Predictive Analysis and Simulation Tools

Deliverable D3.2

Project information	
Project Acronym	COLOURS
Project Title	Collaborative On-cloud Lab for the conservation and digital restoration of ColOUred heritage collectionS
Grant Agreement Number	101233413
Call and Topic Identifier	HORIZON-CL2-2024-HERITAGE-ECCCH-01/ HORIZON-CL2-2024-HERITAGE-ECCCH-01-05
Funding Scheme	HORIZON Innovation Actions
Project Duration	36 Months [1 October 2025 – 30 September 2028]

Funded by the European Union

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

Deliverable Information			
Work Package No	3		
Del. Number	3.2	Del. Title	Predictive Analysis and Simulation Tools
Type	Report		
Dissemination Level	PU - Public		
Due Date	M12	Delivery Date	30.09.2026
Lead Beneficiary	2 - FhG		
Contributors	Paul Julius Kühn (FhG), Bruno Fanini (CNR), Daniele Ferdani (CNR), Giorgio Trumpy (NTNU), Sophia Sotiropoulou (FORTH), Marios Pitikakis (FORTH), Panagiotis Siozos (FORTH), Ioanna Stefanopoulou (FORTH), Anastasia Pentari (FORTH), Tanzima Habib (NTNU), Andreas Zapf (FhG)		

Document History

Version	Release	Summary of changes	Authors - Institutions
V0	29.07.2026	First release. Specifications of all digital tools defined, including feature lists, implementation plans, risk assessment, evaluation criteria and ECCCH integration strategy	FhG, CNR, FORTH and NTNU
V1.0	09.10.2026	Final version for submission; consortium review comments addressed.	FhG, CNR, FORTH and NTNU

Funded by the European Union

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

Project Context

Project start date: 1 October 2025

Project Duration: 36 months

Project website: https://colours-eccch.eu/

Project Coordinator: Daniele Ferdani - CNR ISPC daniele.ferdani@cnr.it

COLOURS Consortium

No	Short Name	Institution Name	Country
1	CNR	CONSIGLIO NAZIONALE DELLE RICERCHE	Italy
2	FhG	FRAUNHOFER GESELLSCHAFT ZUR FORDERUNG DER ANGEWANDTEN FORSCHUNG EV	Germany
3	FORTH	IDRYMA TECHNOLOGIAS KAI EREVNAS	Greece
4	NTNU	NORGES TEKNISK-NATURVITENSKAPELIGE UNIVERSITET NTNU	Norway
5	FBA	FUNDINGBOX ACCELERATOR SP ZOO	Poland
6	FBC	FUNDINGBOX COMMUNITIES SL	Spain
7	ECE	ERASMUS CENTRE FOR ENTREPRENEURSHIP BV	Netherlands
8	ZVKDS	JAVNI ZAVOD REPUBLIKE SLOVENIJE ZA VARSTVO KULTURNE DEDISCINE	Slovenia
9	BCS	Brakebush Börries	Germany
10	GMB	GALERIA MESTA BRATISLAVY	Slovakia
11	MIC	MINISTERO DELLA CULTURA	Italy

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

Executive Summary

This deliverable defines the base technologies and tools for the predictive analysis and simulation of colour change in cultural heritage (CH) artefacts, taking into account the requirements defined in WP2 (D2.1, D2.2) and the colour-management work of WP5. Alignment with the specification of the Collaborative Simulation Environment (Task T4.1, D4.1 due at M14) will be addressed in the second iteration. It specifies four simulation tools — doLCE 2.0, the Light Damage Estimator, VarnishLab and Neg2Pos — and the Degreening Tool, together with StyleShade3D, an AI-based semantification tool, and ATON and MuLaX, the inspection and visualisation tools through which prediction outputs are explored and compared against reference and diagnostic data, that will be adapted, partially integrated and enhanced. The current tools address light-induced fading, discolouration and darkening, and the visual effect of varnish removal and re-varnishing. Other phenomena listed in Task T3.2, such as loss, bleeding and spotting, are not yet modelled and will be reassessed in the second iteration.

For each tool, the deliverable describes its current state of development — including, where inherited from the PERCEIVE project, its technology readiness level — its planned new features and enhancements, the implementation approach, and the criteria by which its success will be assessed. Across the toolset, the Colour Knowledge Repository (CKR) is established as the central workflow hub through which tools exchange reference data, calibration information and simulation outputs, with the toolset positioned for deployment as a single Vertical Application within the ECHOES-ECCCH platform.

The toolset described builds substantially on tools developed within the preceding PERCEIVE project, whilst introducing targeted adaptations and new components identified through requirements gathering with domain experts and pilot activities. The specifications presented in this version reflect the state of knowledge available at the time of submission; a subsequent iteration is planned at Month 24 to incorporate more detailed insights on data availability, model performance, and workflow integration emerging from ongoing experimentation, workshops, and pilot activities.

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

Table of contents

Introduction **6**

Predictive technologies **8**

doLCE 2.0 8

Light Damage Estimator (LDE) 8

VarnishLab 13

Neg2Pos 14

Degreening Tool 14

Semantification Tools **15**

StyleShade3D: AI-Assisted Semantification of 3D Cultural Heritage Assets 15

Inspection/Visualization Tools **22**

ATON 22

MuLaX 24

Conclusions **26**

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

Introduction

Purpose

This deliverable specifies the predictive analysis and simulation technologies developed and adapted within Work Package 3 of the COLOURS project. It is the main output of Task 3.2 "Predictive analysis and simulation" (lead FORTH; contributors FhG, NTNU, CNR; M1–M36), which adapts, partially integrates and enhances base technologies for the prediction and simulation of colour change in cultural heritage (CH) artefacts — fading, loss, discolouration, darkening, bleeding and spotting — in coordination with Task 4.1, which specifies the Hybrid Collaborative Simulation Environment, and with the colour-management protocols of Work Package 5. FhG, as lead beneficiary of this deliverable, coordinates its drafting and submission on behalf of the contributing partners. Together with D2.1, D2.2, D3.1, D5.1, D7.1, D7.2, D14.1 and D14.2, this deliverable contributes to Milestone MS2 "Requirements, tech adaptation and exploitation strategy" (M12).

