



CRAEFT

care, judgment, dexterity

D3.2 Advanced digitisation technologies

Project Acronym	Craeft
Project Title	Craft Understanding, Education, Training, and Preservation for Posterity and Prosperity
Project Number	
Deliverable Number	D3.2
Deliverable Title	Deliverable
Work Package	
Authors	Xenophon Zabulis, Nikolaos Partarakis, Panagiotis Koutlemanis, David Arnaud, Lymberis Lymberiadis
Number of pages	172



This project has received funding from the European Commission, under the Horizon Europe research and innovation programme, Grant Agreement No 101094349.

<http://www.craeft.eu/>

Executive summary

This deliverable (D3.2) reports a set of cost-efficient and operationally practical methods for the digitisation of tangible cultural heritage, with an emphasis on high-fidelity capture and inspection of surface geometry and visual appearance. The work complements the Craeft digitisation palette by addressing cases that are commonly challenging in practice: wide-area surface documentation, surface relief (2½D) digitisation, garments on mannequins, transparent bodies, and controlled inspection of difficult materials.

The deliverable consolidates and extends two prior foundations: (i) a surface digitisation capability originating in Mingei, here expanded for wide-area acquisition and 2½D surface digitisation, and (ii) a transparent-object strand informed by earlier Transparent3D work, here adapted into an apparatus-oriented workflow with explicit applicability conditions and documented failure modes. In addition, the deliverable introduces controlled material-aware inspection assets for challenging materials, including porcelain body state comparisons (T1–T8) supported by an interactive viewer and associated spectral reference information.

Key outputs include: (1) a configurable surface scanning apparatus with acquisition planning, calibration, and performance characterisation for high-resolution 2D mosaics and 2½D relief reconstructions; (2) a multi-modal workflow for garment digitisation on mannequins; (3) a practical approach for digitising transparent bodies under constraints relevant to heritage handling; and (4) controlled comparative visualisation tools enabling traceable inspection across motion modes and material states. The annexes provide reproducibility and capacity-building assets, including apparatus descriptions, calibration notes, parameter tables, representative results, and viewer documentation.

The reported methods primarily capture geometry and texture, which is sufficient for a wide range of documentation and inspection tasks, but does not fully represent materiality in the broader sensory sense, nor does it by itself model material-dependent light interaction. Future work will therefore (i) extend surface digitisation outputs toward haptic rendering by leveraging 2½D surface relief, and (ii) integrate acquired assets with physics-based rendering, in alignment with the material-specific simulation work reported in D3.1, to improve visual presentation of translucent and transparent artefacts.

Document history

Date	Author	Affiliation	Comment
2024-06-15	Xenophon Zabulis	ICS/FORTH	Created the basic format of the deliverable
2024-07-01	Xenophon Zabulis	ICS/FORTH	Created the first draft and added references
2024-07-15	Ioanna Demeridou	ICS/FORTH	Inserted figures & tables.
2024-07-20	Xenophon Zabulis	ICS/FORTH	Added results, references, and bibliography.
2024-07-26	Ioanna Demeridou	ICS/FORTH	Added an executive summary
2024-07-28	Xenophon Zabulis	ICS/FORTH	Proofread and made corrections.
2024-07-29	Xenophon Zabulis	ICS/FORTH	Read the review and adjustments.
2024-08-10	Xenophon Zabulis	ICS/FORTH	Update for reviewer comments.
2024-08-20	Xenophon Zabulis	ICS/FORTH	Finalising document (v1).
2025-11-20	Xenophon Zabulis	ICS/FORTH	Drafting document (v2).
2025-12-11	Xenophon Zabulis	ICS/FORTH	Drafting document (v2).
2026-01-07	Xenophon Zabulis	ICS/FORTH	Finalising document (v2).
2026-02-07	Xenophon Zabulis	ICS/FORTH	Revised document (v2).

Abbreviations

C3Dp	Cartesian 3D printer	C2Ds	Cartesian 2D scanner
FoV	Field of View	GPS	Global Positioning System
JPEG	Joint Photographic Experts Group	USD	United States Dollar
bAh	milliamp Hour	c	Cent
G	Giga	hrs	hours
mm	millimetre	µm	micrometre
cm	centimetre	kg	kilogram
p	pixel	pt	point
ppi	points per inch	T	tera
V	Volt	W	Watt
AC-RANSAC	A Contrario Random Sample Consensus	AR	Augmented Reality
BSDF	Bidirectional Scattering Distribution Function	CNC	Computer Numerical Control
CT	Computed Tomography	GA	Grant Agreement
HDRI	High Dynamic Range Image	ICP	Iterative Closest Point
IOR	Index of Refraction	IR	Infrared
LCD	Liquid Crystal Display	LED	Light Emitting Diode
LiDAR	Light Detection and Ranging	MP	Megapixel
MR	Magnetic Resonance	MVS	Multi-View Stereo
NeRF	Neural Radiance Fields	OCT	Optical Computed Tomography
OPT	Optical Projection Tomography	PBR	Physically Based Rendering
PCT	Patent Cooperation Treaty	PLY	Polygon File Format
PNG	Portable Network Graphic	RANSAC	Random Sample Consensus
SIFT	Scale-Invariant Feature Transform	SLAM	Simultaneous Localisation and Mapping
SPP	Samples Per Pixel	THz	Terahertz
UV	Ultraviolet	VR	Virtual Reality
XCT	X-ray Computed Tomography		



Table of contents

Executive summary	2
Document history	3
Abbreviations	4
Table of contents	5
1 Introduction	12
1.1 Scope and contributions	12
1.2 Continuity with prior work and novelty in Craeft	13
1.2.1 Surface digitisation.....	13
1.2.2 Transparent objects	14
1.2.3 Other outcomes	14
1.2.4 Complementarity with D3.1.....	14
1.3 Reader’s guide.....	15
2 Surface digitisation.....	16
2.1 Introduction	16
2.1.1 Novelty	16
2.1.2 Why accuracy is important	17
2.1.3 Legacy and novel apparatus.....	17
2.1.4 Overview	20
2.2 Background	21
2.2.1 2D surface scanning	21
2.2.2 2½D surface scanning	23
2.3 Proposed surface digitisation approaches.....	26



- 2.3.1 Hierarchical image acquisition 26
- 2.3.2 Point correspondence 28
- 2.3.3 2D surface reconstruction 31
- 2.3.4 2½D surface reconstruction 34
- 2.3.5 Evaluation 37
- 2.3.6 Figure-ground segmentation 43
- 2.4 Wide-area use cases 44
 - 2.4.1 Wide-area 2D scans 45
 - 2.4.2 Wide-area 2½D scans 49
- 2.5 Fine-detail use cases 51
 - 2.5.1 2D results 52
 - 2.5.2 2½D results 54
- 2.6 Conclusions 58
- 3 Garment digitisation 60
 - 3.1 Data acquisition 61
 - 3.2 Geometrical reconstruction 61
 - 3.3 Texture synthesis 63
 - 3.4 Results 63
- 4 Digitisation of transparent bodies 66
 - 4.1 Background 66
 - 4.2 Method 67
 - 4.3 Materials 68
 - 4.3.1 Optical apparatus 69
 - 4.3.2 Computational infrastructure 69



- 4.4 Calibration 69
 - 4.4.1 Intrinsic camera calibration 70
 - 4.4.2 Extrinsic camera–turntable calibration..... 70
 - 4.4.3 Motor step calibration 70
 - 4.4.4 Metric scale calibration..... 70
- 4.5 Reconstruction 71
- 4.6 Results 73
 - 4.6.1 Setup 74
 - 4.6.2 Evaluation 75
- 4.7 Comparison with view synthesis..... 79
 - 4.7.1 Theoretical expectations..... 80
 - 4.7.2 Experimental assessment 81
 - 4.7.3 Conclusions 83
- 4.8 Conclusions 83
 - 4.8.1 Technical 83
 - 4.8.2 Patent application 84
- 5 Digitisation of challenging materials..... 85
 - 5.1 Introduction 85
 - 5.1.1 Professional significance 85
 - 5.1.2 Ecological motivation..... 86
 - 5.1.3 Definitions 87
 - 5.2 Background 87
 - 5.2.1 Digitisation 88
 - 5.2.2 Visualisation 89



5.2.3 Spectrometry and spectroscopy	89
5.3 Transparent and translucent materials.....	90
5.3.1 Introduction	91
5.3.2 Theory	91
5.3.3 Software parameterisation.....	92
5.3.4 Material representation.....	93
5.3.5 Geometric representation	93
5.3.6 Validation experiments.....	94
5.4 Applications.....	98
5.4.1 Solids	98
5.4.2 Glass containers	98
5.4.3 Porcelain bodies.....	101
5.4.4 Discussion.....	109
6 Conclusions	111
6.1 Surface digitisation.....	111
6.2 Garment digitisation	111
6.3 Transparent bodies	111
6.4 Challenging materials and material-aware presentation	112
6.5 Limitations and next steps	112
Annex A. Surface scanner	113
A.1 Apparatus.....	113
A.1.1 Configuration for close-range 2D surface scanning.....	113
A.1.2 Configuration for close-range 2½D surface digitisation	114
A.2 Surface scanning and digitisation apparatus.....	115



A.2.1 Off-the-net and off-the-shelf components..... 115

A.2.2 Motor and motion 116

A.2.3 Imaging 116

A.2.4 Communication and control 117

A.2.5 Printed components 118

A.3 Calibration and image acquisition 118

 A.3.1 Calibration..... 118

 A.3.2 Image acquisition..... 119

A.4 Market survey..... 121

Annex B. 2D surface scanning legacy results 123

 B.1 Qualitative..... 123

 B.1.1 Paintings..... 124

 B.1.2 Paper and canvas 125

 B.1.3 Repetitive patterns 126

 B.1.4 Fabrics 126

 B.1.5 Fiona..... 127

 B.1.6 Approximately planar surfaces 127

 B.2 Quantitative 128

Annex C. Wide-area 2D scans 131

 C.1 Mosaics 131

 C.1.1 Gallery 131

 C.1.2 Dataset 131

 C.1.3 Foreground segmentation 134

 C.2 Batch scans 136



C.2.1 Gallery 136

C.2.2 Foreground segmentation 137

Annex D. Garment digitisation..... 140

D.1 Acquiring 3D data 140

D.1.1 Step 1: Preparation..... 140

D.1.2 Step 2: Handheld scan 140

D.1.3 Step 3: Laser scanning 140

D.1.4 Step 4: Photogrammetry 141

D.1.5 Step 5: Quality-control checklist..... 141

D.2 Cretan textiles collection 141

Annex E. OCT..... 146

E.1 Principles 146

E.1.1 Measuring visible light ray beam shift for small sidewall thicknesses..... 146

E.1.2 Light rays propagating through hollow circular objects 148

E.2 Patent application documents 149

Annex F. Porcelain material states..... 156

Annex G. Spectral measurements..... 157

G.1 CERFAV study..... 157

G.2 Reference samples 158

G.2.1 UG11 - “Invisible” UV glass..... 158

G.2.2 BG3 - saturated blue-green bandpass 158

G.2.3 BG38 - pale cyan IR-cut glass..... 159

G.2.4 VG20 - rich green, visible-pass / NIR-cut 159

G.2.5 GG495 - yellow longpass (green-yellow cut-on) 159



G.2.6 OG530 - orange longpass 159

G.2.7 RG610 - deep red / visible-to-IR separator..... 159

Notes..... 160

Annex H. Interactive viewer..... 161

 H.1 Top bar elements..... 161

 H.2 Four-panel layout and status footer..... 161

 H.3 Motion semantics and comparative use 161

 H.1 Top bar elements..... 162

 H.2 Footer 162

 H.3 Comparative purpose 162

Annex I. NeRFs 164

References 165

1 Introduction

That “traditional craftsmanship is perhaps the most tangible manifestation of intangible cultural heritage”¹ underscores a practical challenge for digital documentation: craft outcomes are experienced through materiality. In this deliverable, we prioritise methods that capture and communicate shape and visual detail with high fidelity.

1.1 Scope and contributions

The democratisation of digital imaging and 3D reconstruction has increased expectations for digital inspection of objects, including handmade products and heritage artefacts. However, high-quality digitisation is not always simple or economical when fine surface detail, accurate scale, and consistent appearance are required. Traditional photographic workflows demand expertise that many practitioners do not possess, and high-resolution digitisation often requires specialised equipment and careful acquisition planning. In addition, many craft outputs employ materials that are challenging for standard acquisition methods: specular, glossy, translucent, and transparent surfaces violate the approximate Lambertian assumptions that underlie many photogrammetric and active scanning pipelines, and specialised solutions can be expensive or unsuitable for sensitive objects.

Against this background, D3.2 reports a set of digitisation methods that are intended to be cost-efficient, configurable, and operationally practical, and that complement the Craeft palette of approaches for reconstructing tangible heritage. The key contributions are:

- Surface digitisation at very high resolution: a scanning approach and apparatus supporting both 2D surface mosaics and 2½D surface digitisation, including scaling from a small prototype to wide-area scanning regimes.
- Garment digitisation on mannequins: a practical workflow combining complementary acquisition modalities to obtain metrically meaningful surfaces with usable visual appearance.
- Transparent-object digitisation: a scanning method and apparatus targeting transparent bodies under constraints relevant to cultural heritage handling.
- Material-aware inspection and presentation for challenging materials: controlled rendering studies and an interactive viewer enabling traceable comparison of material states (e.g., porcelain bodies) under consistent motion and illumination conditions.
- Capacity-building and reproducibility assets: annex material documenting apparatus, calibration, operating procedures, parameterisations, and inspection tools to support reuse and comparative assessment.

Two cross-cutting design axes guide these contributions. First, cost-efficiency is pursued through low-cost sensors, printable or readily available hardware components, and acquisition plans that explicitly manage memory, storage, and computation constraints. Second, ease of installation and operation is pursued so that non-expert users can reproduce results without relying on ad-hoc expertise or fragile manual steps.

¹ <https://ich.unesco.org/en/traditional-craftsmanship-00057>

Part of the reported work builds on earlier FORTH developments and is extended here in the Craeft context. In particular, a close-range surface scanning capability initially developed in prior work is expanded through wide-area scaling and 2½D digitisation, and the transparent-object strand extends and adapts insights from earlier research toward workflows suitable for the Craeft objectives. Table 1 summarises the relationship between prior foundations and the extensions reported in this deliverable.

1.2 Continuity with prior work and novelty in Craeft

We distinguish between (i) capabilities inherited from prior projects and publications and (ii) contributions developed or substantively extended within Craeft. D3.2 therefore reports both continuity and novelty across two main foundations: the Mingei surface digitisation line and the Transparent3D strand for challenging transparent objects.

Table 1 overviews the contributions per project, as described below in this subsection.

Table 1. Contributions per project.

Project	Result	Target	Material	Scale
Mingei	2D scan	Planar surfaces	Opaque, matte or shiny	Small
Craeft	2D scan	Planar surfaces		Large
Craeft	2½D digitisation	Undulating surfaces		Large
Transparent3D	Equipment & Theory	Solids	Semi-transparent	Any
Craeft	3D digitisation	Thin-walled solid		Small

1.2.1 Surface digitisation

Continuity from Mingei

The close-range surface scanning concept and its core acquisition rationale were established in the Mingei project and documented in Zabulis X, Koutlemanis P, Stivaktakis N, Partarakis N. A Low-Cost Contactless Overhead Micrometre Surface Scanner. Applied Sciences. 2021; 11(14):6274. [doi.org:10.3390/app11146274](https://doi.org/10.3390/app11146274). In that earlier work, the emphasis was on cost-efficient imaging, repeatable acquisition plans, and the construction of high-resolution surface mosaics supported by deliberate overlap and multi-layer capture. Craeft inherits these foundations, including the general approach to scan planning, acquisition orchestration, and evidence-driven evaluation.

Novelty and extensions in Craeft

Craeft extends the Mingei foundations in three concrete ways. First, the scanner is re-engineered and validated for wide-area operation, including the practical use of multiple laterally overlapping scan “pyramids” to increase coverage while maintaining registration stability (Section 2; Annex A). Second, the pipeline is extended from 2D mosaicking toward 2½D surface digitisation, producing metrically meaningful relief geometry together with high-resolution texture for material-sensitive surfaces (Section 2; Annex A). Third, Craeft strengthens reproducibility assets: calibrated parameter sets, configuration tables, and operational guidance are consolidated into annex material intended for reuse and inspection. This led to the following publication: Koutlemanis P, Zabulis X, Stivaktakis N, Partarakis N, Zidianakis E and Demeridou

I (2024) A low-cost close-range photogrammetric surface scanner. *Front. Imaging*. 3:1341343. [doi:10.3389/fimag.2024.1341343](https://doi.org/10.3389/fimag.2024.1341343).

1.2.2 Transparent objects

Continuity from Transparent3D

Transparent-object digitisation in D3.2 builds on concepts previously introduced in Transparent3D and documented in Stavroulakis, P.I.; Ganetsos, T.; Zabulis, X. Large Scale Optical Projection Tomography without the Use of Refractive-Index-Matching Liquid. *Sensors* 2023, 23, 9814. [doi:10.3390/s23249814](https://doi.org/10.3390/s23249814). That strand established that transparent bodies require specialised handling because refraction and specular transport invalidate the assumptions of standard photogrammetry and many active scanning approaches.

Novelty and adaptation in Craeft

Craeft adapts the transparent-object approach to the constraints typical of heritage handling and craft documentation, emphasising (i) practical apparatus configurations, (ii) repeatable acquisition, and (iii) an evidence-led description of applicability conditions and characteristic failure modes (Section 4; Annex E). This positions the method as a usable component within a broader digitisation palette rather than a one-size-fits-all solution. The integrated technology and apparatus (setup) exhibit industrial interest in glass and plastic manufacturing and quality control. Thus, we initiated the process of applying for a patent (patent prosecution). Currently, this process has reached the PCT (Patent Cooperation Treaty) stage, often referred to as the 'international phase' or 'PCT phase' of patent prosecution. The PCT ID of this application is PCT/GR2023/000051.

1.2.3 Other outcomes

The work in Section 3 on garment digitisation has been published in Xhako, A., Katzourakis, A., Evdaimon, T., Zidianakis, E., Partarakis, N., & Zabulis, X. (2024). Reviving Antiquity in the Digital Era: Digitization, Semantic Curation, and VR Exhibition of Contemporary Dresses. *Computers*, 13(3), 57. <https://doi.org/10.3390/computers13030057>.

The viewer in Section 5 featuring the comparative visualisations of porcelain appearance rendering below can be downloaded at: <https://doi.org/10.5281/zenodo.18190038>. Reference: Zabulis, X. (2026). Interactive viewer for porcelain materiality and lineage assessment. Zenodo. doi: [10.5281/zenodo.18190038](https://doi.org/10.5281/zenodo.18190038)

1.2.4 Complementarity with D3.1

Finally, the deliverable explicitly separates *digitisation* from *material-aware visualisation*: D3.2 focuses on the acquisition of geometry and texture together with controlled inspection assets, while the material-specific simulation of light interaction that is required for more faithful appearance reproduction is studied in D3.1 (craft-specific action simulations).



1.3 Reader's guide

Section 2 presents the surface digitisation approach (2D and 2½D) and its wide-area configurations. Section 3 reports the garment digitisation workflow. Section 4 reports the method for transparent objects. Section 5 addresses digitisation and presentation challenges for difficult materials and introduces the controlled viewer-based comparative studies (including porcelain material states and associated measurement/inspection assets). Section 6 concludes with limitations and next steps. The annexes provide apparatus details, calibration and operating parameters, and supporting materials intended for inspection, capacity building, and reuse.

2 Surface digitisation

Section 2 at a glance:

- **Purpose.** This section presents the surface digitisation pipeline for high-resolution 2D and 2½D capture of craft-relevant surfaces, with emphasis on accuracy and repeatability.
- **Inputs and setup.** Imaging apparatus and acquisition configurations (see Annex A); representative legacy and wide-area datasets (see Annexes B–C).
- **Method summary.** Hierarchical acquisition → correspondence estimation → 2D reconstruction / 2½D reconstruction → evaluation.
- **Outputs & evidence.** Quantitative/qualitative evaluation (Section 2.3.5) and application galleries (Sections 2.4–2.5; Annexes B–C).

2.1 Introduction

We propose a low-cost method for digitising nominally planar or quasi-planar surfaces, such as textiles, painted panels, and shallow reliefs, in both 2D (appearance) and 2½D (surface shape). The core idea is to trade expensive optics for sampling density: we acquire a large number of close-range images so that each photograph resolves fine detail over a small footprint, and the collection jointly covers the full area.

This design choice yields very high resolution and strong geometric fidelity using modest mechanical and optical components, making the approach suitable for workshops and FabLabs. An open-source, open-design, DIY ethos further supports substitution of components (camera, lenses, motion stages) according to local availability, without changing the underlying method.

Dense imaging, however, creates two coupled challenges: (i) accurate image registration (“stitching”) at scale, and (ii) practical management of large data volumes. Our acquisition strategy and processing pipeline are built to address both. The method also scales predictably with resources: improved photographic equipment increases image definition, while more capable computing increases reconstructed area and/or reduces processing time.

2.1.1 Novelty

The main contribution is the scanning strategy that interleaves proximal (close-up) and distal (farther) images. Close-up images provide detail but limited context; if stitched purely in a serial, neighbour-to-neighbour manner, small alignment errors accumulate and eventually become visible as drift or local discontinuities. Distal images, while lower resolution, capture a larger portion of the surface and therefore supply long-range geometric constraints. With sufficient distance, a single distal view can even span the entire target, providing a global reference frame.

We exploit these properties in a coarse-to-fine, hierarchical registration scheme in which distal images act as pivots: they anchor the mosaic, limit drift, and guide how local registrations are composed. Conventional global optimisation is still used, but primarily as a final refinement step and not as the sole

mechanism for 'spreading' error after it has accumulated. In effect, the method aims to prevent error build-up by introducing global constraints early, and then to polish the result with a standard global adjustment.

To reduce human effort and operator-induced variability, image capture is automated by coordinating a CNC Cartesian stage with the camera's acquisition interface. Beyond automation, the CNC provides proprioceptive cues about camera position, which serve as informative priors for registration and improve digitisation accuracy.

2.1.2 Why accuracy is important

Accuracy is not a luxury feature here; it directly enables cost reduction. By leveraging auxiliary (distal) imagery and proprioceptive cues, the pipeline can compensate for mechanical imperfections, tolerate greater jitter, and remain robust under less stringent hardware specifications. This relaxes requirements on motors, bearings, and structural materials, allowing the use of lighter, lower-cost components and, consequently, lower-power actuation.

Finally, the acquisition pattern is adapted to the reconstruction target. For 2D mosaicking, overlap must support reliable global alignment. For 2½D reconstruction, overlap is increased to strengthen stereo correspondences across neighbouring views and stabilise surface recovery. In both settings, the same principle applies: combine dense close-up coverage with occasional distal constraints and motion-stage priors to obtain high-resolution, wide-area digitisation without high-end instrumentation.

2.1.3 Legacy and novel apparatus

Legacy apparatus

A mechanical apparatus automates the labour-intensive stage of image acquisition. It consists of a CNC Cartesian mechanism that provides three degrees of freedom, allowing the camera to be positioned frontally to the target surface at controlled offsets and at a range of stand-off distances. Motion and triggering are coordinated via the device's software interface, enabling repeatable capture protocols.

The resulting photographs are organised conceptually as layers of an image pyramid, indexed by distance to the surface. This layered structure is not merely descriptive; it is central to the method. Images acquired at larger distances are treated as auxiliary views: they are not used for surface reconstruction at a fine scale, but instead provide long-range consistency constraints that stabilise registration and improve overall accuracy.

Figure 1 illustrates this idea. The left panel abstracts the motion capabilities of the hardware, while the right panel depicts the sampling density implied by the layered strategy. The base of the pyramid corresponds to the final high-resolution mosaic, assembled from the densest close-range sampling. Each higher layer represents a mosaic formed at a greater distance, culminating in one or a few overview images that span most or all of the scene. Registration is constrained hierarchically: mosaics at lower layers must remain geometrically consistent with those above. In the limiting case where the top layer contains a single global view, it acts as a direct measurement that all derived mosaics should agree with.




Figure 1. Abstraction of C3Dp motion (left, from [1]) and proposed scanning locations (right).

During acquisition, the system records proprioceptive data (motor readings) from the CNC rig. These readings provide approximate camera poses that are injected into the processing pipeline to reduce the correspondence search space and to support global geometric optimisation.

Annex A summarises prior work on the hardware platform, the image-acquisition driver software, and calibration procedures, and includes a brief market survey of comparable devices.

Craeft extensions

Within Craeft, we extend the legacy system along two complementary axes:

1. **Algorithmic extension (from 2D to 2½D):** We move from appearance-only mosaicking to surface reconstruction by introducing stereoscopic processing. This requires denser sampling to increase the number and diversity of stereo correspondences.
2. **Scale-up (from 30 × 40 cm² to 3 × 4 m²):** We expand the scanning envelope by roughly two orders of magnitude in area, which introduces new constraints in optics, structural rigidity, and motion control.

In Craeft, we performed hardware upgrades only to create the wide-area digitisation apparatus. For the fine-detail capture, the hardware remained the same, and only the algorithm was extended. The two devices are shown in Figure 2.




Figure 2. Surface scanners for fine-detail capture (left) and wide-area digitisation (right).

Hardware and system upgrades

To demonstrate that the acquisition and registration strategy generalises across spatial scales, we extended the original fine-detail implementation into a wide-area digitisation platform. The fine-scale configuration targets micrometre-level surface detail, whereas the extended apparatus supports metre-scale capture of large textiles, maps, and related planar artefacts. During evaluation, a further use case emerged: batch scanning, where multiple items are digitised sequentially under a consistent, automated capture protocol.




Figure 3. Open-source CNC imaging platform.

The platform follows an open-source CNC design philosophy, with targeted modifications to improve structural stiffness, positioning accuracy, wiring reliability, and payload capacity for professional imaging equipment. The resulting system is intended for research-grade documentation and repeatable automated capture, where both geometric consistency and operational robustness are required.

Base design and overall dimensions

The system adopts a Cartesian architecture derived from a largely 3D-printed open-source CNC design. This baseline was selected for modularity, cost efficiency, and ease of adaptation. The final



implementation provides a large operational envelope with overall dimensions of approximately $4 \times 3 \times 1.8 \text{ m}^3$, enabling wide-area coverage while maintaining controlled linear motion along all axes. This geometry supports applications such as large-format object documentation, surface scanning, and structured image acquisition over extended scenes.

Structural reinforcement and stability enhancements

Scaling the platform to metre dimensions requires explicit measures to control flex, vibration, and long-term drift. The printed structure was therefore reinforced with an auxiliary framework of metallic iron rods mechanically coupled to the feet of the rig. This support skeleton (i) reduces torsional deformation and vibration during motion, (ii) improves load distribution across the base, and (iii) increases structural durability. The hybrid construction, 3D-printed components paired with metallic reinforcement, preserves the design's modularity while substantially improving mechanical rigidity.

Motion control electronics and precision improvements

To increase positioning accuracy and repeatability across the full workspace, we developed a custom motion controller incorporating stepper motor drivers with advanced microstepping. Relative to typical off-the-shelf solutions, the upgraded electronics provide (i) finer effective motion resolution, (ii) smoother acceleration and deceleration profiles, and (iii) reduced resonance and audible noise. These characteristics are particularly relevant when capturing dense image sequences, where small pose deviations can propagate into registration error.

Wiring design and cable management

Reliable operation over large travel distances depends on cable routing that tolerates repeated motion cycles without introducing strain-induced failures. We redesigned wiring paths and cable management to (i) minimise bending fatigue and strain, (ii) prevent interference with moving components, and (iii) maintain stable power delivery and signal integrity. This reduces the likelihood of intermittent faults during long scanning runs.

Payload support and camera mount design

To support professional imaging equipment, we designed and fabricated a custom 3D-printed camera mount capable of carrying a professional camera body and a close-range macro lens. The mount was engineered for rigidity, vibration damping, and precise alignment with the platform's motion axes. This mechanical stability is critical for macro imaging and for slow, continuous capture trajectories, where even small oscillations can degrade sharpness and compromise correspondence estimation.

Detailed apparatus configurations, components, and calibration/acquisition procedures are provided in Annex A.

2.1.4 Overview

The contribution is based on the legacy image acquisition strategy designed to aid the computational steps that follow. That is, images are captured in a (pyramidal) frustum arrangement, comprising multiple layers above the scanned surface. The top layers serve an auxiliary role; they are used to establish global

consistency constraints rather than for texturing the final reconstruction. This strategy provides lateral overlap for images on the same layer and medial overlap for images across layers.

To support the advanced requirements stemming from the goal of 2½D digitisation, the computational pipeline is updated to integrate hardware-assisted image acquisition with a robust computer vision processing chain.

1. **Feature Detection.** Once images are acquired, the software processes them to identify distinct features. The method detects keypoint features with content descriptors in all images. The baseline implementation uses the Scale-Invariant Feature Transform (SIFT), though other algorithms can be substituted.
2. **Correspondence.** The core computational task involves establishing matches between points in neighbouring images. This process is optimised through several constraints to handle the large data volume and ensure accuracy, including the reduction of search space based on the camera motion plan, as well as conventional techniques borrowed from the literature.
3. **Feature Tracks.** To increase robustness, additional geometric constraints are gained by extending pairwise correspondences into feature tracks, which are sequences of the same physical point identified across three or more images laterally or medially.
4. **Global Optimisation.** A conventional global optimisation refines the image registration and camera pose estimation results.

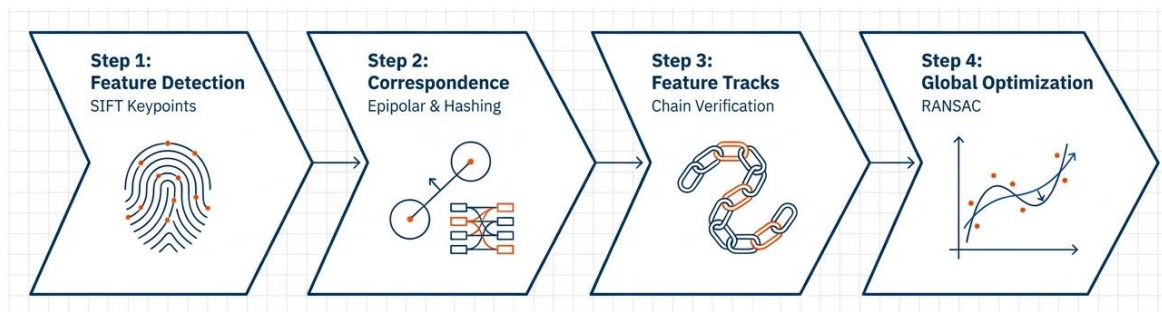


Figure 4. The computational pipeline.

2.2 Background

This review of 2D and 2½D surface scanning methods underscores the novel aspects of this work.

2.2.1 2D surface scanning

Generating image mosaics out of partial images of a surface is useful in many applications, because mosaics cover larger surface portions than single images. When a mapping between pixel and metric coordinates is calibrated, world measurements are performed using the mosaic, much like in cartographic maps. Image registration upon general surfaces enables photorealistic maps for Geographical Information Systems, photo-panoramas [2], but also specialised mosaics from usually unseen surfaces such as the inside of pipes [3], the gastrointestinal tract [4], and the human retina [5]. At the core of all methods for image mosaic generation is the problem of image registration.

A straightforward case of mosaic generation is imaging a planar surface by multiple and conveniently tessellated overlapping frontal views. This case is common in remote sensing [6], document scanning [7], bioinformatics [8, 9], and art [10]. Approaches based purely on visual cues need to estimate camera motion, and though they are continuously making progress, their precision is eventually limited by pixel quantisation; that is, for a large number of images, even minute errors accumulate, leading to distorted results. Applications that compute large mosaics utilise independent sources of information about the location of the camera. Notably, stitching of remote sensing and aerial images is supported by GPS measurements. We use this principle in the context of overhead scanning, where approximate location measurements are available from the motion mechanism of the scanner.

Visual image registration

The problem of computational image registration dates back at least four decades of study. Methods in the literature are typically referred to as 'local' if they utilise point feature correspondences, or 'global' if they employ overall image similarity [11, 12]. When combining images in a mosaic, these images are required to have some lateral 'overlap', and the registration task is called stitching. Due to restricted overlap, stitching is more accurate when local methods are used.

Given point correspondences across two images of a planar surface, robust registration solutions have been established and are now a standard piece of knowledge [13, 14]. When multiple images are registered sequentially to create a mosaic, errors accumulate, causing distortion. A solution involves employing a 'global alignment' or 'bundle adjustment' step, which either yields a more precise solution or, at the very least, distributes errors to preserve the shape of the scanned area. Although this enhances the overall outcome, the errors from registering multiple images appear as local, noticeable distortions, often referred to as 'seams'. For a small number of images, these distortions are minimal and effectively managed by methods that minimise their visual impact. However, when the number of images increases by two orders of magnitude, we observe these distortions becoming substantial and noticeable upon close and macroscopic inspection.

Proprioceptive image registration

Another way to register images in mosaics is by scanners, which use mechanisms to drive the sensor to locations where the acquired images would precisely match. The utilised sensors are most often line cameras with intense illumination and less often photographic cameras. A market survey of pertinent solutions can be found in Appendix B.

Contact-based, flatbed scanners provide up to 1000 ppi at a significantly high cost. Large-format scanners provide resolution up to 1200 ppi and are almost contactless. However, the scanned material should be less than a thickness threshold, i.e. 3 cm, to pass through the scanning slit and also incur a high cost. Film scanners exhibit higher resolution, but require contact and material transparency, and are limited to the frame size of photographic film.

Large-format scanners are contactless and designed for sensitive documents, but are also used for scanning textiles and other similar materials. They reach up to A0 scanning size. The precision required for this mechanical task elevates the cost of the scanning hardware.

Book scanners are contactless and exhibit resolution in the range of 600 – 1200 ppi. Their cost ranges from low to very high, though in many cases, the elevated cost is due to the mechanics for the automation of page-turning. It ought to be noted, though, that V-shaped, as opposed to flatbed, book scanners are unsuitable in the case of deformable materials, such as fabrics or sand.

Recently, the need for realistic textures has given rise to flatbed, contactless surface scanners, known as 'material scanners'. They use camera photography and exhibit resolution up to 1000 ppi. Their effective scanning surface measures approximately $30 \times 40 \text{ cm}^2$. Lower-cost material scanners utilise the sensor of the mobile phone [15]; however, they exhibit lower resolution and definition compared to the aforementioned solutions.

2.2.2 2½D surface scanning

The approaches which achieve accurate and reliable surface digitisation in large detail are as follows. X-ray Computed Tomography (XCT) uses X-rays to provide volumetric scans of objects and their surfaces, and has been used for the digitisation of small objects [16]. Magnetic Resonance (MR) tomography has also been used for detailed volumetric scans of small objects [17]. Laser scanning uses the time-of-flight principle to measure distances and reconstruct surfaces. Structured light scanners project light patterns onto surfaces, capture the deformed patterns using a camera, and, from these deformations, estimate surface structure. Finally, as this work proposes, photogrammetric reconstruction of small artefacts has also been employed, but with significant constraints as to the scanning area and hardware requirements.

Albeit highly accurate, tomographic methods (CT, MR) exhibit severe costs and can be applied only in specialised laboratories. Moreover, they provide volumetric and accurate structure reconstruction but do not reconstruct surface appearance (texture). Laser scanning and structured light scanning capture surface appearance and are marginally less accurate, but are more cost-efficient because they require less equipment. Photogrammetry is even more affordable and widely available, as it relies on standard cameras only. Moreover, photogrammetry captures the appearance of surfaces because it images surfaces without the effects of active illumination. Photogrammetry produces high-resolution and highly detailed textured models that accurately represent visual appearance. Compared to laser and structured light methods, photogrammetry can be less accurate, but it is passive in that it does not require the projection of energy in the form of radiation (e.g., light, X-rays, etc.) upon the scanned surface. This is important in light-sensitive materials encountered in historical, archaeological, biological, and cultural contexts.

The literature is first reviewed in terms of the applications where close-range photogrammetric systems are employed, as well as the optics and illumination methods utilised for their implementation. Thereafter, data acquisition quantities and reconstruction methods are reviewed. Last but not least, photogrammetry creates reconstructions 'up to scale', that is, the resultant reconstruction is not in metric units but in arbitrary ones. As the goals of this work include surface measurement, the presentation of the ways that the scale factor is estimated for close-range photogrammetric methods concludes this section.

Applications

Photogrammetry has been applied in a wide variety of contexts over the last three decades, mainly referring to large-scale reconstruction targets in the range of hundreds to a few meters, such as built and terrain structures captured from aerial views, indoor environments, and cultural heritage objects. More



recently, close-range photogrammetry has found applications in the high-resolution reconstruction of structures of much smaller scale, of industrial [18, 19, 20], biological [21], archaeological [22, 23], anthropological [24], and cultural interest [25, 26], in the range of millimetres (see [27] for a comprehensive review of close-range photogrammetry applications).

Motorisation

When using a single camera, camera motion is required to acquire 3D information. Motorisation of camera movement has been used in photogrammetric reconstruction to alleviate user effort and acquire images at numerous, prescribed viewpoints. The main strategies of motion are either circular around the target, using a turntable, or in a Cartesian lattice of viewpoints [28]. Cartesian approaches exhibit the advantages of being unconstrained by the turntable size, and the camera can be moved arbitrarily close to the scanning target. This work follows the latter approach to scan wider surface areas, avoiding the occurrence of shadows and illumination artefacts. The most relevant scanning apparatus for this work is [29], which uses a CNC to move the camera.

Optics

In the millimetre and sub-millimetre range, zoom and microscopic, or 'tube', lenses have been used [30, 31, 28], employing tedious calibration procedures and relatively inaccurate results [32]. Instead, macro lenses are more widely used in the photogrammetric reconstruction of such small structures [33]. However, macro lenses exhibit a very limited depth of field. This limitation makes photographs acquired with a macro lens out of focus in the periphery of the image. To compensate for this effect, focus stacking [34] is often employed in macro photography [35].

Illumination

Illumination is necessary for photogrammetric methods so that the surfaces are visible to the camera. Photogrammetry typically operates upon environmental illumination. Ambient illumination is ideal for photogrammetric methods because it prevents the formation of shadows, which confound the visual documentation of the reconstructed surfaces. Most works use setups that insulate the target object from environmental illumination and use a specific light source in conjunction with illumination diffusers to prevent the formation of shadows. In closer relation to this work, some systems use light sources that move along with the camera [36, 37, 38].

Purposefully designed illumination is used in photogrammetric methods to facilitate the establishment of more stereo correspondences. This technique is known as 'active illumination', 'active photogrammetry', or 'structured light' [37, 39, 40]. Pertinent methods use active illumination to support the reconstruction of the geometry of the reconstructed surface, by artificially creating reference points on the surface which can be matched across images. The disadvantage is that the projected light alters the visual appearance of the reconstructed surfaces. As such, this work does not use active illumination as it strives to realistically capture the visual appearance of the target surfaces.

Data

The data acquired from the close-range photogrammetric methods vary depending on the hardware and optical apparatus employed.

