^{*\text{proj}}(t) - L_{\text{mean}}^{*\text{proj}}(t). \quad (17)$$

Here, D_{ref} represents the normal-induced shading contrast measured from the screen reference, while $D_{\text{proj}}(t)$ represents the corresponding contrast measured from the projected result at the current iteration.

Instead of directly editing the geometry of the physical surface or applying a post-rendered brightness correction to the final image, the workflow updated the normal map that generated the shading response. Based on the ratio between the reference and projected contrasts, a normal-map update factor was computed as

$$\eta(t) = \left(\frac{D_{\text{ref}}}{\max(D_{\text{proj}}(t), \epsilon)} \right)^{\beta}, \quad (18)$$

where β is a relaxation factor controlling the update magnitude and ϵ is a small lower bound used to avoid numerical instability when $D_{\text{proj}}(t)$ is close to zero. In the experiments, $\beta = 0.5$ and $\epsilon = 10^{-9}$, as reported in Appendix A.

The update factor was then applied to the tangential components of the current normal map. For a unit normal vector

$$\mathbf{n}^{(t)} = \left(n_x^{(t)}, n_y^{(t)}, n_z^{(t)} \right),$$

the updated tangential components were obtained by

$$n_x^{(t+1)} = \eta(t)n_x^{(t)}, \quad n_y^{(t+1)} = \eta(t)n_y^{(t)}. \quad (19)$$

The corresponding z -component was then reconstructed so that the vector remained a valid unit normal:

$$n_z^{(t+1)} = \sqrt{1 - \left(n_x^{(t+1)} \right)^2 - \left(n_y^{(t+1)} \right)^2}. \quad (20)$$

The resulting normal vector was renormalised and encoded back into RGB form to generate the normal map for the next Unity rendering iteration. Under the normal-map convention used in the implementation, RGB values were decoded into normal-vector components before update and re-encoded after update. This ensured that the optimisation acted on the material-level normal representation rather than only on a temporary scalar rendering parameter.

After each update, the Unity scene was rendered again using the updated normal map, and the projector-side flat and mean patches were regenerated and remeasured. The residual between the screen-side and projector-side shading contrast descriptors was defined as

$$r_D(t) = |D_{\text{proj}}(t) - D_{\text{ref}}|. \quad (21)$$

The optimisation process continued until

$$r_D(t) < \tau_D, \quad (22)$$

where τ_D is the stopping tolerance in the CIE L^* domain.

The threshold τ_D was set to $1.0 L^*$. This value should be interpreted differently from the ΔE_{00} threshold used in the albedo branch. It is not claimed here as a universal psychophysical threshold for projected shading perception in all P-SAR conditions. Instead, it is used as a conservative lightness-domain operational tolerance for the specific shading contrast descriptor defined in Equations (16) and (17). The residual $r_D(t)$ does not represent a raw RGB difference, a physical luminance difference, or a pixel-wise colour difference. Rather, it quantifies the remaining mismatch between two CIE L^* -based contrast descriptors measured under a calibrated uniform-albedo condition.

This lightness-domain formulation was adopted because chromatic texture variation is suppressed in the shading branch. Under the uniform-albedo condition, the remaining appearance mismatch is mainly associated with normal-induced lightness modulation. A tolerance expressed in CIE L^* is therefore more appropriate than a tolerance expressed in 8-bit RGB grey levels, since adjacent RGB code values do not correspond to uniform physical luminance or uniform perceptual lightness steps. By contrast, CIE L^* provides a perceptually motivated lightness scale from 0 for black to 100 for the reference white. This use of a perceptual lightness-domain descriptor is also consistent with grey-scale evaluation literature, where neutral grey-scale appearance is treated as a perceptual problem rather than as a direct function of device code values [46,47].

The strictness of $1.0 L^*$ can also be related to CIEDE2000 under neutral-grey conditions. When two colours are neutral greys, namely $a^* = b^* = 0$, the chroma and hue terms in CIEDE2000 vanish, and the colour difference is governed by the lightness-weighted term:

$$\Delta E_{00,\text{grey}} = \frac{|\Delta L^*|}{S_L}, \quad (23)$$

where S_L is the CIEDE2000 lightness weighting term [39,40]. Since $S_L \geq 1$ over the valid L^* range, a residual difference of $1.0 L^*$ corresponds to a CIEDE2000 value not exceeding approximately 1.0 under neutral-grey conditions. Therefore, although $\tau_D = 1.0 L^*$ is not treated as a universal perceptual boundary for projected shading, it is a deliberately conservative tolerance when interpreted in the neutral-grey lightness domain used by this branch.

4.3. Comparative Baseline Methods and Scope of Comparison

To provide objective comparative evidence, the proposed optimisation procedures were evaluated against simpler compensation baselines implemented under the same screen-referenced, colourimeter-based measurement protocol. These baselines were not intended to represent the full state of the art in projector compensation. Instead, they were included to test whether the proposed material-component-aware updates provide advantages over global correction strategies that do not separately model dominant albedo colour groups or normal-induced shading response.

Full projector compensation methods such as CompenNet and CompenNet++ are highly relevant to the broader projector compensation literature, but they were not used as direct baselines in this study. These methods are generally formulated as camera-based image-level or full projector compensation frameworks, where the projector input image is optimised so that a camera-observed projection matches a desired target image. A fair implementation would require a different experimental pipeline, including dense projector-camera acquisition, spatial registration, camera-based training data generation, network implementation or retraining, and image-domain evaluation under the same physical setup. The present study instead uses point-based colourimeter measurements and material-component-related optimisation variables. Therefore, the comparison was

limited to baselines that could be implemented fairly within the same measurement framework. The aim was not to claim superiority over full projector compensation networks, but to compare the proposed method with simpler global strategies under equivalent experimental conditions.

4.3.1. Albedo Compensation Reference Condition and Baseline Methods

For the albedo branch, the proposed group-wise optimisation was compared with two global baseline methods: global neutral-grey correction and global Lab-offset correction. In addition, the initial no-compensation condition was retained as the uncorrected reference condition for calculating the relative improvement. The initial condition, the two baseline methods, and the proposed method all used the same source textures, dominant-colour grouping, screen-side reference measurements, and group-wise ΔE_{00} evaluation criterion.

In the initial no-compensation condition, the source group colours were used directly:

$$\mathbf{C}_k^{\text{out}} = \mathbf{C}_k^{\text{src}}. \quad (24)$$

This condition represents the direct screen–projection colour mismatch before any albedo compensation is applied.

The first baseline method was a global neutral-grey correction. A single neutral grey patch was measured on the screen and under projection, and the resulting Lab residual was applied uniformly to all colour groups:

$$\Delta \mathbf{C}_{\text{grey}} = \mathbf{C}_{\text{grey}}^{\text{ref}} - \mathbf{C}_{\text{grey}}^{\text{proj}}, \quad \mathbf{C}_k^{\text{out}} = \mathbf{C}_k^{\text{src}} + \Delta \mathbf{C}_{\text{grey}}. \quad (25)$$

This represents a single-shot display-level grey-balance correction.

The second baseline method was a global Lab-offset correction. For each iteration, the measured residuals of all dominant colour groups were combined into one area-weighted Lab offset:

$$\Delta \mathbf{C}_{\text{global}}(t) = \sum_{k=1}^K a_k (\mathbf{C}_k^{\text{ref}} - \mathbf{C}_k^{\text{proj}}(t)), \quad \mathbf{C}_k^{\text{out}}(t+1) = \mathbf{C}_k^{\text{out}}(t) + \alpha \Delta \mathbf{C}_{\text{global}}(t). \quad (26)$$

Unlike the proposed group-wise method, this baseline method applies the same correction to all colour groups and therefore cannot independently compensate different chromatic regions.

4.3.2. Normal-Map-Based Shading Compensation Reference Condition and Baseline Method

For the normal-map-based shading branch, the proposed contrast-guided normal-map update was compared with one image-domain baseline method: global luminance-contrast compensation. In addition, the initial no shading-specific compensation condition was retained as the uncorrected reference condition after uniform-albedo baseline calibration. The initial condition, the baseline method, and the proposed method all used the same uniform-albedo calibration, flat/mean patch construction, and residual metric $r_D = |D_{\text{proj}} - D_{\text{ref}}|$. The purpose of this comparison was to distinguish the proposed material-level normal-map update from an output-level image-domain correction.

In the initial no shading-specific compensation condition, the original normal map was used directly after the uniform-albedo calibration:

$$N_{\text{out}} = N_{\text{src}}. \quad (27)$$

This condition represents the residual shading mismatch that remains after baseline luminance alignment but before any normal-map correction is applied.