Works that report scanning areas fall in the range of $\approx [0.14, 158.7] \text{cm}^2$. Specifically, the maximum area reported from these works is (approximately) as follows: 0.196cm^2 [30], 10cm^2 [39], 15cm^2 [37], 22.5cm^2 [23], 25cm^2 [41], 32cm^2 [29], 40cm^2 [36], 57.6cm^2 [42], 60cm^2 [22], 91.5cm^2 [31], 93.4cm^2 [30], and 158.7cm^2 [42]. Out of these works, only [30] reports the achieved resolution, which is 3745p/mm^2 .

Works that report the number of images and pixels processed fall in the range of $[9.6, 3590.0] \text{MP}$. The data load is calculated as the number of utilised images times the resolution of each image. Specifically, the maximum number of pixels reported from these works is (approximately), in MP, as follows: $16 \times 0.6 = 9.6$ [31], $7 \times 10.1 = 70.7$ [38], $12 \times 10.1 = 121.2$ [41], $30 \times 24 = 720.0$ [40], $72 \times 12.3 = 885.6$ [29], $40 \times 36.3 = 1452.0$ [23], $245 \times 10.2 = 2,499.0$ [21], and $359 \times 10 = 3590.0$ [36].

Reconstruction

Several works use commercial photogrammetry software to reconstruct the imaged scene, with the most popular being the Pix4D and AgiSoft suites. These software suites provide less accurate results, as they are agnostic to the image acquisition strategy employed in this work and do not utilise feature tracking.

Some of the reviewed works perform partial reconstructions, which result in individual point clouds and then merge them using point cloud registration methods. These methods are mainly based on the Iterative Closest Point algorithm [43, 44], either implemented by the authors or provided by software utilities, such as CloudCompare and MeshLab. The disadvantages of merging partial reconstructions are the error of the registration algorithm and the duplication of surface points that are reconstructed by more than one view. These disadvantages impact the accuracy and efficacy of the reconstruction result, respectively.

Scale

Several ways to generate metric reconstructions have been proposed. One way is to place markers at known distances so that when they are reconstructed, they yield estimates of the absolute scale. However, this method is prone to the localisation accuracy of said markers. An improvement to this approach comes from [38], which uses the ratio of the reconstructed objects over their actual size, which is manually measured. However, it requires the careful selection of the reference points to estimate the size of the object both in the real and the reconstructed objects and is, thus, prone to human error. This can be a difficult task for free-form artworks as they may not exhibit well-defined reference points, i.e. as opposed to industrially manufactured objects. Therefore, the accuracy of this approach is dependent on the accuracy of manual measurements of the actual object and its reconstruction.

Another way to achieve metric reconstruction is to include objects of known size in the reconstructed scene, such as printed 2D or 3D markers [38, 36, 41]. This method does not involve human interaction and is not related to the structure of the reconstructed objects. The main disadvantage of this approach is the production of these markers, as printing markers at a micrometre scale is not achievable by off-the-shelf 2D or 3D printers.

Finally, a way to estimate the scale factor is based solely on the reprojection error of the correspondences of a stereo pair [45]. This approach is independent of the shape of the target object and does not require human interaction. However, this approach is formulated for a stereo pair and not for a single camera and is also, is highly dependent on the calibration accuracy of this pair.

2.3 Proposed surface digitisation approaches

A different reconstruction method is used to generate the 2D and 2½D surface digitisation. Surface scanning generates an image mosaic and surface digitisation a 3D textured mesh of triangles. These methods are described below. Their common characteristics, which are pronounced in image acquisition and pointy correspondence steps below, while underscoring the differences per the implementation of each, 2D or 2½D, case.

2.3.1 Hierarchical image acquisition

Section 2.3.1 at a glance

Inputs. A nominally planar/quasi-planar target surface; a user-specified acquisition volume (frustum extent) and number of layers (stand-off distances); a within-layer grid of viewpoints; CNC stage motion commands and reported motor readings/acknowledgements providing approximate camera pose priors.

Outputs. A multi-layer image set with (i) lateral overlap within each layer, (ii) medial overlap across layers, and (iii) per-image coarse pose metadata (from CNC priors). A parent–child organisation of viewpoints across layers (often a quadtree when $\tau_m = 25\%$), enabling coarse-to-fine registration.

Method. Images are acquired at multiple elevations (layers) so that each higher layer covers a larger ground footprint, forming a frustum-shaped acquisition volume. Within each layer, viewpoints are placed on a nominal grid to ensure sufficient lateral overlap between neighbouring images. Across layers, viewpoints are configured hierarchically: each 'parent' view at a higher layer is positioned so that its footprint overlaps (medially) with the union of several 'child' views at the layer below by a prescribed fraction (τ_m). In parallel, CNC-derived pose cues (commanded/reported positions) are recorded and used as coarse priors that constrain matching, reduce the correspondence search space, and stabilise global optimisation by discouraging implausible alignments.

Key parameters.

- **Number of layers and stand-off distances** (upper extent of the frustum; chosen by the user per required global context and workspace).
- **Within-layer grid spacing** (controls lateral overlap and boundary coverage).
- **Lateral overlap target** between neighbouring views within a layer (sufficient overlap for reliable matches).
- **Medial overlap fraction** (τ_m) (parent footprint overlap with child-footprint union; ($\tau_m = 25\%$) naturally yields a quadtree).
- **CNC prior type** (commanded positions, reported positions, acknowledgements) and how it is recorded per image.
- **Expected stage uncertainty** drivers (e.g., stepper resolution, backlash, structural compliance) that bound how 'tight' the pose prior can be used.

Implementation note. The acquisition software issues motion commands in G-code (e.g., 'go to x, y, z'), and records the corresponding motor readings/acknowledgements alongside each captured image; these pose priors are then used downstream to accelerate matching and to regularise global optimisation.

Our integrated hardware and software acquisition method rests on two mechanisms:

1. Multi-layer imaging with medial field-of-view (FoV) overlap. Images are acquired at multiple stand-off distances, creating overlap across depth (medial rather than purely lateral). This enables correspondences to be established between images taken at different elevations.
2. Proprioceptive priors from the CNC stage. Motor readings (and, more generally, commanded and reported positions) provide approximate camera localisation that constrains matching and supports downstream optimisation.

Within each layer, camera viewpoints are arranged on a nominal grid (Figure 5, left), selected to ensure adequate lateral overlap between neighbouring images at layer boundaries. Across layers, the set of viewpoints forms a frustum-shaped volume above the surface: increasing elevation expands the ground footprint of each image, and therefore induces medial overlap between views at different distances.

Crucially, cross-layer viewpoints are configured hierarchically. Each “parent” viewpoint at a higher layer is positioned so that its image footprint overlaps medially with the union of several “child” viewpoints at the lower layer by a prescribed fraction, denoted τ_m . Figure 5 (right) illustrates the construction: viewpoint A covers the region imaged by A1–A4. When $\tau_m = 25\%$, the hierarchy naturally instantiates a quadtree, which provides a convenient organisational structure for coarse-to-fine registration. The number of layers, and therefore the upper extent of the frustum, is chosen by the user according to the required global context and available workspace.

We refer to this cross-layer overlap as medial overlap, and to the resulting cross-layer matches as median correspondences, to distinguish them from the conventional lateral correspondences formed between neighbouring views within the same layer.




Figure 5. Image acquisition across layers.

The second feature is the use of CNC-derived cues to constrain registration. The acquisition software issues motion commands expressed in G-code (e.g., 'go to 50.1, 40.2, 101.7'), and the system records the corresponding motor readings and/or execution acknowledgements. These values are not metrically exact, as stepper resolution, backlash, structural compliance, and other uncertainties limit absolute accuracy, and in any case, the stage is not specified to pixel-level precision. Nevertheless, the reported pose provides a coarse but highly informative prior. In practice, it reduces the correspondence search space, accelerates matching, and stabilises global optimisation by discouraging implausible alignments inconsistent with the commanded motion.

2.3.2 Point correspondence

Section 2.3.2 at a glance

Purpose. Establish reliable point correspondences across the image set so that 2D mosaicking and 2½D reconstruction have stable geometric constraints.

Inputs. Overlapping images acquired within each layer and across layers, together with CNC motion commands and recorded motion cues that provide a coarse pose prior for each view.

What we mean by 'correspondence'. Two points correspond when they are projections of the same physical surface point. In our setting, the surface is static, so both stereo-style correspondences and sequence-style correspondences can be used, but they are interpreted as registration constraints rather than motion evidence.

Outputs.

- **Lateral correspondences** within a layer, formed between neighbouring views.
- **Median correspondences** across layers, used to stabilise registration across scales.
- **Feature tracks** that link the same physical point across three or more images and provide additional geometric redundancy.

Method summary.

1. **Detect keypoints and descriptors** in every image, with SIFT used as a baseline and other detectors allowed.
2. **Constrain candidate search** using a motion-predicted neighbourhood derived from the commanded displacement.
3. **Filter geometrically** by enforcing epipolar consistency with a user-set tolerance, rejecting candidates whose epipolar error exceeds that bound.
4. **Select matches by descriptor similarity**, accelerated using cascade hashing to prune unlikely candidates early.
5. **Enforce mutual consistency** by accepting a match only if it is recovered in both matching directions.
6. **Build feature tracks** to strengthen robustness beyond isolated pairwise matches.



Practical note on overlap. The 2½D stage is substantially more sensitive than 2D mosaicking. In practice, dense overlap is treated as necessary for stable 2½D recovery, whereas 2D mosaics can often tolerate moderate overlap.

Two points in two different images are said to correspond if they are projections of the same physical point on the surface. A putative match, therefore, defines an image correspondence, which may be correct or spurious. Establishing reliable correspondences is a central and actively researched problem in computer vision.

Point correspondences are typically exploited in two canonical settings:

- **Stereo (spatial) correspondences.** When images are taken from different viewpoints, matched points constrain epipolar geometry and can be used to estimate depth. In standard stereo acquisition, the baseline is predominantly horizontal, and correspondences are often discussed as lateral; i.e., across the image plane.
- **Temporal correspondences.** When images are taken at different times, matched points can be interpreted as evidence of motion (optical flow, tracking). Here, overlap is usually near-superimposed between consecutive frames, and correspondences are local in extent; in that sense, they are frequently described as medial in the acquisition sequence rather than as wide-baseline matches.

In our setting, the surface is assumed static, so both correspondence types are available but interpreted differently. We use lateral correspondences within a layer in the usual way, and we also exploit cross-layer medial (median) correspondences as constraints for registration, not for motion estimation.

This distinction matters because 2½D reconstruction is substantially more sensitive than 2D mosaicking. A 2D mosaic can often tolerate moderate overlap, e.g., ~50%, because the objective is a consistent appearance map. By contrast, reliable 2½D reconstruction requires dense, redundant multi-view constraints; in professional photogrammetry, overlaps on the order of 90% are commonly treated as a practical minimum for stable surface recovery.

Key points

Image correspondence is established via keypoint features. We use the Scale-Invariant Feature Transform (SIFT) [48] as a baseline detector–descriptor, but the pipeline is not SIFT-specific: any suitable alternative (e.g., more recent learned features) may be substituted without changing the overall logic. For each image, keypoints are detected and associated with local descriptors that encode appearance in a compact, matchable form.

Stereo correspondences

Correspondences are formed by matching keypoint descriptors across image pairs. Because the relative camera pose is approximately known from the scanner motion, matching can be constrained geometrically, reducing both computational cost and the rate of spurious matches. Descriptor similarity then provides the discriminative signal required to select the correct partner within the reduced search region.

Epipolar constraint

We first restrict candidate matches using a motion-predicted neighbourhood: for each keypoint, potential partners are sought only within a circular region consistent with the commanded displacement of the scanner. Within this region, we enforce the epipolar geometry. Let x and x' denote the homogeneous image coordinates of a putative correspondence in images i and j . A match is accepted only if it satisfies the epipolar constraint $x'^T F_{i,j} x < \tau_c$, where $F_{i,j}$ is the estimated fundamental matrix, and τ_c is a user-defined tolerance. The threshold τ_c absorbs practical imaging factors and is interpreted as an upper bound on admissible geometric error. Operationally, this converts the search from a 2D region to a 1D locus and rejects candidates whose reprojection error exceeds τ_c .

Keypoint similarity

Within the geometrically admissible set, descriptor comparison determines the best matches. To accelerate and stabilise this step, we employ cascade hashing [49], which prunes unlikely candidates early using inexpensive, coarse tests before applying more discriminative comparisons. The method constructs a hierarchy of hash tables: upper levels encode coarse descriptor structure to eliminate most mismatches quickly, while lower levels refine the candidate set using progressively more detailed information. This coarse-to-fine filtering reduces computation and improves robustness under large candidate pools.

Keypoint matching

Correspondence establishment is enforced symmetrically, following the standard mutual consistency principle [50]. A match between keypoints in I_i and I_j is accepted only if it is recovered in both directions: matching $I_i \rightarrow I_j$ and independently $I_j \rightarrow I_i$, often referred to as a left–right check [51]. This bidirectional agreement is a simple but effective filter for eliminating many-to-one matches and reducing the proportion of erroneous correspondences.

Feature tracks

To increase robustness, we do not treat correspondences as isolated matches between image pairs. Instead, we track keypoints across multiple images, forming feature tracks that link the same physical surface point through three or more views. Multi-view tracks provide stronger evidence than binocular matches: they allow geometric consistency to be evaluated repeatedly, and they propagate constraints that help reject spurious correspondences and retain true ones. The required constraints are derived from the fundamental matrices of the image pairs implied by the track.

For each image pair, conformity to the epipolar constraint is quantified using the corresponding fundamental matrix. When a track spans multiple images, the epipolar check is extended across all image pairs induced by the track, from its first to its last image. As in standard binocular filtering, if any projected feature deviates from its expected epipolar-consistent location by more than a tolerance τ_p , the track and, therefore, all matches that constitute it, are rejected.

Tracks are not restricted to laterally neighbouring images within a single acquisition layer. They are also formed across layers, which is essential to global consistency: distant views constrain the admissible spatial arrangement of features in close-up views, thereby limiting drift and reducing the probability that local errors propagate into the final reconstruction.

Intuitively, consider a keypoint in one image that is matched to keypoints in two other images (see Figure 6). These three images become linked by two correspondence pairs, which together define a minimal track. In general, a feature track contains three or more matched observations of the same putative surface point.




Figure 6. Three images linked through the same keypoint in the middle one.

Formally, let a track contain N matched keypoints observed in N images, with image locations u_i for $i \in [1, N]$. The set $\{u_i\}$ defines a feature track. If the track contains spurious correspondences, then at least one u_i will not be the projection of the same physical point; no particular ordering of track elements is assumed.

We evaluate each track via a composite epipolar-consistency test involving all track correspondences. Let $F_{1, N'} = F_{1,2} \cdot F_{2,3} \cdot \dots \cdot F_{N-1, N}$ be a composite fundamental mapping between the first and last observations in the track. Note that images I_1 and I_N need not be neighbours; even if they are, we prefer $F_{1, N'}$ over a directly estimated $F_{1, N}$ because the composite construction forces the entire chain of correspondences to be jointly consistent. This increases sensitivity to a single erroneous link within the track.

Using $F_{1, N'}$, we project the first observation u_1 into image I_N as $u_p = F_{1, N'} \cdot u_1$; the distance of these two points is $\delta = |u_p - u_N|$. If the track is correct, δ should be small, limited primarily by calibration error and residual pose uncertainty. If the track contains an incorrect correspondence, the inconsistency typically becomes large, yielding $\delta > \tau_p$. We therefore discard the entire track when δ exceeds the threshold τ_p .

The memory footprint of subsequent stages scales with the number of retained tracks and can become substantial for large image sets. When memory pressure is a concern, tracks may be downsampled at this stage. A pragmatic strategy is to discard the shortest tracks first, because they provide the weakest multi-view constraints and are usually the most numerous.

2.3.3 2D surface reconstruction

Section 2.3.3 at a glance

Purpose. Register all acquired images into a common mosaic coordinate frame and reconstruct a globally consistent 2D surface map, while using scanner pose priors to regularise the solution.

Inputs. Proprioceptive estimates of camera centres from the scanner motion plan, and established point correspondences between laterally and medially adjacent images across all pyramid layers.

Outputs. A projective homography H_i for each image I_i mapping image coordinates to mosaic coordinates, and the final 2D mosaic, optionally exported as tiles for scalable viewing.

Method summary.

1. **Define the map.** Represent the imaged surface in pixel coordinates in the mosaic frame.
2. **Initial homography estimation.** Estimate homographies by minimising a joint objective that enforces agreement of matched features across neighbouring images and promotes compliance with scanner-derived pose priors.
3. **Exploit pyramid adjacency.** Use the acquisition hierarchy to form a topological graph whose nodes are scan-plan locations and whose edges encode adjacency, constraining the optimisation search space.
4. **Graph-based least-squares optimisation.** Solve the resulting graph-structured least-squares problem using a SLAM or bundle-adjustment style framework, which also tolerates missing or discarded homography estimates.
5. **Global optimisation and outlier rejection.** For each neighbouring pair, estimate the fundamental matrix with RANSAC and least squares to reject outlier correspondences; keep inliers to refine geometric consistency.
6. **Render the mosaic.** Store the mosaic as a single large image, or export tiles and use a tile-based viewer to support interactive exploration.

Practical note. Tile-based export avoids loading a massive mosaic at once and supports smooth navigation at high resolution.

We call a map the imaged surface in pixel coordinates, in the coordinate frame of the mosaic to be created. The input to image registration is the proprioceptive estimates of the camera centres and the established point correspondences across laterally or medially adjacent images. The output is a set of projective homography transforms H_i , estimated for each image I_i , where i enumerates the images across all pyramid layers. These homographies associate image locations in each I_i with the corresponding locations in the mosaic.

Homography estimation

World points C_i are the proprioceptively obtained coordinates for these locations in 3D space. Image points c_i are the image centres of images I_i . Initially, projective homography H_g is estimated across this map and the 3D grid locations C_i , using least-squares. For each pair of adjacent images I_i and I_j , we enumerate the correspondences between them using k and denote their locations in I_i and I_j as f_{ki} and f_{kj} , respectively. The computation estimates the homographies by optimising the following objective function

$$\sum_i \sum_j (H_i \cdot f_{ki} - H_j \cdot f_{kj})^2 + \sum_i (H_g \cdot C_i - H_i \cdot c_i)^2. \quad (1)$$

The first term is the conventional reprojection error metric for point correspondences. In that term, j enumerates the neighbours of i . The second term promotes compliance with the scanner coordinates. In Figure 7, the notation is illustrated.




Figure 7. Illustration of objective function notation.

The projective homography has 8 free variables and, thus, the optimised variables are 8 times the number of images. The optimisation capitalises on the adjacency information contained in the pyramid data structure to create a topological graph, as the one in Figure 1, right. This graph has points C_i as nodes and, as vertices, their adjacency relations. These relations constrain the search space of the optimisation. We employed the work in [53], which is a framework for least-squares optimisation of an error function that can be represented by a graph and has been specifically designed for Simultaneous Localisation and Mapping (SLAM) or bundle adjustment problems.

A benefit of using the aforementioned graph-based method is the robustness of 'missing estimates'. Such a case is encountered when we discard unreliable homography estimates. Another is when the apex of the pyramid cannot be reached by the hardware. The latter case is encountered when covering wider areas, using multiple, laterally overlapping pyramids.

Global optimisation

Once candidate correspondences have been established, a global optimisation step discards potentially spurious and inaccurate correspondences and refines camera pose estimates, as follows. The pairs of neighbouring images, let l_i and l_j , within and across layers, are found in the scan plan. The fundamental matrices, F , for all pairs of neighbouring images are computed. Matrix $F_{i,j}$ denotes the fundamental matrix for image pair (i, j) . Matrix $F_{i,j}$ defines the geometric constraints that exist between points in image i and their corresponding epipolar lines in image j . Matrices $F_{i,j}$ are estimated from the available correspondences between image pairs (i, j) . The estimate of $F_{i,j}$ is obtained using Random Sample Consensus (RANSAC) [52] and least squares fitting as follows to reject spurious correspondences, or 'outliers', between images i and j . The remaining, 'inlier', correspondences are used to estimate $F_{i,j}$, using least squares.

Output

For surface scanning, the result is an image mosaic. This is typically stored in a single image. As this image can be very large in dimensions, tile-based rendering is supported by exporting the result in smaller images

('tiles'). An appropriate tile-based viewer dynamically loads higher-resolution tiles for the specific area viewed, allowing for smooth and detailed exploration without loading a massive image.

2.3.4 2½D surface reconstruction

Section 2.3.4 at a glance

Purpose. Reconstruct a metrically scaled 2½D surface as a textured triangle mesh, using large-scale multi-view geometry while controlling drift and texture discontinuities.

Inputs. Camera calibration, the scan plan with approximate camera poses in metric units, pairwise geometry-consistent correspondences and multi-view feature tracks, and cross-layer links that provide long-range constraints.

Outputs. A sparse point cloud and refined camera poses, followed by a dense point cloud, a refined triangle mesh, and a textured mesh stored as a binary PLY file together with a texture image and per-vertex texture coordinates.

Method summary.

1. **Prune the view set.** Build a connectivity graph from feature tracks and keep only the largest connected component; discard the rest to avoid unconstrained fragments and reduce memory demands.
2. **Choose an initial pair.** Select a starting image pair that balances wide baseline with strong feature support, using a baseline-weighted correspondence score.
3. **Sparse reconstruction to refine poses.** Run an incremental Structure from Motion pipeline to estimate an initial sparse point cloud and improve camera extrinsics, using robust estimation to reject outliers.
4. **Bundle adjustment with constrained variables.** Refine primarily lens distortion and extrinsics, using scan-plan estimates to ease convergence and reduce instability in very large image sets.
5. **Preserve metric scale.** Use the fact that scan-plan camera locations are already in metric units, so refined extrinsics remain metrically scaled without an additional external scale object.
6. **Dense reconstruction.** Compute per-view depth maps, fuse them into a dense point cloud, generate a mesh surface, and apply mesh refinement.
7. **Visibility-aware texturing.** Compute triangle index maps, combine multiple views to reduce seams and discontinuities, and pack the result into a texture atlas suitable for very large numbers of images.

Practical note. With tens of thousands of views, small registration errors can accumulate into global distortions and texture seams; cross-layer constraints and multi-view texturing are used to reduce these artefacts.

Implementation note. The sparse stage uses an incremental SfM pipeline with robust estimation and bundle adjustment, and the dense stage uses a multi-view stereo library to produce depth maps, fused geometry, meshing, and multi-view texturing.



The surface is reconstructed as a textured mesh of triangles. This mesh is comprised of two lists. The first is a list of 3D, floating-point locations that represent the mesh nodes. The second is a list of integer triplets that contain indices to the first list and indicate the formation of triangles through the represented nodes. A texture image accompanies the mesh, along with a third list of 2D coordinates that represent the texture coordinates of each node; thus, the third list has the same length as the first one.

Initialisation

Using the obtained feature tracks as a connectivity relation, a connected component labelling of the input images is performed. The largest connected component is selected, and the remaining images are discarded. These components are discarded because they correspond to groups of images that are not linked through fundamental matrices to the rest. As such, they cannot contribute to the main reconstruction. They are, thus, discarded to reduce memory capacity requirements.

The reconstruction method requires an image pair, let (i', j') , as the basis for the surface reconstruction. This pair is selected to exhibit a wide baseline to reduce reconstruction uncertainty and error [55]. The reason is that wider baselines result in more separation between the optical rays, reducing ambiguity in the 3D reconstruction. At the same time, the reliability of this pair depends on the number of feature correspondences existing in this pair. Thus, the initial pair is selected as the product of the baseline with the number of correspondences, or $(i', j') = \operatorname{argmax}_{i,j} (b_{i,j} \cdot k_{i,j})$ where $b_{i,j}$ is the baseline of pair (i, j) and $k_{i,j}$ is the number of feature correspondences between I_i and I_j .

Sparse, point-based reconstruction

A sparse reconstruction is first performed, using the camera poses in the scan plan. The purpose of this reconstruction is to refine these pose estimates. The Open Multiple View Geometry [56] is utilised to obtain a sparse point cloud, which includes implementations of the Incremental Structure from Motion pipeline [57] and 'A Contrario' RANSAC [58].

The large number of images, i.e. tens of thousands, required to cover wide surfaces in detail, poses accuracy problems that are not pronounced when reconstruction involves a few hundred images. In the context of thousands of images, even small registration errors may accumulate, leading to globally inconsistent surface structures. As correspondences include pairs of points across layers, the reconstruction process is guided to produce a structure that is consistent with far-range views. In the experiments, it is observed that these additional constraints reduce global distortion errors in the final result.

A bundle adjustment, using [59], is then performed in the end. This operation is adapted to optimise only the lens distortion and the extrinsic parameters for each image. The reason is to ease the convergence of the bundle adjustment optimisation.

The method is formulated in Algorithm 1.

Algorithm 1 Incremental Structure from Motion

- **Require:** internal camera calibration (matrix K)

- **Require:** pairwise geometry consistent point correspondences
- **Ensure:** 3D point cloud
- **Ensure:** the camera poses
 - compute correspondence tracks t
 - compute connectivity graph G (1 node per view, 1 edge when enough matches)
 - pick an edge e in G with a sufficient baseline (compare F and H)
 - robustly estimate the essential matrix from images of e (AC-RANSAC)
 - triangulate $t \cap e$, which provides an initial reconstruction
 - contract edge e
- **while** G contains an edge, **do**
 - pick edge e in G that maximizes $\text{track}(e) \cap \{3D \text{ points}\}$
 - robustly estimate pose (external orientation/resection) (AC-RANSAC)
 - triangulate new tracks
 - contract edge e
 - perform bundle adjustment (uses our initial estimation from the scan plan)
- **end while**

Finally, the scale factor is estimated at this step. In contrast to the methods in Section 2.2, this work estimates scale factors using the extrinsic parameters of the camera. This accuracy is relatively high because it is founded on the accurate motorisation of CNC devices. The initial estimates of camera locations obtained from the scan plan are already in metric units and, thus, so are the refined ones.

Dense, textured mesh reconstruction

A mesh of triangles is computed based on the obtained sparse point cloud. In the implementation, a Multi-View Stereo library, specifically the OpenMVS [60], is utilised.

The 'Z-buffering technique' [61] is employed to obtain depth maps, D_i . These maps are images that have the same dimensions as I_i , imaging the scene from the same viewpoint as I_i and with the same intrinsic parameters. In $D_i(u)$, each depth map stores the distance of the surface point imaged at $I_i(u)$ from the optical centre where I_i was acquired. Given the camera extrinsic parameters, the pixels of the depth map are converted into a point cloud in world coordinates. This way, depth maps D_i are aggregated in a dense point cloud, using [62]. Utilising the work in [63] a mesh surface that interpolates the dense point cloud is generated. Afterwards, this mesh is refined via the variational method in [64].

Next, index maps T_i are computed, which store at each pixel the id of the mesh triangle, imaged at that pixel, that is in $T_i(u)$, and encoded is the id of the triangle that is imaged at u . Using maps D_i and T_i , texturing the mesh considers the visibility of the triangle in each I_i . When a triangle is imaged in multiple views, then several choices can be made as to which view to select to acquire the texture or how to combine these multiple images of the same surface regions into a better texture [65, 66].

Despite the accuracy improvements, the mesh and the camera pose still contain residual errors. Albeit these errors are relatively small, they are well noticed by the human visual system as texture discontinuities. The phenomenon is particularly pronounced when a large number of images is utilised.

To address these inaccuracies, multiple views are combined using [67], which is a method designed for large numbers of images. The generated texture is efficiently packed in a 'texture atlas' as in [68].

Output

In surface digitisation, the resultant reconstruction is a textured mesh of triangles. This mesh consists of three lists and one texture image. The first is a list of 3D, floating-point locations that represent the mesh nodes. The second is a list of integer triplets that contain indices to the first list and indicate the formation of triangles through the represented nodes. The third is a list of floating-point 2D coordinates in the texture image, one for each node. This mesh is stored in the Polygon File Format (PLY) format, in two files. The first file contains the mesh representation and uses the binary version of the format to save disk space. The second file contains the texture in either the Joint Photographic Experts Group (JPEG) or Portable Network Graphic (PNG) image file format.

2.3.5 Evaluation

Section 2.3.5 at a glance

Purpose. Evaluate the accuracy and practical limitations of the 2D mosaics and the 2½D reconstructions produced by the pipeline.

Inputs.

- For 2D digitisation: legacy evaluation results from Mingei, with additional evidence referenced in Annex B.1.
- For 2½D digitisation: Euro coin scans, calliper measurements used as ground truth for dimensions, and an external high-accuracy scan used as ground truth for surface-structure comparison.

Outputs. A compact set of quantitative and qualitative indicators, including registration error characterisation, planarity limits, dimensional error statistics, circularity and isotropy indicators from orthographic depth projections, and surface-structure similarity metrics, supported by Figures 8–12 and Tables 2–3.

Method summary.

1. **2D digitisation.** Report mosaic registration accuracy and inspect mosaics for distortions due to departures from the planarity assumption; note photometric and focus-stacking limitations that affect feature availability and appearance consistency.
2. **2½D digitisation: dimensions.** Compare reconstructed diameters and thicknesses against ground truth for a controlled set of reference coins, and report measurement error.
3. **2½D digitisation: global distortions.** Generate orthographic depth projections, extract edges, fit a circle model, and report deviations from circularity and an aspect-ratio-based isotropy indicator.

4. **2½D digitisation: surface structure.** Compare depth maps against an external ground-truth scan using cross-correlation and edge-based dissimilarity, supported by qualitative inspection.

Practical note. There is no established public benchmark for close-range photogrammetry at the micrometre scale, so the experiments use standardised coins and an external scan to provide interpretable reference points.

Evaluation of 2D digitisation

Detailed 2D evaluation is reported in Annex B.1; key figures are summarised below.

The quantitative experiments show that registration errors are on the order of 10 px in mosaics comprising ≈ 4 Tpx. The proposed approach employs auxiliary images to strengthen image registration cues and fuses proprioceptive data to produce mosaics of the scanned surface with a resolution of 19,800 ppi.

The obtained mosaics were inspected for distortions due to departures from the planarity assumption. We found the limitation of the current configuration to be sharp steps of over 3 mm.

We performed no correction for the global optimisation of image intensities. Setting the camera acquisition mode to automatic brightness adjustment adapts the dynamic range of image acquisition to the content of each image. This can be observed in the mosaics of the top layers, where brightness differences across surfaces of the same luminance are observed. Compensation methods tailored for mosaics exist in the literature, e.g. [80, 81].

We did not control the stacking process provided by the sensor. By assigning this control to the embedded system accompanying the sensor, we may be wasting potential sensitivity to depth variations. Control of bracketing techniques would provide better focus and, thus, more image features. In addition, it can be supported even by weak depth cues, such as depth from defocus [82] or stereo vision.

Evaluation of 2½D digitisation

To measure the accuracy of reconstruction, targets of known size and structural features are utilised. To the best of our knowledge, there exists no public benchmark for close-range photogrammetry at the micrometre scale. Therefore, we used coins as reference targets to render this work comparable.

Dimensions

The purpose of the experiment was to measure reconstructed dimensions and compare them with the ground truth.

State-manufactured coins are of standard size and accurately manufactured to avoid counterfeiting. The scanned coins belong to the euro currency. All eight coins from this family were scanned on a planar and textured surface. The nominal dimensions provided by the manufacturer (European Central Bank) were confirmed for the particular coins with an electronic calliper. No larger disparity over 0.01 mm was found in these measurements. These dimensions were considered ground truth, as the coins were used and may

have suffered distortions. The digital reconstructions are in metric units, and their dimensions were compared to the ground truth. Their discrepancy is the measurement error.

Original images are shown in Figure 8. The top two rows show images from the highest elevation layer, and the other two images from the lowest layer. The corresponding reconstructions are shown, in the same order, in Figure 9 and in Figure 10.



Figure 8. Original images of circular coins from the closest layer to the target.






Figure 9. Untextured reconstructions of circular euro coins.



Figure 10. Textured reconstructions of circular euro coins.

The dimensions and errors are reported in Table 2. The first column notes the currency value. Column 'Nominal' reports the dimensions provided by the manufacturer (European Central Bank), column 'Measured' reports our calliper measurements (ground truth), column 'Measurement' reports the dimensions of the digital model, and 'Error' is the percentage error. The thickness of the coins was measured as the distance of the supporting plane to the top face of the reconstruction. The average measurement error is $\approx 2.91\%$. In Tables 2–3, 'c' denotes euro cents.

Table 2. Coin dimensions (diameter \times thickness): nominal, measured, reconstruction, and errors.

Currency	Nominal (mm)	Measured (mm)	Reconstruction (mm)	Error (%) (D \times T)
1c	16.25 \times 1.67	16.25 \times 1.66	16.51 \times 1.70	1.58 \times 2.19
2c	18.75 \times 1.67	18.76 \times 1.66	19.45 \times 1.68	3.69 \times 1.50
5c	21.25 \times 1.67	21.26 \times 1.67	20.21 \times 1.73	4.95 \times 3.58
10c	19.75 \times 1.93	19.76 \times 1.91	19.06 \times 1.96	3.56 \times 2.87
20c	22.25 \times 2.14	22.27 \times 2.14	21.51 \times 2.23	3.40 \times 4.00
50c	24.25 \times 2.38	24.27 \times 2.38	25.41 \times 2.26	4.70 \times 4.97
€1	23.25 \times 2.33	23.32 \times 2.34	22.90 \times 2.31	1.80 \times 1.43
€2	25.75 \times 2.20	25.75 \times 2.19	26.06 \times 2.16	1.19 \times 1.24

Aspect ratio

The purpose of this experiment is to measure any deviation from the circular shape of coin edges to assess global distortions in the reconstruction. To assess global distortions in the reconstruction, the orthoimages of the coin reconstructions were generated. Orthoimages are perpendicular projections of

the reconstruction upon a hypothetical plane parallel to the surface. These images are 'map accurate' in that they do not contain perspective distortions included in the original photographs. Depth maps computed for orthophotos share the same property.

In Figure 11, we performed edge detection on the depth orthophotos. We then superimposed the detected depth edges on the inlier edge set. The average error as a deviation from perfect circularity is ≈ 0.58 px.



Figure 11. Depth edges superimposed on the textured reconstructions of euro coins.

In Table 3, deviations of the detected edges from the fitted circle are reported, as the mean distance of these edges from the fitted circle and their standard deviation. The small deviations and the visual results indicate that the circles were appropriately fitted. The average error is ≈ 0.58 px.

Given the representativeness of the circle, the aspect ratio of the detected points is computed to assess whether the reconstruction is isotropic. To compute this, the leftmost p_1 , rightmost p_2 , top p_3 , and bottom p_4 points of the inlier edge set were found. Then, aspect ratio $|p_2 - p_1|/|p_4 - p_3|$ is an indicator of anisotropy over the horizontal and vertical surface dimensions. The last column of Table 3 reports these ratios, indicating a mean aspect ratio of 0.998, corresponding to a 0.0018% deviation from isotropy.

Table 3. Radius, mean circle fit error, and standard deviation, for each measured coin.

Currency	Radius (px)	Error, std (px)	Aspect ratio
1c	224.58	0.49 (0.38)	0.995
2c	243.45	0.69 (0.41)	1.000
5c	232.66	0.48 (0.35)	1.000
10c	234.63	0.67 (0.41)	1.008
20c	238.50	0.66 (0.43)	0.985
50c	248.21	0.51 (0.37)	0.997
€1	231.52	0.71 (0.40)	0.997

€2	237.88	0.41 (0.33)	1.000
----	--------	-------------	-------

Surface structure

The purpose of this experiment is to assess the accuracy of the reconstruction of surface structure. We used the digitisation of a coin by a scanner that is more accurate than photogrammetry and used that digitisation as ground truth. By comparing this higher accuracy model, we obtain a measure of the accuracy of our method. To quantify the error of the proposed method, we used cross-correlation to measure the similarity between two depth maps.

A higher-quality scan of the 2€ coin is available from the TetraVision company. The authors of the dataset used an elaborate scanning technique that involved binocular (stereo) imaging and structured light, using the 'Atos III Triple Scan' 3D scanner manufactured by the GOM company. This dataset is available in [79].

The depth maps of the three scans were produced, and the top surface region of the coin was isolated to compare the same surface regions. In these maps, the following measurements were acquired. First, we computed the cross-correlation between image pairs (S1, S3) and (S2, S3), which were 99.84% and 86.43%, respectively. Second, we performed Canny edge detection [77] on all three maps, using the same parameters. The spatial arrangements of the obtained edges were compared using the Hausdorff distance for the same image pairs. The results were 73.53p and 110.63p for image pairs (S1, S3) and (S2, S3), respectively. In Figure 12, the textureless reconstructions and the edge detection results are shown.




Figure 12. Top: surface reconstructions of a 2E coin. Bottom: edge detections on the depth maps of the reconstructions. Left: proposed. Right: ground truth.

Qualitatively, the comparison of S1 with S2 offers similar observations: S2 exhibits significant levels of noise. This is also observed by the structure of edges in Figure 12 (bottom, middle). Quantitatively, both the correlation and the Hausdorff measures of dissimilarity indicate the greater accuracy of the proposed method, as it provides more similar results to S3.

2.3.6 Figure-ground segmentation

Section 2.3.6 at a glance

Purpose. Automatically isolate the textile from the floor in the wide-area scanning setup, so that downstream processing operates on the object region without requiring a controlled chroma background.

Inputs. A static overhead camera view, a pre-acquired background reference image of the empty floor, and a live foreground image captured after placing the textile, under stable illumination.

Outputs. A binary mask M that labels each pixel as textile or floor, suitable for cropping, masking, and subsequent wide-area scanning steps.

Method summary.

1. **Acquire a robust background model.** Capture a short burst of empty-floor images and compute a pixelwise median to suppress noise and transient fluctuations, yielding the background reference.
2. **Capture the live frame.** Acquire a foreground image after placing the textile.
3. **Compute a difference image.** Subtract background from foreground using absolute difference over intensity or colour channels.
4. **Threshold to obtain a mask.** Apply a threshold to produce a binary mask, using adaptive local thresholding to compensate for slow illumination gradients.
5. **Use the mask operationally.** Treat mask-true pixels as textile and mask-false pixels as floor.

Practical note. Background subtraction is only reliable under stable imaging conditions; the protocol therefore uses ambient darkness and a dedicated LED ring flash integrated with the overhead camera to reduce time-varying illumination and improve repeatability.

The wide-area scanning architecture uses a fixed, overhead vision system observing a floor area that serves as the scanning platform. To increase processing automation, the textile must be isolated from the underlying surface, i.e., a figure-ground segmentation problem. To keep the scanner broadly deployable, we avoid requiring a controlled chroma background (e.g., a ‘green screen’), which makes simple colour-thresholding unreliable for boundary detection.

We frame segmentation as the estimation of a binary mask by comparing a live scan image I_F (textile on the floor) with a pre-acquired reference image I_B (empty floor). Because the camera is static and viewing conditions are designed to be stable, we apply static background subtraction to obtain a clean separation between the foreground textile (figure) and the background environment (ground). The procedure comprises three main steps.

Background image acquisition

We first acquire a reference image I_B that models the appearance of the scanning surface in its empty state, including texture and illumination. Before placing any textile, the overhead system captures a high-resolution image of the floor. In practice, rather than relying on a single background frame, we record a short burst of background images $\{I_{B1}, I_{B2}, \dots\}$ and define the reference background as their pixelwise median. This suppresses sensor noise and transient illumination fluctuations, yielding a more stable background model.

We then capture the scene containing the object of interest. The textile is placed on the floor, and a second image, I_F , is acquired. This frame contains both the unchanged background and the textile region superimposed on it.

Segmentation

Segmentation is performed by subtracting the background model from the live frame. We compute a difference image $I_D = |I_F(u, v) - I_B(u, v)|$, where (u, v) are the pixel coordinates and the subtraction is applied to intensity or colour channels.

A threshold T is then applied to obtain a binary mask M . Pixels in I_D with a value greater than T correspond to the textile, while values less than T correspond to the floor. The result, M , is true where the segmentation result is $I_D(u, v)$ greater than T , and false otherwise. Mask M delineates the textile from the floor.

To address slow spatial illumination gradients, we employ adaptive (local) thresholding. The image is processed into blocks. Threshold T_i is computed per block, based on local intensity variance. This adaptation improves robustness when residual shading varies across the field of view.

Practicalities

Background subtraction is only as stable as the imaging conditions that support it. We therefore enforce a controlled acquisition protocol. Scans are captured under ambient darkness, with illumination provided exclusively by a dedicated LED ring flash integrated with the overhead camera.

This configuration offers two practical benefits. First, switching off room lights eliminates time-varying ambient contributions, such as varying sunlight, fluorescent flicker, or dynamic shadows and reflections from activity in the space. Second, the ring flash provides a repeatable, spatially consistent illumination profile for each capture.

As a result, the photometric properties of I_B and I_F remain highly correlated, which improves the reliability of subtraction in the presence of small residual variations that would otherwise degrade segmentation.

2.4 Wide-area use cases

Section 2.4 at a glance

Purpose. Demonstrate how the surface digitisation pipeline is deployed at a wide scale, covering both ultra-high-resolution 2D mosaics and wide-area 2½D reconstructions.

Inputs. The wide-area scanning setup, scan plans for large floor coverage, and the figure-ground segmentation mask when batch scanning multiple textiles.

Outputs.

- **Wide-area 2D mosaics** in two operational modes: high-definition single-artefact scans and batch scans of multiple smaller items.
- **Wide-area 2½D reconstructions** of a large metallic printing plate, presented as orthophotos and 3D views with textured and untextured renderings.
- **Evidence and access points** via galleries and datasets in Annex C, and linked interactive or video-based views.

Method summary.

1. Configure the optical setup and scan plan for either maximum resolving power or maximum throughput.
2. Acquire images and reconstruct the mosaic, using segmentation to isolate items in batch mode.
3. For 2½D, reconstruct geometry and produce orthographic and viewpoint-based visualisations.
4. Curate results as galleries and datasets for inspection at multiple scales.

Practical note. The two 2D modes share the same mechanical and computational framework but deliberately trade resolving power against acquisition time, data volume, and operational throughput.

2.4.1 Wide-area 2D scans

The wide-area version of the 2D scanning method was deployed in two distinct operation modes. These modes were developed to address complementary digitisation needs that arise in the documentation of textile surfaces. The first mode prioritises ultra-high-resolution imaging of individual artefacts, maximising spatial fidelity. The second mode optimises efficiency by scanning multiple smaller items simultaneously, thus reducing data acquisition time. Both configurations share the same mechanical and computational framework but differ in optical properties, imaging distances, and scan plans.

The distinction between the two modes is not a minor adjustment of parameters but a deliberate reconfiguration of the optical system. Mode A prioritises resolving power through a narrow field of view, minimal imaging distance, and high numerical aperture. These settings maximise detail but require careful focusing and stable illumination. Mode B is designed for optimal coverage and operational efficiency, featuring a wide field of view, extended depth of field, and rapid exposure settings.

The choice of lens and imaging distance determines not only the resolution but also the achievable signal-to-noise ratio and illumination uniformity. Each configuration represents an equilibrium between optical performance, scanning area, and temporal cost. The ability to switch between them underlines the adaptability of the Craeft scanner as a research and documentation tool.

Mode A. Scanning individual articles

This operational mode centres on generating extremely high-resolution mosaics of individual textile artefacts. The aim is to achieve maximal surface definition, capturing fine morphological detail such as fibre structure, weave regularity, and micro-relief in the yarn topology. For this purpose, we employ long-focal macro-optics and operate the camera at close range, ensuring a shallow depth of field and dense lateral image overlap. The overlap is critical for accurate registration and seamless reconstruction, especially in the presence of local texture periodicity. While this configuration achieves micrometre-scale sampling, it is computationally and temporally intensive. Each scan cycle demands extended acquisition time and larger data volumes, making this mode best suited for unique or research-priority artefacts where the preservation of minute visual and geometric features outweighs considerations of throughput.

Large textiles $3 \times 4 \text{ m}^2$ are scanned in one piece, and at a resolution of $12,000 \times 9,000 \text{ px}$ (108 MP). A geometrically accurate scan enables the viewing of the entire textile at resolutions that allow the study of individual fabric knots (see Figure 13).





Figure 13. Traditional textile high-resolution scan. Video version: <https://youtu.be/NdNW5znwkSw>

In this way, we digitised several wide textiles in the context of the Cretan Textiles RCI-8. The table row below illustrates their representation in our knowledge base. This representation includes the link to the Zenodo repository that hosts the digitisation file, a 'Media Object' entry that represents the digital asset and contains its metadata, as well as a representation of the craft artefact and an image thumbnail for rapid visualisation. The hyperlinks below lead to the Craeft online knowledge base.

Storage	Asset	Product	Resolution	Thumb
set3.png	Media Object	BH: 015	27171 × 12582	

The entire collection is provided in Annex C.1. Each element of this collection is presented in an individual video that guides the observer through the details of each textile article. A playlist of these videos can be found online, on Craeft’s YouTube channel.²

Mode B. Batch scanning

The second operational mode emphasises throughput over absolute resolution. This method facilitates the routine documentation of textile samples. Here, the same scanning platform is used to digitise multiple smaller textile samples arranged spatially in a batch layout. A wider field of view is achieved by increasing the imaging distance and employing shorter focal-length optics. This allows multiple items to be scanned during a single XY traversal. The overlap between adjacent views is reduced compared to the high-resolution mode, reflecting the lower requirement for sub-pixel alignment precision. The reduced number of frames per artefact translates into faster acquisition cycles and more efficient use of computational resources, without loss of perceptual fidelity for general documentation purposes.

The figure-ground segmentation algorithm has a larger role, facilitating the automation of individual textile segmentation in the scans. The example below illustrates this operational mode and the utility of the segmentation algorithm. Figure 14 illustrates the method on a large-scale mosaic that images multiple woven articles.

² https://youtube.com/playlist?list=PLQZAktGyYNimFKU7_g6p7GhUOyZLdn-Zh



Figure 14. Figure ground segmentation on an image mosaic.

Figure 15 shows the original digitisation of the previous example and previews its interactive visualisation zooming into the detail of a particular textile article.



Figure 15. Video version: https://youtu.be/nkvwJ_HA_U4

A collection of smaller textiles was digitised using this approach. The scans are provided in Annex C.2, tabulated in individual textiles as background-subtracted from the images. As in the previous case, they are accompanied by links that lead to their documentation in the Craeft online knowledge base. As in the previous case, the results are presented in the form of videos that overview each scan, at multiple scales. In all videos, the scan is presented orthographically (as a map), and the camera zooms in to reveal details.

2.4.2 Wide-area 2½D scans

The method is demonstrated for a stereotype printing plate (also known as a cliché), a metallic printing form used in the production of the Athenian daily newspaper 'Εστία'. It is a tangible and significant artefact documenting the history of Greek printing technology and political journalism. This plate derives its historical importance from its connection to 'Εστία' and the archives of the Hellenic Parliament: 'Εστία' (Estia), first established as a daily newspaper in 1894. This plate documents a key moment in the history of Greek media production. The stereotyping process allowed the newspaper to transition from slow, flat-bed printing to the faster, high-volume rotary printing required for mass circulation.

The printing plate is a durable, single-piece metallic form used to transfer ink onto paper during a high-speed press run. It replaced the original, fragile assembly of individual movable type. The plate is cast from Type Metal, a specific alloy predominantly consisting of Lead (Pb), alloyed with Antimony (Sb) for hardness and expansion control, and Tin (Sn) for fluidity during casting. The surface features a raised relief structure, which is a mirror image of the printed text and imagery. This relief structure holds the ink used to print the page.

Figure 16 shows part of an original image and the mosaic (orthophoto) of the printing plate.



Figure 16. Part of an original image (left) and frontoparallel view of the obtained mosaic (right).

To evaluate wide-area performance, it is useful to examine both the assembled coverage and the local fidelity of the resulting surface. In particular, wide-area fusion can introduce artefacts at tile boundaries (e.g., small geometric discontinuities or texture misalignments) even when individual tiles reconstruct correctly. Figure 17 shows the reconstructed geometry from a macroscopic view. The rest of its panels show a reconstruction detail as a mesh of triangles, as a shaded surface mesh, and as a textured mesh, to highlight the structure of local surface relief and texture detail in the result.






Figure 17. A macroscopic (top-left) and a close-up view from a wide-area 2½D reconstruction. The close-up view is shown in three alternative renderings: mesh of triangles (top-right), as a shaded surface mesh (bottom-left), and as a textured mesh (bottom-right).

2.5 Fine-detail use cases

Section 2.5 at a glance

Purpose. Present fine-detail results that stress-test the pipeline at close range, covering both legacy 2D mosaics and Craeft 2½D reconstructions on challenging shiny metallic objects.

Inputs.

- For 2D: legacy fine-scale scans and associated media evidence (Annex B.1 and linked videos).
- For 2½D: close-range images acquired with the same hardware as 2D, using the Craeft 2½D pipeline, with calliper measurements as ground truth for dimensional evaluation.

Outputs.

- Fine-scale 2D mosaics illustrating close-range detail capture, with representative examples in Figure 18 and additional evidence in Annex B.1.
- Fine-detail 2½D reconstructions of metallic nuts and screws, with untextured and textured renderings (Figure 21, Figure 22, and Figure 23), dimensional error statistics (Table 4), and a comparative failure case using Pix4D (Figure 24).

Method summary.

1. **2D legacy evidence.** Provide representative examples and link to galleries and video walkthroughs to illustrate fine-scale mosaic quality and registration across depth layers.
2. **2½D on shiny metals.** Acquire close-range imagery where limited texture cues remain available, reconstruct geometry and texture, and assess artefacts in high-curvature specular regions.
3. **Quantitative checking.** Compare reconstructed dimensions against the calliper ground truth and report errors in a consistent order.

4. **Baseline comparison.** Reconstruct the same scene with a conventional photogrammetry tool and compare sparsity and failure modes.


Practical note. Highly reflective surfaces violate the uniqueness constraint across views, which limits correspondence reliability; close-range imaging and the proposed constraints can still recover structure, with artefacts concentrated in specular high-curvature regions.

2.5.1 2D results

The results of the fine-scale 2D scanning mode are legacy and have been presented in Mingei. A library of such legacy results from fine-scale scans demonstrating close-range scanning is provided in Annex B.1. A video playlist of these results is available on Craeft's YouTube channel at: <https://www.youtube.com/playlist?list=PLQZaktGyYNikhZkYeCgZ23TmeW5Rufck>.

Two indicative examples are shown in Figure 18, which demonstrates the scanning of a silk textile and a banknote. The figure is accompanied by a video animation showing the entire scans as well as the registration accuracy across depth layers.








Figure 18. Mosaics of a silk textile (left column) and a banknote (right column). The top row overviews the physical item. The middle row shows 4×4 mosaics at the second closest imaging layer; these mosaics are 6052×4585 px. The bottom row shows segments of the original images at the closest layer. Video version <https://youtu.be/3rhcAUcekkM>

2.5.2 2½D results

In Craeft, we used the same hardware as for 2D scanning with the new algorithm approach for 2½D surface digitisation. The experiments presented in this section provide a dual challenge. First, metallic, shiny surfaces are well-documented to hinder photogrammetric methods because they do not comply with the Lambertian model of surface reflectance. Second, the reconstructed surfaces were selected to exhibit detailed, sharp, and repeatable structures, which are also known to hinder stereo reconstruction methods. The purpose of the experiment is to assess reconstruction quality for metallic surfaces and find the limits of the proposed configuration of this type of surface.

Use case selection rationale

Stereo vision and photogrammetry usually fail to reconstruct even moderately shiny surfaces. The reason is that they reflect different parts of the environment from each viewpoint, and thereby, the 'uniqueness constraint' [73] is not met. The employment of additional algorithmic methods, such as photometric stereo, has countered this incapability [74], which requires additional and high-end hardware and illumination, the application of matting spray [75], or manual editing [76]. A study demonstrating the problems caused by highly reflective surfaces in multiple 3D reconstruction modalities, including photogrammetry, can be found in [100].

Experimental approach

When the imaging range is very close, even shiny surfaces contain some texture.

Shiny metallic nuts and screws were scanned because they feature multiple orientations and curvatures. These structural features are susceptible to illumination artefacts because they reflect light from multiple


directions. Some of these structures are higher than the depth of focus range of the camera. For this reason, coarse sand was used as a substrate, and the targets were partially submerged in it to study the upper surface of these structures. Original images are shown in the top two rows of Figure 21.

The images exhibit specular reflectance, which is reduced with increasing imaging distance. Surface regions that exhibit such reflectance are typically confined to the high-curvature parts of the surface, such as the creases of the bolts and the railings of the screws. In the images, these specularities are expressed as saturated (white) regions of pixels in the high curvature regions. Still, small imperfections, dust particles, and structural features give rise to some feature correspondences.

A practical reason for using multi-distance acquisition is that the *same* metallic surface can appear qualitatively different depending on imaging distance. Figure 18 illustrates this for a washer: from the furthest imaging distance the target appears strongly glossy, with broad specular highlights and limited visible surface texture, whereas closer views of the same region reveal fine-scale roughness (scratches, machining marks, and local texture) and a markedly reduced “mirror-like” appearance. This is important because photogrammetric correspondence estimation benefits from stable, repeatable texture cues; in our setting, these cues become available primarily at close range.

Figure 19 provides representative original images from the washer dataset. The examples highlight (i) how specular highlights concentrate on high-curvature regions and (ii) how microtexture becomes increasingly visible as imaging distance decreases, yielding additional features that can support correspondence estimation.

Data






Figure 19. Original images of the washer across the acquisition distances. Representative views illustrate how apparent gloss and specular highlight extent change with imaging distance, and how fine surface texture becomes visible at close range, providing additional feature support for reconstruction.

The reconstruction outcome is best interpreted by separating geometry formation from appearance. Figure 20 therefore shows three stages of the same result: the sparse point cloud that reflects the distribution of reliable correspondences, the untextured surface mesh that captures the recovered geometry, and the textured mesh that maps the original photographic information onto the reconstructed surface. Presenting these stages side-by-side helps distinguish geometric completeness from texture detail, and makes reconstruction artefacts (when present) easier to localise.




Figure 20. Reconstruction of the washer from the multi-distance dataset. (a) Sparse point cloud, (b) untextured mesh, and (c) textured mesh, shown for the same reconstruction to separate correspondence support, recovered geometry, and mapped appearance.

Results

Metallic nuts and screws were scanned to evaluate reconstruction accuracy. Original images are shown in Figure 21. The corresponding reconstructions are shown in subsequent figures, in the same order, in Figure 22 and Figure 23. The ground truth dimensions of these items were measured using a calliper.

The obtained reconstructions are shown, in the same order, in the four bottom rows of Figure 21. It is observed that the reconstructions do not suffer from gross structural errors. Surface patterns, i.e. screw threads and bolt markings, are reconstructed. When specular reflections are systematic over broad regions of pixels, the reconstruction exhibits artefacts that occur exactly at the high-curvature regions of the surface.



Figure 21. Original images of metallic nuts and screws from the top layer.




Figure 22. Untextured renderings of the obtained reconstructions of the metallic nuts and screws.



Figure 23. Textured renderings of the obtained reconstructions of the metallic nuts and screws.

We compare calliper measurements against reconstruction-derived dimensions in Table 4, in the same order as Figure 21. Dimensions are reported as diameter × thickness, except for screws, where we report length × head diameter. The average error is ≈ 2.8%.

Table 4. Measurements of nut and screw dimensions (diameter × thickness): measured, reconstruction, and errors (mm).

#	Measured (mm)	Reconstruction (mm)	Error (%) (D × T)
1	16.80 × 7.70	17.29 × 7.52	2.89 × 2.36
2	17.80 × 1.50	18.15 × 1.53	1.97 × 2.08
3	12.80 × 6.30	12.63 × 6.54	1.36 × 3.73
4	9.80 × 4.80	10.15 × 4.97	3.59 × 3.54
5	7.80 × 4.80	7.95 × 4.64	1.94 × 3.43
6	6.90 × 3.00	7.18 × 2.89	4.08 × 3.65
7	24.50 × 9.80	25.00 × 10.13	2.03 × 3.33
8	13.00 × 10.00	12.70 × 9.52	2.27 × 4.76

To compare against conventional photogrammetry, we have reconstructed the same scene using the Pix4D software and present the results in Figure 24 (left). As can be observed, very little of the scene is reconstructed. To indicate the difference, the sparse reconstruction using the proposed method is shown, from the same viewpoint, in Figure 24 (right). The reconstruction obtained using Pix4D is poor and manages to reconstruct only a few parts of the scene. The reason is the lack of point correspondences and the establishment of many erroneous correspondences, due to the shiny material of the targets.




Figure 24. Left: reconstruction of a scene with shiny objects using Pix4D. Right: sparse reconstruction of the same scene using the proposed method.

Legacy and fine-detail qualitative and quantitative results are summarised in Annex B.

2.6 Conclusions

The Craeft surface digitisation system now operates as a configurable platform rather than a fixed device. The same mechanical and control foundation supports two complementary 2D operating modes: forensic-level imaging of individual artefacts at ultra-high spatial resolution, and higher-throughput batch documentation of multiple items in a single scan. This flexibility makes the platform suitable across different phases of heritage digitisation workflows, from detailed research capture to routine collection documentation.



At the method level, we presented a complete surface digitisation pipeline comprising hierarchical image acquisition, constrained correspondence estimation, 2D mosaicking via globally consistent registration, optional 2½D reconstruction, and evidence-driven evaluation. For wide-area batch scanning, automation is strengthened by a figure-ground segmentation step based on static background subtraction under controlled illumination, enabling reliable isolation of textiles without requiring a chroma background. The wide-area and fine-detail galleries and datasets (Annexes B–C) further demonstrate that the same acquisition rationale can be deployed across scale, from large textiles to small, detail-rich objects.

Using the same acquisition rationale at coarse and fine scales, we extended the software pipeline to support 2½D surface reconstruction as a scanning modality. Similar to the 2D case, the 2½D approach employs multi-distance acquisition and feature tracking to strengthen geometric constraints, and it remains usable in the presence of specularities that hinder conventional photogrammetry. In the fine-detail experiments on metallic nuts and screws, reconstruction-derived dimensions agree with calliper measurements with average errors on the order of a few per cent, with artefacts concentrated in high-curvature regions dominated by specular highlights.

The current 2½D configuration is limited primarily by the restricted depth of focus of macro imaging, which constrains where stacked images provide maximal sharpness and feature support. Two improvements are envisaged. First, a two-pass scan could coarsely reconstruct the surface to estimate an elevation map, then refine camera motion to follow the surface more closely so that subsequent captures remain within the depth range where the surface is known to occur. Second, stack acquisition can provide an additional cue, depth-from-focus [83], which could contribute to depth estimation in both passes. We already use stacked images, but we do not yet exploit the additional information produced by the stacking operation. The optimal combination of these cues with scan planning and camera motion remains a clear topic for further study.

3 Garment digitisation

Section 3 at a glance

Purpose. Digitise garments on mannequins despite non-rigid cloth drift during acquisition, by combining complementary modalities whose strengths are fused computationally.

Inputs. A multi-modal acquisition set comprising (i) multi-view photographic images, (ii) static multi-station laser scans, and (iii) a quick handheld structured-light scan with IMU used as a geometric control model and world reference frame.

Outputs. A metrically consistent 3D garment model with (i) geometry supported by high-confidence laser regions merged into a unified surface, and (ii) a high-detail photographic texture projected and blended onto the final mesh; supported by overview and result figures (Figure 25, Figure 26, and Figure 27) and capacity-building instructions (Annex D.1).

Method summary.

1. Acquire the three datasets under protocols designed to minimise drift and maximise complementary information.
2. Reconstruct a textured photogrammetric mesh for visual fidelity.
3. Filter and mesh laser scans to retain high-confidence geometry.
4. Register all partial surfaces in the handheld scan coordinate frame, fuse in a volumetric representation, and re-mesh.
5. Synthesise a unified texture by projecting photographic colour onto geometry and blending with confidence weighting.

Practical note. Textile compliance means the surface is not rigidly constrained across time; the workflow, therefore, treats the handheld scan as a rapid control model for alignment seeding, drift detection, and deformation monitoring.

The 3D digitisation of garments on mannequins presents challenges due to the mechanical compliance of textiles. Small environmental perturbations, such as air movements or the operator's wake flow when moving around the object, mean the surface geometry is not rigidly constrained. These perturbations induce subtle non-rigid shape changes over the duration of a scan. These millimetre-scale drifts disrupt spatiotemporal consistency: multi-view datasets no longer correspond to a single stable geometric state. For this reason, the workflow adopts three complementary acquisition modalities whose information is later fused computationally.

A comprehensive review of 3D scanning technologies and approaches can be found in [The Mingei Handbook on Heritage Craft representation and preservation](#), at Step 3 'Craft recording', in Section 3.2, 'Digitisation of enduring assets'.

Conventionally, multi-view photographic capture provides the foundational dataset for camera pose solving and delivers the highest visual fidelity for texture (Figure 25, left). Laser scanning provides dense

point-based geometry with sub-millimetre accuracy. On the other hand, because repositioning the scanner is tedious and the data has a large capacity, usually fewer viewpoints are used in Light Detection and Ranging (LiDAR) reconstructions (Figure 25, middle). This results in more limited coverage per station. Furthermore, colour information is usually lower due to wider FOV optics used in LiDAR scanners. For the same reason, colour is much less consistent across adjacent views. Handheld structured light with an Inertial Measurement Unit produces a rapid, medium-quality mesh and texture (Figure 25, right). Though lacking the geometric precision and optical quality for high-fidelity reconstruction, its rapid data acquisition protocol provides a geometric control model for early registration, drift detection, and subsequent monitoring of deformation during the longer-term acquisition of the other datasets.




Figure 25. Overview of the multi-modal 3D digitisation workflow.

3.1 Data acquisition

Photographic capture is performed by acquiring overlapping images around the vertical axis of the artefact. A rotational circular coverage is used with ~50% image overlap, repeated at multiple heights with the same vertical overlap. Lighting and exposure remain fixed for photometric consistency across the image sequences.

Static laser scans are acquired from multiple stations around the object, sampling near-normal incidence regions of the textile where range noise is minimised. Laser scanning is a time-consuming process. Therefore, the aim is to capture high-quality depth data across the garment's surface, minimising the acquisition of unnecessary geometry. Each station produces a coloured point cloud, which is filtered to retain only the high-confidence geometry (near-normal incidence).

The handheld structured-light device is used for a single quick orbit around the mannequin to generate a medium-resolution mesh with a globally consistent scale. This mesh defines the world coordinate frame for all subsequent registrations and enables alignment seeding and quality control; it reveals whether cloth drift occurred during the static scan phase.

A capacity-building instruction sheet for this acquisition protocol is provided in Annex D.1.

3.2 Geometrical reconstruction

Section 3.2 at a glance

Purpose. Fuse complementary 3D datasets into a single metrically consistent garment surface by combining photogrammetric coverage and texture fidelity with laser-derived geometric confidence, anchored in a handheld reference frame.

Inputs. Photographic image set for photogrammetry, multi-station laser point clouds, and a handheld structured-light scan used as the world reference model.

Outputs. A topologically unified garment mesh in a common world frame, with geometry informed by high-confidence laser regions and prepared for downstream texture synthesis.

Method summary.

1. Reconstruct a photogrammetric mesh to obtain dense coverage and high-fidelity texture support.
2. Filter each laser scan to retain near-normal incidence regions, then mesh to produce partial high-confidence surfaces.
3. Rigidly register each laser-derived surface to the handheld reference (coarse feature alignment, then point-to-plane ICP).
4. Fuse registered surfaces in a volumetric representation and remesh to obtain a single unified surface.
5. Simplify the mesh with an anisotropic, field-aligned decimation that preserves principal curvature directions.

Practical note. The photogrammetric mesh is used primarily for coverage and texture support, while laser regions are filtered for geometric reliability before fusion.

The reconstruction fuses the input data, combining the advantages of each, as follows.

Initially, we perform two operations. Photogrammetry reconstructs a textured surface mesh with high-fidelity texture; the resulting geometry is known to be less accurate and susceptible to fine-scale structural errors. Laser data analysis retains only regions that are near-normal to the scanning direction, and this filtered point cloud is meshed, producing multiple partial meshes with high geometric confidence.

Data	Geometry	Texture	Time	Coverage
Photogrammetry	Low	High (high resolution, high accuracy)	Medium	Dense
Laser scanning	High	Low (lower fidelity, inconsistent)	Slow	Sparse
Handheld device	Medium	Medium-High	Fast	Medium

These partial meshes are merged, using the third dataset acquired from the handheld device as the world reference frame, as follows. Each laser scan is rigidly registered to this reference model; initially via feature-based coarse alignment, then refined by the point-to-plane Iterative Closest Point (ICP) algorithm. We then transform the registered partial meshes into a common volumetric representation. This volume is re-meshed to yield a single topologically unified surface. Finally, mesh simplification is performed using

an anisotropic, field-aligned algorithm to reduce polygon count while preserving principal curvature directions.

3.3 Texture synthesis

Section 3.3 at a glance

Purpose. Produce a uniform, high-detail colour texture for the unified garment mesh by transferring photogrammetric colour onto geometrically reliable regions and blending contributions across scans.

Inputs. The unified mesh from Section 3.2, rigid registrations between datasets, photogrammetric colour observations, and the per-scan angular mask (alpha) derived from near-normal incidence filtering.

Outputs. A single high-resolution albedo texture is baked onto the unified mesh, with seam reduction achieved by confidence-weighted blending.

Method summary.

1. Project photogrammetric colour onto each registered partial mesh and use the angular alpha mask to gate unreliable regions.
2. Bake the per-scan textures onto the unified mesh.
3. Blend textures into a single albedo using confidence weights and produce the final texture map.

Practical note. Weighting should favour close-range, near-normal observations to reduce blur, colour inconsistency, and seam artefacts.

Using the registration result, a colour texture is extracted for the entire model as follows.

1. For each partial laser scan mesh, the corresponding photogrammetric colour data is projected onto its geometry. The angular alpha mask from geometric filtering is reused to produce an alpha channel for texture selection. These intermediate textures are then baked onto the unified mesh.
2. A weighted average image stacking procedure combines all per-scan textures into a single high-detail albedo. The weights represent sample confidence, taken to increase with near-normal incidence and decrease with distance (for example, proportional to the cosine of the incidence angle and inversely proportional to distance). This yields a uniform, high-resolution photographic texture applied to the metrically accurate unified geometry.

3.4 Results

Section 3.4 at a glance

Purpose. Present representative outputs of the garment digitisation workflow and indicate how results are accessed and inspected.

Inputs. The unified textured garment models produced by Sections 3.2–3.3 and their publication in the Craeft online knowledge base.

Outputs. Example visualisations of a digitised dress (Figure 26 and Figure 27) and a catalogue of additional garment models with repository and knowledge-base links (Annex D.2).

Practical note. Mannequins may be removed to avoid biasing visual interpretation, but can be retained when anatomical morphology is relevant to garment design.

The digitised dress, as displayed in our knowledge base’s online viewer, is illustrated in Figure 26. The top row provides front and back views of the garment. The bottom row offers a magnified inspection of its geometric (left) and textural (right) attributes.



Figure 26. The 3D digitisation of a dress.

In post-processing, the mannequin is usually removed so that its style does not influence visual interpretation. In some cases, it is retained when its anatomical morphology is relevant to the garment design; when retained, the mannequin is headless (Figure 27).



Figure 27. The 3D digitisation of a dress worn by a mannequin.

The list of 3D garment models acquired in Craeft is reported in Annex D.2, with hyperlinks to the Zenodo repository for inspection and to the corresponding entries in the Craeft online knowledge base. The visualisations in Figure 26 and Figure 27 are screen captures from the embedded 3D viewer.

4 Digitisation of transparent bodies

The need to create 3D models of objects made from semi-transparent materials is found in various industries and academic fields. Accurate 3D reconstruction of transparent objects is challenging due to light refraction and reflection, which distort the final image. Conventional methods involve covering the object with a substance to reduce transparency or immersing it in a liquid that matches its refractive index. These methods aim to avoid problems caused by light reflection and refraction rather than directly addressing them. However, these approaches add complexity, can be impractical, and might not be suitable for delicate objects, like those of archaeological or cultural significance, due to the risk of contamination or damage. More specific methods for this problem involve capturing 2D images and processing them into a 3D model, such as in Optical Computed Tomography (OCT), which uses light to create a digital volumetric model of an object.

4.1 Background

Section 4.1 at a glance

Purpose. Motivate why transparent-object digitisation is difficult for conventional optical metrology and justify the need for a sprayless, non-contact method suitable for heritage contexts.

Key challenge. Transparent objects exhibit transmission, reflection, and refraction, which violate the diffuse-reflectance assumptions that underpin many practical scanning methods.

Conventional workaround. Opaque coating (e.g., matting spray) enables standard scanners but is costly, time-consuming, and often unacceptable for sensitive heritage artefacts.

Positioning in this deliverable. We investigate Optical Projection Tomography (OPT) for thin-walled transparent objects, targeting a cost-efficient and automatable workflow that avoids refractive-index matching liquids and opaque coatings.

A review of Optical Projection Tomography (OPT) and its imaging principles is provided in Annex E.1.

The accurate and practical optical 3D reconstruction of transparent objects has been an open challenge for the field of optical metrology [84, 85, 86]. The main difficulty in using conventional practical optical metrology tools, such as structured light scanning, laser scanning, and photogrammetry, to reconstruct transparent objects is the combination of transmission, reflection, and refraction exhibited by transparent objects [87]. These effects violate the diffuse-reflectance assumptions required by many practical tools, which typically rely on stable diffuse surface reflection to establish reliable measurements [88].

The typical way of getting around this limitation is to render the transparent object's surface opaque by spray-coating it with a diffuse reflection layer [89]. The transparent object can then be reconstructed with conventional visible-spectrum optical metrology tools. Spraying the object is time-consuming, increases the cost of the reconstruction process, and is not permitted for sensitive heritage objects.

There have been multiple attempts to create alternative sprayless digitisation techniques for transparent objects in the past [84, 85]. The proposed techniques include the use of OPT with a refractive index-matching liquid [90], the fringe projection for detection of pattern differences [91], infrared (IR) to detect induced surface heating [92, 93], infrared digital holography [94], ultraviolet (UV) fluorescence [95, 96], a combination of X-ray tomography and photogrammetry [97], shape from polarisation [98], a combination of polarisation imaging and inverse rendering [99], shape from interaction [100], the visual hull technique [101], various AI-based image processing methods [102, 103], and terahertz (THz) imaging and tomography [104, 105], passive single-pixel imaging [106], and edge estimation computer vision techniques [107].

In this deliverable, we investigate the possibility of achieving the 3D reconstruction of thin-walled transparent objects for quality control and digital preservation purposes in cultural heritage applications.

The proposed method is cost-efficient, non-contact, and simple to automate. It retains the characteristics of conventional optical metrology tools, but can additionally operate without the need for an opaque coating. To achieve this, the use of OPT [108] without a refractive index-matching liquid was investigated.

4.2 Method

Section 4.2 at a glance

Purpose. Acquire projection images of a thin-walled transparent object under controlled illumination and reconstruct a 3D model from multiple viewpoints.

Inputs. The transparent object, a controlled optical enclosure, a motorised turntable (viewpoint control), an LCD screen used as a programmable illumination source, and a high-resolution camera.

Outputs. A multi-view image set (projection views) and a reconstructed 3D model produced by computational processing of these views.

Method summary.

1. Place the object in a controlled, diffuse enclosure and align it on a motorised turntable.
2. Illuminate the object using an LCD screen and acquire images at multiple rotation angles with a fixed camera.
3. Process the resulting projection images computationally to reconstruct a 3D model.

Practical note. The method is designed to be low-contact and automatable, and to avoid opaque coatings that are unsuitable for sensitive heritage objects.

The proposed method illuminates the object, acquires photographs from multiple viewpoints, and processes these images to produce a 3D model.

The novelty lies in a scanning geometry inspired by cone-beam X-ray computed tomography (XCT): a broad-area source and a 2D detector are arranged so that each captured image acts as a projection view for tomographic reconstruction (see Annex E.1).

The proposed scanning apparatus has the following core components: a white chamber, a motorised turntable, a Liquid Crystal Display (LCD) light source, and a high-resolution camera, configured as in Figure 28.

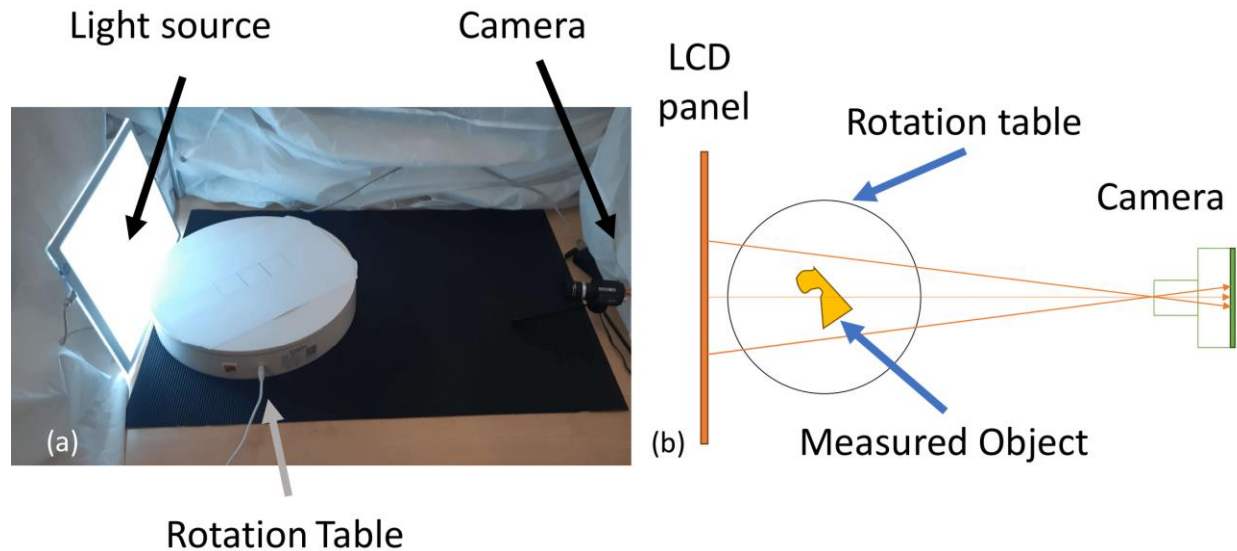


Figure 28. Acquisition setup for projection imaging: (a) photograph from inside the white chamber, showing the chamber walls, LCD screen, and motorised turntable; (b) schematic of the setup (LCD screen, turntable, camera, and enclosure).

4.3 Materials

Section 4.3 at a glance

Purpose. Specify the optical apparatus and computational setup used to acquire projection images and run the OPT reconstructions.

Inputs. White enclosure (to stabilise illumination), LCD screen used as a backlight source, camera, motorised turntable, and the reconstruction geometry and sampling settings passed to the ASTRA Toolbox.

Outputs. A reproducible acquisition setup and a declared compute configuration, including reconstruction runtime characteristics.

Method summary.

1. Use a controlled enclosure to suppress ambient light and stabilise photometry.
2. Place the object on a motorised turntable and acquire evenly spaced projection images over 360°.
3. Reconstruct using ASTRA with an explicit cone-beam geometry and a declared number of views and detector sampling.

Practical note. Reconstruction time is dominated by compute mode; enabling GPU acceleration can substantially reduce runtime relative to CPU-only execution.

4.3.1 Optical apparatus

Opposite the camera, an LCD screen serves as the illumination source. The screen provides stable backlighting so that light transmitted through the object forms projection images that emphasise contours and thickness variations. The camera is aligned with the screen and captures images of the object while it rotates on the turntable.

Acquisition takes place inside a white chamber to suppress external illumination and stabilise the light field. A motorised turntable at the base of the chamber supports the object and provides smooth, controlled rotation. The turntable is driven in fixed angular increments and, at each predefined angle, the camera captures one image. The sequence continues until a full 360° set of views is acquired.

Tomographic geometry and sampling (ASTRA Toolbox). In the cone-beam configuration used for reconstruction, the following parameters are set in the ASTRA Toolbox: detector pixel spacing in x and y (set to 1, in normalised units), detector rows (512), detector columns (512), projection angles (64 views), source-to-centre distance (70 cm), and centre-to-detector distance (20 cm).

4.3.2 Computational infrastructure

Reconstruction was performed on a single laptop with an Intel Core i7-1065G7 CPU and an NVIDIA GeForce MX250 GPU with 2 GB GDDR5 memory. For the ASTRA Toolbox runs reported here, CPU mode was used without GPU acceleration. With this configuration, reconstruction required approximately 2–5 min per object. Runtime can be reduced by enabling GPU acceleration and by using a higher-performance CPU and GPU.

4.4 Calibration

Section 4.4 at a glance

Purpose. Calibrate the imaging and motion geometry so that projection views are associated with accurate intrinsics, known rotation angles, and a metrically meaningful scale for reconstruction.

Inputs. Camera, checkerboard calibration target mounted on the turntable, motor step control, and a reference sphere of known diameter used for scale estimation.

Outputs. Calibrated intrinsic parameters, a consistent camera–turntable geometry and rotation-axis reference, verified per-step angular increment (projection angles), and a global scale factor.

Method summary.

1. Estimate camera intrinsics from checkerboard images (standard calibration).

2. Estimate extrinsic relations between the camera and turntable coordinate frame using the same target while rotating the turntable.
3. Verify and refine the motor's angular increment so that each acquired projection is assigned an accurate rotation angle.
4. Establish a metric scale using a sphere of known diameter reconstructed from its silhouettes.

Practical note. For tomographic reconstruction, the critical product of calibration is a consistent mapping from each projection view to its acquisition angle and geometry; small angular or axis misalignments can produce streaking or shape bias.