The baseline method was a global luminance-contrast compensation. In this case, the normal map was not modified. Instead, the rendered shading image was adjusted in the image domain by globally scaling the lightness deviation around the projector-side flat-patch lightness:

$$L_{\text{img,proj}}^{*'}(x, y) = L_{\text{flat,proj}}^* + \gamma \left[L_{\text{img,proj}}^*(x, y) - L_{\text{flat,proj}}^* \right], \quad \gamma = \frac{D_{\text{ref}}}{\max(D_{\text{proj}}^0, \epsilon)}. \quad (28)$$

Here, D_{proj}^0 denotes the projector-side shading descriptor before image-domain compensation, and ϵ is a small positive constant used to avoid division by zero. This baseline method represents a simple output-level contrast correction. It modifies the final rendered lightness distribution, whereas the proposed method modifies the normal map that generates the shading response. Therefore, the comparison separates the initial no-update condition, the image-domain luminance-contrast baseline method, and the proposed material-level normal-map update.

5. Experimental Results and Practical Evaluation

This section presents the experimental evaluation of the proposed screen-referenced appearance optimisation workflow. The results are reported according to the two material appearance components addressed by the method. Section 5.1 evaluates the albedo appearance optimisation in terms of group-wise convergence, comparison with global albedo compensation baselines, and sensitivity to the number of dominant colour groups. Section 5.2 evaluates the normal-map-based shading appearance optimisation in terms of convergence behaviour and comparison with a global luminance-contrast baseline. Section 5.3 presents qualitative integrated appearance illustrations using photographic examples of representative material cases. Section 5.4 discusses practical feasibility in terms of execution time, operational requirements, and potential automation. Together, these analyses examine the objective effectiveness, comparative value, practical feasibility, and evaluation scope of the proposed workflow under controlled P-SAR conditions.

5.1. Albedo Appearance Optimisation Results

The albedo appearance optimisation results are first presented for the proposed group-wise method. The aim of this subsection is to evaluate whether the proposed method can reduce the screen-to-projection colour mismatch of the dominant albedo colour groups below the adopted CIEDE2000 threshold. Comparative baseline results and the sensitivity of the dominant colour-group number K are reported separately in Sections 5.1.2 and 5.1.3.

5.1.1. Group-Wise Convergence of the Proposed Method

The ten tested textures comprised two source categories: Textures 01–05 were reconstructed from scans of real physical materials, whereas Textures 06–10 were digitally authored textures obtained from online material libraries. The dataset covered a range of surface types, including woven fabrics, leather and velvet surfaces, a soil-and-moss ground texture, wood, a painted wall, rusted metal, and brick. Across all textures, a total of 23 dominant colour groups were identified and optimised, with each texture containing between one and three groups.

As shown in Figure 8, all group trajectories exhibited a clear decreasing trend in ΔE_{00} , demonstrating that the proposed group-wise optimisation consistently reduced the colour mismatch between the projected appearance and the screen-side reference. Importantly, convergence was evaluated at the group level rather than only by using a texture-level mean value. All 23 dominant colour groups eventually satisfied the adopted threshold of

$\Delta E_{00} < 2.3$, indicating that the optimisation did not rely on averaging across colour regions to mask poorly corrected groups.

A common pattern visible in Figure 8 is that the largest decrease usually occurred between Iteration 0 and Iteration 1, while subsequent updates mainly refined the remaining residual error. This behaviour suggests that the dominant colour bias was corrected at an early stage, whereas later iterations were primarily used for fine adjustment. In addition, the trajectories remained monotonic or near-monotonic in all tested cases, and no strong oscillatory behaviour was observed. This indicates that the Lab-space update rule was numerically stable under the tested screen–projection conditions.

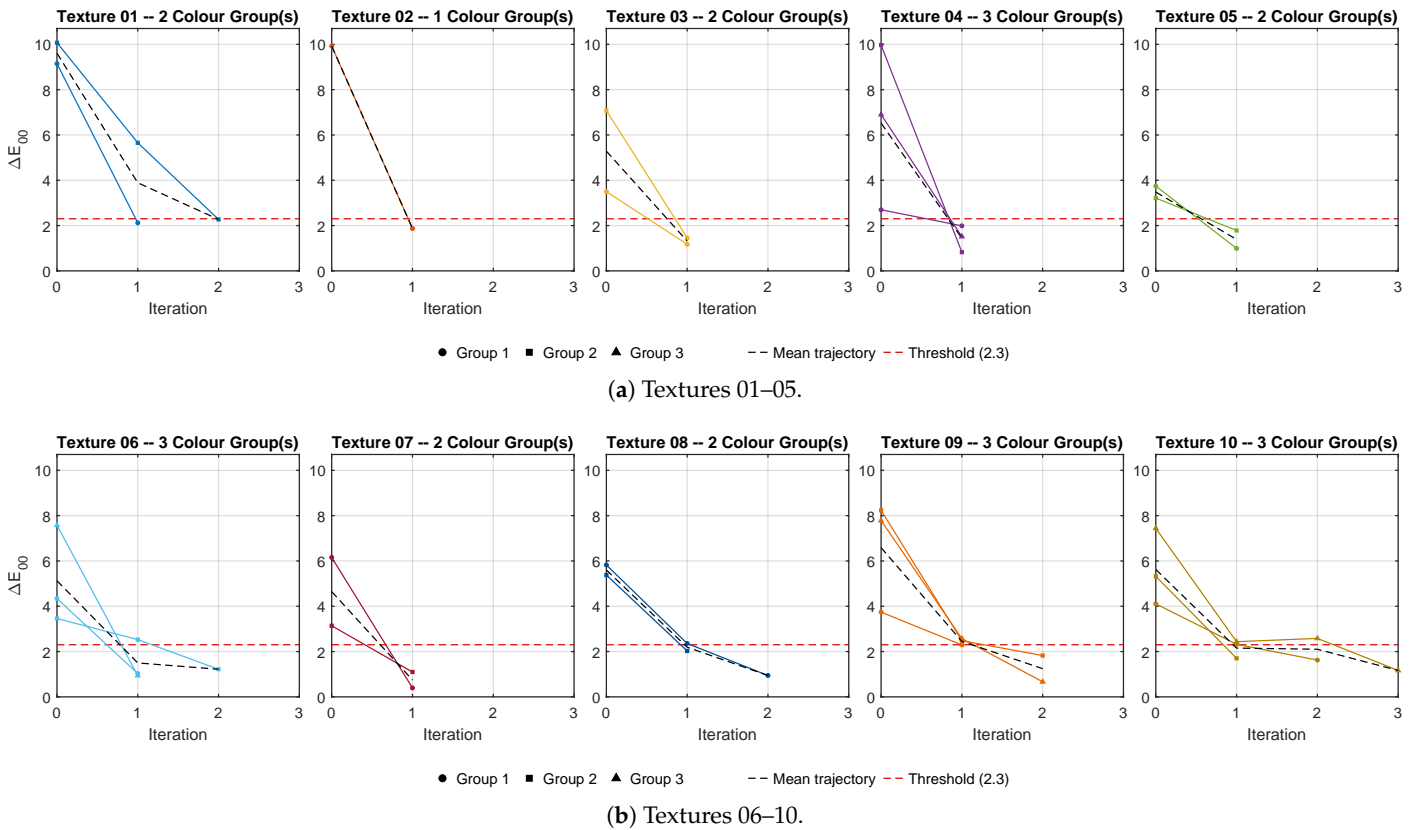


Figure 8. Group-wise ΔE_{00} convergence during the proposed albedo optimisation. Iteration 0 denotes the initial projection state. Coloured lines show individual colour groups, the black dashed line indicates the mean trajectory, and the red dashed line marks the adopted threshold ($\Delta E_{00} = 2.3$). (a) Textures ID 01–05. (b) Textures ID 06–10.

The quantitative summary in Table 1 confirms the same tendency at the texture level. Before optimisation, the initial mean ΔE_{00} values ranged from 3.48 for Texture 05 to 9.93 for Texture 02, with an overall mean of 6.24. After optimisation, the final mean ΔE_{00} values were reduced to a narrower range of 0.75–2.27, with an overall mean of 1.36. This corresponds to an average reduction of 77.43% across the ten textures. Although the table reports texture-level mean values for compactness, the convergence requirement was applied at the group level: a texture was considered fully converged only when all of its dominant colour groups satisfied $\Delta E_{00} < 2.3$.

Table 1. Summary of albedo appearance optimisation results for all tested textures using the proposed group-wise method.

Albedo Map ID	No. of Colour Groups	No. of Updates	Mean Initial ΔE_{00}	Mean Final ΔE_{00}	Mean Reduction (%)
01	2	2	9.61	2.27	76.36
02	1	1	9.93	1.87	81.16
03	2	1	5.29	1.32	75.11
04	3	1	6.52	1.44	77.89
05	2	1	3.48	1.39	60.09
06	3	2	5.12	1.22	76.21
07	2	1	4.65	0.75	83.94
08	2	2	5.60	0.94	83.15
09	3	2	6.59	1.25	81.10
10	3	3	5.62	1.17	79.27
Mean	2.3	1.6	6.24	1.36	77.43

The required number of updates was limited. Five textures, namely, Textures 02, 03, 04, 05, and 07, converged after a single optimisation update from the initial condition. Four textures, namely Textures 01, 06, 08, and 09, required two updates, while only Texture 10 required three updates. The mean number of optimisation updates was therefore 1.6 per texture. From a practical perspective, this indicates that the proposed albedo branch achieved substantial colour improvement without requiring a prolonged iterative process.

Although the overall convergence trend was consistent, some differences among textures can still be observed. Texture 02 presented the largest initial mean error, with an initial mean ΔE_{00} of 9.93, but was reduced to 1.87 after a single update, corresponding to an 81.16% reduction. This indicates that a large initial mismatch did not necessarily imply a difficult optimisation process. In contrast, Texture 10 required the largest number of updates, even though its initial mean ΔE_{00} of 5.62 was only moderate relative to the other samples. Its final mean ΔE_{00} was nevertheless reduced to 1.17, yielding a total reduction of 79.27%. This case suggests that textures with more colour groups or more uneven group-wise convergence may require additional refinement even when the initial average error is not the highest.

The best final result was obtained for Texture 07, whose mean ΔE_{00} decreased from 4.65 to 0.75, corresponding to the largest reduction ratio of 83.94%. Texture 08 also showed a similarly strong improvement, decreasing from 5.60 to 0.94 with a reduction of 83.15%. By contrast, Texture 05 showed the smallest relative reduction, from 3.48 to 1.39, corresponding to 60.09%. However, this case also had the lowest initial mean error among all textures, which means that the optimisation had less room for numerical improvement from the outset. A smaller reduction ratio in this context therefore does not indicate failure, but reflects a milder initial mismatch.

At the group level, all 23 dominant colour groups were eventually reduced below the threshold of $\Delta E_{00} = 2.3$. The largest initial group error was observed in Texture 01, Group 2, with an initial ΔE_{00} of 10.07, which was reduced to 2.27 at the final state. Although this remained the highest final residual among all groups, it still fell slightly below the adopted threshold. Several groups achieved much lower final residuals. For example, Texture 07, Group 1 decreased from 6.15 to 0.39, Texture 04, Group 2 decreased from 9.97 to 0.82, and Texture 09, Group 3 decreased from 7.78 to 0.67. These results indicate that the proposed optimisation was not only effective on average, but was also capable of substantially correcting the most difficult local colour regions.

Representative examples of the resulting albedo texture correction are shown in Figure 9. Texture 02 was selected as the case with the largest initial texture-level mismatch, Texture 07 as the case with the lowest final residual, and Texture 10 as the case requiring the largest number of optimisation updates. The examples show that the optimisation modified the overall colour tendency of the albedo maps while preserving the spatial structure of the source textures. This behaviour is consistent with the correction-field reconstruction described in Section 4.1.3, where group-wise Lab corrections are propagated through soft spatial masks rather than replacing the texture with discrete colour regions.



Figure 9. Representative examples of digital albedo textures before and after optimisation. Texture 02 represents the case with the largest initial mismatch, Texture 07 the case with the lowest final residual, and Texture 10 a representative case requiring the largest number of optimisation updates.

Taken together, these results confirm that the proposed group-wise albedo optimisation effectively improved projected base-colour matching across all tested textures. All dominant colour groups satisfied the adopted $\Delta E_{00} < 2.3$ criterion, and the texture-level mean colour difference was reduced substantially. The next subsection compares this proposed group-wise method with simpler global albedo compensation baselines.

5.1.2. Comparison with Global Albedo Compensation Baselines

To examine whether the improvement was specifically associated with the proposed group-wise formulation, the albedo optimisation was compared with two simpler global baseline methods. The initial no-compensation condition was also retained as the uncorrected reference condition for calculating the relative reductions, but it was not considered a compensation method.

The first baseline method was a global neutral-grey correction, in which a single neutral-grey patch was used to estimate a screen–projector colour offset and the same offset was applied once to all colour groups. This represents a simple display-level grey-balance correction. The second baseline method was a global Lab-offset correction, in which the residuals of all dominant colour groups in one texture were combined into a single texture-level Lab residual and iteratively applied to all groups. This represents a stronger texture-level feedback correction, but it does not allow different colour groups to be updated independently.

Figure 10 compares the final mean ΔE_{00} values obtained by the two global baseline methods and the proposed group-wise optimisation for the ten tested textures. The global neutral-grey correction reduced the initial colour mismatch but remained above the threshold of $\Delta E_{00} = 2.3$ for all textures, indicating that a single neutral-grey correction can only compensate for general projection colour bias. The global Lab-offset correction produced a larger improvement, but several textures still retained residual errors above the threshold. In contrast, the proposed group-wise optimisation consistently achieved lower final colour differences across the tested textures.

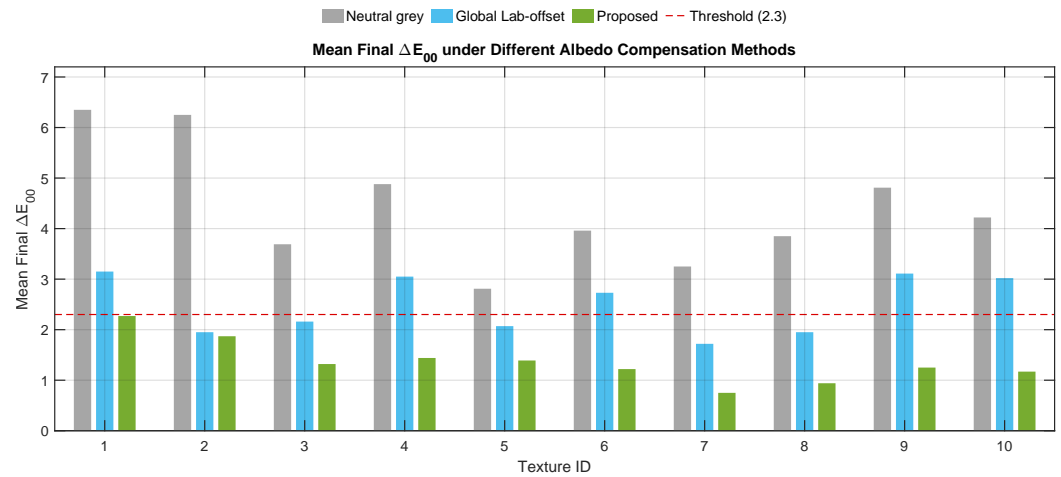


Figure 10. Comparison of final mean ΔE_{00} values under different albedo compensation strategies. Each grouped bar set corresponds to one texture, and the dashed horizontal line indicates the threshold of $\Delta E_{00} = 2.3$.

The method-level summary is reported in Table 2. The uncorrected initial condition had a mean ΔE_{00} of 6.24. The global neutral-grey correction reduced this value to 4.41, corresponding to an average texture-wise reduction of 28.16%, but no texture reached full convergence. The global Lab-offset correction further reduced the mean final error to 2.49, with an average reduction of 57.44%, but only five of the ten textures fully converged. By contrast, the proposed method achieved the lowest mean final error of 1.36, the highest reduction of 77.43%, and full convergence for all ten textures.

Table 2. Comparison of the uncorrected initial condition, global albedo compensation baselines, and the proposed group-wise correction. Mean reduction was calculated as the average of texture-wise percentage reductions. N/A: Not Applicable.

Condition/Strategy	Mean Final ΔE_{00}	Mean Reduction (%)	Mean No. of Correction Steps	Fully Converged Textures
Initial no-compensation condition	6.24	N/A	N/A	0/10
Global neutral-grey correction	4.41	28.16	1.0	0/10
Global Lab-offset correction	2.49	57.44	6.5	5/10
Proposed group-wise correction	1.36	77.43	1.6	10/10

The comparison also explains why the group-wise formulation is necessary. When a texture is represented by only one dominant colour group, the global Lab-offset baseline is almost equivalent to the proposed method. However, as the number of dominant groups increases, a single global offset becomes less reliable because different colour regions may require different correction directions in Lab space. This was reflected in the four textures represented by three colour groups, none of which fully converged under the global Lab-offset baseline, whereas all converged under the proposed group-wise method. The benefit of the proposed method is therefore not only a lower average colour difference, but also more reliable convergence across textures with different chromatic complexity.