4.4.1 Intrinsic camera calibration

Intrinsic camera parameters are estimated using a conventional checkerboard calibration procedure. A planar checkerboard target is placed on the turntable and imaged from the fixed camera. A set of calibration images is acquired while ensuring the checkerboard covers a range of positions and orientations within the field of view (for example, by modestly tilting or translating the target between captures). From detected corner correspondences, the intrinsic matrix and lens distortion parameters are estimated using a standard pinhole-with-distortion model.

4.4.2 Extrinsic camera–turntable calibration

To relate projection views to a rotation model, the checkerboard target is placed on the turntable and aligned as close as practicable to the turntable's rotation axis. The turntable is then rotated in controlled increments, and the camera records an image at each step. Using the known checkerboard geometry and the detected corner features, the pose of the target is estimated for each step. These poses define the relative geometry between the fixed camera and the turntable coordinate frame and provide an empirical estimate of the turntable rotation axis in the camera frame.

4.4.3 Motor step calibration

The step motor is calibrated to verify that rotational increments are accurate and consistent. The expected angular displacement per motor step is compared against the observed angular change implied by the checkerboard poses across successive captures. Deviations from the nominal increment (e.g., due to mechanical backlash, step rounding, or slip) are used to correct the angle assignment for each view. The result is a set of explicit projection angles that are then passed to the reconstruction pipeline, consistent with the acquisition order.

4.4.4 Metric scale calibration

A metric scale factor is established using a polished glass sphere ('lens ball') of known diameter. In the setup used here, a sphere of 80 mm diameter is imaged under backlighting (Figure 29). Backlighting yields a high-contrast silhouette that can be segmented reliably. Because the sphere is convex and its silhouette is stable across views, its 3D shape can be reconstructed from multiple silhouettes using a visual-hull surface approximation method [111]. The reconstructed sphere is then compared to the known physical diameter to derive a global scale factor, which is propagated to subsequent reconstructions.




Figure 29. Back-illuminated images (a) captured by the setup camera are used to extract a reference silhouette (b), from which the sphere is reconstructed via the visual hull technique to establish metric scale.

4.5 Reconstruction

Section 4.5 at a glance

Purpose. Reconstruct the 3D shape of thin-walled transparent objects from projection images using an OPT-inspired acquisition and a CT-style reconstruction pipeline, without refractive-index matching liquids.

Inputs. Calibrated camera intrinsics and projection angles (Section 4.4), projection image set acquired over one full rotation (64 views), and a global metric scale factor (Section 4.4.4).

Outputs. A reconstructed 3D density volume, a surface estimate extracted from that volume, and a scaled 3D point cloud of the object surface.

Method summary.

1. Acquire projection images while rotating the object on the turntable under LCD back-illumination.
2. Pre-process images (normalisation and intensity inversion) to obtain a consistent “absorption-like” signal.
3. Reconstruct a 3D density volume with CT-style filtered backprojection (FBP) in Astra Toolbox.
4. Threshold the volume to remove air and isolate the object region.
5. Extract the thin sidewall surface by peak detection in each slice, then export the resulting 3D points.
6. Apply the metric scale factor from the reference sphere calibration.

Practical note. When refractive-index matching is not used, the reconstructed ‘density’ is an operational quantity driven by the imaging and pre-processing choices. The objective here is consistent localisation of the sidewall surface rather than physically interpretable attenuation coefficients.

To measure the shape of thin-walled transparent items using OPT principles without a refractive-index matching fluid, we use the following experimental and processing sequence:

1. **Setup.** Configure the rotation stage, camera, and LCD field light source as in the apparatus of Figure 30.
2. **Projection imaging.** Illuminate the object with the LCD panel and capture transmission images with the camera.
3. **Multi-view acquisition.** Place the object at the centre of rotation and acquire 64 evenly spaced views over a full 360° rotation.
4. **Pre-processing.** Normalise and invert the monochrome images so that regions corresponding to higher apparent absorption appear brighter and lower absorption appears darker.
5. **Tomographic reconstruction.** Use the Astra Toolbox CT [109] reconstruction pipeline to compute a 3D volume from the projection images.
6. **Background removal.** Threshold the reconstructed volume to remove low-density voxels corresponding to air, retaining the higher-density region associated with the object.
7. **Sidewall localisation.** To extract a single surface for a thin-walled object, process the thresholded volume slice-by-slice and line-by-line, and retain the peak-density location(s) along each line profile as the sidewall estimate (Figure 30).
8. **Export and scaling.** Convert the extracted peak locations into a 3D point cloud and apply the global scale factor established in Section 4.4.4 before saving.

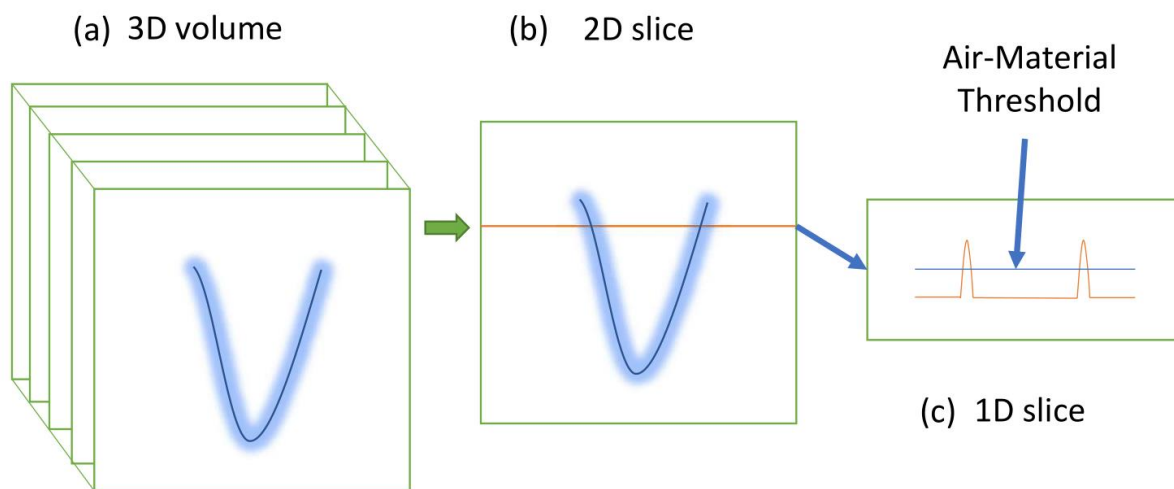


Figure 30. Process of thresholding (a) the reconstructed 3D voxel density to isolate the object region. Each thresholded vertical slice (b) is processed line-by-line (c) to extract absorption peaks that indicate the thin sidewall surface.

Filtered backprojection (FBP) is used to reconstruct a volumetric field from the set of projection images across angles [110]. In essence, each horizontal line in a projection image is treated as a 1D measurement whose frequency content is filtered and then backprojected across the corresponding reconstruction slice. In our implementation, the explicit projection angles are passed to Astra Toolbox so that each view is associated with its calibrated acquisition angle.

4.6 Results

Section 4.6 at a glance

Purpose. Characterise which classes of transparent objects are well suited to the OPT-based reconstruction, quantify external-surface accuracy against an industrial structured-light reference, and compare performance against two camera-based baselines (MVS and NeRF).

Inputs. Six transparent test objects (three polymer bottles and three glass drinking vessels), OPT reconstructions from Section 4.5, reference external-surface scans acquired with a structured-light scanner after opaque coating, and baseline reconstructions produced with MVS and NeRF.

Outputs. A qualitative overview of objects and reconstructions (Figure 31), distance-to-reference visualisations (Figure 33, Figure 36, and Figure 37), representative failure modes for refractive regions (Figure 34), a quantitative error summary (Table 5), and a qualitative side-by-side comparison (Figure 38).

Method summary.

1. Select representative hollow transparent objects and classify them as shelled (approximately uniform wall thickness) or non-shelled (spatially varying thickness).
2. Acquire reference external geometry using a structured-light scanner with opaque spray coating.
3. Reconstruct each object with the OPT pipeline (Section 4.5) and export a scaled point cloud.
4. Align each OPT point cloud to the reference and compute closest-point distances; visualise distances as colour maps and report mean distances (Table 5).
5. Repeat the same comparison for MVS and NeRF baselines.

Practical note. Non-shelled glass objects introduce strong, spatially varying refraction (for example, at the base, rim, and embossed features), which can produce local distortions and ‘ghost material’ artefacts. Mean distance errors should therefore be interpreted together with the qualitative failure cases (Figure 34).

The types of objects to which this measurement principle is most suited are hollow, thin-walled, cylindrically symmetric objects, which do not induce considerable refraction as light traverses through them.

Hollow objects can be further subdivided into two categories, namely ‘shelled’ (objects whose internal and external surfaces are identical in shape and one is scaled down relative to the other by the size of the wall thickness), usually made of plastic. The other type is that of ‘non-shelled’ hollow objects (objects whose internal and external surfaces are not identical in shape and therefore have variations of thickness around the object), which are commonly made of glass.

Both types are of interest in manufacturing and cultural heritage applications, which this investigation aims to focus on.

4.6.1 Setup

To test the category of 'shelled' objects that are commonly used in the beverage industry (e.g., soda and water bottles), we selected a soda bottle and two water bottles (Figure 31, a–c). To test the category of hollow 'non-shelled' objects, applicable mainly to cultural heritage and drink and food containers made of glass, contemporary glass cups (two liqueur and one wine glass) with and without embossed features were measured. (Figure 31, d–f).

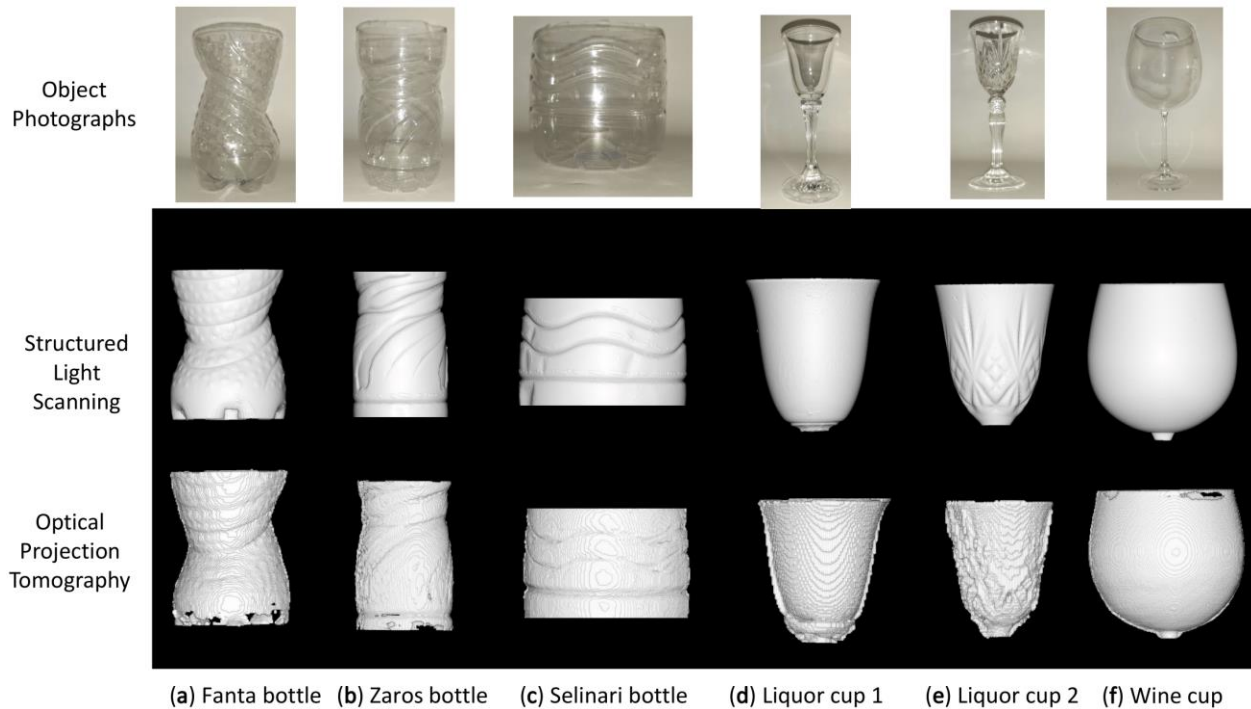


Figure 31. Photographs of transparent objects measured (top row), 3D reconstructions acquired from the Shining 3D Einscan Pro 2X industrial 3D scanner (middle row), and 3D reconstructions by our OPT method (bottom row). (a) Fanta bottle (b), Zaros bottle (c), Selinari bottle (d), Liquor cup 1 (e), Liquor cup 2 (f), Wine cup.

However, 'non-shelled' hollow objects are more difficult to reconstruct since the thickness of the material is not consistent around the whole object and therefore of varying refraction. They contain areas where the light passing through the object encounters thick layers of material and therefore gets refracted (neck, base, and cup bottom). Additionally, many cultural heritage items also contain embossed features, which further add to the variation of material thickness in specific areas of the object. We nevertheless tested such items to test the limits of the suggested OPT method.

To obtain reference 3D reconstruction results for the outside shapes of the measured items, we scanned the objects with a conventional white-light structured light optical scanner used for industrial purposes, called the Shining 3D Einscan Pro 2X scanner (Figure 32). The reference scanner was calibrated to an accuracy of $\pm 22\mu\text{m}$ using the calibration plates provided by Shining 3D. To use this scanner, the transparent objects needed to be coated with the 'AESUB blue' opaque spray coating.



Figure 32. A static structured-light experimental setup used for the acquisition of the reference 3D reconstructions.

4.6.2 Evaluation

The evaluation of 3D requires volumetric error measures to quantify the deviation between the actual 3D structure of the object and its reconstructed version.

Accuracy

Then, using the CloudCompare [112] point cloud software, the reference point cloud reconstructions and the point cloud reconstructions created by the OPT process were aligned. To extract the dimensional error, the residual point cloud distances were calculated. The point cloud errors are depicted as colour textures on the OPT point clouds in Figure 33, and the numerical average of the point cloud distances for each object is reported in Table 5.

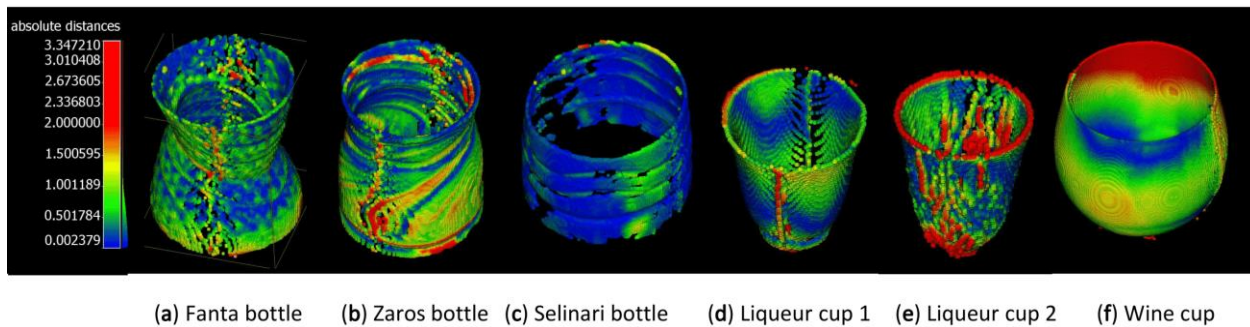


Figure 33. Use of colour texturing on the OPT reconstructed point clouds of the objects in Figure 31, which visually represent the value of the closest distances of each point in the OPT reconstruction from the reference point cloud.

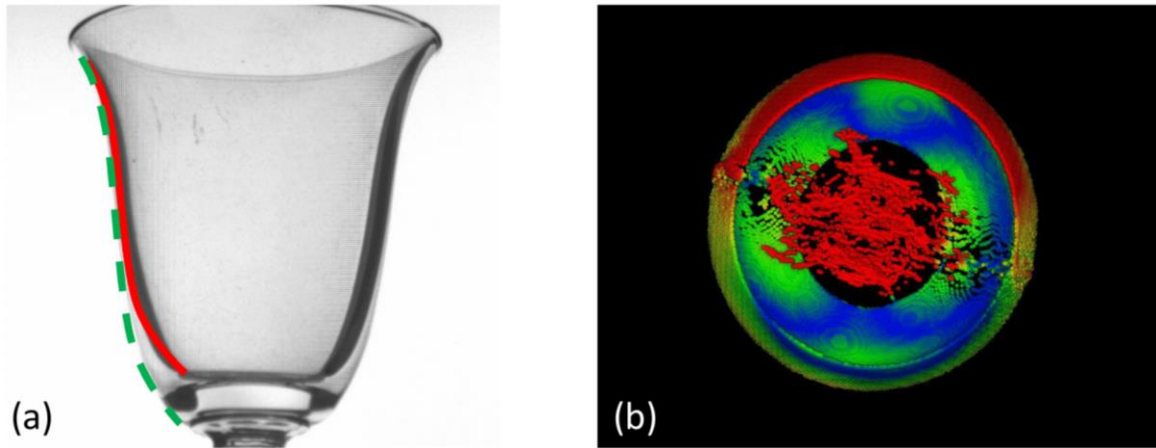


Figure 34. (a) Image of the object in Figure 31d showing a severely refracted liqueur-glass outer edge (red line) compared to the true surface edge (dashed green line), producing erroneous density and hence an erroneous surface estimate. (b) Top view of the wine-glass reconstruction in Figure 31f showing “ghost material” partially filling the cavity.

Table 5. Results of comparing the point clouds acquired by MVS, NeRF, and OPT on the transparent objects shown in Figure 31 to the reference measurements acquired from an industrial-grade 3D scanner.

Object	MVS (mm)	NeRF (mm)	OPT/Ours (mm)
Fanta soda bottle	6.1	3.4	0.3
Zaros water bottle	6.8	5.8	0.4
Selinari large water bottle	8.6	3.2	0.3
Liquor glass 1	5.9	3.7	0.5
Liquor glass 2	3.1	2.2	0.8
Wine glass	10.4	5.6	1.5

There are multiple methods of comparing point clouds [113, 114, 115], using point-to-point, point-to-mesh, and mesh-to-mesh strategies. We opted for using the closest point-to-point distance, rather than comparing point-to-mesh or mesh-to-mesh, because in our case, the reference point clouds created by the structured light scanner were extremely dense, and therefore, it was not necessary to create a mesh surface to accurately compare the point clouds, as suggested in [114].

For hollow non-shelled objects, the average distance errors (0.93 mm), as expected, were higher on average than those measured for the hollow 'shelled' objects (0.34 mm). This fact alone, however, was not reflective of the much wider type of errors experienced on these objects due to intense refraction effects, which manifested as distorted embossed shapes (Figure 33-e), artificial 'ghost material' partially filling up the hollow areas, and the reduction of the object's size.

A simple precision bound can be estimated from sampling: dividing the field of view by the detector width gives an effective sampling of about 0.5 mm per pixel. This implies a quantisation limit of about 0.25 mm (half a pixel) on lateral dimensions, which is consistent with the lowest mean errors observed in Table 5.

Comparison with visual hull reconstruction

In this work, we use OPT to extract the shape of thin-walled objects as a single surface, and it is, therefore, worth qualitatively comparing it to another technique, the visual hull, which is very similar and used during

calibration. The visual hull technique can operate in the visible spectrum without the use of spray coatings. It is well known that it can extract only the external convex shape via the use of silhouettes [111], which is why it was used to measure our calibration sphere.

For hollow non-convex objects, however, it is known that this technique cannot be used, as it produces a solid convex 3D shell around the object. For example, when this technique was used in [101] to reconstruct a wine glass, the opening of the hollow end was covered. Similarly, it is known that any convex cavities (small craters) around the external surface of the object are 'filled up' due to the nature of the visual hull technique. OPT, being a tomography technique similar to X-ray CT, does not have these drawbacks, as it can reconstruct hollow objects and also deal with convex surface structures.

Moreover, OPT can also measure internal surfaces. To demonstrate OPT's ability to measure internally, we placed two cut-offs of plastic bottles, one inside the other, and the reconstructed result is shown in Figure 35. However, since we could not perform a reference measurement for internal surfaces (e.g., using an X-ray CT machine), it was not possible to confirm the achievable accuracy.




Figure 35. Demonstration of OPT's ability to reconstruct internal surface structures, photo of a mixed bottle structure (a), and 3D reconstruction sliced in half (b), showing the internal structure.

Comparison with multi-view stereo photogrammetry

Multi-view stereo (MVS) requires a surface texture to operate, which is why it is well known to perform very poorly on transparent objects [100]. We attempted to reconstruct the objects in this study using MVS, as it is one of the most commonly used camera-based reconstruction techniques today, to contrast its reconstruction quality to that achieved by the OPT technique.

In Figure 36, the prohibitive mean closest-point distances involved in reconstructing transparent objects using MVS, when compared to our reference 3D reconstruction performed via structured light, can be noticed. In Figure 36, these large errors are visualised by aligning the point clouds to the reference reconstructions, and colouring each point of the MVS reconstructed point cloud with its minimum distance to the reference point cloud.

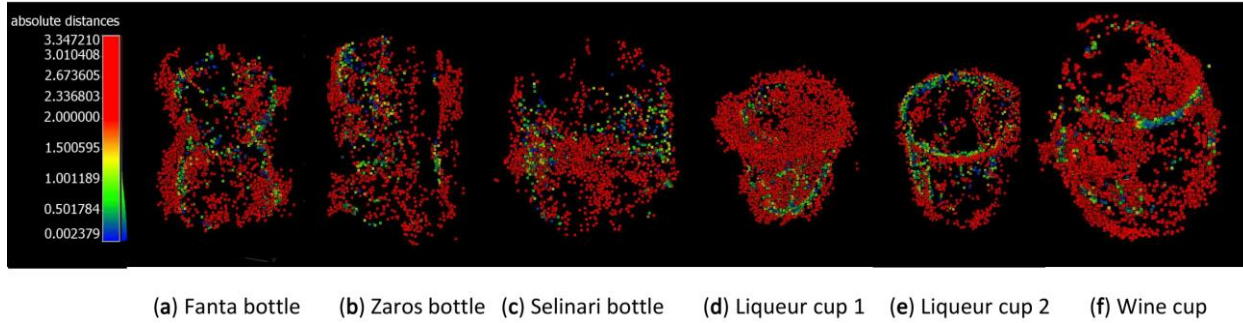


Figure 36. Use of colour texturing on the MVS reconstructed point clouds, which visually represent the value of the closest distances of each point in the MVS reconstruction from the reference point cloud.

Comparison with neural radiance fields

Neural radiance fields (NeRF) are a relatively new reconstruction technique [126]. It uses artificial intelligence to build a non-linear relationship between the input, which is a single continuous 5D coordinate (the spatial location (x, y, z) and viewing direction (θ, ϕ)), and the output, which is the volume density and view-dependent emitted radiance at that spatial location.

It is primarily used for rendering purposes, but it can also recover the voxelised 3D shape of the object. Due to the complication of light transport between views, the reconstruction of transparent objects is not fully successful. It is, however, more successful than multi-view stereo.

In this section, the point clouds created via NeRF with the data acquired by the reference reconstructions in Figure 31 are compared. It can be observed clearly in Figure 37 that NeRF performs better than MVS but worse than OPT. The numerical averages of the errors confirm this observation.





Figure 37. Use of colour texturing on the NeRF reconstructed point clouds, which visually represent the value of the closest distances of each point in the NeRF reconstruction from the reference point cloud.

Discussion

We observe that OPT is much more accurate in reconstructing the external surface of these transparent objects most of the time, with its accuracy being an order of magnitude better than the other two.

If we perform a qualitative study on the 3D reconstruction quality of the same reconstructed object, shown in Figure 38, we notice that OPT retains the most surface details and also has the highest level of

reconstruction completeness. Secondly, we notice NeRF with an acceptable level of completeness but without the ability to reconstruct any of the surface details, and lastly, MVS, which has both very poor completeness and reconstruction fidelity.



(a) Fanta bottle (b) Structured light (with spray) (c) Ours (d) NeRF (e) MVS

Figure 38. Qualitative comparison of different reconstruction methods results from left to right: (a) Object photograph, (b) structured light reconstruction, (c) OPT reconstruction (ours), (d) NeRF reconstruction, (e) MVS reconstruction.

Compared to other reconstruction methods, which have been suggested for the reconstruction of transparent objects mentioned in the introduction, the cost of OPT is minimal, as the only two things required are a field illumination source, such as a large LED, a means of rotation, and a black-and-white camera. The speed of the method has to be divided into acquisition speed and data processing speed, which can be done asynchronously if required. Since the acquisition is performed by cameras, it can potentially be performed on the level of milliseconds. The data processing speed demonstrated here can also be improved many times over by the use of parallel GPUs and with the use of more professional hardware. Regarding the potential for complete automation, it can be completely automated either by adding a robotic arm or via the use of a conveyor belt.

When it comes to the specific class of transparent objects considered in this study, OPT therefore does seem to have the potential to provide near real-time 3D reconstruction at an incredibly low price and with much higher accuracy than any of the techniques that have preceded it so far.

4.7 Comparison with view synthesis

Section 4.7 at a glance

Purpose. Assess whether view-synthesis methods can substitute for metrological 3D reconstruction when digitising transparent or translucent craft artefacts for documentation and presentation.

Inputs. Multi-view image datasets of representative matte and transparent objects acquired under controlled illumination, and reconstructions produced with Neural Radiance Fields (NeRFs) and Gaussian Splatting.

Outputs. A theory-led expectation of failure modes for transparent materials, a comparative experimental demonstration (Figure 39 and Figure 40), and a workflow conclusion: geometry should be acquired by scanning, while appearance should be presented via physically based rendering.

Method summary.

1. Explain why view-synthesis assumptions align better with matte (near-Lambertian) surfaces than with transparent materials.
2. Evaluate reconstruction and view quality on matte vs transparent objects under comparable acquisition conditions.
3. Repeat the transparent-object experiment with a substantially denser image set to test whether additional data reduces artefacts.
4. Conclude when view synthesis is suitable for visualisation and when metrological geometry is still required.

Practical note. For transparent objects, refraction, specular reflection, and multi-path light transport break the straight-ray, single-path rendering assumptions used by standard NeRF-style models; failures should be interpreted as a mismatch between physics and model class, not as poor implementation.

We explored recent technologies that replace explicit geometric reconstruction with view synthesis, namely Neural Radiance Fields (NeRFs) and Gaussian Splatting. These methods learn to synthesise novel views of a scene from 2D images. Although they can yield a volumetric or point-based scene representation, they do not primarily target metric geometry and are therefore typically unsuitable for dimensional measurement. Their objective is to approximate the plenoptic function, also known as the light field. Using this approximation, they can produce realistic views beyond those directly captured. For presentation purposes, the result is often more photorealistic than conventional renderings of a metrologically accurate reconstruction.

4.7.1 Theoretical expectations

NeRFs reconstruct matte objects more reliably than transparent ones because the core assumptions behind NeRF training are closer to how matte surfaces behave in images. A typical NeRF learns a function that maps a 3D location and a viewing direction to a density and a colour, and it renders each pixel by integrating colour and density along a straight ray from the camera into the scene. This straight-ray, single-path model works well when the appearance of a surface point is consistent across different viewpoints.

Gaussian Splatting is also a view-synthesis method trained from multi-view imagery and rendered along straight camera rays; therefore, transparent objects challenge it for the same underlying reason: refraction and multi-path transport break the single-path, straight-ray image-formation model used at render time.

Matte objects are close to Lambertian: a point on the surface tends to have a similar colour from many viewing directions, and changes in appearance are relatively smooth. That gives the training process strong multi-view constraints. The same 3D point is observed repeatedly with similar colour, so the network can use photometric consistency with parallax to determine where surfaces must be in 3D.



Transparent objects violate these conditions in several ways. First, refraction bends light paths. The colour recorded at a pixel when looking through glass or plastic is produced by light that has been deflected by the object, so it does not correspond to the straight camera ray that the NeRF renderer assumes. As the viewpoint changes, the bent path changes too, which means that the same nominal ray direction can sample different parts of the background or the environment. Second, much of what the camera sees through a transparent object is not the object's intrinsic colour, but the background and surrounding scene distorted by the object's shape and refractive index. This distortion is view-dependent, so the training signal that would normally tie a stable appearance to a stable 3D location becomes weak.

Third, transparent objects commonly produce strong specular reflections and Fresnel effects. Highlights and reflections move rapidly with the viewpoint and can be very sharp. Although NeRFs can represent some view dependence, these effects are often too concentrated and too entangled with transmitted background content, so the model may explain them in ways that do not correspond to the true surface geometry. Fourth, real transparent materials introduce complex light transport such as internal reflections, multiple bounces, and caustics. Standard NeRF rendering models a single straight-ray accumulation and cannot capture these multi-path phenomena. When the model cannot represent the correct physics, it tends to place spurious density or colour in the volume to match the observed pixels. Finally, many transparent objects have low or ambiguous texture, so classical correspondence cues across views are weak. In addition, transparent objects often provide weak and view-dependent cues for pose estimation (for example, SfM), because refraction and specularities disrupt the image correspondences that these pipelines rely on. In practice, these issues often manifest as cloudy reconstructions, incorrect shapes, and spurious 'blobs' that float in space.

4.7.2 Experimental assessment

We tested the capacity of these technologies for the 3D visualisation of transparent and translucent objects for digitising craft products and presenting them in 3D.

The 3D reconstruction of matte objects in the lab is a well-studied problem, but exhibits artefacts when scanning transparent or glossy objects. Thus, we tested whether neural view synthesis would help digitise the appearance of objects composed of such materials.

As expected, this case is more challenging due to the lack of Lambertian surfaces and, consequently, the lack of rigid correspondences between views. To compensate, we ensured abundant, stable visual features in the scene (for example, a textured background and well-controlled illumination) to improve pose estimation and multi-view constraints. The experiment revealed that the problem remains challenging even with favourable imaging conditions, highlighting the ongoing need for 3D reconstruction of transparent and translucent artefacts. The results for a piece of pottery and a drinking glass are shown below. We observe that although the view synthesis is satisfactory for matte objects, the quality is insufficient for transparent items in our use case.





Figure 39. View-synthesis comparison under comparable acquisition conditions: a matte pottery object (left) and a transparent drinking glass (right). The matte object yields stable novel-view rendering, whereas the transparent object exhibits view-dependent artefacts consistent with refraction and specular transport.

The video presentations of this comparative experiment are provided online.³

To better understand the phenomenon, we repeated the experiment under the same conditions, this time acquiring ten times more input images than usual for a conventional NeRF reconstruction. The rationale for using this large amount of data was to force the learning of the plenoptic function. The results are shown below.



³ Video 1: https://youtu.be/_cKOK6NX3EQ and Video 2: <https://youtu.be/9XqQqX6X6XQ>.

Figure 40. NeRF result for the transparent object using an order-of-magnitude denser input image set than the baseline experiment. Increased sampling improves background-through-object appearance, but residual volumetric artefacts remain, limiting suitability for faithful presentation of aesthetically rich transparent artefacts.

In general, this approach improves the result. As opposed to the first example, the result reconstructs transparency, meaning that the observer can see the background through the object. However, the results still contain artefacts that render them unsuitable for the realistic presentation of aesthetically rich objects.

4.7.3 Conclusions

NeRF-style methods do not directly produce an explicit surface model suitable for conservation workflows without additional post-processing, and their reconstructions are not designed to be metrically reliable. Moreover, their performance is not sufficient for the realistic presentation of transparent objects. As such, our proposed workflow is the reconstruction of geometric structure using scanning and the visualisation of the result using physically based rendering.

Still, neural radiance technologies are quite useful for capturing the appearance of conventional scenes and targets with modest means, proving sufficient for many users. To demonstrate this, we created a walkthrough of an archaeological site featuring marble works and stone masonry from multiple historical eras. The result is provided in Annex I.

As discussed in Annex E.1, the light propagation conditions required by OPT are approximately satisfied for the thin-walled objects considered here. Some errors are still expected due to the minimal but non-zero refraction induced by the object's sidewalls.

4.8 Conclusions

4.8.1 Technical

The study demonstrates that an OPT-inspired workflow can reconstruct the 3D shape of a specific class of thin-walled transparent objects without refractive-index matching liquids. Compared to conventional X-ray CT, the approach is non-ionising, lower cost, simpler to deploy in a lab or heritage setting, and easier to automate with camera-based acquisition. At the same time, unlike standard optical metrology (structured light, laser scanning, photogrammetry), it can operate without applying opaque spray coatings, which are often unacceptable for sensitive artefacts.

The experiments indicate that the method is best suited to hollow, thin-walled, near-cylindrically symmetric objects that induce limited refraction along most optical paths. In our dataset, this corresponds to shelled objects with approximately uniform wall thickness (e.g., polymer bottles), for which the mean surface distance to the structured-light reference scan was approximately 0.34 mm. For non-shelled objects with spatially varying thickness (e.g., glass vessels), refraction becomes strongly location dependent—particularly at thick regions such as the base, rim transitions, stem–vessel joints, and embossed features—and the mean distance increased to approximately 0.92 mm.

Importantly, the modest difference between these mean values does not capture the qualitative gap in failure modes. For non-shelled glass objects, the reconstructions may exhibit local distortions of embossed details, apparent shrinkage, and “ghost material” partially filling hollow regions (see Figure 34). These artefacts are consistent with the decision to avoid index matching: thick optical paths produce stronger refraction and violate the straight-ray assumption that underpins CT-style filtered backprojection under visible-light transmission (see Annex E.1 for the analytical context).

Compared with two camera-based baselines (MVS and NeRF), the OPT pipeline yields substantially improved completeness and retention of surface structure for the object classes where refraction remains limited. In this sense, OPT provides a practical compromise: it does not replace X-ray CT for arbitrary transparent artefacts, but it does enable non-contact, sprayless external (and in some cases internal) surface recovery for a narrow yet industrially and culturally relevant class of objects.

In summary, the work supports the claim that OPT-style reconstruction without refractive-index matching is feasible at the centimetre scale when the object’s sidewalls are sufficiently thin and uniform to avoid significant refraction. This makes the technique promising for dimensional quality assurance of hollow-shelled objects (for example, polymer packaging). For glass heritage vessels and other non-shelled objects with substantial thickness variations, the method is substantially more error-prone and should be used cautiously, with quantitative metrics interpreted alongside visible failure modes.

4.8.2 Patent application

As part of the project’s innovation objectives, a patent application has been initiated to protect the intellectual property associated with the proposed transparent-object digitisation workflow.

A PCT application was filed through the World Intellectual Property Organisation to establish an international route for protection. A European patent attorney was engaged to prepare the application, including the claims defining the scope of protection.

Following filing, an International Search Report (ISR) and Written Opinion of the International Searching Authority (WOISA) were received. The ISR identified relevant prior art, while the WOISA provided an assessment of novelty, inventive step, and industrial applicability. The feedback also included formal recommendations (e.g., claim formatting/numbering and language consistency) and substantive recommendations (e.g., claim refinements to better distinguish from the cited prior art and to strengthen support in the description).

In response, the claims were amended, and a reply to the observations was prepared in collaboration with the patent attorney. The updated submission was filed with the European Patent Office on 10 January 2025. Application identifiers: PCT/GR2023/000051, Submission number 13933600, EP application number 23800947.6; Receiving Office: European Patent Office, The Hague.

No charges have been made to the Craeft budget for the above. The Technology Transfer Office ([PRAXI](#)) of Foundation for Research and Technology - Hellas covers the patent attorney fees. All invoicing and business transactions complied with the PEPPOL BIS 3.0 specification.

Supporting documentation is provided in Annex E.2.

5 Digitisation of challenging materials

5.1 Introduction

Section 5.1 at a glance

Purpose. Introduce *material digitisation* as a complement to conventional 2D/3D digitisation, motivating why “challenging” materials (transparent, translucent, glossy) require measured optical data to support faithful visualisation and analysis.

Inputs. (i) Conventional geometry/appearance capture outputs (meshes, textures, photographs), and (ii) measured optical properties (e.g., spectral transmittance/reflectance, refractive index) suitable for simulation.

Outputs. A conceptual framing for Section 5: why appearance cannot be treated as a fixed texture under fixed lighting, and operational definitions of transparency/translucency/opacity used throughout the section.

Method summary.

1. Distinguish conventional digitisation (geometry + appearance *under a specific capture illumination*) from material digitisation (optical parameters enabling *re-rendering under controlled illumination*).
2. Explain why thickness-dependent light transport matters for craft materials (especially glass).
3. Provide definitions that will be used consistently in the remainder of the section.

Practical note. Photographs are not “wrong”, but they are *capture-condition dependent* (illumination spectrum, geometry, exposure, white balance). Material digitisation supports controlled comparison by predicting appearance under *specified* lighting conditions.

Advanced digitisation goes beyond photographic capture and conventional 3D scanning. It targets material composition and optical behaviour—the properties required to create digital twins whose appearance can be *predicted* under specified illumination conditions rather than merely *recorded* under the lighting present during acquisition.

Conventional digitisation typically captures surface geometry and an apparent colour/texture that is entangled with the illumination at the time of capture. By contrast, material digitisation encodes the information required to simulate light–material interaction. In practice, this supports *controlled re-visualisation*: the same geometry can be rendered under different, defined lighting setups, allowing principled comparisons across materials and thicknesses. Physically based rendering (PBR) uses such data to simulate reflection, transmission, refraction, and absorption, thereby predicting appearance under arbitrary, specified illumination conditions.

5.1.1 Professional significance

For practitioners and designers, material appearance is not defined only by “colour recipes” or surface finish, but by how light travels through the material. Glass is a canonical example: two samples can appear similar when thin, yet diverge substantially once shaped into thicker objects, layered elements, or architectural components, because absorption and internal transport effects accumulate with optical path length.

Material digitisation makes such differences inspectable *before* committing to physical production. By comparing candidate materials against familiar references under matched geometry and lighting, it becomes easier to anticipate how perceived colour intensity, brightness, and warmth will change with thickness. This supports concrete design decisions, such as selecting glass types for vessel bodies, decorative panels, or lighting elements.

For practitioners working with recycled or alternative raw materials, this is particularly important. Novel compositions raise uncertainty: will colour darken excessively with thickness? Will the object become visually dull? Will it shift towards an undesired warmth? The approach demonstrated here provides an early-stage tool for reasoning about these issues, reducing trial-and-error, saving material, and supporting more sustainable production practices. The goal is not to replace hands-on experimentation, but to complement it: to help craftspeople reason about material behaviour, communicate choices, and integrate innovative circular glass materials with greater confidence.

5.1.2 Ecological motivation

This work sits at the intersection of two complementary efforts: CRAEFT, which explores how advanced digital tools can support sustainability and the transmission of craft knowledge, and the Colver national research project, led by CERFAV, which investigates glass colouration through agro-mineral and circular processes. The collaboration offered a concrete setting in which material innovation in glassmaking can be documented, analysed, and communicated in ways that respect both accurate measurement and practitioner understanding.

Within CRAEFT, a key objective is to represent material behaviour in a form intelligible to practitioners, designers, conservators, and researchers. Glass is a particularly challenging case: its identity is defined not only by shape and surface finish, but by how it transmits, absorbs, refracts, and reflects light—often in thickness-dependent ways that are hard to infer from tabulated data alone.

The Colver work provides a timely use case. By developing formulations grounded in circular-economy principles, it raises practical questions that matter to makers and designers: how does the new glass compare to standard references; how does colour emerge with increasing thickness; how stable is its visual behaviour under neutral illumination? Addressing these questions requires more than isolated chemical or optical measurements: it requires a comparative visual framework grounded in measured data.

Using spectrophotometric and refractometric measurements alongside reference datasets, we can construct controlled comparative studies in which geometry is held constant while material properties vary. By rendering identical bodies across materials and thicknesses, the studies isolate perceptual consequences of absorption, reflectance, and refraction, enabling differences to be attributed to material behaviour rather than to form, lighting variability, or surface decoration. This aligns with circular-economy

principles by reducing uncertainty and costly trial-and-error when adopting recycled or alternative materials.

The work also resonates with the ambitions of the New European Bauhaus, particularly its emphasis on practices that integrate sustainability, aesthetics, and inclusion. Here, simulation is treated as a mediating tool: it translates measured optical data into visual evidence that can be read and discussed by non-specialists, enabling a common discourse among scientists, artisans, and designers.

5.1.3 Definitions

The terms below are used *operationally* in this deliverable. In practice, materials lie on a continuum; here we distinguish categories by the degree of scattering and the preservation of image coherence.

Transparency

- **Definition:** A state in which a material permits the passage of light with negligible scattering, such that image structure is largely preserved.
- **Context:** The material may absorb selected wavelengths (tint) and may refract rays, but it does not substantially randomise photon directions. Consequently, an observer can discern geometric detail of objects behind the material.

Translucency

- **Definition:** A state in which a material permits the passage of light but induces significant subsurface scattering, thereby destroying image coherence.
- **Context:** Light is transmitted, but internal microstructure diffuses directions, so spatial information is scrambled. An observer perceives illumination, shading, and sometimes coarse silhouettes, but not crisp geometric detail.

Opacity

- **Definition:** A state in which a material prevents transmission of light by dominant absorption and/or reflection, forming a barrier to optical propagation.
- **Context:** Incident light energy is absorbed (often converted to heat) and/or reflected into the environment; negligible light exits the rear side, so no image is visible through the material.

5.2 Background

Purpose. Position “challenging materials” (shiny, transparent, translucent) as a gap between conventional digitisation pipelines (geometry + RGB texture) and physically faithful appearance simulation; motivate why spectral material capture matters.

Inputs. Prior work on 3D digitisation (photogrammetry/laser scanning), established visualisation approaches (PBR, real-time rendering, global illumination, image-based lighting), and spectral measurement concepts.

Outputs. A concise framing of (i) digitisation vs visualisation, (ii) the evolution from texture baking to physically based material models, and (iii) how spectrometry/spectroscopy supports material-aware digital twins.

Practical note. For non-Lambertian materials, “good geometry + RGB texture” is rarely sufficient: appearance depends strongly on illumination, view direction, and wavelength-dependent transport.

5.2.1 Digitisation

To date, much cultural heritage research has prioritised the digitisation and interactive visualisation of existing artefacts—i.e., the construction of geometric and appearance models from acquired data (in computer vision parlance, reconstruction). Scanning and modelling technologies have therefore become central to photographic and 3D documentation workflows. High-resolution laser scanning and photogrammetry are widely used to capture detailed geometry and to acquire high-fidelity textures suitable for inspection and dissemination. A comprehensive review of 3D scanning approaches is provided in [The Mingei Handbook on Heritage Craft Representation and Preservation](#), Step 3 (“Craft recording”), Section 3.2 (“Digitisation of enduring assets”).

In parallel, material-oriented digitisation methods aim to reveal properties that are weakly visible or invisible in conventional RGB capture. These approaches rely on imaging beyond the visible spectrum (e.g., IR and UV) and on multispectral or hyperspectral acquisition, enabling the identification of materials and supporting conservation analysis. For example, [131] reviews advances in multispectral and hyperspectral imaging for archaeology and art conservation and highlights their relevance to material identification and preservation.

However, many practical pipelines still implicitly assume predominantly Lambertian (matte) behaviour: they acquire geometry plus an RGB texture that visually “looks right” under the capture lighting. When artefacts are shiny, transparent, or translucent, appearance varies strongly with illumination and viewpoint, and the model must explicitly represent reflectance/transmittance behaviour rather than only a baked colour texture.

The table below presents the advances in material digitisation and visual representation (rendering) chronologically.

<p>Photographic and photogrammetric</p>	<p>Records the appearance under a single, controlled lighting condition. Photographic quality texture is acquired as a Diffuse Map or Colour Map and applied to the 3D model. The texture has baked-in shadows, highlights, and ambient occlusion. This approach lacks versatility and interactivity. It produces contradictory and unrealistic results when the virtual light source in the scene does not match the light source in the photograph.</p>
--	---

<p>Geometrical and Empirical Maps (early PBR)</p>	<p>In an interim step, workers in computer graphics approximated light interaction without physics by separating the colour from structural details. Graphics shaders used maps. Diffuse maps store the base colour. Specular maps control surface reflection intensity. Normal maps render surface detail, without geometrically reconstructing these details. The process often violates the energy conservation principle and lacks a model for how reflections should change with the viewing angle. Artists empirically adjust these maps for a particular lighting setup. This enables interactive 3D viewing, allowing the object to be rotated and viewed from any angle. Surface quality and reflections remain visually inconsistent when the light source is moved.</p>
<p>Physically Based Rendering based on Analytical Methods</p>	<p>PBR is a comprehensive shading and material system where all texture inputs correspond to physically measurable properties of a real-world surface. The advancement is the shift from capturing how the object looks to what the object is. This guarantees consistent and accurate material behaviour under any lighting conditions. Digitisation separates the colour, geometry, and reflection characteristics. This makes PBR versatile, realistic, and highly consistent across different applications and lighting environments.</p>

5.2.2 Visualisation

The visualisation of heritage artefacts has advanced substantially with real-time 3D graphics and immersive technologies (AR/VR), enabling interactive educational experiences and broader public access. Real-time rendering remains essential for immersive applications and typically relies on GPU acceleration to deliver interactive frame rates [123].

For photorealistic depiction, global illumination methods model complex light transport in a scene [125], while image-based lighting uses captured real-world illumination to improve realism and material plausibility [121].

For material-specific appearance, physically based rendering (PBR) provides a principled framework in which reflectance and transmission are simulated rather than “painted in” [125]. Modern rendering infrastructures can model light–material interaction with high fidelity and produce photorealistic imagery appropriate for documentation, presentation, and design tasks [122].

More recently, Neural Radiance Fields (NeRFs) have attracted attention as view-synthesis methods that generate novel views from multi-view imagery by learning a volumetric scene representation [126]. While such methods can produce compelling visualisations, their representations are not primarily designed for metric measurement, and they are typically tied to the capture conditions (notably the scene’s illumination and surrounding context). This limits their suitability when the goal is to create editable, relightable digital surrogates for documentation and conservation workflows.

5.2.3 Spectrometry and spectroscopy

Spectrometry refers to measuring light as a function of wavelength—i.e., quantifying spectral intensity across a range of wavelengths. The device is a spectrometer (or, in many material workflows, a spectrophotometer), and it can measure wavelength-dependent reflectance (for opaque materials) and

transmittance/absorption (for transparent materials). Conventional cameras and many scanners sample a small portion of the visible spectrum through broad RGB filters, whereas spectral measurement records a much denser wavelength profile.

In this deliverable, spectral measurements are incorporated to support more faithful material simulation. For transparent and translucent materials, wavelength-dependent transmittance and absorption directly influence perceived colour and depth effects, while refractive behaviour affects sparkle, brilliance, and view-dependent appearance. Integrating spectral data into rendering workflows supports more reliable “digital twins” for design, documentation, and presentation, reducing reliance on repeated physical prototyping and enabling consistent predictions under different illuminations and viewing geometries.

5.3 Transparent and translucent materials

Section 5.3 at a glance

Purpose. Enable faithful visual reproduction of transparent and translucent artefacts by representing light–material interaction at (i) interfaces, considering reflection and refraction, and (ii) volumes, considering absorption and scattering. In both cases, using wavelength-dependent (spectral) parameters.

Inputs. (i) Geometry: solid or hollow; (ii) spectral material data: refractive index, transmittance, and absorption; and (iii) rendering configuration: integrator, SPP, and environment illumination.

Outputs. Spectral and RGB renderings whose appearance differences arise from measured/standard optical behaviour (dispersion, selective attenuation, Beer–Lambert absorption), rather than from ad hoc RGB approximations.

Practical note. For transparent media, texture and geometry are insufficient: appearance depends strongly on view direction, illumination, and wavelength-dependent transport.

Transparent and translucent materials are characterised by three coupled behaviours: reflection at interfaces (specular appearance), refraction (ray bending), and attenuation within the medium (absorption and, for translucent media, internal scattering). Our primary goal is faithful visual reproduction, i.e., predicting how a material absorbs, reflects, and transmits light under arbitrary illumination and viewing conditions. To achieve this, we treat light transport in two regimes:

1. Interface regime (surface boundary). Reflection and refraction are governed by Fresnel behaviour and Snell’s law. Spectral dispersion arises because the refractive index varies with wavelength, i.e., $n(\lambda)$.
2. Volume regime (interior medium). Light is attenuated and potentially scattered as it travels through the medium. Spectral absorption (and optionally scattering) is modelled with Beer–Lambert attenuation and a volume scattering model.

By using wavelength-dependent optical constants, the rendered appearance of each glass variant becomes sensitive to dispersion and selective attenuation, allowing visual differences among materials to emerge naturally under the same illumination and geometry.

5.3.1 Introduction

Light transport describes how radiant energy propagates through an environment and interacts with surfaces and media: light is emitted, reflected, transmitted, absorbed, and scattered. For transparent and translucent artefacts, the decisive phenomena are (i) reflection/refraction at boundaries and (ii) absorption/scattering within the volume. In practice, these effects create a view- and illumination-dependent appearance that cannot be captured reliably by a single RGB texture.

5.3.2 Theory

Geometry and materiality as independent dimensions

We treat geometry and materiality as independent dimensions. This allows controlled comparisons: the same shape can be rendered across different materials and environments, and differences in appearance can be attributed to optical properties rather than to geometry.

Interface behaviour: reflection and refraction

An interface is the boundary between two media with different optical densities (e.g., air and glass). At each interface, incident light is partitioned into reflected and transmitted components. In rendering terms, this is modelled by a BSDF that includes both reflection and transmission lobes. Macroscopically, refraction follows Snell’s law, and the reflection/transmission ratio follows the Fresnel equations.

We denote the refractive indices of the two media as spectral functions:

- $n_{\text{ext}}(\lambda)$: refractive index of the *exterior* medium (typically air),
- $n_{\text{int}}(\lambda)$: refractive index of the *interior* medium (the transparent material).

Because rays can arrive from either side of the surface, the roles of “interior” and “exterior” swap depending on ray direction relative to the surface normal; the renderer therefore assigns the appropriate pair $(n_{\text{ext}}(\lambda), n_{\text{int}}(\lambda))$ per interaction.

Microscopically, surfaces exhibit roughness, which broadens highlights and changes reflection sharpness. We model roughness with a microfacet distribution: Trowbridge–Reitz [128], implemented as GGX [129]. The user-facing roughness parameter is mapped to the microfacet parameter (α) such that the perceptual response is approximately linear; in our implementation, user inputs are squared internally, i.e., $(\alpha = r^2)$.

Volume behaviour: absorption and scattering

Between interfaces, rays travel through a participating medium (the glass volume). Volume effects are governed by:

- Absorption coefficient $a(\lambda)$: selective removal of energy by wavelength,
- Scattering coefficient $\sigma_s(\lambda)$: redirection of energy within the medium,
- Extinction $\sigma_t(\lambda) = \sigma_a(\lambda) + \sigma_s(\lambda)$.

Beer–Lambert attenuation over path length (d) is: $T(\lambda, d) = \exp(-\sigma_t(\lambda), d)$.

For clear homogeneous glass, we typically model pure absorption (no volumetric scattering), i.e., $\sigma_s(\lambda)=0$. For translucent materials, $\sigma_s(\lambda)$ becomes non-zero and must be paired with a phase function (e.g., Henyey–Greenstein) to model directional scattering.

Notation used in this section

Symbol	Meaning
λ	Wavelength (nm)
$n_{\text{ext}}(\lambda)$	Exterior refractive index spectrum (air)
$n_{\text{int}}(\lambda)$	Interior refractive index spectrum (material)
$\sigma_a(\lambda)$	Spectral absorption coefficient
$\sigma_s(\lambda)$	Spectral scattering coefficient
$\sigma_t(\lambda)$	Spectral extinction coefficient ($\sigma_a + \sigma_s$)
$T(\lambda, d)$	Spectral transmittance over distance d
α	Microfacet roughness parameter (GGX)

5.3.3 Software parameterisation

This work uses the light-transport visualisation toolbox reported in D3.2. Here we detail the parameters relevant to material representation.

Sampling quality (SPP). Image quality is primarily determined by samples per pixel (SPP), i.e., the number of Monte Carlo light paths evaluated per pixel. Larger SPP reduces variance (noise) and improves convergence of view-dependent phenomena such as caustics, sharp specular reflections, and transmission through complex geometry. Computational cost scales approximately linearly with SPP.

Integrator and medium model.

We use a volumetric path tracer to account for:

- interface refraction/reflection (dielectric BSDF),
- volumetric attenuation along interior paths (Beer–Lambert),
- optional microfacet roughness (GGX).

For homogeneous, clear media, we typically set the medium to pure absorption (equivalently, volumetric albedo (= 0)), which suppresses scattering while retaining selective attenuation.

Outputs

Results are generated from a virtual perspective camera. The output is dual: a conventional RGB image and a spectral representation (used for faithful wavelength-dependent simulation and downstream conversion).

Scene assumption

All bodies are immersed in air. For hollow bodies, the cavity is also air, at the same pressure and refractive index as the surrounding environment.

5.3.4 Material representation

Faithful simulation requires measurements of material optical properties. For transparent media, both reflected light (interface) and transmitted light (volume) are relevant; spectroscopy and spectrophotometry provide the necessary inputs.

Spectral data sources

The input data are spectrophotometric measurements, either custom-measured or obtained from manufacturer references:

- Annex G.1: measurements for a new glass type (custom dataset),
- Annex G.2: manufacturer data (SCHOTT), including BK7 (a widely used optical glass).

These datasets provide wavelength-dependent transmittance/attenuation measurements over broad spectral ranges; for rendering, we filter and resample to the visible domain used by the spectral pipeline.

Spectral refractive index (dispersion)

To capture dispersion, the refractive index is defined as a function of wavelength rather than a scalar. For materials with Sellmeier coefficients (e.g., BK7), we reconstruct: $n^2(\lambda) = 1 + (B_1\lambda^2)/(\lambda^2-C_1) + (B_2\lambda^2)/(\lambda^2-C_2) + (B_3\lambda^2)/(\lambda^2-C_3)$, where λ is in micrometres and (B_i, C_i) are material coefficients from optical catalogues. We then evaluate $n(\lambda)$ at discrete wavelengths (e.g., 380, 390, ..., 780 nm) for spectral sampling.

Absorption/extinction spectrum.

For each dataset, tabulated attenuation values are parsed as pairs $(\lambda_i, \sigma_{t,i})$, filtered to 380–780 nm, and interpolated onto a uniform grid (e.g., 10 nm spacing). These spectra drive Beer–Lambert attenuation: $T(\lambda, d) = \exp(-\sigma_t(\lambda), d)$.

Because computers cannot represent continuous functions exactly, both refractive index and attenuation spectra are stored as discrete vectors sampled over wavelength.

5.3.5 Geometric representation

Solid bodies

Solid bodies are represented as watertight triangle meshes with outward-facing normals. These surfaces define a single dielectric boundary between air and the material medium.

Hollow bodies

A hollow glass body is represented by two watertight meshes:

- Outer surface: normals point outward towards the external air.
- Inner surface: normals point inward towards the cavity (air).

The region between the two surfaces is the material medium. When rays cross the outer interface, they enter the medium; when they cross the inner interface, they leave the medium and re-enter air. The optical path length inside the medium is therefore the geometric distance between the two intersections, defined by the two meshes. This wiring mirrors the physical situation of a hollow glass shell.

Stratified bodies

Stratified objects (e.g., glazed ceramics) combine layered materials and therefore require modelling multiple sequential interactions: light must (i) refract through a clear specular coating (glaze), (ii) attenuate within it, (iii) scatter from the rough diffuse substrate (clay), and (iv) traverse the glaze again on exit. This produces characteristic optical depth and colour enrichment.

To simulate such behaviour, the model requires:

- Absorption of the coating (possibly tinted), affecting light on entry and exit (double pass),
- Reflection and albedo of the diffuse substrate,
- Layer thickness/path length, which modulates attenuation and perceived depth.

5.3.6 Validation experiments

Data

The spectral inputs are from Annex G.1 (custom) and Annex G.2 (SCHOTT reference), resampled for the visible band. In all cases, the refractive index for air is set to 1.0.

Setup

The illumination is an HDRI environment map ([Studio Small 08](#) from Poly Haven), shown in Figure 41.



Figure 41. The HDRI environment was used for illumination in the validation experiments.

Subjects

We use spheres as test objects because they exhibit all surface orientations. Two glass types are tested:

- BK7 (reference optical glass),
- CLV (custom/measured glass type).

For hollow bodies, inner and outer surfaces are identical in shape; the region between them is glass, attenuated using Beer–Lambert with wavelength-dependent spectra [130]. The hollow geometry is a spherical shell formed by two nested spheres:

- outer radius R ,
- inner radius $R - \tau$,

where τ is the wall thickness. In the experiments, ($R = 5$) cm and ($\tau = 1$) mm.

Results

Figure 42 supports scene understanding: it compares hollow spheres made of BK7 (left) and CLV (right) against a calibration substrate. The upper and lower hemispheres can be inspected comparatively: the upper region is dominated by refraction/transmission, while the lower region shows stronger reflection of the environment. In this example, SPP = 256 and resolution = 512×512 ; the camera axis is oriented at 45° relative to the substrate.




Figure 42. Comparative scene understanding of hollow spheres (BK7 vs CLV) showing refraction and reflection patterns against a calibration substrate.

Figure 43 introduces local surface perturbations (“scratches”) by adding deformations on the outer surface to demonstrate the interaction of light with more realistic surface geometry.




Figure 43. Rendering of a glass object with added surface deformations (“scratches”).

Figure 44 uses solid spheres to compare BK7 and CLV under matched illumination, highlighting interface effects (Fresnel reflection/refraction and microfacet roughness) and, where enabled, volumetric attenuation. The individual panels are self-labelled; interpretation should follow those per-panel labels rather than any fixed row/column ordering.




Figure 44. Solid-sphere comparisons for BK7 and CLV under identical illumination and camera settings. Panels are self-labelled (material and parameter settings); visual differences reflect interface behaviour (Fresnel + roughness) and, where enabled, volumetric effects.




Figure 45. Experiments with the measured glass. Top: the Stanford bunny benchmark rendered as hollow and solid glass. Bottom: two glass bodies designed by CERFAV rendered as frosted glass.

5.4 Applications

Section 5.4 at a glance

Purpose. Demonstrate how the material digitisation pipeline supports (i) solids, (ii) hollow glass containers, and (iii) stratified/participating materials such as porcelain across making states.

Inputs. Digitised spectral material descriptions (Annex G.1–G.2 and Colver measurements), canonical test geometries (spheres, Stanford bunny), and representative container meshes with controlled thickness profiles.

Outputs. Comparative renderings under matched camera and illumination settings for (i) solids, (ii) hollow glass artefacts in contextual environments (Figures 43–45), and (iii) porcelain making states with static, animated, and interactive visualisation (Figures 46–47).

Practical note. Throughout, comparisons hold geometry, illumination, and camera fixed wherever possible, so that visible differences can be attributed to material parameters rather than scene variation.

5.4.1 Solids

We demonstrate solid transparent objects using the Stanford bunny as a canonical test geometry. The bunny provides nontrivial curvature, self-occlusion, and edge silhouettes, making it a useful stress test for refraction, Fresnel reflection, and caustic-like transport effects under controlled illumination.

Material parameters are instantiated from the spectral reference measurements in Annex G.2 (including BK7) and the Colver project measurements. All solid-body experiments use watertight meshes with outward-facing normals and are rendered under matched camera and illumination settings to isolate material-driven appearance differences.

5.4.2 Glass containers

We next apply the same material models to hollow containers rendered in visually plausible contexts. To expose view-dependent behaviour (reflection/refraction) under everyday conditions, we use four HDRI backgrounds from Poly Haven: studio, office, forest, and park (left to right in Figures 43–44).

Two container geometries are tested: (i) a bottle with approximately uniform wall thickness (~ 2 mm), and (ii) a vase with spatially varying wall thickness (~ 3.5 – 4.5 mm). Unless stated otherwise, the surface micro-roughness parameter is set to $\alpha = 0.06$, corresponding to a well-polished but non-pristine surface (i.e., not an optical lens finish).

Figure 46 and Figure 47 compare BK7 (top row) against Colver glass (bottom row) under identical camera settings. The Colver tint in these renderings is selected to be consistent with photographs of comparable objects provided in Annex G.1.




Figure 46. Bottle renderings under contextual HDRI illumination. Material: BK7 (top) and Colver glass (bottom). Backgrounds (left to right): studio, office, park.




Figure 47. Vase renderings under contextual HDRI illumination. Material: BK7 (top) and Colver glass (bottom). Backgrounds (left to right): studio, office, park.

The glass reference measurements in Annex G.2 are further used to render the vase in a neutral studio environment across a set of SCHOTT glass variants (Figure 48). Hyperlinks in the figure panels point to the corresponding video renders; the associated playlist is provided for convenience.




Figure 48. Vase rendered with SCHOTT reference glass variants under a neutral studio environment. Individual panels are self-labelled by glass type; embedded hyperlinks point to the corresponding YouTube video renders. Playlist: <https://www.youtube.com/playlist?list=PLQZAktGyYNinJQ6Chg50LVuWF3MLIZMnw>

5.4.3 Porcelain bodies

We apply our rendering capabilities to illustrate its utility in rendering various qualities and processing stages of porcelain. This study employs physically based light transport simulation to visualise and analyse the optical appearance of porcelain surfaces representative of different making states. Eight canonical configurations were defined to model the progressive transformation of the surface and subsurface structure during manufacture.

Context and lineage

For the trained eye, porcelain is never simply white and smooth; it is the visible trace of both what it is made from and how it is made. The material refers to the geological origin of the clay. Much of the finest porcelain begins with Limoges kaolin, a naturally bright, low-iron clay found in central France. It gives a

potential for whiteness, purity, and translucency. The lineage is what makers do with that material: the recipes, firing, glaze, and finishing methods that transform clay into recognisable qualities. Lineages often become associated with ‘house’ styles.

The Limoges lineage aims for optical clarity. The body is hard-paste and densely vitrified, the glaze sits glass-flat, and highlights stay tight and fine, producing a cold, blue-white tone and knife-edge rim brightness. A neutral lineage, by contrast, favours warmth and tactility. Its body, though still refined, is slightly more open in texture. Under identical illumination, highlights broaden and soften. Edges bloom rather than gleam; the light seems to breathe rather than to snap. In Annex F, we provide the colour values for plain and Limoges porcelain.

Through the making states, these qualities unfold. In the unglazed body (biscuit), one already senses the distinction: Limoges holds light tautly, while the neutral porcelain diffuses it gently along its surface. After glazing, Limoges hardens into brilliance, whereas the neutral lineage rounds into warmth. In satin-matte and celadon variants, the same contrast plays out. In overglaze decoration, the Limoges style ensures motifs remain sharply defined beneath a perfect glaze; conversely, neutral surfaces tend to blend decoration more gently into the surface.

Eight examples demonstrate the impact of subtle alterations in the top layer on our perception. Starting with bare porcelain, they progress through clear coatings of varying sheens to glazes that either tint or cloud the underlying body.

The table below summarises the physical mechanisms that determine the phenomena mentioned above.

Category	Light transport mechanism	Visual sensation
Unglazed (T1–T2)	Subsurface scattering in the body; weak surface specular	Matte to semi-translucent; body-colour dominant
Transparent glaze, surface model (T3–T4)	Dielectric surface reflection (Fresnel) + roughness-controlled specular lobe	Gloss level changes; highlights sharpen/broaden
Explicit glaze film, clear (T5)	Explicit film + multiple internal bounces + body return	Deeper gloss; thickness-dependent depth cues
Coloured transparent glaze (T6, T8)	Absorption in the glaze film (participating medium)	Green tint with optical depth; saturation increases with absorption/thickness
Opaque scattering glaze (T7)	Strong volumetric scattering in glaze film	Milky/creamy appearance; suppressed detail beneath

The appearance of a surface, ranging from chalky and matte to deep, glossy, and luminous, is determined by how it influences the probability distribution of photon path lengths and the angular redistribution of emergent radiance. This concept is demonstrated across eight states, which progressively simulate the radiative transport space. These states vary from a purely diffusing medium, such as unglazed ceramic, to a strongly refractive and absorptive film, like a coloured glaze.

Rendering study design and reading guide

The objective of this subsection is to demonstrate how the eight making states (T1–T8) manifest perceptually under *identical* geometry, illumination, and camera settings, and to isolate the effect of the

porcelain body palette (Neutral vs Limoges-type). We therefore treat the rendering study as a controlled experiment: for each condition, all parameters are held fixed except the material state and, in the Neutral/Limoges comparison, the body colour palette.


Static figure layout

Figure 49 is presented as a compact 2×4 grid for rapid comparison. In each image (Neutral and Limoges), the top row (left → right) shows T1, T2, T3, T4, and the bottom row (left → right) shows T5, T6, T7, T8. This grid layout is used only for the static figure to keep all eight states visible at once.

Neutral vs Limoges-type comparison

Figure 49 and the paired animations compare two body palettes under identical settings. The Neutral palette is an achromatic baseline, while the Limoges-type palette is slightly cooler (blue-biased) to emulate higher purity kaolin with reduced long-wavelength absorption. Importantly, the glaze parameters are held constant; any systematic difference between the two sets should therefore be interpreted as body-driven, except in states where the glaze itself strongly absorbs or scatters (notably T6–T8).

In the example below, we visualise treatments T1–T8 for Neutral (top) and Limoges (bottom) porcelain. In both cases, the top row (left→right) shows T1, T2, T3, T4, and the bottom row (left→right) shows T5, T6, T7, T8. Both images use identical lighting, camera, geometry, and glaze settings. The only change is the body palette. The Neutral image uses an achromatic white body (equal RGB), while the Limoges-type image uses a slightly cooler (blue-biased) body to emulate higher-purity kaolin with reduced long-wavelength absorption. Any visible difference between the two images therefore arises from the body material, not from illumination or glaze parameters.






Figure 49. Treatments T1–T8 rendered on neutral (top) and Limoges-type (bottom) porcelain bodies

Preliminary renderings and key observations

Figure 49 provides a controlled, static comparison of the eight making states (T1–T8) under identical geometry, camera, and illumination. The figure serves as a compact reading guide for the richer studies that follow (contextual renderings, animations, and interactive inspection). It isolates (i) differences attributable to the making state, and (ii) within each state, differences attributable to the porcelain body palette (Neutral vs Limoges-type).

In each image (Neutral and Limoges-type), the top row (left→right) shows T1–T4 and the bottom row (left→right) shows T5–T8. All rendering parameters are held constant across the two images except the body palette; non-body parameters remain unchanged between Neutral and Limoges-type for the same state.

Our key observations are:

- Body-driven differences are clearest when the glaze is absent or optically clear (T1–T5). Limoges-type bodies appear slightly cooler/cleaner in the unglazed states (T1–T2) and maintain that cooler base tone under clear glazes (T3–T5).
- Glaze-driven differences dominate when the glaze absorbs or scatters strongly (T6–T8). In the celadon states (T6, T8), the green appearance is primarily determined by absorption within the glaze layer; the underlying body modulates the result secondarily. In the opaque state (T7), strong scattering within the glaze reduces sensitivity to body colour.
- Specular behaviour distinguishes surface-only from explicit-film modelling. The surface-model glazes (T3–T4) mainly affect highlight sharpness via roughness, whereas explicit-film states (T5–T8) additionally express thickness-dependent cues arising from absorption and/or scattering within the glaze volume.

Realistic renderings

Whereas Figure 49 uses a 2×4 grid for compact comparison, the contextual studies below use a circular ring layout to support animation and interactive inspection.

To place the material differences in a more natural viewing context, we render the eight states on identical spheres positioned on a diffuse ground plane in a circular layout. The spheres are evenly spaced on a ring (radius 0.8 m), each sphere has a radius of 0.12 m, and the scene is illuminated by an HDRI environment map to provide both high-frequency specular cues and stable diffuse shading under consistent conditions.

Unless otherwise stated, renderings use a perspective camera with a 35° field of view at 960 × 540 px and 256–512 samples per pixel (SPP). These controls ensure that differences across states are attributable to the material models rather than to lighting, geometry, or camera variability.

Animated visualisation

Static images convey the overall appearance of each state, but animations make differences easier to perceive by revealing how highlights, reflections, and subsurface cues evolve with viewpoint. We therefore provide two complementary motion regimes:

1. Turntable mode. The camera and HDRI remain fixed while the entire ring of spheres rotates as a rigid group through 360°. This emphasises specular dynamics (highlight travel, edge brightening, and roughness-dependent lobe width), and is particularly informative for differentiating glaze states.
2. Orbit mode. The objects remain fixed while the camera moves around them through 360°. This supports cleaner observation of body tone and translucency under changing viewpoints, with reduced confounding from changes in illumination direction relative to the objects.

Each animation samples one full revolution at 0.5° increments (720 frames). The eight spheres are ordered clockwise, starting from the sphere nearest the camera, as: T3 → T4 → T5 → T6 → T7 → T8 → T1 → T2 (see Section 5.4.3.3).

Presentation conditions. To stress progressively more demanding appearance cues, we provide three conditions:

- Plain (CMP): no decoration; isolates intrinsic material appearance.
- Gold rings (RNG): adds a narrow metallic equatorial band; stresses highlight behaviour and gloss contrast.
- Patterned (PTN): adds a surface pattern; stresses visibility of fine texture through glazes and under scattering.

Video link table. The table lists the animation datasets used to compare porcelain-making states across body palettes and motion regimes. Columns separate the porcelain body palettes (Neutral vs Limoges-type) under identical glaze parameters. Rows group three presentation conditions—Plain (CMP), Gold Rings (RNG), and Patterned (PTN)—each shown in two motion regimes: Turntable (camera and illumination fixed; objects rotate) and Orbit (objects fixed; camera moves). Each cell links to the corresponding video sequence.

Condition	Motion	Neutral	Limoges-type
Plain (CMP)	Turntable	https://youtu.be/ETse0vCMWNs	https://youtu.be/Zg8OfUyPiMQ
	Orbit	https://youtu.be/GJBy7W_JCOM	https://youtu.be/tUfdboQ-kHs
Gold Rings (RNG)	Turntable	https://youtu.be/Rt698abrEtQ	https://youtu.be/TX5NtuvBdXQ
	Orbit	https://youtu.be/28w8jAHDKEg	https://youtu.be/HvEG_DkVgmA
Patterned (PTN)	Turntable	https://youtu.be/nW1dyWevf4M	https://youtu.be/uTa7DSA4xEU
	Orbit	https://youtu.be/mq3uJkpycgl	https://youtu.be/bly2ISkbHx8

Interactive visualisation

While the videos provide complete coverage of the rendered frames, interactive control (pause, scrubbing, reversal, and side-by-side comparison) supports more precise inspection of subtle differences—especially gloss–diffuse separation, view-dependent highlight behaviour, and the visibility of surface patterns through transparent and scattering layers. For this reason, we also provide a lightweight, browser-based HTML viewer that assembles rendered frame sequences into an interactive layout.

The viewer featuring the interactive examples below can be downloaded at: <https://doi.org/10.5281/zenodo.18190038>. Reference: Zabulis, X. (2026). Interactive viewer for porcelain materiality and lineage assessment. Zenodo. doi: [10.5281/zenodo.18190038](https://doi.org/10.5281/zenodo.18190038)

The viewer mirrors the comparison logic used in this subsection:

- Horizontal comparison: Neutral vs Limoges-type bodies under identical conditions.
- Vertical comparison: turntable vs orbit motion regimes.

The viewer presents these conditions as three tabs labelled Plain, Rings, and Pattern (corresponding to CMP, RNG, and PTN).

The viewer is provided with the deliverable evidence bundle as a self-contained HTML page alongside the corresponding rendered-frame folders.

Screenshots of the viewer for the three conditions are shown in Figure 50, Figure 51, and Figure 52, and the user controls are described in Annex H.

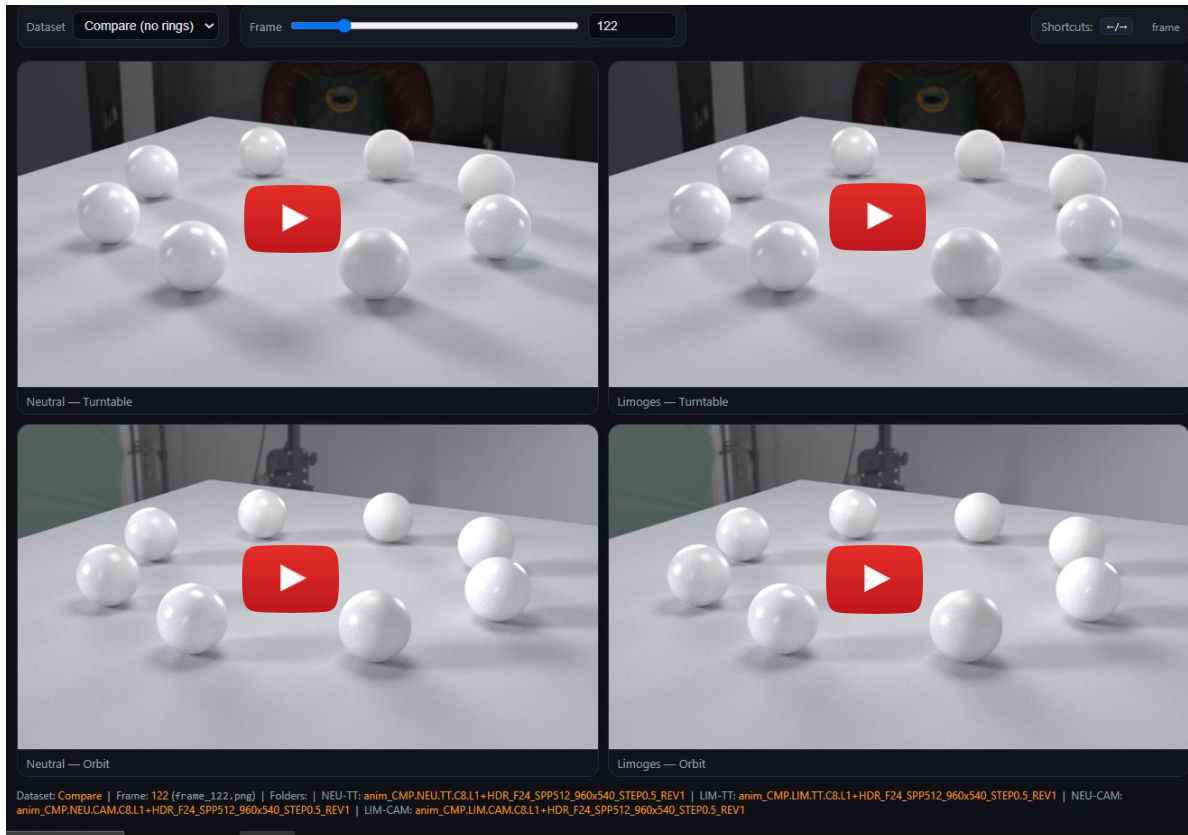


Figure 50. HTML viewer UI for the Plain (CMP) condition. The viewer shows a 2x2 comparison matrix with columns Neutral (left) vs Limoges-type (right) and rows Turntable (top) vs Orbit (bottom).

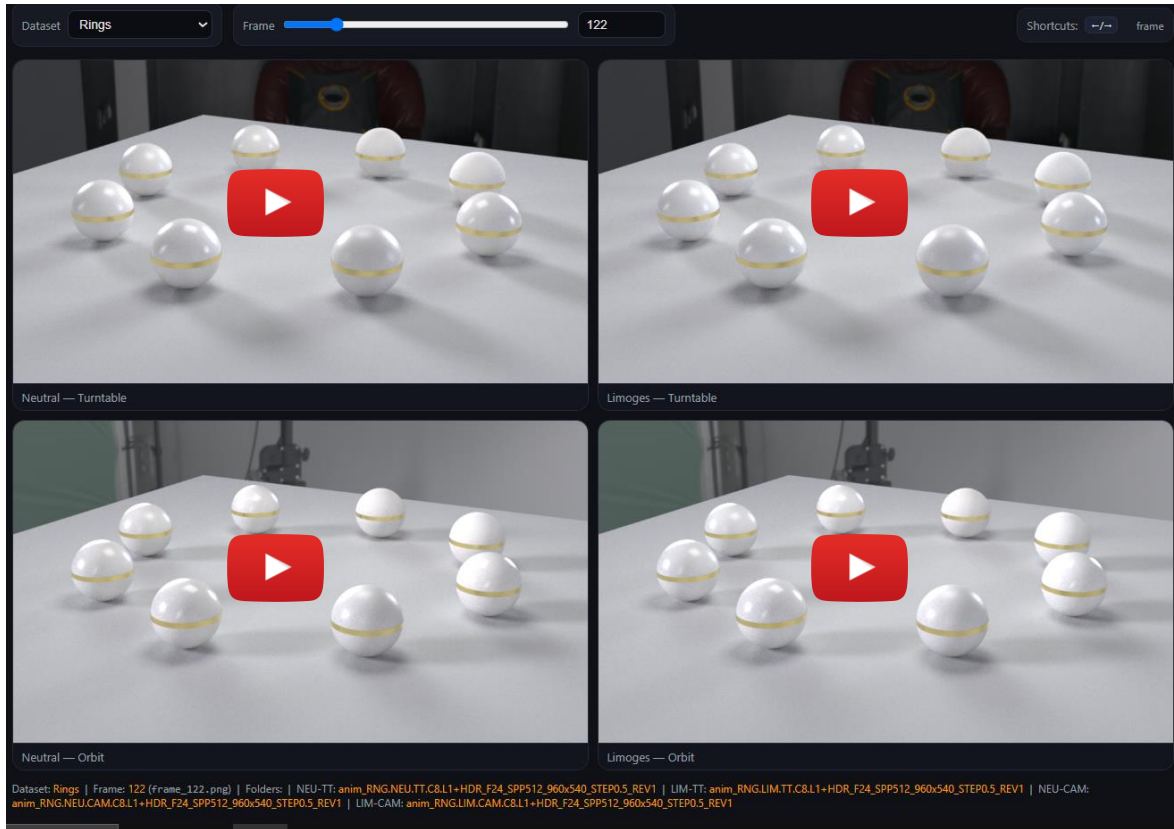


Figure 51. HTML viewer UI for the Gold rings (RNG) condition. The viewer shows a 2x2 comparison matrix with columns Neutral (left) vs Limoges-type (right) and rows Turntable (top) vs Orbit (bottom).