5.1.3. Sensitivity to the Number of Dominant Colour Groups

Since the proposed albedo optimisation relies on dominant-colour decomposition, the number of colour groups K determines how the chromatic structure of a texture is represented. In the main experiment, K was manually selected between one and three according to the apparent colour complexity of each texture. To examine the influence of this parameter, a sensitivity analysis was conducted using three representative textures: Texture

02, a near-uniform brown leather texture; Texture 01, a striped woven fabric with two visually dominant colour regions; and Texture 09, a rusted metal texture with more complex chromatic variation. For each texture, the optimisation was repeated using $K = 1$, $K = 2$, and $K = 3$, while keeping the remaining measurement and update settings unchanged.

Table 3 shows that the influence of K depends on texture complexity. For the near-uniform leather texture, all three evaluated K values converged after one update and produced very similar final errors, indicating that $K = 1$ was already sufficient. Increasing K therefore mainly introduced redundant colour groups rather than a meaningful improvement in correction accuracy.

Table 3. Sensitivity of albedo optimisation results to the number of dominant colour groups K .

Texture ID	Texture Type	Visually Inspected K	Evaluated K	No. of Updates	Converged	Mean Final ΔE_{00}	Reduction (%)
02	Near-uniform leather	1	1	1	Yes	1.87	81.16
02	Near-uniform leather	1	2	1	Yes	1.81	81.77
02	Near-uniform leather	1	3	1	Yes	1.82	81.67
01	Striped woven fabric	2	1	10	No	5.29	44.95
01	Striped woven fabric	2	2	2	Yes	2.27	76.36
01	Striped woven fabric	2	3	2	Yes	2.19	77.21
09	Rusted metal	3	1	10	No	5.12	22.31
09	Rusted metal	3	2	10	No	2.66	59.64
09	Rusted metal	3	3	2	Yes	1.25	81.10

For the striped woven fabric, $K = 1$ failed to converge even after ten updates because the texture was reduced to a single averaged colour representation. Using $K = 2$ separated the two main chromatic regions and achieved convergence, while increasing K to 3 produced only a small additional reduction in the final mean ΔE_{00} . This suggests that $K = 2$ provided an appropriate balance between representational accuracy and measurement effort for this texture.

For the rusted metal texture, the sensitivity to K was more pronounced. Both $K = 1$ and $K = 2$ failed to reach convergence, although $K = 2$ reduced the final error substantially compared with $K = 1$. Full convergence was achieved only with $K = 3$, indicating that the main chromatic structure of this complex texture required three dominant groups to be represented adequately.

These results indicate that K should be understood as a representational parameter rather than a purely numerical clustering parameter. For simple or near-uniform textures, a small K is sufficient and avoids unnecessary measurements. For textures with multiple visually distinct colour regions, however, too small a K can under-represent the chromatic structure and lead to non-convergence or an averaged correction target. In the current implementation, K was selected through clustering results and manual inspection; automatic or semi-automatic selection of K remains an important direction for future work.

5.2. Normal-Map-Based Shading Appearance Optimisation Results

The normal-map-based shading appearance optimisation results are first presented for the proposed contrast-guided normal-map update. The aim of this subsection is to evaluate whether the proposed method can reduce the screen-to-projection mismatch in the lightness-domain shading contrast descriptor. Comparative results against a shading compensation baseline are reported separately in Section 5.2.2.

5.2.1. Convergence of the Proposed Normal-Map Update

The shading appearance mismatch was evaluated using the residual

$$r_D = |D_{\text{proj}} - D_{\text{ref}}|, \quad (29)$$

where $D = L_{\text{flat}}^* - L_{\text{mean}}^*$. In this definition, D_{ref} denotes the shading contrast descriptor measured from the screen-side reference, while D_{proj} denotes the corresponding descriptor measured from the projected result. Since the descriptor was computed under a uniform-albedo condition, the residual mainly reflects the mismatch in normal-induced lightness contrast rather than spatial albedo colour variation. The adopted operational tolerance was $r_D < 1.0 L^*$. The ten tested normal maps corresponded directly to the ten albedo textures analysed in Section 5.1, so that each material was evaluated as a paired albedo–normal case.

As shown in Figure 11, six of the ten normal maps, namely n01, n03, n04, n05, n07 and n08, already satisfied the tolerance at the initial evaluation. The remaining four cases, namely n02, n06, n09 and n10, exceeded the tolerance initially but converged after one normal-map update. All tested cases therefore reached the target criterion with either zero or one update. This indicates that the proposed shading optimisation functioned mainly as a targeted refinement of normal-induced shading strength rather than as a prolonged iterative correction process.




Figure 11. Initial-to-final shading residuals for all tested normal maps. Each pair of points represents the initial residual and the final accepted residual for one normal map, and the red dashed line marks the operational tolerance of $r_D = |D_{\text{proj}} - D_{\text{ref}}| = 1.0 L^*$.

The quantitative summary in Table 4 confirms the same tendency. Across all ten normal maps, the mean residual decreased from 1.164 to 0.264 L^* , corresponding to an average reduction of 77.32%. The final residuals were confined to a narrow range of 0.039–0.502 L^* , and all cases remained below the adopted tolerance. The mean number of updates was only 0.4 per normal map. From a practical perspective, this indicates that, after the uniform-albedo baseline had been established, the proposed shading branch required only limited additional measurement and re-rendering.

Table 4. Summary of shading appearance optimisation results for all tested normal maps using the proposed normal-map update. The reduction in the mean row was calculated from the mean initial and mean final residuals, rather than as the arithmetic mean of the per-case reductions.

Normal Map ID	No. of Updates	Mean Initial r_D (L^*)	Mean Final r_D (L^*)	Reduction (%)	Final Scale Factor
01	0	0.407	0.407	0.00	1.000
02	1	1.857	0.081	95.65	0.840
03	0	0.407	0.407	0.00	1.000
04	0	0.175	0.175	0.00	1.000
05	0	0.343	0.343	0.00	1.000
06	1	3.462	0.354	89.77	0.765
07	0	0.502	0.502	0.00	1.000
08	0	0.226	0.226	0.00	1.000
09	1	2.219	0.039	98.22	0.765
10	1	2.040	0.105	94.84	0.783
Mean	0.4	1.164	0.264	77.32	–

A clear distinction can be observed between the direct-pass cases and the update-required cases. For the six normal maps that already satisfied the tolerance at the initial evaluation, the residuals ranged from 0.175 to 0.502 L^* , with a mean value of approximately 0.343 L^* . This suggests that, once the uniform-albedo baseline calibration had reduced broader luminance bias, the projected shading contrast of these cases was already close to the screen-side reference. Therefore, no further normal-map modification was applied to these samples.

For the four update-required cases, the initial residuals were considerably larger, ranging from 1.857 to 3.462 L^* , with a mean value of approximately 2.395 L^* . After one update, these residuals were reduced to 0.039–0.354 L^* , with a mean value of approximately 0.145 L^* . The largest initial mismatch was observed for n06, whose residual decreased from 3.462 to 0.354 L^* . The lowest final residual was obtained for n09, which decreased from 2.219 to 0.039 L^* . These results show that, when a clear shading mismatch was present, the contrast-guided update rule was able to remove most of the residual error within a single correction step.

The final scale factors further indicate the direction of the required correction. All four update-required cases converged with scale factors smaller than 1.0, ranging from 0.765 to 0.840. This indicates that, under the tested oblique lighting configuration, the projected rendering tended to produce stronger-than-reference shading contrast in these cases. The optimisation therefore primarily attenuated the effective normal strength by scaling the tangential normal components, rather than increasing the shading amplitude. This behaviour is consistent with the intended role of the shading branch: after the uniform-albedo baseline had been aligned, the remaining mismatch was mainly associated with the amplitude of normal-induced shading.

A representative example is shown in Figure 12 using Texture/Normal Map 06, corresponding to the soil-and-moss ground texture. This case was selected because it presented the largest initial shading residual and required one update. As shown in Figure 12a, the normal-map modification was visually subtle and preserved the overall normal pattern. The main effect is instead observed in the rendered shading: the optimised projection-side shading in Figure 12c shows reduced excessive contrast and becomes closer to the corresponding screen-side reference in Figure 12b.




Figure 12. Representative example of shading appearance optimisation for Texture/Normal Map 06 (soil-and-moss ground texture). (a) Comparison of the original and optimised normal maps. (b) The corresponding screen-side reference shading renderings. (c) The corresponding projection-side shading renderings.

Taken together, these results show that the proposed normal-map update was selective, efficient and stable under the tested conditions. It did not unnecessarily modify cases that already satisfied the operational tolerance, while the cases with larger initial residuals converged after only one update. The shading branch therefore served as a targeted material-level refinement of normal-induced shading strength, complementing the preceding albedo appearance optimisation. The next subsection compares this proposed normal-map update with a simpler image-domain shading compensation baseline.

5.2.2. Comparison with an Image-Domain Shading Compensation Baseline

To examine whether the improvement was specifically associated with the proposed material-level normal-map update, the shading optimisation was compared with a simpler image-domain baseline method. The initial no shading-specific compensation condition was also retained as the uncorrected reference condition after the uniform-albedo baseline calibration, but it was not considered a shading compensation method. In this condition, the original normal map was used directly, and the remaining residual represented the shading mismatch before any normal-map update.