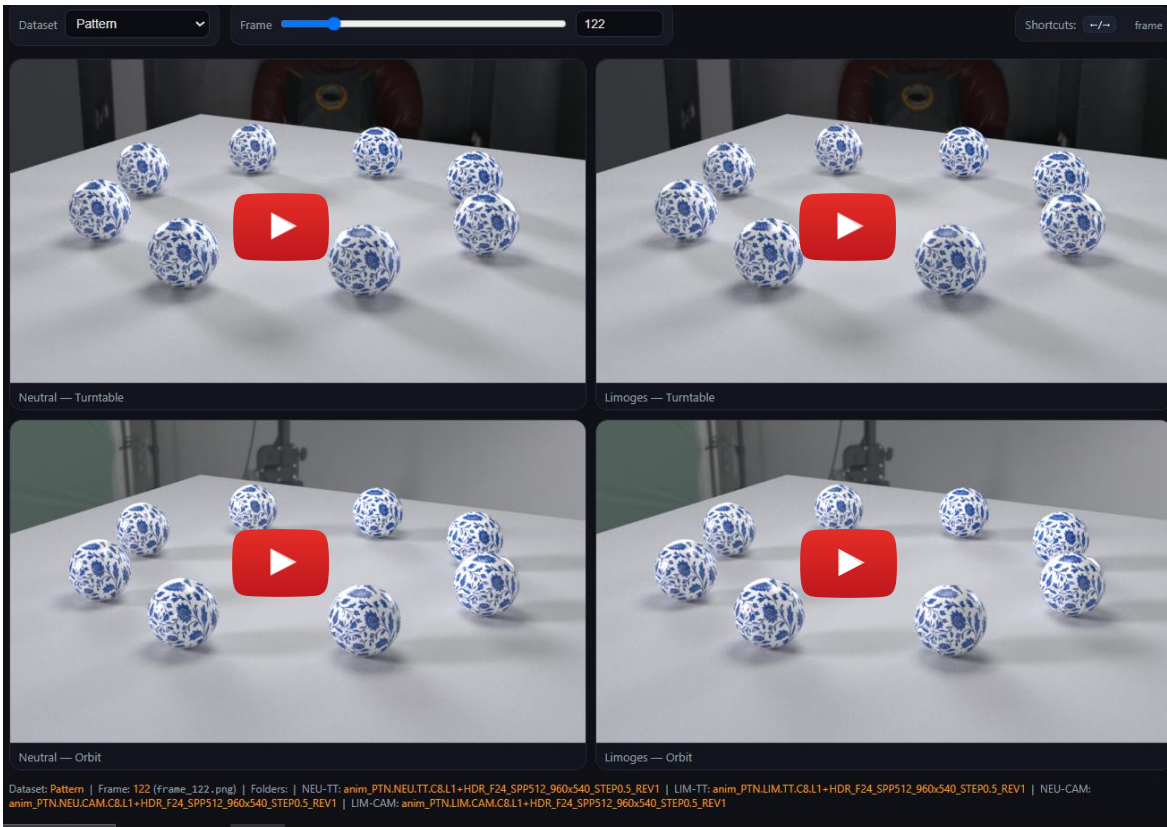


Figure 52. HTML viewer UI for the Patterned (PTN) condition. The viewer shows a 2×2 comparison matrix with columns Neutral (left) vs Limoges-type (right) and rows Turntable (top) vs Orbit (bottom).

5.4.4 Discussion

This digitisation capability has dual use. First, it supports computer-aided design and material exploration: by keeping the object geometry, illumination, and viewing conditions controlled, render studies provide a stable basis for comparing how alternative porcelain body states affect perceived whiteness, gloss, and translucency. Second, it supports communication and evaluation: by varying only the material state (T1–T8) while holding all other factors constant, the resulting videos and interactive viewer enable traceable side-by-side comparison that can be reviewed by craft experts and non-experts without conflating material effects with changes in camera motion or lighting.

From the Neutral vs Limoges comparisons, we observe the following:

- In the Limoges rendition, both biscuit and Parian (unglazed) bodies appear marginally cooler and “cleaner”. Interpreting this perceptually, the overall white point shifts slightly towards the blue end of the spectrum, which becomes noticeable because subsurface scattering integrates small spectral differences over longer path lengths.
- The glaze is colourless in the clear-glaze states, so the appearance beneath the specular highlight is governed primarily by the diffuse body. The two palettes produce comparable highlight structure, while the Limoges-type body yields a slightly cooler base tone beneath the same



specular response. This is expected for a transparent glaze whose dominant effect is attenuation rather than recolouring.

- In the celadon states, the green hue is dominated by absorption within the glaze layer. The underlying body palette modulates the perceived cleanliness of the green, but the dominant chromatic effect originates in the glaze.
- In the opaque state, strong volumetric scattering within the glaze reduces sensitivity to the underlying body palette by repeatedly redirecting light before emergence. Body-driven differences are therefore attenuated relative to the clear-glaze cases.

In summary, replacing a neutral body with a cooler, Limoges-type body shifts the base tone most clearly when the glaze is absent or optically clear, whereas glaze-driven effects dominate when the glaze itself absorbs or scatters strongly (such as celadon/opaque). This separation of body-driven versus glaze-driven contributions is precisely what makes the controlled rendering workflow useful both for material exploration and for faithful, reproducible presentation of scanned artefacts.

6 Conclusions

This deliverable consolidates a set of cost-effective, reproducible, and practically deployable digitisation methods for heritage-related artefacts, covering (i) 2D and 2½D surface digitisation, (ii) garment digitisation on mannequins, (iii) transparent-object reconstruction, and (iv) material-aware visualisation studies for challenging materials. In these areas, the guiding principle remains the same: capture shape and appearance data under controlled conditions, and organise results so they are traceable, comparable, and reusable in subsequent documentation and presentation workflows.

6.1 Surface digitisation

The Craeft surface digitisation system is reported as a configurable platform rather than a fixed device. A shared mechanical and control foundation supports complementary operating modes: (a) ultra-high resolution 2D imaging (mosaicked capture for detailed documentation) and (b) higher-throughput batch documentation (wide-area capture of multiple items). The corresponding pipeline combines hierarchical multi-distance acquisition, constrained correspondence estimation, globally consistent registration for mosaicking, optional 2½D reconstruction, and evidence-driven evaluation. Wide-area use cases demonstrate how the same acquisition rationale scales across object sizes and materials, while fine-detail experiments illustrate both the strengths of the approach and where artefacts tend to concentrate (notably in high-curvature, specularly-dominated regions). The annexes provide the apparatus description, calibration, legacy and wide-area galleries, and supporting datasets to enable inspection and reuse.

6.2 Garment digitisation

The deliverable reports a pragmatic workflow for digitising garments on mannequins in the presence of textile compliance and drift. The approach is explicitly multi-modal: photographic capture supports visual fidelity, while structured-light and laser scanning contribute geometric control and high-confidence surface regions. Registration and fusion are performed in a common reference frame to obtain a metrically consistent garment surface, with a texture synthesis step that projects and blends photographic colour onto the fused geometry. The annex material provides operational guidance intended for repeatable acquisition and quality control.

6.3 Transparent bodies

For transparent artefacts, the deliverable reports an OPT-inspired reconstruction workflow that targets thin-walled, hollow objects without requiring refractive-index matching liquids and without applying opaque spray coatings. The results show that the method can provide useful reconstructions for a narrow but relevant object class (near-cylindrically symmetric, approximately uniform wall thickness), while also documenting characteristic failure modes when refraction becomes strongly location-dependent (e.g., thick bases, embossed features, and non-uniform thickness). The conclusions are therefore limited: the method is presented as a practical compromise—useful under specified geometric and material conditions and to be interpreted alongside visible artefact patterns when those conditions are not met.

6.4 Challenging materials and material-aware presentation

The deliverable also addresses the gap between “a textured mesh” and “a faithful visual appearance”, by separating (i) geometry acquisition from (ii) material/appearance modelling and then using controlled rendering studies to support comparative inspection. The porcelain body studies exemplify this approach: material states (T1–T8) are defined explicitly (including intentional replicates), and an interactive viewer enables four-way comparison across base type and motion mode under controlled conditions. The discussion emphasises that such controlled studies support both material exploration (understanding the perceptual effect of parameter changes) and faithful presentation (maintaining consistency across viewpoint and illumination), provided that the assumptions and parameterisations are stated clearly.

6.5 Limitations and next steps

Across the reported digitisation methods, the dominant outputs remain geometry first and appearance second, typically as texture mapped to the recovered surface. This is an appropriate and useful baseline for documentation, but it does not fully capture materiality in the broader sense (e.g., acoustic and haptic properties), and it does not, by itself, model the interaction between light and matter that drives many perceptual cues—especially for transparent and translucent artefacts. Two concrete directions follow from this observation. First, richer materiality capture can be pursued by using the 2½D surface digitiser as a source of physical surface texture that can be mapped to haptic rendering pipelines (planned for the next iteration of the deliverable). Second, a more realistic visualisation of acquired assets can be achieved by integrating the digitised geometry and measured/parameterised material properties into physics-based rendering workflows, in alignment with the material-specific light transport simulation work reported in D3.1 Craft-specific action simulations.

Annex A. Surface scanner

A.1 Apparatus

The components required to construct the proposed piece of hardware are cost-efficient, and several of them are 3D printable.

The specifications of the computer used in the experiments were as follows: CPU ×64 Intel i7 8-core 3GHz, RAM 64Gb, GPU Nvidia RTX 8Gb RAM (RTX2060 SUPER), SSD 256Gb, HDD 2Tb. The critical parameters are CPU and GPU RAM, as they determine the number of correspondences that can be processed and, therefore, the area that can be reconstructed.

The flash is mounted on and moves along with the camera. A ring flash is employed, which is a circular illumination device that fits around the camera lens, creating a ring of light around the subject. Ring flashes provide even and shadow-free illumination for close-up shots of small subjects and evenly light the subject from all angles. Sensor brightness, contrast, and colour balance were set to automatic. The sensor hardware and firmware implemented focused stacking.

A textured substrate is employed to aid reconstruction, as it yields more feature point correspondences. The data are stored on a 128 GB SD card. The storage capacity of this card determines the maximum number of images that the system can acquire. The average size of the image file is 2.4 MB.

The visual sensor was an Olympus Tough TG-5 configured for macro acquisition (4,000 × 3,000 px, JPEG; FoV 16.07° × 12.09°). The closest focusing distance was approximately 2 mm (as configured). Individual exposures have limited depth of field; therefore, in-device focus stacking was used to obtain an effective depth of field of approximately 1 cm under the acquisition settings used. Illumination was produced by the sensor's flash. Here, "closest focusing distance" refers to the minimum focusable object distance, while "depth of field" refers to the object-space range appearing sharp, increased via focus stacking.

The scanner operates in a Cartesian frame where x and y span the scan plane, and z is the elevation (stand-off distance). Lateral motion in x – y is achieved by moving the camera/gantry (printer mechanics), while changes in z are achieved by setting the camera elevation relative to the substrate. This elevation determines the height of a hypothetical square frustum. The doubling of the elevation for each layer is to create a medial overlap of ≈ 4 . In this way, an image at a higher layer oversees 4 images from a lower layer.

A.1.1 Configuration for close-range 2D surface scanning

The number of layers was 4, and they were configured as noted in Table 6. In this table, E is the elevation, n is the total number of acquired images, and n_x and n_y are the numbers of images acquired in the horizontal and vertical dimensions, respectively. Moreover, s_x and s_y are the lengths of the steps that the camera is moved to acquire an image in the horizontal and vertical dimensions, respectively. Finally, s_p is the length of the side of the square surface region that is imaged by an image pixel, at the specific layer.

Table 6. Spatial arrangement of camera locations in a square frustum structure.

Layer	E (mm)	n (#)	n _x (#)	n _y (#)	s _x (mm)	s _y (mm)	s _p (μm)
1	160.0	50	5	10	4.51	3.38	11.296
2	80.0	551	19	29	2.25	1.69	5.648
3	40.0	3102	47	66	1.12	0.84	2.824
4	20.0	14420	103	140	0.56	0.42	1.412

This level of medial overlap was sufficient for the samples we scanned. Denser or sparser configurations are treated in the same way. For the utilised sensor, the base layer of this pyramid is $5 \times 5 \text{ cm}^2$ and is covered by $25 \times 19 = 475$ images. The inequality between the n_x and n_y steps is by preference. Given the rectangular camera FoV, this configuration results in a square scanned area.

We report an optical resolution of 19754 horizontal and 19820 vertical ppi, without interpolation. Mosaic pixels, thus, deviate by a factor of $\varphi = 0.00334$ from squareness. The measurement was obtained using the banknote, which is of known dimensions. This means that mosaics are linearly scaled by a factor of φ in the vertical direction. If needed, the mosaic can be resampled to feature square pixels. We report the horizontal as the scanner resolution, thus 19754 ppi. The benchmark was obtained using a banknote and mm-grade graph paper.

The use of time and computational resources, as a function of the scanned area, is reported in Table 7.

Table 7. Temporal and computational requirements.

Area (cm ²)	Scan time (hrs)	RAM (Gb)	Storage (Gb)	Computation time (hrs)
25	2	3.4	1.4	3.5
91	8	12	5.5	11.5
176	15.3	25	10.5	38.5

We tested the utilisation of multiple, laterally overlapping pyramids in samples of larger areas. The targets were a $21 \times 7 \text{ cm}^2$ piece of industrially woven, patterned silk fabric and a $12 \times 6.2 \text{ cm}^2$ banknote. The results are shown in Figure 53. The pyramids used were X and Y for these cases, respectively. For the fabric, an arrangement of 5×2 pyramids was used; thus, the top layer was a mosaic of 10 images. For the banknote, a 2×3 arrangement was used; thus, the top layer comprised 6 images. In this configuration of RAM and scanning resolutions, the maximum scan size is $42 \times 14 = 588 \text{ cm}^2$.



Figure 53. Larger scans.

A.1.2 Configuration for close-range 2½D surface digitisation

For the reconstruction of a 2.5 cm^2 surface area, 3,348 images were acquired, whose dimensions were $4,000 \times 3,000$ pixels. In these images, $7.52 \cdot 10^7$ key point features were detected. The computation lasted

3.55 hours, and the amounts of RAM utilised were 21 Gb and 3.29 Gb for the CPU and the GPU, respectively. The result consisted of a mesh with 371,589 nodes and 741,885 triangles and a texture map of $16,384 \times 16,384$ pixels. The obtained resolution for the geometry of the reconstruction (mesh nodes) is $\approx 714 \text{ p/mm}^2$ or $\approx 679 \text{ ppi}$, while its texture is $\approx 257 \text{ Kp/mm}^2$ or $\approx 13,000 \text{ dpi}$. These measurements can be verified in the euro coin reconstructions.

Scanning area limitations stem from the memory capacity of the computer. The maximum scanning area achieved is 5cm^2 . The following parameter values are utilised: $\tau_o = 90\%$, $\tau_m = 25\%$, $\tau_p = 50$ pixels, and $\tau_c = 25$ pixels. Moreover, four (4) elevation layers are used. Given that medial overlap is $\tau_m = 25\%$ and that the smallest elevation is at the minimum focus distance of the camera (2 mm), the elevations were 2 mm, 4 mm, 8 mm, and 16 mm.

An example of the scanning dimensions is provided for a $24 \times 28 \text{ cm}^2$ artwork. In Figure 54, the artwork is (left) and is marked with a rectangle, which indicates the scanned area. The four images right to it are original images, one from each elevation layer; starting from top to bottom and shown in left-to-right order, respectively.



Figure 54. Artwork and original images, one from each elevation layer.

A.2 Surface scanning and digitisation apparatus

The proposed approach is implemented using the following materials.

A.2.1 Off-the-net and off-the-shelf components

The visual sensor was an Olympus Tough TG-5, which has a minimum focus distance of 1cm, $4,000 \times 3,000$ px resolution, and a FoV of $16^\circ \times 12^\circ$.

The proposed Cartesian 2D scanner (C2Ds) is a device that is attached next to the printing head of a C3Dp. The C3Dp is not otherwise modified; the attachment can be removed without affecting its operation. The motion mechanism belongs to the C3Dp. This mechanism moves the printing plate laterally, in two dimensions, and the printing head only vertically. The C3Dp is commanded to reach the imaging locations by a microcontroller.

The C2Ds were built on top of an adaptation of the Prusa i3 series C3Dp, which were chosen due to their wide adoption, low cost, and ease of construction. The operating volume is $24.89 \times 21.08 \times 6.86 \text{ cm}^3$. The parts for the C3Dp are cited in A.2.5 below.

The motor is controlled by the Marlin open-source firmware. Marlin is widely used and runs on the cost-efficient 8-bit Atmel AVR microcontrollers. The reference platform for Marlin is the Arduino Mega 2560 with RAMPS 1.4, which is directly compatible with the equipment used for implementing the printer. This

firmware runs on the motherboard and manages real-time controls for heaters, steppers, sensors, lights, LCDs, buttons, etc. The control language is a derivative of G-code.

The power supplies shipped with 3D printers usually generate up to 350 W on a 12 V output. In our implementation, a more robust solution was preferred to accommodate the power requirements of the visual sensor. To this end, a 650 W ATX power supply was used. The electronics and the sensor are connected to the 5 V output, while the stepper motors are connected to the 12 V output.

A.2.2 Motor and motion

The implementation of a C3Dp contains some free variables, such as the quality of materials, the torque of motors, etc. To reduce the effect of vibrations and increase motion accuracy, the device was implemented as follows.

Aluminium frames of 40 × 40 mm thickness were used for the truss. The backside of the print bed was enhanced with an aluminium frame to increase its weight. High-quality, heat-hardened steel rods of 12 mm thickness and high-quality linear bearings were used for the motorised part of the printing bed. Printable components of the apparatus were printed using PET-G and a 60% infill rate to enhance their stiffness and reduce the possibility of heating deformation due to intensive use.

To provide enough torque for this implementation, motors were standard Nema 17-sized high-torque stepper motors. The motor's motion was transmitted via 6mm non-elastic timing belts, integrated with steel threads for enhanced stiffness. Motors are driven by the Texas Instruments DRV8825 Stepper Motor Controller ICs. The controller supports up to 1/32 micro-stepping. The device is operated through a microcontroller built on top of the Arduino Mega 2560. For the wiring of the C2Ds, the RAMPS 1.4 Arduino Mega Pololu Shield was used.

On account of the achieved mechanical robustness, the printing bed was increased by a factor of 4.764 from its specification to 50 × 60 cm². The bed was coated with a 5 mm-thick aluminium sheet to ensure a flat slide for the placement of samples.

A.2.3 Imaging

The camera faces the imaged surface perpendicularly. To mount the camera, a sensor base was designed using the TinkerCad software. The design, shown in Figure 55 (left), was exported in STL format and printed on the C3Dp; see Figure 55 (right). The design of the mount is compatible with the print head and is placed on its backside, thus allowing both heads to be mounted concurrently.




Figure 55. Design of the camera base (left) and photograph from its printing on the C3Dp (right).

The time required to capture the required number of images exceeds the duration of typical consumer-grade batteries, i.e. ≈ 1300 mAh. To avoid interrupting the scan for changing and the consequent sensor displacements, power was continuously provided as follows. A printed case emulating the battery was wired to a power supply of the appropriate voltage and current. The second component in Figure 56 guides and stabilises wiring. The models can be found in the supplementary material.



Figure 56. Designs of a battery emulator and photographs of its implementation.

A.2.4 Communication and control

The scanner control software runs on a personal computer wirelessly connected to the controller. This computer runs software mediating the image acquisition process. The software executes a scanning plan, containing the locations of image acquisition, encoded as 3D coordinates.

Specifically, the print bed and sensor are drawn to a designated relative position due to a G-Code command to the controller. The signal for image acquisition is then sent.

The command sequence is sent to trigger the acquisition of an image. Image acquisition failures and delays are treated as follows. The software checks if the image has indeed been acquired. The command sequence acquires the list of stored images to be compared with a previously collected one. These two sequences are repeated until the image is acquired. Each filename is stored to later conveniently rename the acquired images.

When all the pictures are acquired, the files are manually copied from the memory card of the sensor and automatically renamed by the software. The renaming includes the coordinates of image acquisition in the filename, in the form of Z-Y-X.jpg. These 3D coordinates are only readings of the C3Dp controllers;

they are not regarded as absolute measurements but are combined and fused with visual cues later on in the processing pipeline.

A.2.5 Printed components

Table 8. Printed components.

Name	Source
Z-axis leadscrew	https://www.thingiverse.com/thing:519391
Controller case	https://www.thingiverse.com/thing:2047732
Bowden extruder	https://www.thingiverse.com/thing:2243325
x-carriage	https://www.thingiverse.com/thing:2514659
z-axis	https://www.thingiverse.com/thing:1692666
y-axis belt holder	https://www.thingiverse.com/thing:1030200
y-belt tensioner	https://www.thingiverse.com/thing:3404464
y-axis motor holder	https://www.thingiverse.com/thing:2808408

A.3 Calibration and image acquisition

A.3.1 Calibration

Our sensor is active and composite. Its active part is a 3D CNC that translates the camera to a prescribed location and distance from the target surface. The sensor is a macro camera. Both components are controlled and orchestrated by a PC using the appropriate drivers. The two components of the sensor are calibrated so that each acquired image is associated with the location and distance where it was prescribed to be acquired. The calibration of the sensor for 2D and 2½D scanning is the same.

Camera

The proposed approach requires estimates of intrinsic and extrinsic camera parameters. Intrinsic parameters regard the resolution and optical properties of the visual sensor. Extrinsic parameters regard its position and orientation in space and represent the estimation of camera motion relative to the target surface.

The camera is calibrated to estimate its intrinsic parameters and its lens distortion, using the methods in [46, 47]. The intrinsic parameters represent the location of the optical centre, the skew of the optical axis, and the focal length of the camera. The FoV of the camera is also derived from these parameters. In addition, lens distortion is also estimated in this step. The aforementioned parameters are independent of the camera location and, thus, are estimated once before mounting the camera.

Extrinsic parameters represent the orientation and the location of the camera, as a rotation and translation of the camera with respect to some world coordinate system. An initial estimate of extrinsic camera parameters is provided by the CNC device, specifically from the readings of the stepper motor controllers that move the camera. The area of the scan table and the elevation range are measured to estimate the motor step length per dimension, given that these motors produce equal motion steps. The number of horizontal and vertical steps is denoted as n_x and n_y , while the length of the horizontal and

vertical steps is s_x and s_y , respectively. Step s_x and s_y are defined in metric units. Thereby, the obtained initial estimate of the extrinsic parameters is also in metric units. This initial estimate of camera locations is refined at a later stage of the method.

CNC motion calibration

Motion calibration establishes the mapping from commanded steps to metric displacement in x , y , and z , and verifies repeatability over the scan envelope. We (i) measured effective step length per axis using a known-dimension reference and repeated moves, (ii) verified that commanded returns to the same coordinate yield consistent camera placement within the mechanical tolerance of the platform, and (iii) checked alignment between the scan plane and the camera optical axis to minimise systematic drift during mosaicking. The CNC readings provide an initial estimate of camera extrinsic parameters, which is subsequently refined using image correspondences.

Discussion

The benefit of calibrating the camera with respect to the CNC measurements is that each acquired image is associated with a coarse estimate of its extrinsic parameters. In terms of delivering a “plug and play” device, these calibration steps are done once at manufacturing (or installation) time. As a device, the sensor receives a “scan plan” with image locations and provides a dataset with the image acquired at those locations and appropriately indexed by location.

A.3.2 Image acquisition

A simple utility GUI facilitates the generation of a scan plan given the surface region to be scanned.

Given the sensor’s FoV and the proportions of lateral or medial overlap, camera locations are pre-computed and saved in the scan plan file. When acquiring datasets for surface scanning or digitisation, the planned image acquisition locations are depth-first serialised and converted to scanner coordinates.

A corresponding segment of G-code is accordingly generated. This code is transmitted to the CNC, and images are acquired. A driver program orchestrates the operation in the computer, halting motion when images are acquired and moving the camera across the prescribed locations. In the output dataset, image filenames are associated with the camera locations in the scan plan.

Single image acquisition

At each location, a focus-stacked image is acquired. That is, several photographs are taken with the focus point shifted slightly in each image to cover the entire depth of the subject. Using software, the images are aligned to ensure they overlay correctly, and the in-focus parts of each image are combined to produce a single image with a greater depth of field than any of the individual images. In our case, this operation is embedded in the sensor; however, it can be performed at the computer as well. The acquired images are not compensated for lens distortion at this stage. Due to the focused stacking operation that is applied individually on each image, lens distortion is estimated at a later stage.

Dataset acquisition

Although the CNC apparatus acquires all images in one pass, the acquired images are conceptually classified into layers modulated by imaging distance from the surface of interest. The software interface drives the device to acquire all images and delivers them organised per layer.

Within each layer, viewpoints are organised in a lattice, as illustrated in the example Figure 57 (left). In the example, $n_x = 4$ and $n_y = 5$. Solid lines mark the surface regions imaged by the camera. Dots mark camera centre locations. The horizontal and vertical distances between camera centres are denoted as s_x and s_y , respectively. Transparent regions indicate the lateral overlap between neighbouring images of the same layer. The FoV of the camera and the required amounts of lateral and medial overlaps determine the values of s_x , s_y , and distances between layers. The top layer, indicating the upper end of this frustum, is also user-determined. We use the same amount of lateral overlap for the horizontal and vertical dimensions of each layer, let τ_o . Layer distances are configured so that a “parent” viewpoint images four times the area of its “child” viewpoints; this leads to a quad-tree arrangement of viewpoints and the doubling of the elevation per layer.




Figure 57. Illustrations of CNC motion parameterisation for one layer.

The purpose of lateral overlap is the establishment of lateral (stereo) correspondences across neighbouring images from the same layer. These correspondences are the basis of every photogrammetric surface reconstruction method. The purpose of medial overlap is the anchoring of error accumulation when registering images and reconstruction results from the bottom layers. That is, the proposed approach constrains potential distortions in the final result by requiring that the surface reconstruction is consistent with the appearance of the surface at larger distances.

The difference in configuration for surface scanning and surface digitisation is that for surface scanning and digitisation, 50% and 80% overlap are used, respectively. Naturally, this has direct implications for the number of images our memory can handle and, consequently, the amount of area we can scan or digitise at a given resolution.

A.4 Market survey

All prices are approximate and estimated on the M18 of Craeft. Unreported prices require asking for a quote, and all exhibit a larger price than the others in the same table.

Table 9. Flatbed A0 scanners.

Name	Optical resolution (ppi)	Price (USD)
Kurabo K-IS-A0FW	1000	50K
Microtek LS-4600	600	60K

Table 10. Film scanners.

Name	Optical resolution (ppi)	Price (USD)
Plustek OpticFilm 8100	7200	400
Epson Perfection V550 Photo	12800	600
UScan+ HD LTE	2400	Quote

Table 11. Large format scanners.

Name	Optical resolution (ppi)	Price (USD)
Colortrac SmartLF Sci	1200	5-12K
Colortrac SmartLF SC 42 Xpress	1200	7K
Contex IQ Quattro X - 44"	1200	7K
Contex IQ Quattro 4450 / 4490	1200	7-8K
Image Access WideTEK 48CL	1200	6K
ROWE 850i - 55"	1200	22K
Image Access WideTEK 60CL	1200	12K
CRUSE ST Light 300	300	30K
CRUSE ST Light 600	600	30K
CRUSE Synchron Table (ST)	830	30K
CRUSE CS 85/145 ST-T	600	30K
CRUSE CS 82 ST-T 2450	600	30K

Table 12. Book scanners.

Name	Optical resolution (ppi)	Price (USD)
Suprascan double A0	600	Quote
Suprascan Quartz A0 LED HD	600	Quote
Sma Scanmaster 0 3650	600	Quote
book2net Hornet	400	Quote
Czur ET16	275	355
Fujitsu Scansnap SV600	280	500
SMA ScanMaster 2	1200	10K
SMA RoboScan 2	600	10K
IID Bookeye 5 V3	600	10K
Bookeye 4 V3 Kiosk	600	10K

Bookeye 4 V2 Semiautomatic	600	10K
Bookeye 4 V2 Professional Archive	600	10K
Zeutschel OS Q1	600	10K
Zeutschel OS HQ	1000	10K
Zeutschel OS Q0	600	10K

Table 13. Material scanners.

Name	Optical resolution (ppi)	Price (USD)
Vizoo xTex	1000	Quote
xrite TAC7	385	Quote

Table 14. Scanning services.

Name	Optical resolution (ppi)	Price (USD per sample)
Materialcapture	600	100 – 300
Muravision	920	100 – 300
Overnight scanning	600	100 – 300

Annex B. 2D surface scanning legacy results

B.1 Qualitative

The goal of the experiments was to assess mosaic registration accuracy and to explore tolerance to departures from the assumptions of Lambertian reflectance and surface planarity. Indicative samples were drawn from applications in sensitive materials, from art, biology, document analysis, and textiles. The selected surfaces exhibit variability in their reflectance properties and their 3D surface texture. We included shiny and rough materials in the samples. We did not include highly transparent, highly specular, or highly reflective materials.

Macroscopic images of the surfaces, acquired by a conventional camera, are shown in Figure 58. In the figure, from left to right and top to bottom, samples 1 – 5 are paintings; sample 6 a blank piece of cotton canvas; sample 7 is a scarcely handwritten A4 page; sample 8 is a stamped and signed passport; samples 9 and 10 are blank and printed graph paper, respectively; samples 11-13 are pieces of silk fabric; sample 14 is a leaf; samples 15 and 16 are fine and coarse-grained sand, respectively; sample 17 is an assortment of coins; sample 18 is a banknote. Samples 13 and 18 were scanned entirely. For the rest, a $5 \times 5 \text{ cm}^2$ region was scanned. In Figure 59, $2,048 \times 2,048 \text{ px}$ regions from images of the finest layer are shown, in the same order as in Figure 58.



Figure 58. Scanned surfaces.





Figure 59. Samples from the closest range where photographs were acquired.

Our primary investigation regarded the translational component of the estimated homographies. The estimated homographies lead to image shifts that are always less than 4 pixels. In other words, none of the estimated homographies suggests an update of C_i that would cause an image shift not larger than 4 pixels. In turn, this suggests that no homography estimate is in gross contradiction with the proprioceptive readings. Second, we observe that the method is robust to the occurrences of missing information that were encountered in the experiments. The information missing was either a single pyramid apex in the use of multiple pyramids or registration failures due to a lack of reliable point correspondences. In all cases, a complete mosaic is provided for the layer of the greatest detail.

B.1.1 Paintings

To study paintings, we acquired painted samples on canvas, Canson paper, and regular printing paper. The samples exhibited various degrees of surface roughness. The colours were made from soft pastel or oil. The results are shown in Figure 60. To better investigate the effect of height variations, in the example of the top row, the impasto painting technique [69, p. 100] was utilised. This technique involves painting in overlapping layers and gives rise to surface anomalies. The average height step of these anomalies was 0.75mm. The example is centred upon a 1.5 mm protrusion. In the remaining rows of Figure 60, the examples are sorted in the declining level of surface roughness.






Figure 60. Paintings.

B.1.2 Paper and canvas

A white sheet of paper and a white piece of canvas were utilised to test the method in documents and plain fabrics. Canvas is a textile with a repeated structure, but only macroscopically, as at the imaging resolution, cotton plies are in the range of 12 – 20 μm . The results are shown in Figure 61. Under investigation were potential effects due to blank surface space. No such effects were observed, as the paper, when closely imaged, reveals a rich texture. The same was the case for cotton canvas used for painting. Still, in the third layer from the top, a homography estimate was discarded.




Figure 61. Canvas, handwritten paper, and a passport with a stamp and handwriting.

B.1.3 Repetitive patterns

To test the sensitivity of feature-based image registration to repetitive patterns, conventional graph paper was used. The results are shown in Figure 62. As in the case of the canvas, fine layers reveal sufficient uniqueness cues to abstain from gross misregistration errors. Nevertheless, a failure is observed in the auto-focus function of the sensor. At the second layer from the top, some images were out of focus, possibly due to the sensor’s auto-focus mechanism’s sensitivity to repeated texture. The homography estimate was considered unreliable homographies and discarded. We repeated the experiment this time using printed text, using the ‘Liberation Serif – Regular’ font, at 6 pt. In this condition, the effect did not occur.

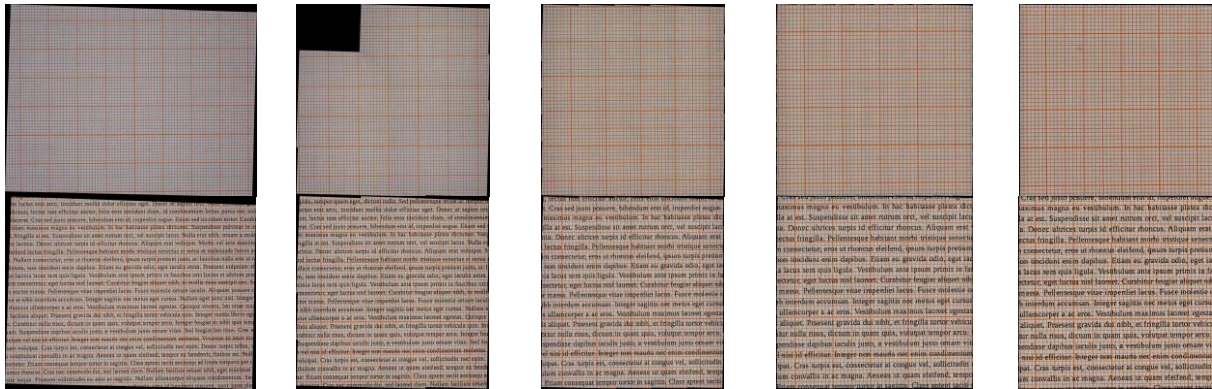


Figure 62. Repetitive patterns.

B.1.4 Fabrics

Though the study of fabrics is related to heritage [70] and industrial applications, ways to digitise textiles and fabrics are constrained in the products reviewed in Appendix C. These data approaches do not scan the fabric in sufficient resolution to reveal the fine crafting of some textiles. We chose patterned silk fabrics handwoven on a Jacquard loom because this type of weaving allows for intricate patterns on the fabric. We chose silk as the most challenging material because its fibrils are sleek, reflecting light from many angles, attributing it to its characteristic sheen. Also, silk is one of the finest plies. Silk fibres from the Bombyx Mori, as in the example, are in the range of 5 – 10 µm (a human hair is ≈ 50 µm). We scanned two samples woven with the same two-colour pattern but with alternating colours. The results are shown in Figure 63.





Figure 63. Jacquard woven silk textile.

B.1.5 Fiona

To image biological tissue, a leaf was scanned. The result is shown in Figure 64.




Figure 64. Leaf.

B.1.6 Approximately planar surfaces

To observe the effects of surface roughness, we scanned two types of sand and an assortment of coins. The results are shown in Figure 65, in increasing order of surface roughness. Coin thickness ranged from 1.67 to 2.33 mm. For the coarse-grained sand, the height steps between adjacent grains exceeded 3 mm.




Figure 65. Approximately planar surfaces.

We did not detect artefacts from fine-grained sand or coins. Though not easily found, the case of coarse-grained sand exhibits some tractable mismatches, as the image combination method cannot compensate for the lack of accurate registration between the overlapping image regions. They are shown in Figure 66, where each image shows a region of $\approx 1 \text{ cm}^2$.



Figure 66. Examples of registration failure.

B.2 Quantitative

Ideally, when viewed at the same resolution, mosaic layers that image the same surface region should be identical. The similarity between mosaic layers was quantified by cross-correlation, in the domain of $[-1, 1]$. We computed this metric between consecutive layers, as well as between the top and bottom layers. In Table 15, we report the correlation values. The first row reports the correlation coefficient between the coarsest and the finest layer. The remaining columns show the correlation values for consecutive layers.

Table 15. Correlation coefficients between mosaic layers.

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
0-4	.5	.3	.1	.2	.6	.6	.5	.6	.9	.9	.2	.5	.6	.3	.7	.8	.3	.5
0-1	.9	.8	.9	.9	.9	.8	.7	.8	.5	1.0	.9	.8	.9	.9	.9	.9	.5	.8
1-2	.8	.7	.7	.9	.9	.4	.6	.9	.5	1.0	.9	.9	.9	.9	.9	.9	.6	.9
2-3	.6	.3	.6	.8	.9	.5	.6	.9	1.0	1.0	.8	.8	.9	.8	.9	.9	.6	.9
3-4	.7	.5	.4	.4	.8	.7	.7	.7	.9	1.0	.2	.7	.8	.7	.8	.8	.7	.9

To measure systematic distortions, we used images of coins, which are circular structures. We performed Canny edge detection [77] in the finest mosaic layer and selected the edges corresponding to the circular creases of the 2c and a 5c coin. The selected edges were used to fit circles, using least squares without a robust selection of inliers. The selected edges are shown on the left image pair of Figure 67. In Table 16, we report deviations of the detected edges from the fitted circle.




Figure 67. Edges belonging to the inner and outer creases of two coins (left) and edge detection details (right).

Table 16. Mean circle fit error and standard deviation.

Circle	Radius (p)	Error (p)
2c outer	3464.8	3.05 (2.62)
2c inner	3216.8	4.09 (3.50)
5c outer	3936.4	4.86 (3.17)
5c inner	3630.8	5.20 (3.36)

This example facilitates observations at the locations of lateral image overlap, where seams are typically observed. In general, due to accurate registration and the effect of the method in [54], seams are usually easily observable in fine scales. In the experiment, the focus distance was automatically adjusted. As the coins are elevated from the background, near the coin boundary, the sensor focuses either on the paper background or the coin. When the focus is placed on the background, the image region where the coin appears is blurry; instead, the background is focused. Though the structure distortion is minute, the difference in the focus of the blended images is observable when image edges are detected, as in the right pair of images in Figure 67. Figure 67. Edges belonging to the inner and outer creases of two coins (left) and edge detection details (right).

In turn, the different amounts of image blur at the boundaries of the blended image give rise to edge dislocations. In Figure 68, a mosaic of 2×2 images is shown for each sample in the experiments (minified for document scale). In the supplementary material, these images can be found in their true resolution. We observe that although no highly frequent seams are observed, a global brightness difference is observed between stitched images.





Figure 68. Mosaics of 2 × 2 original images.

Annex C. Wide-area 2D scans

C.1 Mosaics

This subsection presents results obtained from scanning large traditional textiles as single items.

C.1.1 Gallery

The figure below overviews the results.



Figure 69. Gallery of scanned textiles with hyperlinks to videos that present them individually.




C.1.2 Dataset

The textile mosaics are available in our knowledge base as well as a separate Zenodo dataset: <https://doi.org/10.5281/zenodo.14732270>. Each scan is accompanied by a video that illustrates the result. A playlist of these videos is available at: <https://www.youtube.com/playlist?list=PLQZAktGyYNikhZkYeCgZ23TmeW5Rufrcck>.

The table below overviews this dataset, which is not constrained to the digital assets but represents scanned objects as cultural heritage entities. Per column and from left to right, the figure tabulates the link to digital assets as stored on Zenodo, the link to the knowledge entity representing the digital asset in the Craeft online knowledge base, the link to the knowledge entity representing the cultural object as a heritage craft product, the resolution of the digital asset, and its thumbnail.

Table 17. A dataset of wide-area textiles individually scanned.

File	Asset	Product	Resolution	Thumbnail
set1_1.png	Media Object	BH: 001	6593 × 11214	
set1_2.png	Media Object	BH: 002	6815 × 7584	
set1_3.png	Media Object	BH: 003	8581 × 8540	
set1_4.png	Media Object	BH: 004	6917 × 6880	
set1_5.png	Media Object	BH: 005	13852 × 6594	
set1_6.png	Media Object	BH: 006	9856 × 13789	

set1_7.png	Media Object	BH: 007	7505 × 6767	
set1_8.png	Media Object	BH: 008	6841 × 6127	
set1_9.png	Media Object	BH: 009	8368 × 4512	
set2_1.png	Media Object	BH: 010	12068 × 10630	
set2_2.png	Media Object	BH: 011	11199 × 11003	
set2_3.png	Media Object	BH: 012	8077 × 11620	
set2_4.png	Media Object	BH: 013	15663 × 6944	

[set2_5.png](#)

[Media Object](#)

[BH: 014](#)

14375 × 7857



[set3.png](#)

[Media Object](#)

[BH: 015](#)

27171 × 12582



C.1.3 Foreground segmentation

Results of the foreground segmentation method on wide-area textiles. Shown in the figure below are the individual textiles extracted from the images.







Figure 70. Foreground segmented wide-area textiles from high-resolution mosaics.

C.2 Batch scans

This subsection contains results of batch scans where multiple textiles were concurrently scanned, similarly to the digitisation of multiple analogue photographs on a flatbed scanner and their subsequent segmentation into individual snapshots.

C.2.1 Gallery

The figure below overviews the acquired 2S scans.






Figure 71. Gallery of scanned textile vatches with hyperlinks to videos that present them individually.

C.2.2 Foreground segmentation

Results of the foreground segmentation method on batch textile scans. Shown in the figure below are the individual textiles extracted from the images.







Figure 72. Foreground segmented wide-area textiles from high-resolution mosaics.

Annex D. Garment digitisation

D.1 Acquiring 3D data

This annex trains the operator to acquire data that can be fused reliably. Its purpose is not to teach reconstruction, but to correct field behaviour that secures stable, measurable input.

This guide trains the operator to acquire multi-modal data (photography, static laser scans, handheld structured-light) in a way that preserves geometric stability and maximises fidelity. The operator is primarily responsible for preventing cloth movement and acquiring three coherent datasets for later fusion.

D.1.1 Step 1: Preparation

1. Close all windows and doors. Eliminate air drafts.
2. Remove unnecessary personnel from the room. Movement creates micro-airflow.
3. Turn off any fans, ventilation, or heating nozzles that are aimed directly at the object.
4. Allow the garment to “settle” for at least 2 minutes before starting acquisition.
5. Verify that the mannequin is stable and cannot rock or rotate.
6. Place your fiducials/calibration aids if used.

Operator principle: The garment must remain motionless and maintain its original shape throughout the entire process.

D.1.2 Step 2: Handheld scan

This step always comes first.

1. Walk a smooth orbit around the mannequin with the handheld device.
2. Capture a complete pass (not multiple partial ones).
3. Check that the device reports a consistent scale.

Inspect the preview mesh. Confirm there are no “floating islands” or broken track sections.

Outcome of this stage: You now have a control mesh that defines the reference world coordinate system. This mesh will be used later to check whether the cloth stayed still.

D.1.3 Step 3: Laser scanning

These produce the highest-accuracy geometry.

1. Position the scanner at the planned station number (e.g., around the mannequin at regularly spaced angles).
2. Before pressing 'scan', wait 5-10 seconds. Let any disturbed air settle.

3. Capture the scan.
4. After every scan, align it to the control mesh and check the numerical fit.
5. Rule of thumb: Discard and repeat the scan if the alignment RMS deviates by more than approximately 1-2 mm in rigid mannequin regions. This should be done once the garment has settled again.

Proceed to the next station only after each scan has passed the deviation check.

D.1.4 Step 4: Photogrammetry

These images will produce the final high-detail texture.

1. Fix exposure, white balance, and aperture.
2. Walk a consistent circular path, keeping a constant distance from the mannequin.
3. Maintain ~50% overlap between images.
4. Do not cast shadows onto the garment.

Avoid microflow caused by body motion by keeping your body and clothing clear of the air volume around the dress.





D.1.5 Step 5: Quality-control checklist

Item	
The garment remained still from the beginning to the end.	<input type="checkbox"/>
The handheld scan successfully generated a complete and stable control mesh.	<input type="checkbox"/>
All laser scans aligned to the control mesh within acceptable RMS	<input type="checkbox"/>
No scan was kept if the RMS was too high.	<input type="checkbox"/>
Photography done with fixed exposure and high overlap.	<input type="checkbox"/>

If any of the above are “No”, repeat the relevant stage.

D.2 Cretan textiles collection

The garments are provided as a separate Zenodo dataset: <https://doi.org/10.5281/zenodo.14732270>. List of individual garments and accessories digitised. The assets are also available through our knowledge base, as well as Media Objects.

			
<p>dress_01.glb Media Object Golden silk draped garment</p>	<p>dress_02.glb Media Object Electronic loom</p>	<p>dress_03.glb Media Object Silk garment 1</p>	<p>dress_04.glb Media Object Net, Dress with prints</p>
			
<p>dress_05.glb Media Object Embroidered garment 1</p>	<p>dress_06.glb Media Object The Snake Goddess</p>	<p>dress_07.glb Media Object Garment woven on loom 1</p>	<p>dress_08.glb Media Object Minoan Lily</p>




[dress_09.glb](#)
Media Object
[Circe](#)

[dress_10.glb](#)
Media Object
[Silk garment 2](#)

[model_01.glb](#)
Media Object
[Lanvin by Alber Elbaz](#)

[model_02.glb](#)
Media Object
[Escada](#)



[model_04.glb](#)
Media Object
[model_04](#)

[model_05.glb](#)
Media Object
[Pleats Please](#)

[model_06.glb](#)
Media Object
[Sophia Kokosalaki BH 3](#)

[model_07.glb](#)
Media Object
[Versace](#)




[model_09.glb](#)
[Media Object](#)
[model_09](#)

[model_10.glb](#)
[Media Object](#)
Alberta Ferretti

[model_11.glb](#)
[Media Object](#)
[model_11](#)

[model_12.glb](#)
[Media Object](#)
[model_12](#)



[model_13.glb](#)

[model_15.glb](#)

[model_19.glb](#)

[Forema_01.glb](#)

<p>Media Object Alexander Wang</p>	<p>Media Object Chimaera Dress</p>	<p>Media Object Garment woven on a loom</p>	<p>Media Object Skirt</p>
<p>Forema_02.glb Media Object Vionnet</p>	<p>Forema_02.glb Media Object Carven</p>	<p>Forema_03.glb Media Object Helmut Lang</p>	<p>Forema_04.glb Media Object Sophia Kokosalaki BH 4</p>
<p>Forema_05.glb Media Object Hermès</p>	<p>Forema_06.glb Media Object Halston Heritage 1</p>	<p>Forema_07.glb Media Object Halston Heritage 2</p>	<p>Forema_08.glb Media Object Halston</p>

Annex E. OCT

E.1 Principles

Optical Projection Tomography (OPT) [108, 116], as the name suggests, is a tomographic method whose principle of operation is similar to that of XCT. It reconstructs objects in 3D by measuring light absorption through an object from different points of view, with the difference being that X-ray CT uses the X-ray spectrum of light to perform the reconstruction, and OPT uses the visible spectrum of light.

For visible light to replicate the functionality of X-ray light in X-ray CT, two principles need to be adhered to: the first is that light must be able to pass through the object and be detected on the other end (i.e., not absorbed completely so that the absorption is measurable), and the second is that the light that passes through the object is not ‘bent’ or refracted as it passes through the object [108].

The two conditions above are normally achieved in OPT for biological imaging, where the measured object is very small in size (1–10 mm in thickness) so that visible light can shine through it [108], and by the use of refractive index-matching liquid so that its refractive index (RI) matches that of the environment. No refraction occurs as light passes through [117]. This refractive index matching process is called ‘optical or tissue clearing’ because it has the effect of making the biological tissue less refractive to visible light and easier to photograph clearly.

For light rays in the visible spectrum, glass material used for creating drinking cups found in cultural heritage collections nominally has an RI between ≈ 1.3 and 1.7 [118]. Polyethylene terephthalate (PET), on the other hand, a very common plastic material, has a nominal RI of ≈ 1.58 [118]. Both materials therefore have an RI that is very different to that of the surrounding medium of air ($RI \approx 1$), and therefore we would normally expect considerable refraction when visible light rays propagate through these materials.

However, for very thin-walled, cylindrically-symmetric objects, such as plastic bottles and wine glasses, two effects come into play, which minimize refraction: the first is that the wall thickness is so thin that the refraction effect is expected to be insignificant (Section E.1.1), and the second is the fact that the objects are mostly cylindrical, which means that parallel light rays will experience two, almost opposite, beam shifts as they pass through, ‘self-correcting’ their path (Section E.1.2). Therefore, for these types of objects, it is valid to assume that OPT can be achieved directly in air, without the use of index-matching liquid. We quantify these phenomena via simulation in the following subsections.

E.1.1 Measuring visible light ray beam shift for small sidewall thicknesses

It is assumed in this work that the refraction occurring at any part of a transparent object can be approximated by that of a flat, thin slab of transparent material. This is a valid approximation because the sidewalls are very thin and the surface texture contours are much larger (≈ 1 mm–10 cm) than the wavelength of light used (400–600 nm). Therefore, at the scale of the photon, every light ray striking an object will essentially experience the surface as a thin flat slab of material. To measure the parallel beam-shift effect of a light ray through a thin slab (Figure 73), we use Snell’s law of refraction [87] both on the incoming and also on the outgoing surfaces of each sidewall: $n_1 \sin(\theta_1) = n_2 \sin(\theta_2)$.

Using the flat slab approximation and Snell’s law, it is possible to calculate the amount of parallel beam-shift for various angles of incidence and slab thickness. The calculations were performed for the glass and plastic materials used in the objects we have selected to measure, with an RI of 1.51 for BK7 glass (Figure 74) and 1.58 for PET (Figure 75). As can be seen in Figure 74 and Figure 75, the beam shift increases with the angle of incidence and reaches a maximum that is almost equal to the slab’s thickness at an angle of incidence of 89 degrees from the surface normal.




Figure 73. Effect of parallel beam shift predicted by Snell’s Law of refraction, when a beam is incident on a slab of material of RI n_2 in a medium of RI n_1 .

The systematic analysis, which was carried out, therefore shows that, for the glass and plastic materials that were selected, the beam shift will be smaller than that of the sidewall thickness at any angle of incidence and therefore at any point on the object. For the hollow objects, the sidewalls (0.15–2 mm) are small compared to their diameter (52–80 mm), so refraction is not expected to affect the reconstructed shape significantly. Especially so, when combined with the effect described in Section E.1.2.




Figure 74. Beam shift distance d as the angle of incidence θ_1 increases between 1 and 89° as calculated by Snell’s law, for PET of RI 1.58 and different thicknesses h between 0.2 and 2 mm.




Figure 75. Beam shift distance d as the angle of incidence θ_1 increases between 1 and 89° as calculated by Snell's law, for BK7 glass of RI of 1.51 and different thicknesses h between 0.2 and 2 mm.

E.1.2 Light rays propagating through hollow circular objects

The second effect that makes hollow cylindrical objects especially measurable using OPT is that parallel light rays experience two opposing beam shifts as they propagate through these objects: a first parallel beam shift towards the centre of the object as the ray enters the hollow transparent object, and a second beam shift away from the centre of the object as the ray exits the object. This effect is simulated for multiple parallel beams entering a hollow thin-walled circular disk made of PET with an RI of 1.58.

When the camera is placed sufficiently far away, the rays that reach the lens are those which are mostly initiated in parallel. This can be seen in the performed ray tracing simulations, where it can be seen that parallel beams become slightly convergent after passing through the object. This is advantageous as these rays can be collected by a single camera placed 'far away' from the object, without the need to 'stitch an image,' such as in other large-scale OPT approaches [119]. In our experiment, we placed our camera 70 cm from the object, which, compared to the object diameters of 5–8 cm, is at least a $\approx 9:1$ distance-to-size ratio. The field light source selected was a white light LED panel, and it was also placed as far back as possible, whilst concurrently being able to illuminate the whole object (in our case, 20 cm away). So the total distance between the light source and the camera was 90 cm.

In non-hollow cylindrically-symmetric objects filled with PET, though, the light beams pass through a lot more optical material, which causes larger amounts of refraction. The overall beam paths are therefore changed by a far greater amount, resulting in an intense lensing effect that does not enable easy reconstruction using OPT.

The operational principle is the following. In 2D, the available tomographic configurations are a 'fan-beam' and a 'parallel beam' setup. For 3D scanning, they correspond to 'cone' and 'parallel' setups. A diagram illustrating these configuration modes is shown in Figure 76. In this work, the 'cone' configuration was used by replacing the X-ray source with an optical camera and the X-ray detectors with a field light source (LCD Panel). It is assumed that the camera used can be modelled by the 'pinhole camera model', and therefore, we essentially use the 3D 'cone' beam geometry in reverse.

equipment. Hence, a polished glass ball 80 mm in diameter (also known as a lens ball in scientific photography) was used for calibration. Since the reference ball needs to be visible and has a convex shape, its outer surface was reconstructed by manually selecting its silhouette from each of the 30 images and using the built-in 3D reconstruction method. The sphere reconstruction was then loaded as a reference point cloud processing reference (CloudCompare) and its diameter was measured. The measured diameter was then divided by the ball 80 mm apparent diameter to acquire the object's calibrated scaling factor.

[0042] In the Adobe Tachyon software, there was used the filtered back projection algorithm which is typically used in X-ray CT to reconstruct the measured volume density from the projection data at each rotation angle. The density of each voxel $\rho(x, y, z)$ was calculated per line l at a particular θ angle in the image I as follows:

$$I_{\theta}(l) = \int \rho(x \cos \theta + y \sin \theta) dl \quad (1)$$

[0043] Where θ is the rotation angle's angle, x and y are the particular voxel's x and y coordinates and l is the filtered Fourier transform of the distance image described in Eq. 2.

$$\rho(x, y) = \int F_{\theta}(l) e^{-i(l(x \cos \theta + y \sin \theta))} dl \quad (2)$$

[0044] Where θ is the frequency in the Fourier domain, the spatial dimension of the 3D absorption measurement of each slice (the size of pixels of the projection acquired at each rotation).

[0045] In the conducted experiments, for measuring the effect of the red-film transfer method with a polished end-sphere of the present invention without the use of reference light-making fluid, the following experimental sequence was followed:

1. A rotation stage, 1 camera and a light-guide were setup as in the apparatus illustrated in Figure 4.
2. Light from a led source (LED Panel) was projected through the object and reflected at the camera which was a black and white camera.

objects measured, compared to the plastic objects, there was more intense reflection effects which resulted in appearance of shadow-like distorted shapes (Fig. 10), which should be removed, probably, by the hollow resin, and the reduced light intensity.

[0046] The maximum position required for the specific OPT setup used in the present experiment, in general, was calculated by checking the available camera pixels to the half of area and was fixed to be 0.12 mm per pixel. Besides, the distance displacement error required in the lateral and vertical directions is half the value, i.e. 0.06 mm. The half-pixel shift is the size of the calibration-measurement table (2). The measurement with the least error achieved was an average point cloud distance of 0.020 mm between the OPT and the reference reconstruction for object 1 "large water bottle" (Table 1).

[0047] In the conducted experiments the inverted surface object is particularly complex because of the present invention, with the "resin half" technique which was used during calibration. The result technique may operate in the visible spectra without the use of any special lighting and/or control only the external camera object in use of photographs, which is why it is used to measure the calibration sphere as described further above. For future reference objects however, it is known that the "resin half" technique cannot be used as a reference to build common 3D data around the object. Secondly, it is known that, when using the resin half technique, some camera small errors around the visible surface of the object may appear as "faded" due to the nature of the resin half technique. However, the technique method according to the present invention, which may be considered as a half of solid transparent hemisphere that has been embedded in a dry OPT, does not have these drawbacks as it can reconstruct hollow objects and can also deal with convex surface structures.

[0048] However, the method according to the present invention can also measure internal structures. To demonstrate the present method's ability for measuring features found in the interior of the objects, the reference point cloud of glass bottles was used to build the other, and the reconstructed model, i.e., the 3D representation generated with the method of present invention, is shown in Fig. 9.

[0049] The use of present invention method instead of X-ray CT may result as a dry OPT reconstruction of common objects, also, expensive and present health risk to operators. One of the main drawbacks of conventional OPT is the necessary use of a matching liquid. The present invention however provides a new type of OPT without the need to use any matching liquid. This new type of OPT can be used for successfully reconstructing i.e.,

3. The object was placed in the middle of the rotation stage and was rotated to acquire 30 calibrated views around the object.

4. The images obtained by the lens ball and white camera were located on the same of high absorption glass light and areas of low absorption appear darker.

5. The images were processed using the Adobe Tachyon X-ray CT reconstruction software with a 50% opacity.

6. Transferring the views of the 3D density volume to remove the density of air, what was left was the final density volume of the object.

7. To reconstruct a single surface from the reconstructed density volume, then the object density volume was split into horizontal slices and the slice z , by selecting the peak distance on each row of the image plane, as shown in Fig. 8. The identified points represent the areas with the most dense material and hence those of the silhouette.

8. The peak locations were then calculating the calibrated scaling factor which was calculated using the polished glass ball as measured further above.

[0050] The types of objects measured in the particular experiments are hollow, thin-walled, cylindrical transparent objects, which are common consumer articles and/or light fixtures through them. Plastic objects of the type are lightweight, and are commonly found and used in the consumer industry, e.g., water bottles. These 6 objects measured were a soda bottle and two water bottles which are indicated as a in Fig. 7. Another category of objects consisted of being more transparent optical lighting devices. The present invention is also suitable for measuring these types of objects for which the use of special lighting should be avoided due to the "hollow" of the particular objects. However, measuring these objects may be particularly challenging task, because in these objects the thickness of the material is not constant around the whole object. These objects contain areas where light passing through the object encounters large amounts of material and therefore gets refracted differently, hence, these features of top view. A lot of optical lighting items also contain irregular features around the object, but each individual feature adds the measurement in specific area. A number of thin-walled hollow cylindrical glass objects similar to those found in cultural heritage collections such as wine and liquor bottles, were included in the experiments performed, and are indicated as b in Fig. 6.

measuring the 3D representation of all plastic objects, especially hollow and thin-walled objects. The inventors can thus use type of OPT, "large OPT" or "dry OPT".

[0051] An article from the experiments described further above, the use of OPT can be used for specific, low-cost use. The reconstruction of plastic bottles in the context of a manufacturing quality control scenario, and the reconstruction of glass objects in the context of a cultural heritage object preservation scenario. Representative plastic and glass objects for these cases were selected and reconstructed using the P-PORT.

[0052] It is shown that the plastic bottles measured an average accuracy of 0.34 mm when compared with the setup used. The fact that in the conducted experiments the best point cloud accuracy achieved is 0.020 mm is due to the theoretical precision of the setup 0.12 mm, indicates that with an even more precise setup, a higher accuracy could potentially be achieved.

[0053] For glass objects in the context of cultural heritage on the other hand, which typically have thinner silhouettes, and also typically have some areas with thin optical paths, considerable attention is produced. Hence, the average error rate of the glass objects measured was higher than that of plastic objects, at 0.32 mm.

[0054] The experiments showed that the P-PORT can be used for measuring hollow-walled objects as a replacement for conventional optical photography using bright light and camera, photogrammetry, since it does not require the use of special coating and a high degree of absorption, camera angles and long travel time, and high data acquisition, high degree of reproduction fidelity and accuracy. It is also an unobstructed internal structure, simplifying construction and from laborious.

[0055] Hence, the invention, i.e., the method of the present invention, can be used for dimensional quality assurance of plastic bottles in manufacturing. Also, the conducted experiments showed that the present invention can also be potentially used for the reconstruction of glass cultural heritage objects.

[0056] It is noted that described further above, the present invention is particularly easy to implement and was exemplarily well by the measurement of the hollow objects because in other, there is not significant extension of the light that is used for illumination and when passing through the objects. This can be understood considering the following related to leads of view:

[0057] Before reconstructing the plastic bottles selected for the experiments, described in the below, several test objects were measured. The test and final object features were accurately measured using optical collation, and the respective measurements are shown in Table 1. The second column is that the objects could fit in the half of area of camera used in the setup (the apparent view), which could only measure objects of about 150 mm in height. To measure the hollow areas, i.e., the top view, of the selected glass bottle, some water and paper glasses, the glass objects were calibrated using the standard table with the hollow side used for the setup and being pointing up. Only the hollow parts of the glass items were measured, and the items used had been chosen so that external noise was ignored in this particular experimental study.

Table 1: Thickness measurements of the object's silhouettes.

Object	Average of eight silhouettes thickness measurements measured to object's rim
a. Soda bottle	0.25 mm
b. Water bottle	0.14 mm
c. Large water bottle	0.14 mm
d. Liquor glass 1	2.1 mm
e. Liquor glass 2	2.3 mm
f. Wine glass	0.7 mm

[0058] In order to obtain reference 3D reconstruction results for the outside shapes of the measured items, we scanned the objects with a conventional white light structured light projector scanner that is typically used for industrial purposes. The reference scanner was calibrated in an average of 20 mm using calibration plates which were provided with the scanner. For rapid this purpose, the transparent objects needed to be coated with an opaque grey coating. Then, using the aforementioned point cloud software, both the reference point cloud reconstructions and the point cloud reconstructions created by using the method of the present invention, which can be considered as being a new type of optical projection

of thickness of less than 3 mm hollow transparent object can be approximated to that of the total of transparent matter. This is a valid approximation, because the silhouette of the measured items, we scanned the objects with a conventional white light structured light projector scanner that is typically used for industrial purposes. The reference scanner was calibrated in an average of 20 mm using calibration plates which were provided with the scanner. For rapid this purpose, the transparent objects needed to be coated with an opaque grey coating. Then, using the aforementioned point cloud software, both the reference point cloud reconstructions and the point cloud reconstructions created by using the method of the present invention, which can be considered as being a new type of optical projection

$$n_1 \sin(\theta_1) = n_2 \sin(\theta_2) \quad (3)$$

[0059] Using the first approximation and Snell's law is possible to calculate the amount of glare that occurs for various angles of incidence and side thickness. The calculations were performed for the glass and plastic materials of the objects we have selected to measure, with a refractive index (RI) of 1.51 for a type of glass (Fig. 11) and 1.50 for PET (Polyethylene terephthalate plastic) (Fig. 12). As can be seen in Fig. 11 and Fig. 12, the beam and increases with the angle of incidence and reaches a maximum when it is almost parallel to the side thickness at an angle of incidence of 89 degrees from surface normal.

[0060] Besides, the silhouette analysis which was carried out above, that for the glass and plastic materials which were selected, the amount of glare that occurs from the side thickness at an angle of incidence and therefore at any point on the object. For the hollow objects measured, the available 10-20mm are reconstructed their diameter (D) object, as indicated in not expanded or distorted completely the reconstructed shape.

[0061] The second effect that makes hollow cylindrical objects especially measurable using the OPT according to the present invention, is that parallel light rays experience less scattering after the ray propagation through these objects. A flat parallel beam light towards the center of the object as the ray enters the hollow transparent object, and a beam will pass through the center of the object as the ray exits the object. This is not the case for solid opaque parallel beams entering a hollow thin-walled circular disk made of PET material with a RI of 1.50 as shown in Fig. 4.

[0062] An object's surface with a concave and a convex surface for any part of the object, the ray which mostly incident parallel match the lens. This can be derived from the ray

transparency (OPT) technique, which was aligned to the hand and then aligned more accurately to the camera. Instead of using OPT in an empty laboratory of 10'x10', in order to obtain the dimensional error, the residual point cloud distances were calculated. The point cloud errors were specified as the distances on the OPT point cloud distance in Fig. 6, and the normalized error of the point cloud distance for each object is reported in Table 1. Fig. 8, 7 shows the normalized 3D representation of the measured objects. Hence, in view of the results shown in Fig. 7, it can be considered that the present invention can be used for acquiring 3D representations of the measured objects.

Table 2: Results of the comparison of the point clouds acquired by applying the present invention and reference measurement from an industrial-grade 3D scanner.

Object	Average cloud point distance to reference 3D point cloud
a. Soda bottle	0.22 mm
b. Water bottle	0.14 mm
c. Large water bottle	0.14 mm
d. Liquor glass 1	2.1 mm
e. Liquor glass 2	2.3 mm
f. Wine glass	0.8 mm

[0063] There is multiple known methods of comparing point clouds, using point-to-point, point-to-mesh, mesh-to-mesh strategies. For the specific experiments shown above, the method used for comparing the point clouds was point-to-point. Other than comparing point-to-point to mesh, because of this particular case, the reference point cloud required by the structure factor scanner were extremely dense, and therefore it was not required to create a mesh surface to accurately compare the point clouds.

[0064] For glass objects, the average distance error is 0.020 mm as expected, which is an average error that measured for the plastic objects (0.34 mm) measured in the glass

being structures shown in Fig. 13 from where it can be seen that parallel beams become highly convergent after passing through the hollow object (a in Fig. 13). This is characteristic of Fresnel lens, can be called as single camera placed far away from the object, without the need to shift an image, both in an other large scale conventional OPT applications. In the experiments conducted and described further above, the camera was placed 20 cm from the object which, compared to the measured object diameter of 5.8 cm is at least an 8x distance to use. Thus, the high light source emitted was a white light LED panel with the glass ball, said ball shown in Fig. 10, is used the Snell's law refraction (Eq. 3) both on the incoming and also on the outgoing surface of each material.


[0065] In the context of the present disclosure, the term "silhouette" should be understood as including values very near to those which accompany the aforementioned term. That is to say, a silhouette with reasonable error from an exact value should be accepted, because a slight error in the area of computation that is a deviation from the value indicates a deviation due to measurement uncertainties.

[0066] Additionally, and within the context of the present disclosure, the term "silhouette" should be understood as including values very near to those which accompany the aforementioned term. That is to say, a silhouette with reasonable error from an exact value should be accepted, because a slight error in the area of computation that is a deviation from the value indicates a deviation due to measurement uncertainties.


[0067] Additionally, and within the context of the present disclosure, the term "silhouette" should be understood as including values very near to those which accompany the aforementioned term. That is to say, a silhouette with reasonable error from an exact value should be accepted, because a slight error in the area of computation that is a deviation from the value indicates a deviation due to measurement uncertainties.

[0068] Additionally, and within the context of the present disclosure, the term "silhouette" should be understood as including values very near to those which accompany the aforementioned term. That is to say, a silhouette with reasonable error from an exact value should be accepted, because a slight error in the area of computation that is a deviation from the value indicates a deviation due to measurement uncertainties.


[0069] Additionally, and within the context of the present disclosure, the term "silhouette" should be understood as including values very near to those which accompany the aforementioned term. That is to say, a silhouette with reasonable error from an exact value should be accepted, because a slight error in the area of computation that is a deviation from the value indicates a deviation due to measurement uncertainties.




17




27




37




47




57





67



77



87



97

Figure 77. Patent documents

Annex F. Porcelain material states

This annex defines the eight porcelain body material states (T1–T8) used consistently across the static grids, videos, and interactive viewer in Section 5.4. Each state specifies (a) the Mitsuba BSDF model, (b) the GGX roughness parameter α , where applicable, and (c) the base reflectance colour for the Limoges-type (LIM) and Neutral (NEU) bases.

Replicates. Some states are intentionally repeated as replicates (same BSDF and parameters) to verify that observed differences in appearance are attributable to the material parameters themselves rather than to animation, Monte Carlo sampling variability, or post-processing. Concretely, T3 is replicated by T5, and T6 is replicated by T8.

Table 18. Porcelain material states (T1–T8).

T#	BSDF	α (ggx)	Base RGB (LIM)	Base RGB (NEU)
T1	diffuse	-	0.72, 0.70, 0.68	0.72, 0.70, 0.68
T2	diffuse	-	0.80, 0.79, 0.78	0.78, 0.76, 0.74
T3	roughplastic	0.05	0.85, 0.84, 0.83	0.82, 0.80, 0.78
T4	roughplastic	0.14	0.85, 0.84, 0.83	0.82, 0.80, 0.78
T5	roughplastic	0.05	0.85, 0.84, 0.83	0.82, 0.80, 0.78
T6	roughplastic	0.07	0.85, 0.84, 0.83	0.82, 0.80, 0.78
T7	roughplastic	0.09	0.85, 0.84, 0.83	0.82, 0.80, 0.78
T8	roughplastic	0.07	0.85, 0.84, 0.83	0.82, 0.80, 0.78

Notes.

- The diffuse states (T1–T2) isolate changes in base reflectance without specular/gloss effects.
- The roughplastic states (T3–T8) hold base RGB fixed within each base type and vary α to modulate surface micro-roughness and, consequently, the perceived gloss/clarity of reflections.

Annex G. Spectral measurements

This annex documents the spectral transmittance measurements used in the deliverable as (i) empirical input to material modelling where applicable, and (ii) reference spectra to support controlled rendering studies and cross-material comparisons. Measurements are reported as wavelength-dependent transmittance curves and should be interpreted with the corresponding sample thickness and measurement geometry, as specified in the associated study notes and datasheets.

G.1 CERFAV study

This deliverable includes a set of spectral transmittance measurements acquired in the context of the CERFAV study. The resulting spectra are presented in Figure 78 and serve as the primary measurement evidence referenced in the relevant sections.

Cerfav CRT - VERRE

Analysis/study/research report:
Technical data Colver glass

Recipients:

Cerfav
Craeft Project

For the attention of:

Mr. David Arnaud
david.arnaud@cerfav.fr

Subject:

Report compiling technical data on Colver glass

July 23, 2025

Written by: Sarah Kirack



I – Colver glass

Colver glass was analyzed to determine all of its optical characteristics.




Figure 1: Colver vase a) sandblasted, b) raw




Figure 2: Sample analyzed

II – Spectrophotometric analyses


1) Transmission

Transmission was performed over a range of 200 to 2500 nm. The exact values can be found in the attached Excel document.



2) Reflection

Reflection was measured over a range of 80 to 780 nm. The exact values can be found in the attached Excel document.



3) Color

The color was measured under a D65 illuminant at an angle of 10° over a wavelength range of 250 to 780 nm. This analysis provides the Cielab color coordinates of the glass, from which we were able to obtain the actual color in RGB coordinates.

Actual thickness	L	b	b	Cab (chroma)	Hab (hue angle)	Actual RGB color
3.9	82.6	2.15	8.228	8.504	75.36	216/204/191

III – Refractometry

The analysis was performed three times on three different samples to obtain an average.

Echantillon	Analysis 1	Analysis 2	Analysis 3	Moyenne
1	1.5144	1.5177	1.5153	1.5158
2	1.5153	1.5153	1.5152	1.5153
3	1.5146	1.5154	1.5165	1.5155
				1.5156

The refractive index n is therefore 1.5155.

IV – Density

The analysis was performed on three different pieces of the same sample to obtain an average.

# d'echantillon	masse air (g)	masse eau (g)	Densité mesurée (g/cm ³)
1	0.7675	0.8331	2.4781
2	0.4472	0.0994	2.4573
3	0.4481	0.0881	2.4998
			Moyenne (Densité (g/cm ³))
			(Ecart type (g/cm ³))

The density d is therefore 2.4786 g/cm³.

Figure 78. Spectral measurements.

G.2 Reference samples

The following reference samples are included to cover a range of representative spectral behaviours (bandpass, longpass, and NIR-cut characteristics). For each sample, the “Visual impression” is a qualitative description of the appearance of a typical plate under broadband illumination and depends on thickness and lighting; the “Spectral role” summarises the transmittance behaviour indicated by the manufacturer’s datasheet.

G.2.1 UG11 - “Invisible” UV glass

- Visual impression: Very dark, almost black glass with a slight violet cast.
- Spectral role: Classic UV bandpass filter - passes UV (around 250-380 nm depending on thickness) while strongly blocking most of the visible and NIR. Widely used when you want to “see” UV fluorescence without visible leak-through.
- Datasheet: [UG11 datasheet & downloads \(SCHOTT\)](#)

G.2.2 BG3 - saturated blue-green bandpass

- Visual impression: Strong blue-green/cyan glass; looks very coloured even at modest thickness.
- Spectral role: Bandpass / shortpass type; transmits in the blue-green region while attenuating the red and much of the NIR. Classified as an ionically coloured bandpass filter / shortpass filter, UV-capable in the data sheet.
- Datasheet: [BG3 datasheet & downloads \(SCHOTT\)](#)

Nice for showing a vivid “cool” glass that still has significant UV/blue transmission.

G.2.3 BG38 - pale cyan IR-cut glass

- Visual impression: Pale cyan / slightly greenish; less saturated than BG3 but clearly coloured.
- Spectral role: A bandpass / shortpass NIR-cut filter: high transmission across most of the visible spectrum, with strong absorption in the near-infrared region. This family (BG18, BG38, BG39, etc.) is marketed as NIR-cut glasses with steep IR edges and high visible transmission.
- Datasheet: [BG38 datasheet & downloads \(SCHOTT\)](#)

BG38 is a great “default” IR-cut piece for your renders: physically plausible and visually relatable (it looks like the lightly tinted glass in some camera filters).

G.2.4 VG20 - rich green, visible-pass / NIR-cut

- Visual impression: Deep green glass.
- Spectral role: A green bandpass / NIR-cut filter: high transmission across the visible with very strong absorption in the near-infrared; the data sheet classifies it as an ionically coloured bandpass filter and NIR cut-off glass.
- Datasheet: [VG20 datasheet & downloads \(SCHOTT\)](#)

VG20 pairs nicely with BG38: both cut NIR, but one reads “cool cyan” and the other “pure green”.

G.2.5 GG495 - yellow longpass (green-yellow cut-on)

- Visual impression: Warm yellow glass.
- Spectral role: A longpass filter with 50 % internal transmittance around 495 nm for a 3 mm plate; it blocks much of the blue and passes green-red wavelengths.
- Datasheet: [GG495 datasheet & downloads \(SCHOTT\)](#)

Perceptually, this is your “classic yellow” filter - good for showing how cutting the blue end shifts the perceived colour of transmitted light.

G.2.6 OG530 - orange longpass

- Visual impression: Strong orange glass, especially at 2-3 mm thickness.
- Spectral role: Longpass filter with 50 % point near 530 nm at 3 mm, shifting the cutoff slightly deeper into the green; the datasheet classifies it explicitly as a longpass filter.
- Datasheet: [OG530 datasheet & downloads \(SCHOTT\)](#)

In a rendered lineup, GG495 → OG530 gives you a nice yellow → orange progression with well-defined spectral cut-ons.

G.2.7 RG610 - deep red / visible-to-IR separator



- Visual impression: Very deep red, almost ruby; looks quite dark at modest thickness.
- Spectral role: Longpass filter with cut-on around 610 nm, used to separate visible light from near-infrared - passing red and NIR while strongly attenuating shorter wavelengths.
- Datasheet: [RG610 product page & datasheet link \(SCHOTT\)](#)

This is a “deep red” anchor and a natural endpoint of the visible palette before sliding into IR.

Notes

- Wavelength ranges and cut-on points are **indicative** and should be taken from the datasheet for the **specific thickness** used.
- The reference set is intended to support **comparative** reasoning (e.g., “bandpass vs longpass vs NIR-cut”) rather than to imply a single “correct” spectrum across all thicknesses and measurement conditions.

Annex H. Interactive viewer

This annex documents the interactive HTML viewer used in Section 5.4 to inspect the porcelain body rendering datasets under controlled viewing and illumination conditions. The viewer provides a four-panel layout enabling direct comparison across **base type** (NEU vs LIM) and **motion mode** (Turntable vs Orbit) for the same material state and dataset.

H.1 Top bar elements

Control	Description
<i>Base Path</i>	Path to the renders/ directory containing the image sequences. Normally left as "." when the viewer HTML file is placed inside renders/. If the HTML file is moved, set a relative path to the renders/ folder.
<i>Dataset Menu</i>	Selects the dataset variant to display: CMP (Compare) – plain surfaces for neutral comparison; RNG (Rings) – spheres with gold equatorial bands; PTN (Patterned) – patterned surface texture.
<i>Frame Slider</i>	Scrubs through animation frames interactively. The current frame index is displayed in the adjacent number box.
<i>Keyboard shortcuts</i>	← / → move one frame backwards/forward.

H.2 Four-panel layout and status footer

The viewer shows the same dataset and frame in four panels corresponding to:

- **NEU-TT**: Neutral base, Turntable mode
- **LIM-TT**: Limoges-type base, Turntable mode
- **NEU-CAM**: Neutral base, Orbit mode
- **LIM-CAM**: Limoges-type base, Orbit mode

A status footer (orange) summarises the active dataset (**CMP/RNG/PTN**), the current frame number, and the source folders used for each panel.

H.3 Motion semantics and comparative use

Motion semantics (as used throughout Section 5.4):

- **Turntable (TT)**: camera and illumination are fixed; the object/ring group rotates as a rigid unit.
- **Orbit (CAM)**: the object/ring group is fixed; the camera moves.

Comparative use. The viewer supports controlled comparison by holding the dataset and frame constant while varying:

1. **Base type:** NEU vs LIM (material lineage), and
2. **Motion mode:** Turntable vs Orbit (viewing geometry).

Within any selected dataset (CMP/RNG/PTN), users may compare corresponding panels to assess how the eight material states (T1–T8, arranged as T1–T4 top row and T5–T8 bottom row in the static grids) differ in *appearance cues* such as surface highlight sharpness/spread and the relative visibility of subsurface light transport, without conflating these effects with changes in dataset selection or frame index.

H.1 Top bar elements

Control	Description
Base Path	Normally left as ".", assuming the viewer resides inside the renders/ directory. If you move the HTML file elsewhere, specify the relative path to the renders/ folder.
Dataset Menu	Selects which set of renderings to display: <ul style="list-style-type: none"> • Rings (RNG) - Spheres with gold equatorial bands. • Patterned (PTN) - Patterned surface texture. • Compare (CMP) - Plain surfaces for neutral comparison.
Frame Slider	Runs through animation frames interactively. The frame number is displayed in the number box.
Keyboard Shortcuts	← / → move one frame backwards/forward.

H.2 Footer

At the bottom, an orange status footer summarises:

- Dataset (Rings / Pattern / Compare)
- Current frame number
- Source folders for each of the four panels:
 - NEU-TT (Neutral Turntable)
 - LIM-TT (Limoges Turntable)
 - NEU-CAM (Neutral Orbit)
 - LIM-CAM (Limoges Orbit)

H.3 Comparative purpose

The four views reveal how material lineage and motion mode interact:



D3.2 Advanced digitisation technologies



- The turntable view effectively demonstrates the movement of illumination across a stationary camera, highlighting specular reflections and the glossy quality of the surface.
- The orbit view illustrates camera movement with consistent lighting, highlighting geometric curvature and material thickness.