The baseline method was a global luminance-contrast compensation, in which the normal map was not modified. Instead, the final rendered lightness distribution was adjusted in the image domain by globally scaling the lightness deviation around the flat-patch baseline. This represents a simple output-level contrast correction, but it does not modify the material normal map that generates the shading response.

All conditions were evaluated using the same residual $r_D = |D_{proj} - D_{ref}|$ and the same operational tolerance of $1.0 L^*$. The method-level summary is reported in Table 5. Under the initial no shading-specific compensation condition, the mean residual remained $1.164 L^*$, with a maximum residual of $3.462 L^*$. Six of the ten normal maps already satisfied the tolerance after the uniform-albedo calibration, but four cases still required additional shading correction.

Table 5. Comparison of the initial no shading-specific compensation condition, the image-domain shading compensation baseline, and the proposed normal-map update. N/A: Not Applicable.

Condition/Strategy	Mean Final r_D (L^*)	Max Final r_D (L^*)	Cases Below $1.0 L^*$	Correction Level
Initial no shading compensation condition	1.164	3.462	6/10	N/A
Global luminance-contrast compensation	0.521	1.108	9/10	Image-domain lightness correction
Proposed normal-map update	0.264	0.502	10/10	Material-level normal-map update

The global luminance-contrast compensation reduced the mean residual from 1.164 to 0.521 L^* , and nine of the ten cases satisfied the adopted tolerance. This confirms that a simple image-domain contrast correction can partially compensate for the mismatch in shading strength. However, one case still remained above the tolerance, with a maximum residual of 1.108 L^* , indicating that a global output-level adjustment does not guarantee convergence for all normal maps.

By contrast, the proposed normal-map update achieved the lowest mean residual of 0.264 L^* , the lowest maximum residual of 0.502 L^* , and convergence in all ten cases. More importantly, the proposed method operates at a different correction level. The baseline modifies the final rendered lightness distribution as an image-domain post-processing operation, whereas the proposed method modifies the tangential components of the normal map itself. The corrected shading response is therefore embedded in the material representation rather than only applied to the final output image.

This comparison supports the use of a material-level normal-map update for the shading branch. Although the image-domain baseline reduced the residual in most cases, the proposed method provided more reliable convergence while remaining consistent with the material-component-aware structure of the workflow.

5.3. Integrated Appearance Illustration

After separately evaluating albedo colour matching and normal-map-based shading consistency, the complete workflow was further examined at the integrated appearance level. In practical P-SAR viewing, base colour and normal-induced shading appear together rather than as isolated components. Therefore, representative material samples were photographed to illustrate the overall appearance tendency after applying the complete two-stage optimisation workflow.

The examples in Figure 13 are intended as qualitative visual illustrations rather than quantitative colour evidence. This limitation arises because the screen reference is self-emissive, whereas the projected sample is observed as reflected light from the projection surface. In addition, the RGB camera used for documentation has a device-dependent spectral response and cannot fully reproduce cross-device colour appearance. For this reason, the quantitative evidence for colour and shading improvement is provided by the colourimetric and lightness-domain results in Sections 5.1 and 5.2, while the photographs are used only to support visual interpretation of the integrated outcome.

Within this limitation, the optimised projections show a closer qualitative visual tendency to the screen-side references than the original projections. In the striped woven fabric, the brightness balance becomes more consistent with the reference. In the leopard print leather, the excessive warm bias is reduced. In the brick wall example, the balance between base tone and relief-related shading becomes closer to the reference. These examples provide an integrated visual summary of the optimisation effect under the tested conditions. However, they should be interpreted only as qualitative illustrations; the main evidence for the proposed workflow remains the objective colourimetric and lightness-domain evaluation reported in the preceding sections.

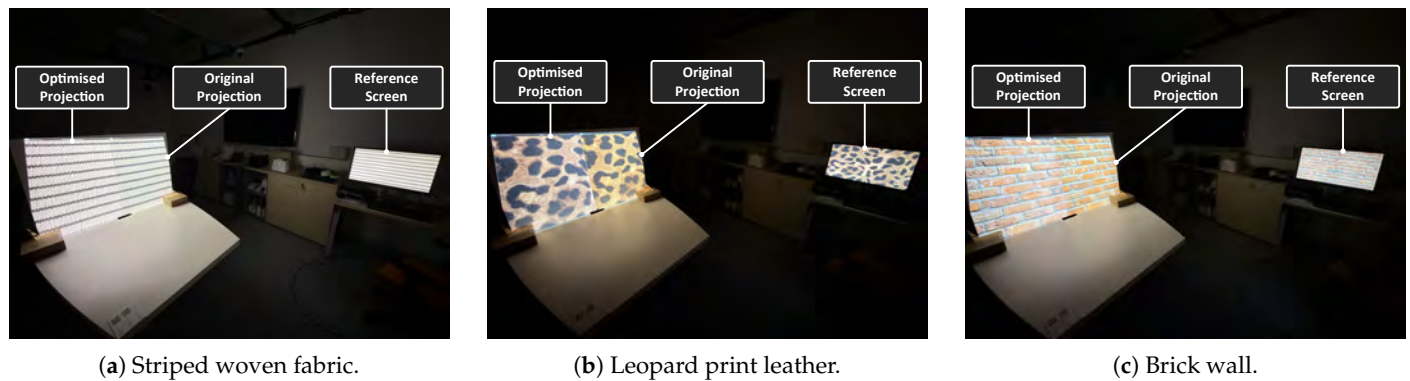


Figure 13. Qualitative photographic illustration of final integrated appearance for three representative materials. Each case compares the screen-side reference, the original projection, and the optimised projection.

5.4. Practicality and Operational Feasibility Evaluation

In addition to the objective appearance results, the practical burden of the proposed workflow was evaluated. Since the method combines controlled rendering, colourimetric measurement, MATLAB-based optimisation and manual result verification, its applicability depends not only on the final residual errors, but also on the time and operational effort required to complete one material case. Table 6 summarises the approximate time consumption of the current implementation. One-off procedures, such as device warm-up, initial monitor calibration and projector setup, are not included.

Table 6. Approximate time breakdown of the proposed workflow per material case. The values correspond to the current manual implementation using one operator and one colourimeter. One-off preparation procedures, such as device warm-up and initial system calibration, are not included.

Operation Type	Main Operation	Time (min)
Albedo rendering/preparation	Dominant-colour grouping and group-patch generation	~3
Albedo colourimetric measurement	Screen/projector measurements of dominant colour-group patches	~7
Albedo computation/inspection	MATLAB update, albedo reconstruction, convergence check and result inspection	~2
Shading baseline preparation	Uniform-albedo baseline calibration	~4
Shading rendering/preparation	Flat/mean patch generation for screen and projection branches	~2
Shading measurement/update	Flat/mean patch measurement, contrast-residual calculation, normal-map update and inspection	~8
Total	Complete albedo and shading appearance optimisation workflow	~26

The breakdown indicates that the dominant time cost comes from physical colourimetric measurement and operator interaction, particularly the repeated screen/projector measurements of colour-group patches and flat/mean shading patches. By contrast, Unity rendering, patch generation, MATLAB update calculation, texture reconstruction and convergence checking are comparatively lightweight. Therefore, the reported duration mainly reflects the manually operated measurement workflow rather than the computational cost of the optimisation algorithms.

The reported time should also be interpreted as a conservative estimate. In the present experiment, one operator and one colourimeter were used to complete all screen and projection measurements sequentially. With multiple operators, or with two calibrated

measurement devices assigned separately to the screen and projection branches, part of the measurement procedure could be performed in parallel. Thus, the elapsed time could be reduced even without introducing a fully automated acquisition system.

The current implementation is not intended as a real-time compensation method to be repeated from scratch after every minor design modification. It is better understood as a material-specific or setup-specific appearance preparation procedure for controlled P-SAR design review. Once a material has been optimised under a given projector, surface, lighting condition and rendering configuration, the compensated albedo map and normal-related parameters can be reused under the same conditions.

Further reduction of the operational burden is possible through automation. Patch generation, Unity rendering/export, MATLAB updates, convergence checking and result logging can be scripted directly. The main remaining bottleneck is manual colourimeter operation, which could be reduced using scripted ArgyllCMS acquisition, a fixed measurement rig, motorised positioning, or integrated projector–camera feedback. Automatic or semi-automatic selection of the dominant colour-group number K would also reduce manual inspection in the albedo branch. These improvements would make the workflow more practical for repeated material preparation and small-batch P-SAR design-review applications.