Comparing corresponding panels allows distinguishing:

- Subtle colour temperature and translucency differences between Neutral and Limoges porcelains.
- Changes in gloss, highlight spread, and internal scattering among the eight material states.

Annex I. NeRFs

We collaborated with a spin-off of FORTH, [SceneOpsis](#), and created a walkthrough of the archaeological site featuring marble works and stone masonry from multiple historical eras. Traditional craftsmanship is seen in the materials, paving techniques, and construction methods. An immersive digital experience of the site highlights the potential of scenes using modest digitisation means (i.e., a mobile device) to capture the “sense” of a location and workspace. This experiment offers a straightforward method for practitioners to record their created scenes and for ethnographers to document their research environments.

The top row of Figure 79 illustrates an interactive walkthrough of a marble and stone archaeological site on a PC, captured using a mobile phone. The bottom row presents the same content viewed on a mobile phone in AR mode, virtually integrating the site and allowing for physical navigation within the virtual environment. The two videos can be found on [Venizelou Walk](#) and [Venizelou 3D](#), respectively.




Figure 79. Interactive walkthrough of an archaeological site. Top row: PC-based navigation of marble and stone masonry. Bottom row: Mobile AR mode allowing for physical navigation within the virtual environment.

References

1. B. Schmitt, C. Zirbes, C. Bonin, D. Lohmann, D. Lencina, and A. Netto, "A comparative study of Cartesian and Delta 3D printers on producing PLA parts," *Materials Research*, vol. 20, pp. 883–886, 2017. doi: [10.1590/1980-5373-MR-2016-1039](https://doi.org/10.1590/1980-5373-MR-2016-1039).
2. M. Brown and D. Lowe, "Automatic panoramic image stitching using invariant features," *International Journal of Computer Vision*, vol. 74, no. 1, pp. 59–73, 2007. doi: [10.1007/s11263-006-0002-3](https://doi.org/10.1007/s11263-006-0002-3).
3. Z. Zhu, J. Fu, J. Yang, and X. Zhang, "Panoramic image stitching for arbitrarily shaped tunnel lining inspection," *Computer-Aided Civil and Infrastructure Engineering*, vol. 31, no. 12, pp. 936–953, 2016. doi: [10.1111/mice.12220](https://doi.org/10.1111/mice.12220).
4. R. Carroll and S. Seitz, "Rectified surface mosaics," *International Journal of Computer Vision*, vol. 85, no. 3, pp. 307–315, 2009. doi: [10.1007/s11263-009-0238-z](https://doi.org/10.1007/s11263-009-0238-z).
5. C. Hernandez-Matas, X. Zabulis, A. Triantafyllou, P. Anyfanti, and A. Argyros, "Retinal image registration under the assumption of a spherical eye," *Computerized Medical Imaging and Graphics*, vol. 55, pp. 95–105, 2017. doi: [10.1016/j.compmedimag.2016.06.006](https://doi.org/10.1016/j.compmedimag.2016.06.006).
6. V. Eremeev, A. Kuznetsov, G. Myatov, O. Presnyakov, V. Poshekhonov, and P. Svetelkin, "Image structure restoration from Sputnik with multi-matrix scanners," in *Image and Signal Processing for Remote Sensing XX*, vol. 9244, SPIE, 2014, p. 92440F. doi: [10.1117/12.2066631](https://doi.org/10.1117/12.2066631).
7. M. Pilu and F. Isgro, "A fast and reliable planar registration method with applications to document stitching," in *Proceedings of the British Machine Vision Conference*, 2002, pp. 67.1–67.10. doi: [10.5244/C.16.67](https://doi.org/10.5244/C.16.67).
8. C. Blasse, S. Saalfeld, R. Etournay, A. Sagner, S. Eaton, and E. Myers, "PreMosa: extracting 2D surfaces from 3D microscopy mosaics," *Bioinformatics*, vol. 33, no. 16, pp. 2563–2569, 2017. doi: [10.1093/bioinformatics/btx195](https://doi.org/10.1093/bioinformatics/btx195).
9. T. Levanic, "ATRICS - a new system for image acquisition in dendrochronology," *Tree-Ring Research*, vol. 63, no. 2, pp. 117–122, 2009. doi: [10.3959/1536-1098-63.2.117](https://doi.org/10.3959/1536-1098-63.2.117).
10. D. Conover, J. Delaney, and M. Loew, "Automatic registration and mosaicking of technical images of old master paintings," *Applied Physics A*, vol. 119, pp. 1567–1575, 2015. doi: [10.1007/s00339-015-9095-5](https://doi.org/10.1007/s00339-015-9095-5).
11. Z. Xiong and Y. Zhang, "A critical review of image registration methods," *International Journal of Image and Data Fusion*, vol. 1, no. 2, pp. 137–158, 2010. doi: [10.1080/19479831003802790](https://doi.org/10.1080/19479831003802790).
12. B. Zitová and J. Flusser, "Image registration methods: a survey," *Image and Vision Computing*, vol. 21, no. 11, pp. 977–1000, 2003. doi: [10.1016/S0262-8856\(03\)00137-9](https://doi.org/10.1016/S0262-8856(03)00137-9).
13. R. Szeliski, *Computer Vision: Algorithms and Applications*. Springer, 2011. doi: [10.1007/978-1-84882-935-0](https://doi.org/10.1007/978-1-84882-935-0).
14. R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, 2nd ed. Cambridge University Press, 2003. doi: [10.1017/CBO9780511811685](https://doi.org/10.1017/CBO9780511811685).
15. A. Salvi, "Your smartphone is a material scanner," R&D Blog, 2017.
16. B. Illerhaus, E. Jasiuniene, J. Goebbels, and P. Loethman, "Investigation and image processing of cellular metals with highly resolving 3D microtomography (uCT)," in *Developments in X-Ray Tomography III*, vol. 4503, SPIE, 2002, pp. 201–204. doi: [10.1117/12.452936](https://doi.org/10.1117/12.452936).
17. K. Semendeferi, H. Damasio, R. Frank, and G. Van Hoesen, "The evolution of the frontal lobes: a volumetric analysis based on three-dimensional reconstructions of magnetic resonance scans of human and ape brains," *Journal of Human Evolution*, vol. 32, no. 4, pp. 375–388, 1997. doi: [10.1006/jhev.1996.0099](https://doi.org/10.1006/jhev.1996.0099).

18. T. Luhmann, "Close-range photogrammetry for industrial applications," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 65, no. 6, pp. 558–569, 2010. doi: [10.1016/j.isprsjprs.2010.06.003](https://doi.org/10.1016/j.isprsjprs.2010.06.003).
19. M. Rodríguez-Martín and P. Rodríguez-Gonzálvez, "Learning based on 3D photogrammetry models to evaluate the competencies in visual testing of welds," in *IEEE Global Engineering Education Conference (EDUCON)*, 2018, pp. 1576–1581. doi: [10.1109/EDUCON.2018.8363428](https://doi.org/10.1109/EDUCON.2018.8363428).
20. M. Rodríguez-Martín and P. Rodríguez-Gonzálvez, "Learning methodology based on weld virtual models in the mechanical engineering classroom," *Computer Applications in Engineering Education*, vol. 27, no. 5, pp. 1113–1125, 2019. doi: [10.1002/cae.22141](https://doi.org/10.1002/cae.22141).
21. M. Fau, R. Cornette, and A. Houssaye, "Photogrammetry for 3D digitizing bones of mounted skeletons: Potential and limits," *Comptes Rendus Palevol*, vol. 15, no. 8, pp. 968–977, 2016. doi: [10.1016/j.crvp.2016.08.003](https://doi.org/10.1016/j.crvp.2016.08.003).
22. D. Gajski, A. Solter, and M. Gašparović, "Applications of macro photogrammetry in archaeology," *ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. XLI-B5, pp. 263–266, 2016. doi: [10.5194/isprs-archives-XLI-B5-263-2016](https://doi.org/10.5194/isprs-archives-XLI-B5-263-2016).
23. S. Marziali and G. Dionisio, "Photogrammetry and macro photography: The experience of the MUSINT II project in the 3D digitizing process of small-size archaeological artifacts," *Studies in Digital Heritage*, vol. 1, no. 2, pp. 298–309, 2017. doi: [10.14434/sdh.v1i2.23351](https://doi.org/10.14434/sdh.v1i2.23351).
24. B. Hassett and T. Lewis-Bale, "Comparison of 3D landmark and 3D dense cloud approaches to hominin mandible morphometrics using structure-from-motion," *Archaeometry*, vol. 59, no. 1, pp. 191–203, 2017. doi: [10.1111/arcm.12242](https://doi.org/10.1111/arcm.12242).
25. L. Inzerillo, "Smart SfM: Salinas Archaeological Museum," *International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. 42, pp. 369–374, 2017. doi: [10.5194/isprs-archives-XLII-2-W5-369-2017](https://doi.org/10.5194/isprs-archives-XLII-2-W5-369-2017).
26. J. Fernández-Lozano, G. Gutiérrez-Alonso, M. Á. Ruiz-Tejada, and M. Criado-Valdés, "3D digital documentation and image enhancement integration into schematic rock art analysis and preservation: The Castrocontrigo Neolithic rock art (NW Spain)," *Journal of Cultural Heritage*, vol. 26, pp. 160–166, 2017. doi: [10.1016/j.culher.2017.02.003](https://doi.org/10.1016/j.culher.2017.02.003).
27. P. Lussu and E. Marini, "Ultra close-range digital photogrammetry in skeletal anthropology: A systematic review," *PLOS ONE*, vol. 15, no. 4, p. e0230948, 2020. doi: [10.1371/journal.pone.0230948](https://doi.org/10.1371/journal.pone.0230948).
28. F. Lavecchia, M. Guerra, and L. Galantucci, "The influence of software algorithms on photogrammetric micro-feature measurement's uncertainty," *International Journal of Advanced Manufacturing Technology*, vol. 93, pp. 3991–4005, 2017. doi: [10.1007/s00170-017-0842-z](https://doi.org/10.1007/s00170-017-0842-z).
29. A. Gallo, M. Muzzupappa, and F. Bruno, "3D reconstruction of small sized objects from a sequence of multi-focused images," *Journal of Cultural Heritage*, vol. 15, no. 2, pp. 173–182, 2014. doi: [10.1016/j.culher.2013.04.001](https://doi.org/10.1016/j.culher.2013.04.001).
30. M. González, J. Yravedra, D. González-Aguilera, J. Palomeque-González, and M. Domínguez-Rodrigo, "Micro-photogrammetric characterization of cut marks on bones," *Journal of Archaeological Science*, vol. 62, pp. 128–142, 2015. doi: [10.1016/j.jas.2015.08.004](https://doi.org/10.1016/j.jas.2015.08.004).
31. G. Percoco and A. Salmerón, "Photogrammetric measurement of 3D freeform millimetre-sized objects with micro features: an experimental validation of the close-range camera calibration model for narrow angles of view," *Measurement Science and Technology*, vol. 26, no. 9, p. 095203, 2015. doi: [10.1088/0957-0233/26/9/095203](https://doi.org/10.1088/0957-0233/26/9/095203).
32. M. Shortis, C. Bellman, S. Robson, G. Johnston, and G. Johnson, "Stability of zoom and fixed lenses used with digital SLR cameras," in *ISPRS Symposium of Image Engineering and Vision Metrology*, 2006, pp. 285–290.

33. L. Galantucci, M. Guerra, and F. Lavecchia, "Photogrammetry applied to small and micro-scaled objects: A review," in *International Conference on the Industry 4.0 Model for Advanced Manufacturing*, Springer, 2018, pp. 57–77. doi: [10.1007/978-3-319-89563-5_5](https://doi.org/10.1007/978-3-319-89563-5_5).
34. L. Galantucci, F. Lavecchia, and G. Percoco, "Multistack close-range photogrammetry for low-cost submillimeter metrology," *Journal of Computing and Information Science in Engineering*, vol. 13, no. 4, 2013. doi: [10.1115/1.4025470](https://doi.org/10.1115/1.4025470).
35. A. Mathys and J. Brecko, "Focus Stacking," in *Handbook of best practice and standards for 2D+ and 3D imaging of natural history collections*, Amsterdam University Press, 2018, pp. 213–216.
36. L. Galantucci, M. Pesce, and F. Lavecchia, "A stereo photogrammetry scanning methodology for precise and accurate 3D digitization of small parts with sub-millimeter-sized features," *CIRP Annals*, vol. 64, no. 1, pp. 507–510, 2015. doi: [10.1016/j.cirp.2015.04.144](https://doi.org/10.1016/j.cirp.2015.04.144).
37. G. Percoco, M. Guerra, S. Sanchez, and L. M. Galantucci, "Experimental investigation on camera calibration for 3D photogrammetric scanning of micro-features for micrometric resolution," *International Journal of Advanced Manufacturing Technology*, vol. 91, pp. 2935–2947, 2017. doi: [10.1007/s00170-016-9964-8](https://doi.org/10.1007/s00170-016-9964-8).
38. G. Percoco, F. Modica, and S. Fanelli, "Image analysis for 3D micro-features: A new hybrid measurement method," *Precision Engineering*, vol. 48, pp. 123–132, 2017. doi: [10.1016/j.precisioneng.2016.11.011](https://doi.org/10.1016/j.precisioneng.2016.11.011).
39. D. Sims-Waterhouse, S. Piano, and R. Leach, "Verification of micro-scale photogrammetry for smooth three-dimensional object measurement," *Measurement Science and Technology*, vol. 28, no. 5, p. 055010, 2017. doi: [10.1088/1361-6501/aa6012](https://doi.org/10.1088/1361-6501/aa6012).
40. D. Sims-Waterhouse, P. Bointon, S. Piano, and R. Leach, "Experimental comparison of photogrammetry for additive manufactured parts with and without laser speckle projection," in *Optical Measurement Systems for Industrial Inspection X*, vol. 10329, SPIE, 2017, p. 103290W. doi: [10.1117/12.2270222](https://doi.org/10.1117/12.2270222).
41. F. Lavecchia, M. Guerra, and L. Galantucci, "Performance verification of a photogrammetric scanning system for micro-parts using a three-dimensional artifact: adjustment and calibration," *International Journal of Advanced Manufacturing Technology*, vol. 96, pp. 4267–4279, 2018. doi: [10.1007/s00170-018-1886-4](https://doi.org/10.1007/s00170-018-1886-4).
42. L. Galantucci, M. Pesce, and F. Lavecchia, "A powerful scanning methodology for 3D measurements of small parts with complex surfaces and sub-millimeter-sized features, based on close-range photogrammetry," *Precision Engineering*, vol. 43, pp. 211–219, 2016. doi: [10.1016/j.precisioneng.2015.08.004](https://doi.org/10.1016/j.precisioneng.2015.08.004).
43. P. Besl and N. D. McKay, "A method for registration of 3-D shapes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, no. 2, pp. 239–256, 1992. doi: [10.1109/34.121791](https://doi.org/10.1109/34.121791).
44. C. Yang and G. Medioni, "Object modelling by registration of multiple range images," *Image and Vision Computing*, vol. 10, no. 3, pp. 145–155, 1992. doi: [10.1016/0262-8856\(92\)90066-C](https://doi.org/10.1016/0262-8856(92)90066-C).
45. M. Lourakis and X. Zabulis, "Accurate scale factor estimation in 3D reconstruction," in *Computer Analysis of Images and Patterns*, Springer, 2013, pp. 498–506. doi: [10.1007/978-3-642-40246-3_62](https://doi.org/10.1007/978-3-642-40246-3_62).
46. Z. Zhang, "Flexible camera calibration by viewing a plane from unknown orientations," in *Proceedings of the Seventh IEEE International Conference on Computer Vision*, vol. 1, 1999, pp. 666–673. doi: [10.1109/ICCV.1999.791289](https://doi.org/10.1109/ICCV.1999.791289).
47. J. Heikkila and O. Silven, "A four-step camera calibration procedure with implicit image correction," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 1997, pp. 1106–1112. doi: [10.1109/CVPR.1997.609468](https://doi.org/10.1109/CVPR.1997.609468).
48. D. Lowe, "Distinctive image features from scale-invariant key points," *International Journal of Computer Vision*, vol. 60, no. 2, pp. 91–110, 2004. doi: [10.1023/B:VISI.0000029664.99615.94](https://doi.org/10.1023/B:VISI.0000029664.99615.94).

49. J. Cheng, C. Leng, J. Wu, H. Cui, and H. Lu, "Fast and accurate image matching with cascade hashing for 3D reconstruction," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 1–8. doi: [10.1109/CVPR.2014.266](https://doi.org/10.1109/CVPR.2014.266).
50. N. Ayache, *Artificial Vision for Mobile Robots: Stereo Vision and Multisensory Perception*. MIT Press, 1991.
51. P. Fua, "A parallel stereo algorithm that produces dense depth maps and preserves image features," *Machine Vision and Applications*, vol. 6, no. 1, pp. 35–49, 1993. doi: [10.1007/BF01211638](https://doi.org/10.1007/BF01211638).
52. M. Fischler and R. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, 1981. doi: [10.1145/358669.358692](https://doi.org/10.1145/358669.358692).
53. R. Kuemmerle, G. Grisetti, H. Strasdat, K. Konolige, and W. Burgard, "g2o: A general framework for graph optimization," in *IEEE International Conference on Robotics and Automation*, 2011, pp. 3607–3613. doi: [10.1109/ICRA.2011.5980334](https://doi.org/10.1109/ICRA.2011.5980334).
54. P. Burt and E. Adelson, "A multiresolution spline with application to image mosaics," *ACM Transactions on Graphics*, vol. 2, no. 4, pp. 217–236, 1983. doi: [10.1145/245.247](https://doi.org/10.1145/245.247).
55. G. Csurka, C. Zeller, Z. Zhang, and O. Faugeras, "Characterizing the uncertainty of the fundamental matrix," *Computer Vision and Image Understanding*, vol. 68, no. 1, pp. 18–36, 1997. doi: [10.1006/cviu.1997.0537](https://doi.org/10.1006/cviu.1997.0537).
56. P. Moulon, P. Monasse, R. Perrot, and R. Marlet, "OpenMVG: Open multiple view geometry," in *International Workshop on Reproducible Research in Pattern Recognition*, Springer, 2016, pp. 60–74. doi: [10.1007/978-3-319-56414-2_5](https://doi.org/10.1007/978-3-319-56414-2_5).
57. P. Moulon, P. Monasse, and R. Marlet, "Adaptive structure from motion with a contrario model estimation," in *Asian Conference in Computer Vision*, Springer, 2013, pp. 257–270. doi: [10.1007/978-3-642-37447-0_20](https://doi.org/10.1007/978-3-642-37447-0_20).
58. F. Espuny, P. Monasse, and L. Moisan, "A new a contrario approach for the robust determination of the fundamental matrix," in *Image and Video Technology – PSIVT 2013 Workshops*, Springer, 2014, pp. 181–192. doi: [10.1007/978-3-642-53926-8_17](https://doi.org/10.1007/978-3-642-53926-8_17).
59. M. Lourakis and A. Argyros, "SBA: A software package for generic sparse bundle adjustment," *ACM Transactions on Mathematical Software*, vol. 36, no. 1, pp. 1–30, 2009. doi: [10.1145/1486525.1486527](https://doi.org/10.1145/1486525.1486527).
60. D. Cernea, "OpenMVS: Multi-view stereo reconstruction library," 2021. [Online]. Available: <https://cdcseacave.github.io/openMVS>.
61. E. Catmull, "A subdivision algorithm for computer display of curved surfaces," PhD dissertation, University of Utah, 1974.
62. C. Barnes, E. Shechtman, A. Finkelstein, and D. B. Goldman, "PatchMatch: A randomized correspondence algorithm for structural image editing," *ACM Transactions on Graphics*, vol. 28, no. 3, 2009. doi: [10.1145/1531326.1531330](https://doi.org/10.1145/1531326.1531330).
63. M. Jancosek and T. Pajdla, "Exploiting visibility information in surface reconstruction to preserve weakly supported surfaces," *International Scholarly Research Notices*, vol. 2014, 2014. doi: [10.1155/2014/757269](https://doi.org/10.1155/2014/757269).
64. V. Hiep, P. Labatut, J. Pons, and R. Keriven, "High accuracy and visibility-consistent dense multiview stereo," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 34, no. 5, pp. 889–901, 2011. doi: [10.1109/TPAMI.2011.172](https://doi.org/10.1109/TPAMI.2011.172).
65. R. Gal, Y. Wexler, E. Ofek, H. Hoppe, and D. Cohen-Or, "Seamless montage for texturing models," *Computer Graphics Forum*, vol. 29, no. 2, pp. 479–486, 2010. doi: [10.1111/j.1467-8659.2009.01617.x](https://doi.org/10.1111/j.1467-8659.2009.01617.x).

66. B. Wang, P. Pan, Q. Xiao, L. Luo, X. Ren, R. Jin, and X. Jin, "Seamless color mapping for 3D reconstruction with consumer-grade scanning devices," in *European Conference on Computer Vision Workshops*, 2018, pp. 633–648. doi: [10.1007/978-3-030-11024-6_43](https://doi.org/10.1007/978-3-030-11024-6_43).
67. M. Waechter, N. Moehrle, and M. Goesele, "Let there be color! Large-scale texturing of 3D reconstructions," in *European Conference on Computer Vision*, Springer, 2014, pp. 836–850. doi: [10.1007/978-3-319-10602-1_54](https://doi.org/10.1007/978-3-319-10602-1_54).
68. J. Jylänki, "A thousand ways to pack the bin - a practical approach to two-dimensional rectangle bin packing," 2010.
69. A. Weyer, P. Roig, D. Pop, J. Cassar, A. Ozkose, V. Jean-Marc, and I. Srsa, Eds., *EwaGlos - European Illustrated Glossary of Conservation Terms for Wall Paintings and Architectural Surfaces*. Michael Imhof Verlag, 2011. doi: [10.5165/hawk-hhg/233](https://doi.org/10.5165/hawk-hhg/233).
70. M. Pérez et al., "From historical silk fabrics to their interactive virtual representation and 3D printing," *Sustainability*, vol. 12, no. 18, 2020. doi: [10.3390/su12187533](https://doi.org/10.3390/su12187533).
71. X. Zabulis, P. Koutlemanis, N. Stivaktakis, and N. Partarakis, "A low-cost contactless overhead micrometer surface scanner," *Applied Sciences*, vol. 11, no. 14, p. 6274, 2021. doi: [10.3390/app11146274](https://doi.org/10.3390/app11146274).
72. B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis, "3D Gaussian splatting for real-time radiance field rendering," *ACM Transactions on Graphics*, vol. 42, no. 4, 2023. doi: [10.1145/3592433](https://doi.org/10.1145/3592433).
73. D. Marr and T. Poggio, "Cooperative computation of stereo disparity," *Science*, vol. 194, no. 4262, pp. 283–287, 1976. doi: [10.1126/science.968482](https://doi.org/10.1126/science.968482).
74. A. Karami, F. Menna, and F. Remondino, "Combining photogrammetry and photometric stereo to achieve precise and complete 3D reconstruction," *Sensors*, vol. 22, no. 21, p. 8172, 2022. doi: [10.3390/s22218172](https://doi.org/10.3390/s22218172).
75. A. Petruccioli, F. Gherardini, and F. Leali, "Assessment of close-range photogrammetry for the low-cost development of 3D models of car bodywork components," *International Journal on Interactive Design and Manufacturing (IJDeM)*, vol. 16, no. 2, pp. 703–713, 2022. doi: [10.1007/s12008-022-00862-y](https://doi.org/10.1007/s12008-022-00862-y).
76. C. Nicolae, E. Nocerino, F. Menna, and F. Remondino, "Photogrammetry applied to problematic artefacts," *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. 40, pp. 451–456, 2014. doi: [10.5194/isprsarchives-XL-5-451-2014](https://doi.org/10.5194/isprsarchives-XL-5-451-2014).
77. J. Canny, "A computational approach to edge-detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 8, no. 6, pp. 679–698, 1986. doi: [10.1109/TPAMI.1986.4767851](https://doi.org/10.1109/TPAMI.1986.4767851).
78. F. Hausdorff, *Grundzüge der Mengenlehre*. Leipzig: Viet, 1914.
79. TetraVision, "3D scan of a 2 euro coin," 2024. [Online]. Available: <https://www.tetraction.be/portfolio-item/3d-scan-of-a-2-euro-coin/>.
80. X. Zhou, "Multiple auto-adapting color balancing for a large number of images," *ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. XL-7/W3, pp. 735–742, 2015. doi: [10.5194/isprsarchives-XL-7-W3-735-2015](https://doi.org/10.5194/isprsarchives-XL-7-W3-735-2015).
81. Y. Shin, M. Park, Y. Jeon, Y. Moon, S. Lee, and K. Yoon, "Tone correction with dynamic objects for seamless image mosaic," in *Trends and Topics in Computer Vision*, Springer, 2012, pp. 104–117. doi: [10.1007/978-3-642-35740-4_8](https://doi.org/10.1007/978-3-642-35740-4_8).
82. A. Pentland, "A new sense for depth of field," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 9, no. 4, pp. 523–531, 1987. doi: [10.1109/TPAMI.1987.4767940](https://doi.org/10.1109/TPAMI.1987.4767940).
83. P. Grossmann, "Depth from focus," *Pattern Recognition Letters*, vol. 5, no. 1, pp. 63–69, 1987. doi: [10.1016/0167-8655\(87\)90028-2](https://doi.org/10.1016/0167-8655(87)90028-2).
84. I. Ihrke, K. N. Kutulakos, H. P. A. Lensch, M. Magnor, and W. Heidrich, "Transparent and specular object reconstruction," *Computer Graphics Forum*, vol. 29, no. 8, pp. 2400–2426, 2010. doi: [10.1111/j.1467-8659.2010.01753.x](https://doi.org/10.1111/j.1467-8659.2010.01753.x).

85. F. Meriaudeau, R. Rantoson, K. M. Adal, D. Fofi, and C. Stolz, "Non-conventional imaging systems for 3D digitization of transparent objects: Shape from polarization in the IR and shape from visible fluorescence induced UV," in *3rd International Topical Meeting on Optical Sensing and Artificial Vision: OSAV'2012*, 2013, pp. 34–40.
86. A. Karami, R. Battisti, F. Menna, and F. Remondino, "3D digitization of transparent and glass surfaces: State of the art and analysis of some methods," *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. XLIII-B2-2022, pp. 695–702, 2022. doi: [10.5194/isprs-archives-XLIII-B2-2022-695-2022](https://doi.org/10.5194/isprs-archives-XLIII-B2-2022-695-2022).
87. M. Born and E. Wolf, *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light*, 7th ed. Cambridge University Press, 1999. doi: [10.1017/CBO9781139644181](https://doi.org/10.1017/CBO9781139644181).
88. P. I. Stavroulakis and R. K. Leach, "Invited review article: Review of post-process optical form metrology for industrial-grade metal additive manufactured parts," *Review of Scientific Instruments*, vol. 87, p. 041101, 2016. doi: [10.1063/1.4944968](https://doi.org/10.1063/1.4944968).
89. "Influence of Coating Spray on Surface Measurement Using 3D Optical Scanning Systems," in *ASME 2019 International Manufacturing Science and Engineering Conference*, vol. 1, 2019, p. V001T02A009. doi: [10.1115/MSEC2019-2882](https://doi.org/10.1115/MSEC2019-2882).
90. B. Trifonov, D. Bradley, and W. Heidrich, "Tomographic reconstruction of transparent objects," in *ACM SIGGRAPH 2006 Sketches*, 2006, p. 55. doi: [10.1145/1179849.1179919](https://doi.org/10.1145/1179849.1179919).
91. H. Guo, H. Zhou, and P. P. Banerjee, "Use of structured light in 3D reconstruction of transparent objects," *Applied Optics*, vol. 61, no. 5, pp. B314–B324, 2022. doi: [10.1364/AO.445214](https://doi.org/10.1364/AO.445214).
92. M. Landmann, H. Speck, P. Dietrich, S. Heist, P. Kühmstedt, and G. Notni, "Fast 3D shape measurement of transparent glasses by sequential thermal fringe projection," *EPJ Web of Conferences*, vol. 238, p. 06008, 2020. doi: [10.1051/epjconf/202023806008](https://doi.org/10.1051/epjconf/202023806008).
93. F. Mériaudeau, L. A. S. Secades, G. Eren, A. Ercil, F. Truchetet, O. Aubreton, and D. Fofi, "3-D scanning of nonopaque objects by means of imaging emitted structured infrared patterns," *IEEE Transactions on Instrumentation and Measurement*, vol. 59, no. 11, pp. 2898–2906, 2010. doi: [10.1109/TIM.2010.2046366](https://doi.org/10.1109/TIM.2010.2046366).
94. H. Huang et al., "Free field of view infrared digital holography for mineral crystallization," *Crystal Growth & Design*, vol. 23, pp. 7992–8008, 2023. doi: [10.1021/acs.cgd.3c00780](https://doi.org/10.1021/acs.cgd.3c00780).
95. R. Rantoson, C. Stolz, D. Fofi, and F. Meriaudeau, "Optimization of transparent objects digitization from visible fluorescence ultraviolet-induced," *Optical Engineering*, vol. 51, no. 3, p. 033601, 2012. doi: [10.1117/1.OE.51.3.033601](https://doi.org/10.1117/1.OE.51.3.033601).
96. M. B. Hullin, M. Fuchs, I. Ihrke, H. P. Seidel, and H. P. A. Lensch, "Fluorescent immersion range scanning," *ACM Transactions on Graphics*, vol. 27, no. 3, pp. 1–10, 2008. doi: [10.1145/1360612.1360686](https://doi.org/10.1145/1360612.1360686).
97. P. Fried, J. Woodward, D. Brown, D. Harvell, and J. Hanken, "3D scanning of antique glass by combining photography and computed tomography," *Digital Applications in Archaeology and Cultural Heritage*, vol. 18, p. e00147, 2020. doi: [10.1016/j.daach.2020.e00147](https://doi.org/10.1016/j.daach.2020.e00147).
98. M. Ferraton, C. Stolz, and F. Mériaudeau, "Optimization of a polarization imaging system for 3D measurements of transparent objects," *Optics Express*, vol. 17, no. 23, pp. 21077–21082, 2009. doi: [10.1364/OE.17.021077](https://doi.org/10.1364/OE.17.021077).
99. M. Shao, C. Xia, D. Duan, and X. Wang, "Polarimetric inverse rendering for transparent shapes reconstruction," *IEEE Transactions on Multimedia*, vol. 26, pp. 6140–6151, 2024. doi: [10.1109/TMM.2024.3371792](https://doi.org/10.1109/TMM.2024.3371792).
100. D. Michel, X. Zabulis, and A. Argyros, "Shape from interaction," *Machine Vision and Applications*, vol. 25, no. 3, pp. 1077–1087, 2014. doi: [10.1007/s00138-013-0579-2](https://doi.org/10.1007/s00138-013-0579-2).

101. M. Mikhnevich and D. Laurendeau, "Shape from silhouette in space, time and light domains," in *Proceedings of the 9th International Conference on Computer Vision Theory and Applications*, 2014, pp. 368–377. doi: [10.5220/0004683003680377](https://doi.org/10.5220/0004683003680377).
102. J. Munkberg, J. Hasselgren, T. Shen, J. Gao, W. Chen, A. Evans, T. Mueller, and S. Fidler, "Extracting triangular 3D models, materials, and lighting from images," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 8248–8258. doi: [10.1109/CVPR52688.2022.00810](https://doi.org/10.1109/CVPR52688.2022.00810).
103. Z. Li, Y. Y. Yeh, and M. Chandraker, "Through the looking glass: Neural 3D reconstruction of transparent shapes," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 1259–1268. doi: [10.1109/CVPR42600.2020.00134](https://doi.org/10.1109/CVPR42600.2020.00134).
104. H. Guerboukha, K. Nallappan, and M. Skorobogatiy, "Toward real-time terahertz imaging," *Advances in Optics and Photonics*, vol. 10, p. 843, 2018. doi: [10.1364/AOP.10.000843](https://doi.org/10.1364/AOP.10.000843).
105. B. Recur, A. Younus, S. Salort, P. Mounaix, B. Chassagne, P. Desbarats, J. P. Caumes, and E. Abraham, "Investigation on reconstruction methods applied to 3D terahertz computed tomography," *Optics Express*, vol. 19, no. 6, p. 5105, 2011. doi: [10.1364/OE.19.005105](https://doi.org/10.1364/OE.19.005105).
106. A. Mathai, N. Guo, D. Liu, and X. Wang, "3D transparent object detection and reconstruction based on passive mode single-pixel imaging," *Sensors*, vol. 20, no. 15, p. 4211, 2020. doi: [10.3390/s20154211](https://doi.org/10.3390/s20154211).
107. C. J. Phillips, M. Lecce, and K. Daniilidis, "Seeing glassware: from edge detection to pose estimation and shape recovery," in *Robotics: Science and Systems*, 2016. doi: [10.15607/RSS.2016.XII.028](https://doi.org/10.15607/RSS.2016.XII.028).
108. J. Sharpe, U. Ahlgren, P. Perry, B. Hill, A. Ross, J. Hecksher-Sørensen, R. Baldock, and D. Davidson, "Optical projection tomography as a tool for 3D microscopy and gene expression studies," *Science*, vol. 296, no. 5567, pp. 541–545, 2002. doi: [10.1126/science.1068206](https://doi.org/10.1126/science.1068206).
109. W. van Aarle et al., "Fast and flexible X-ray tomography using the ASTRA toolbox," *Optics Express*, vol. 24, no. 22, pp. 25129–25147, 2016. doi: [10.1364/OE.24.025129](https://doi.org/10.1364/OE.24.025129).
110. L. A. Feldkamp, L. C. Davis, and J. W. Kress, "Practical cone-beam algorithm," *Journal of the Optical Society of America A*, vol. 1, no. 6, pp. 612–619, 1984. doi: [10.1364/JOSAA.1.000612](https://doi.org/10.1364/JOSAA.1.000612).
111. A. Laurentini, "The visual hull concept for silhouette-based image understanding," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 16, no. 2, pp. 150–162, 1994. doi: [10.1109/34.273735](https://doi.org/10.1109/34.273735).
112. "CloudCompare," 2023. [Online]. Available: <https://www.cloudcompare.org/>.
113. P. Helmholz, D. Belton, N. Oliver, J. Hollick, and A. J. Woods, "The influence of the point cloud comparison methods on the verification of point clouds using the Batavia reconstruction as a case study," in *IKUWA6 Shared Heritage: Proceedings of the Sixth International Congress for Underwater Archaeology*, 2020.
114. G. Antova, "Application of areal change detection methods using point cloud data," *IOP Conference Series: Earth and Environmental Science*, vol. 221, no. 1, p. 012082, 2019. doi: [10.1088/1755-1315/221/1/012082](https://doi.org/10.1088/1755-1315/221/1/012082).
115. S. M. Seitz, B. Curless, J. Diebel, D. Scharstein, and R. Szeliski, "A comparison and evaluation of multi-view stereo reconstruction algorithms," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 1, 2006, pp. 519–528. doi: [10.1109/CVPR.2006.19](https://doi.org/10.1109/CVPR.2006.19).
116. E. Figueiras et al., "Optical projection tomography as a tool for 3D imaging of hydrogels," *Biomedical Optics Express*, vol. 5, no. 10, pp. 3443–3449, 2014. doi: [10.1364/BOE.5.003443](https://doi.org/10.1364/BOE.5.003443).
117. A. Ertürk et al., "Three-dimensional imaging of solvent-cleared organs using 3DISCO," *Nature Protocols*, vol. 7, pp. 1983–1995, 2012. doi: [10.1038/nprot.2012.119](https://doi.org/10.1038/nprot.2012.119).
118. "MatWeb Material Properties," 2023. [Online]. Available: <https://www.matweb.com>.
119. K. J. I. Lee et al., "Macro optical projection tomography for large-scale 3D imaging of plant structures and gene activity," *Journal of Experimental Botany*, vol. 68, no. 3, pp. 527–538, 2017. doi: [10.1093/jxb/erw452](https://doi.org/10.1093/jxb/erw452).

120. "Ray simulator for optics," 2023. [Online]. Available: <https://phydemo.app/ray-optics/simulator/>.
121. P. Debevec, "Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography," in *ACM SIGGRAPH 2008 Classes*, 2008, pp. 1–10. doi: [10.1145/1401132.1401175](https://doi.org/10.1145/1401132.1401175)
122. W. Jakob, S. Speierer, N. Roussel, and D. Vicini, "Dr Jit: A just-in-time compiler for differentiable rendering," *ACM Transactions on Graphics*, vol. 41, no. 4, 2022. doi: [10.1145/3528223.3530099](https://doi.org/10.1145/3528223.3530099)
123. T. Aila and S. Laine, "Understanding the efficiency of ray traversal on GPUs," in *Proceedings of the Conference on High Performance Graphics*, 2009, pp. 145–149. doi: [10.1145/1572769.1572792](https://doi.org/10.1145/1572769.1572792)
124. H. W. Jensen, *Realistic Image Synthesis Using Photon Mapping*. A K Peters, 2001.
125. M. Pharr, W. Jakob, and G. Humphreys, *Physically Based Rendering: From Theory to Implementation*, 4th ed. MIT Press, 2023.
126. Mildenhall, B., Srinivasan, P. P., Tancik, M., Barron, J. T., Ramamoorthi, R., & Ng, R. (2020). NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. *Proceedings of the European Conference on Computer Vision (ECCV)*, 405-421. <https://doi.org/10.1145/3503250>.
127. J. J. Park, S. N. Sinha, A. A. Efros, and N. Snavely, "Deformable neural radiance fields," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021, pp. 5791–5800. doi: [10.1109/ICCV48922.2021.00581](https://doi.org/10.1109/ICCV48922.2021.00581)
128. T. S. Trowbridge and K. P. Reitz, "Average irregularity representation of a rough surface for ray reflection of diffuse radar returns," *IEEE Transactions on Antennas and Propagation*, vol. 23, no. 3, pp. 328–334, 1975. doi: [10.1109/TAP.1975.1141079](https://doi.org/10.1109/TAP.1975.1141079).
129. B. Walter, S. R. Marschner, H. Li, and K. E. Torrance, "Microfacet models for refraction through rough surfaces," in *Proceedings of the 18th Eurographics Conference on Rendering Techniques*, 2007, pp. 195–206. doi: [10.2312/EGWR/EGSR07/195-206](https://doi.org/10.2312/EGWR/EGSR07/195-206).
130. IUPAC, "Extinction coefficient," in *Compendium of Chemical Terminology: Gold Book*. IUPAC, 2019. doi: [10.1351/goldbook.E02283](https://doi.org/10.1351/goldbook.E02283)
131. Liang, H. Advances in multispectral and hyperspectral imaging for archaeology and art conservation. *Appl. Phys. A* 106, 309–323 (2012). doi: [10.1007/s00339-011-6689-1](https://doi.org/10.1007/s00339-011-6689-1).