6. Conclusions

This paper presented a screen-referenced, measurement-driven appearance optimisation framework for P-SAR. Rather than treating the problem as projector calibration or generic image-level radiometric compensation alone, the proposed method addressed the more specific challenge of reproducing a screen-defined material appearance on a projected physical mock-up for appearance-oriented design evaluation. The workflow was formulated around the material structure used in the rendering process: albedo appearance was optimised through dominant-colour group correction, while normal-induced shading appearance was refined through a lightness-contrast descriptor derived from flat and normal-modulated renderings. By combining controlled Unity-based rendering, calibrated screen and projector output, colourimetric measurement, D65 CIE Lab analysis, and MATLAB-based iterative update, the framework established an interpretable closed loop for improving projected material appearance in controlled P-SAR conditions.

The experimental results demonstrated the effectiveness of the proposed method under the tested setup. In the albedo branch, 23 dominant colour groups were identified across ten material textures, and all groups converged below the adopted CIEDE2000 threshold of $\Delta E_{00} = 2.3$. At the texture level, the mean colour difference decreased from 6.24 to 1.36, corresponding to an average reduction of 77.43%. The comparative baseline analysis further showed that the proposed group-wise optimisation achieved lower residual errors and more consistent convergence than global neutral-grey correction and global Lab-offset correction. This supports the use of independent dominant-colour group updates, especially for textures containing multiple visually distinct chromatic regions.

In the shading branch, the proposed contrast-guided normal-map update reduced the mean residual $r_D = |D_{\text{proj}} - D_{\text{ref}}|$ from 1.164 to 0.264 L^* units. Six of the ten normal-map cases already satisfied the adopted 1.0 L^* tolerance at the initial evaluation, while the remaining four cases converged after one update. The comparison with a global luminance-contrast baseline indicated that directly updating the normal-map-driven shading response provides a more material-relevant correction than only adjusting the rendered image intensity in the image domain. In addition, the threshold analysis clarified that the 1.0 L^* criterion should be interpreted as a conservative operational tolerance for the proposed lightness-based shading descriptor, rather than as a universal psychophysical boundary.

The additional analyses also clarified the role and limitations of key parameters in the workflow. The sensitivity analysis on the number of dominant colour groups showed that K should be understood as a representational parameter. For near-uniform textures, increasing K provided limited benefit, whereas for textures with more complex chromatic structure, too small a value of K could under-represent visually important colour regions and reduce the meaningfulness of the optimisation target. This result supports the current use of a small number of dominant colour groups, while also indicating that future automatic or semi-automatic K selection would improve reproducibility and reduce operator dependence.

At the integrated appearance level, photographic examples provided a qualitative visual summary of the complete two-stage workflow. These examples showed a closer overall visual tendency between the optimised projections and the screen-side references in representative material cases, including woven fabric, leopard print leather, and brick wall textures. However, the photographs were used only as qualitative illustrations of the integrated outcome. The main evidence for the proposed workflow remains the objective colourimetric and lightness-domain evaluation reported for the albedo and shading branches.

From a practical perspective, the complete workflow required approximately 26 min per material case, excluding one-off preparation steps such as device warm-up and initial calibration. The refined time analysis showed that the main operational burden comes from manual colourimetric measurement and user inspection rather than from Unity rendering or MATLAB computation. The method is therefore not intended as a real-time compensation technique to be repeated from scratch during every design iteration. Instead, it is better suited to controlled or semi-controlled P-SAR material appearance preparation, where a material-specific or setup-specific correction can be generated and reused under stable display, projection, surface, and lighting conditions.

Some limitations remain. First, the current validation was conducted on a planar white matte projection surface in a darkened environment. Extension to curved or non-planar objects would require geometric registration, projection mapping, and local or multi-region appearance measurement. Coloured, textured, or multi-material substrates would further require substrate characterisation and may introduce achievable-colour constraints. Secondly, the method depends on the measurable and reproducible behaviour of the display, projector, projection surface, and colourimeter combination. Different hardware configurations may have different luminance ranges, colour gamuts, spectral power distributions, contrast responses, and short-term stability, and would therefore require equivalent calibration and spectral correction. Thirdly, the present implementation still involves manual decisions, especially in the selection of the number of dominant colour groups and in the physical measurement procedure. Finally, the integrated photographic examples remain qualitative illustrations only, and the validation reported in this study is based on objective colourimetric and lightness-domain measurements under controlled conditions.

Future work will therefore proceed in three main directions. The first is to reduce operator dependence through automatic or semi-automatic dominant-colour grouping, scripted measurement acquisition, fixed measurement rigs, or integrated projector-camera feedback. The second is to extend the workflow beyond the current planar white-surface setup towards curved objects, non-white substrates, mixed-material mock-ups, and more realistic ambient-light conditions. The third is to establish a unified comparison protocol with camera-based full projector compensation methods, such as CompenNet-style approaches, so that image-level compensation and material-component-aware optimisation can be evaluated fairly under shared experimental conditions. These extensions would fur-

ther strengthen the reproducibility, robustness, and practical applicability of the proposed screen-referenced appearance optimisation framework.

Author Contributions: Conceptualisation, F.M. and G.C.; methodology, L.W.; software, L.W.; validation, L.W., F.M., and G.C.; formal analysis, L.W.; investigation, L.W.; resources, F.M. and G.C.; data curation, L.W.; writing—original draft preparation, L.W.; writing—review and editing, F.M. and G.C.; visualisation, L.W.; supervision, F.M. and G.C.; project administration, L.W.; funding acquisition, F.M. and G.C. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable. No human participants were recruited and no participant data were collected in this study.

Informed Consent Statement: Not applicable. No participant data were collected in the present study.

Data Availability Statement: Dataset available on request from the authors.

Conflicts of Interest: The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Abbreviations

The following abbreviations are used in this manuscript:

P-SAR	Projection-based spatial augmented reality
CIE	Commission internationale de l'éclairage
AR	Augmented reality
LUT	Lookup table
RGB	Red–green–blue
ΔE_{00}	CIEDE2000
SSIM	Structural similarity index
cd/m ²	Candela per square metre
K	Kelvin
LCD	Liquid crystal display
W-LED	White light-emitting diode
ICC	International Color Consortium
CCSS	Colorimeter calibration spectral sample
URP	Universal render pipeline
PNG	Portable network graphics

Appendix A. Implementation Details of the Optimisation Workflow

This appendix provides the implementation parameters and pseudo-code used for the two optimisation branches described in Section 4.

Appendix A.1. Implementation Parameters

Table A1 summarises the main implementation parameters used in the optimisation workflow. Parameters that were applied to both branches, such as the measurement repetition number and the colourimetric conversion basis, are reported together with branch-specific parameters for albedo and shading optimisation.

Table A1. Main implementation parameters used in the screen-referenced appearance optimisation workflow.

Parameter	Default Value/Setting	Branch	Purpose
Unity colour space	Linear	Both	Fixed project-level colour-space setting used throughout the rendering workflow.
Measurement repetitions	3	Both	Number of repeated colourimeter readings averaged before colour-space conversion and optimisation.
Measurement format	CIE XYZ	Both	Device-independent tristimulus values recorded by the colourimeter.
Comparison colour space	D65 CIE Lab	Both	Perceptually motivated colour space used for residual calculation and optimisation.
Number of dominant colour groups, K	1–3	Albedo	Controls the dominant-colour representation of each source albedo texture.
Soft-mask smoothing, σ	5.0 pixels	Albedo	Gaussian smoothing parameter used to convert hard clustering masks into soft spatial masks.
Albedo update relaxation, α	0.5	Albedo	Controls the magnitude of each Lab-space group-colour update.
Albedo convergence threshold, τ_A	2.3	Albedo	Group-wise CIEDE2000 threshold used as the albedo stopping criterion.
Uniform grey input	sRGB [119, 119, 119]	Shading	Neutral albedo input used to suppress texture-dependent colour variation.
Normal-map update relaxation, β	0.5	Shading	Controls the magnitude of the contrast-guided normal-map update.
Small contrast lower bound, ϵ	10^{-9}	Shading	Lower bound used in the contrast-ratio calculation to avoid division by values close to zero.
Shading convergence threshold, τ_D	$1.0 L^*$	Shading	Operational tolerance for the residual between screen-side and projection-side shading contrast descriptors.
Maximum albedo iterations, T_A	10	Albedo	Safety stopping condition for the group-wise albedo optimisation loop.
Maximum shading iterations, T_D	5	Shading	Safety stopping condition for the contrast-guided normal-map optimisation loop.

Appendix A.2. Pseudo-Code of the Albedo Appearance Optimisation

Algorithm A1 reports the implementation logic of the proposed group-wise albedo appearance optimisation. The procedure starts from a source albedo map, decomposes it into a small number of dominant colour groups, optimises each group independently using screen-referenced colour measurements, and reconstructs the final albedo map through a soft-mask-based correction field.

Algorithm A1. Proposed group-wise albedo appearance optimisation.

Require: Source albedo map A_{src} ; number of dominant colour groups K ; soft-mask smoothing parameter σ ; relaxation factor α ; convergence threshold τ_A ; maximum iteration number T_A .

Ensure: Optimised albedo map A_{opt} and iteration log.

- 1: Convert A_{src} from RGB to D65 CIE Lab to obtain $\mathbf{C}_{\text{src}}(x, y)$.
- 2: Cluster the Lab pixels into K dominant colour groups using k-means clustering.
- 3: Sort the colour groups by area and generate the corresponding hard masks.
- 4: Smooth the hard masks using a Gaussian filter with parameter σ .
- 5: Normalise the smoothed masks to obtain soft masks $w_k(x, y)$, where $\sum_{k=1}^K w_k(x, y) = 1$.
- 6: Compute the source group centre $\mathbf{C}_k^{\text{src}}$ for each group as the soft-mask-weighted mean Lab value.
- 7: Initialise the commanded output colour of each group:

$$\mathbf{C}_k^{\text{out}}(0) = \mathbf{C}_k^{\text{src}}.$$

- 8: **for** each colour group $k = 1, \dots, K$ **do**
- 9: Generate the screen-side reference colour patch from $\mathbf{C}_k^{\text{src}}$.
- 10: Display the reference patch on the calibrated monitor and measure the averaged Lab value $\mathbf{C}_k^{\text{ref}}$.
- 11: Set $t = 0$.
- 12: **while** $t < T_A$ **do**
- 13: Generate the projection-side colour patch from $\mathbf{C}_k^{\text{out}}(t)$.
- 14: Display the patch through the projector and measure the averaged projected Lab value $\mathbf{C}_k^{\text{proj}}(t)$.
- 15: Compute the colour difference $\Delta E_{00,k}(t)$ between $\mathbf{C}_k^{\text{ref}}$ and $\mathbf{C}_k^{\text{proj}}(t)$.
- 16: **if** $\Delta E_{00,k}(t) < \tau_A$ **then**
- 17: Mark group k as converged.
- 18: **break**
- 19: **end if**
- 20: Compute the Lab residual:

$$\Delta \mathbf{C}_k(t) = \mathbf{C}_k^{\text{ref}} - \mathbf{C}_k^{\text{proj}}(t).$$

- 21: Update the commanded group colour:

$$\mathbf{C}_k^{\text{out}}(t+1) = \mathbf{C}_k^{\text{out}}(t) + \alpha \Delta \mathbf{C}_k(t).$$

- 22: Apply the valid colour-range constraint.
- 23: Set $t = t + 1$.
- 24: **end while**
- 25: **end for**
- 26: Compute the group-wise Lab correction:

$$\Delta \mathbf{C}_k^{\text{out}} = \mathbf{C}_k^{\text{out}} - \mathbf{C}_k^{\text{src}}.$$

- 27: Construct the full-resolution Lab correction field:

$$\Delta \mathbf{C}_{\text{field}}(x, y) = \sum_{k=1}^K w_k(x, y) \Delta \mathbf{C}_k^{\text{out}}.$$

- 28: Reconstruct the optimised Lab texture:

$$\mathbf{C}_{\text{opt}}(x, y) = \mathbf{C}_{\text{src}}(x, y) + \Delta \mathbf{C}_{\text{field}}(x, y).$$

- 29: Convert $\mathbf{C}_{\text{opt}}(x, y)$ back to RGB and export the optimised albedo map A_{opt} .
- 30: Record measured XYZ/Lab values, ΔE_{00} , update values, output file names, and convergence status in the iteration log.

Appendix A.3. Pseudo-Code of the Shading Appearance Optimisation

Algorithm A2 reports the implementation logic of the proposed shading appearance optimisation based on contrast-guided normal-map update. The procedure first establishes a uniform-albedo condition, computes the screen-side reference shading contrast, and then iteratively adjusts the tangential components of the normal map until the projection-side contrast approaches the screen-side reference.

Algorithm A2. Proposed shading appearance optimisation based on contrast-guided normal-map update.

Require: Original normal map N_{src} ; screen-side uniform albedo input G_{ref} ; calibrated projector-side uniform albedo input G_{proj} ; relaxation factor β ; convergence threshold τ_D ; maximum iteration number T_D .

Ensure: Optimised normal map N_{opt} and iteration log.

1: Establish the uniform-albedo condition. Use G_{ref} for the screen-side reference branch and G_{proj} for the projection branch.

2: Construct the screen-side reference flat patch by rendering G_{ref} with a flat normal map.

3: Measure the averaged lightness value $L_{\text{flat}}^{*\text{ref}}$.

4: Construct the screen-side reference mean patch by rendering G_{ref} with the original normal map N_{src} .

5: Compute the perceptual mean of the valid rendered region in D65 CIE Lab.

6: Generate the corresponding uniform mean patch and measure the averaged lightness value $L_{\text{mean}}^{*\text{ref}}$.

7: Compute the screen-side reference shading contrast:

$$D_{\text{ref}} = L_{\text{flat}}^{*\text{ref}} - L_{\text{mean}}^{*\text{ref}}.$$

8: Initialise the current normal map:

$$N(0) = N_{\text{src}}.$$

9: Set $t = 0$.

10: **while** $t < T_D$ **do**

11: Render the projector-side flat patch using G_{proj} and a flat normal map.

12: Measure the averaged lightness value $L_{\text{flat}}^{*\text{proj}}(t)$.

13: Render the projector-side normal-modulated image using G_{proj} and the current normal map $N(t)$.

14: Compute the perceptual mean of the valid rendered region in D65 CIE Lab.

15: Generate the corresponding projector-side mean patch and measure the averaged lightness value $L_{\text{mean}}^{*\text{proj}}(t)$.

16: Compute the projector-side shading contrast:

$$D_{\text{proj}}(t) = L_{\text{flat}}^{*\text{proj}}(t) - L_{\text{mean}}^{*\text{proj}}(t).$$

17: Compute the residual:

$$r_D(t) = |D_{\text{proj}}(t) - D_{\text{ref}}|.$$

18: **if** $r_D(t) < \tau_D$ **then**

19: Mark the normal map as converged.

20: **break**

21: **end if**

22: Compute the normal-map update factor:

$$\eta(t) = \left(\frac{D_{\text{ref}}}{\max(D_{\text{proj}}(t), \epsilon)} \right)^\beta,$$

where ϵ is a small lower bound used to avoid numerical instability.

23: Decode the current normal map $N(t)$ from RGB into unit normal vectors:

$$\mathbf{n}^{(t)} = (n_x^{(t)}, n_y^{(t)}, n_z^{(t)}).$$

24: Scale the tangential components:

$$n_x^{(t+1)} = \eta(t)n_x^{(t)}, \quad n_y^{(t+1)} = \eta(t)n_y^{(t)}.$$

25: Clamp the tangential components if necessary and reconstruct:

$$n_z^{(t+1)} = \sqrt{1 - (n_x^{(t+1)})^2 - (n_y^{(t+1)})^2}.$$

26: Renormalise the updated normal vector.

27: Encode the updated normal vectors back into RGB form to obtain $N(t+1)$.

28: Export the updated normal map and reload it in Unity for the next iteration.

29: Set $t = t + 1$.

30: **end while**

31: Export the final accepted normal map as N_{opt} .

32: Record measured XYZ/Lab values, D_{ref} , D_{proj} , residual r_D , update factor η , output file names, and convergence status in the iteration log.

References

1. Bloch, P.H. Seeking the ideal form: Product design and consumer response. *J. Mark.* **1995**, *59*, 16–29.
2. Crilly, N.; Moultrie, J.; Clarkson, P.J. Seeing things: Consumer response to the visual domain in product design. *Des. Stud.* **2004**, *25*, 547–577.
3. Blijlevens, J.; Creusen, M.E.H.; Schoormans, J.P.L. How consumers perceive product appearance: The identification of three product appearance attributes. *Int. J. Des.* **2009**, *3*, 27–35.
4. Morosi, F.; Carli, I.; Caruso, G.; Cascini, G.; Dekoninck, E.; Boujut, J.-F. Exploring tablet interfaces for product appearance authoring in spatial augmented reality. *Int. J. Hum.-Comput. Stud.* **2021**, *156*, 102719.

5. Ashby, M.F.; Johnson, K. The art of materials selection. *Mater. Today* **2003**, *6*, 24–35.
6. Veelaert, L.; Du Bois, E.; Moons, I.; Karana, E. Experiential characterization of materials in product design: A literature review. *Mater. Des.* **2020**, *190*, 108543.
7. Sauer, J.; Sonderegger, A. The influence of prototype fidelity and aesthetics of design in usability tests: Effects on user behaviour, subjective evaluation and emotion. *Appl. Ergon.* **2009**, *40*, 670–677.
8. Zhou, X.; Rau, P.-L.P. Determining fidelity of mixed prototypes: Effect of media and physical interaction. *Appl. Ergon.* **2019**, *80*, 111–118.
9. Cox, C.; Hicks, B.; Gopsill, J. Improving mixed-reality prototyping through a classification and characterisation of fidelity. *Proc. Des. Soc.* **2022**, *2*, 353–362.
10. Azuma, R.T. A survey of augmented reality. *Presence: Teleoperators Virtual Environ.* **1997**, *6*, 355–385.
11. Raskar, R.; Welch, G.; Cutts, M.; Lake, A.; Stesin, L.; Fuchs, H. The Office of the Future: A Unified Approach to Image-Based Modeling and Spatially Immersive Displays. In *Proceedings of SIGGRAPH '98*; ACM: New York, NY, USA, 1998; pp. 179–188.
12. Raskar, R.; Welch, G.; Low, K.-L.; Bandyopadhyay, D. Shader Lamps: Animating Real Objects with Image-Based Illumination. In *Rendering Techniques 2001*; Gortler, S.J., Myszkowski, K., Eds.; Springer: Vienna, Austria, 2001; pp. 89–102.
13. Verlinden, J.C.; Horváth, I. Analyzing opportunities for using interactive augmented prototyping in design practice. *Artif. Intell. Eng. Des. Anal. Manuf.* **2009**, *23*, 289–303.
14. Kent, L.; Snider, C.; Gopsill, J.; Hicks, B. Mixed reality in design prototyping: A systematic review. *Des. Stud.* **2021**, *77*, 101046.
15. Nam, T.-J.; Lee, W. Integrating hardware and software: Augmented reality based prototyping method for digital products. In *CHI '03 Extended Abstracts on Human Factors in Computing Systems*; ACM: New York, NY, USA, 2003; pp. 956–957.
16. Akaoka, E.; Ginn, T.; Vertegaal, R. DisplayObjects: Prototyping functional physical interfaces on 3D Styrofoam, paper or cardboard models. In *Proceedings of the Fourth International Conference on Tangible, Embedded, and Embodied Interaction*; ACM: New York, NY, USA, 2010; pp. 49–56.
17. Park, M.K.; Lim, K.J.; Seo, M.K.; Jung, S.J.; Lee, K.H. Spatial augmented reality for product appearance design evaluation. *J. Comput. Des. Eng.* **2015**, *2*, 38–46.
18. Cascini, G.; O'Hare, J.A.; Dekoninck, E.; Becattini, N.; Boujut, J.-F.; Ben Guefrache, F.; Carli, I.; Caruso, G.; Giunta, L.; Morosi, F. Exploring the use of AR technology for co-creative product and packaging design. *Comput. Ind.* **2020**, *123*, 103308.
19. Poulin, M.; Masclat, C.; Boujut, J.-F. Investigating the effects of spatial augmented reality on user participation in co-design sessions: A case study. *Comput. Ind.* **2024**, *154*, 104023.
20. Denni, M.; Spallazzo, D.; Ceconello, M. A. Investigating Spatial Augmented Reality Technology in the Cultural Heritage Sector: A Scoping Review. *Electronics* **2026**, *15*, 540.
21. Liaw, M.-J.; Chen, C.-Y.; Shieh, H.-P. D. Color characterization of an LC projection system using multiple-regression matrix and look-up-table with interpolation. In *Projection Displays IV*; Proceedings of SPIE; SPIE: Bellingham, WA, USA, 1998; Volume 3296, pp. 238–247.
22. Bimber, O.; Iwai, D.; Wetzstein, G.; Grundhöfer, A. The visual computing of projector-camera systems. *Comput. Graph. Forum* **2008**, *27*, 2219–2245.
23. Grundhöfer, A.; Iwai, D. Recent advances in projection mapping algorithms, hardware and applications. *Comput. Graph. Forum* **2018**, *37*, 653–675.
24. Nayar, S.K.; Peri, H.; Grossberg, M.D.; Belhumeur, P.N. A projection system with radiometric compensation for screen imperfections. In *Proceedings of the IEEE International Workshop on Projector-Camera Systems (PROCAMS)*; IEEE: Piscataway, NJ, USA, 2003.
25. Grossberg, M.D.; Peri, H.; Nayar, S.K.; Belhumeur, P. N. Making one object look like another: Controlling appearance using a projector-camera system. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*; IEEE: Washington, DC, USA, 2004; Volume 1, pp. I-452–I-459.
26. Grundhöfer, A.; Bimber, O. Real-time adaptive radiometric compensation. *IEEE Trans. Vis. Comput. Graph.* **2008**, *14*, 97–108.
27. Wetzstein, G.; Bimber, O. Radiometric compensation through inverse light transport. In *Proceedings of the 15th Pacific Conference on Computer Graphics and Applications*; IEEE: Maui, HI, USA, 2007; pp. 391–399.
28. Sheng, Y.; Yapo, T.C.; Cutler, B. Global illumination compensation for spatial augmented reality. *Comput. Graph. Forum* **2010**, *29*, 387–396.
29. Menk, C.; Koch, R. Truthful color reproduction in spatial augmented reality applications. *IEEE Trans. Vis. Comput. Graph.* **2013**, *19*, 236–248.
30. Jones, B.R.; Sodhi, R.; Budhiraja, P.; Karsch, K.; Bailey, B.P.; Forsyth, D. Projectibles: Optimizing surface color for projection. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software and Technology*; ACM: New York, NY, USA, 2015; pp. 137–146.
31. Huang, B.; Ling, H. End-to-End Projector Photometric Compensation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*; IEEE: Long Beach, CA, USA, 2019; pp. 6810–6819.

32. Huang, B.; Ling, H. CompenNet++: End-to-End Full Projector Compensation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*; IEEE: Seoul, Republic of Korea, 2019; pp. 7165–7174.
33. Huang, B.; Sun, T.; Ling, H. End-to-End full projector compensation. *IEEE Trans. Pattern Anal. Mach. Intell.* **2022**, *44*, 2953–2967.
34. Li, Y.; Yin, W.; Li, J.; Xie, X. Physics-based efficient full projector compensation using only natural images. *IEEE Trans. Vis. Comput. Graph.* **2024**, *30*, 4968–4982.
35. Wang, S.; You, Z.; Zhang, Y. A Novel Multi-Projection Correction Method Based on Binocular Vision. *Electronics* **2023**, *12*, 910.
36. Hu, B.; Li, L.; Wu, J.; Qian, J. Subjective and objective quality assessment for image restoration: A critical survey. *Signal Process. Image Commun.* **2020**, *85*, 115839.
37. Zhai, G.; Min, X. Perceptual image quality assessment: A survey. *Sci. China Inf. Sci.* **2020**, *63*, 211301.
38. International Organization for Standardization; International Commission on Illumination. *Colorimetry—Part 4: CIE 1976 L*a*b* Colour Space (ISO/CIE 11664-4:2019)*; CIE: Vienna, Austria, 2019.
39. Luo, M.R.; Cui, G.; Rigg, B. The development of the CIE 2000 colour-difference formula: CIEDE2000. *Color Res. Appl.* **2001**, *26*, 340–350.
40. Sharma, G.; Wu, W.; Dalal, E.N. The CIEDE2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations. *Color Res. Appl.* **2005**, *30*, 21–30.
41. Wang, Z.; Bovik, A.C.; Sheikh, H.R.; Simoncelli, E. P. Image quality assessment: From error visibility to structural similarity. *IEEE Trans. Image Process.* **2004**, *13*, 600–612.
42. Liu, H.; Huang, M.; Cui, G.; Luo, M.R.; Melgosa, M. Color-difference evaluation for digital images using a categorical judgment method. *J. Opt. Soc. Am. A* **2013**, *30*, 616–626.
43. Yang, Y.; Ming, J.; Yu, N. Color image quality assessment based on CIEDE2000. *Adv. Multimed.* **2012**, *2012*, 273723.
44. Díaz-Barrancas, F.; Cwierz, H.; Pardo, P.J. Real-Time Application of Computer Graphics Improvement Techniques Using Hyperspectral Textures in a Virtual Reality System. *Electronics* **2021**, *10*, 2852.
45. Morosi, F.; Caruso, G. High-Fidelity Rendering of Physical Colour References for Projected-Based Spatial Augmented Reality Design Applications. *Comput.-Aided Des. Appl.* **2021**, *18*, 343–356.
46. Commission Internationale de l’Eclairage (CIE). *Grey-Scale Calculation for Self-Luminous Devices*; CIE 228:2018; CIE: Vienna, Austria, 2018.
47. Carter, R. Suprathreshold gray scale is implied by thresholds. *Appl. Opt.* **2018**, *57*, 8751–8756.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.