Scope

This deliverable covers the COLOURS tools that predict or simulate colour change and appearance — the Light Damage Estimator, doLCE 2.0, VarnishLab and Neg2Pos, together with the Degreening Tool, StyleShade3D, which maps simulation results onto 3D assets, and ATON and MuLaX, through which results are visualised— covering, for each tool, its current state of development (including the technology readiness level inherited from the PERCEIVE project, where applicable), its expected new features and enhancements, the planned implementation approach, where available associated risks and mitigation strategies, and the criteria by which its success will be assessed.

It further covers the network design and training strategy of the underlying AI core technologies, the integration of the AI core with a 3D visualisation system based on an adapted version of MuLaX incorporating colour management and sensor-fusion strategies (in this first iteration only as outlined in the individual tool sections; their full specification will be reported in the second iteration at M24), the data inputs and formats to be provided by conservators and heritage scientists, and the common integration strategy that deploys the tools as a single Vertical Application within the ECHOES-ECCCH platform.

Validation of the scientific quality of the tools' outputs against pilot data is addressed by the internal and external validation activities of later work packages rather than here; the detailed interoperability and testing requirements for deploying the Vertical Application on the ECHOES-ECCCH platform are specified in

Funded by the European Union

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

D7.1 and D7.2; and the semantic model and HDT mapping of the Colour Knowledge Repository is the subject of D3.3.

Target audience

The COLOURS tool developers and owners (CNR, FhG, FORTH, NTNU); the Work Package 4 team integrating the 3D visualisation and Collaborative Simulation Environment; the Work Package 5 team defining colour protocols; the Work Package 7 team specifying ECHOES-ECCCH integration and testing; the Work Package 10, 12 and 13 pilot and validation leads; internal reviewers; and the European Commission.

Relation to other deliverables

- COLOURS D3.1 Specification of digital restoration tools — baseline tool inventory and specifications this deliverable builds on for the predictive and simulation subset of tools.
- COLOURS D3.3 Knowledge Repository Design and HDT Mapping (M18) — defines the semantic model and HDT mapping of the Colour Knowledge Repository referenced here as the shared data hub.
- COLOURS D4.1 / D4.2 — specification and use cases of the Hybrid Collaborative Simulation Environment that hosts MuLaX.
- COLOURS D5.1 / D5.2 — colour protocols and colour-management information this deliverable's prediction and visualisation tools build on.
- COLOURS D7.1 / D7.2 — interoperability requirements and the integration/test plan for deploying the tools specified here as part of the COLOURS Vertical Application.
- WP4 Task T4.1 — specification of the Hybrid Collaborative Simulation Environment, with which the tools described here will be aligned in the second iteration.

Structure of the deliverable

Following this introduction, the deliverable specifies the predictive technologies (doLCE 2.0, the Light Damage Estimator, VarnishLab and Neg2Pos), the AI-based semantification tools (StyleShade3D), the inspection and visualisation tools (ATON and MuLaX). Each tool is described along a common set of dimensions: its current state and, where inherited from PERCEIVE, its technology readiness level; planned new features and enhancements; the implementation approach; evaluation criteria; and its integration with the predictive-analysis toolchain and the ECHOES-ECCCH platform. The deliverable closes with conclusions.

Funded by the European Union

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

Predictive technologies

doLCE 2.0

doLCE 2.0 reconstructs colour from monochrome scans of historical lenticular films, such as Kodacolor. Its contribution to predictive analysis lies in estimating the lenticular boundaries needed to decode the colour information physically embedded in the film. It supports the recovery of historical colour appearance rather than forecasting future degradation.

The processing pipeline combines convolutional neural networks for lenticule detection with colour reconstruction based on the historical encoding process. Training uses representative scanned frames, with expert corrections supporting iterative fine-tuning for individual films or scanning conditions. Users can adjust detected boundaries and inspect the resulting colour reconstruction. The accompanying project website[1] presents the human-in-the-loop deep-learning framework underlying this approach.

Within COLOURS, development will consolidate this pipeline into a guided workflow integrating dataset management, boundary editing, model fine-tuning and comparative previews. Reconstructed images and processing metadata will support colour-managed visualisation and deployment within ECCCH. Evaluation will assess boundary-detection accuracy, reconstruction consistency across frames, processing time and usability. Limited training data and variability between scans will be addressed through expert-guided refinement and reuse of saved model checkpoints.

Light Damage Estimator (LDE)

State of the Art

The assessment of light-induced colour change in cultural heritage has historically relied on empirical standardisation rather than predictive modelling. The Blue Wool standard (ISO 105-B01:2014; ISO 105-B02:2014)) established a widely adopted reference framework for quantifying light fastness through the controlled exposure of dyed wool samples under standardised conditions, providing a comparative scale for material sensitivity but offering no capacity to predict colour change for specific objects under real exhibition lighting scenarios. A step toward digital prediction was made by the CCI Light Damage Calculator (Michalski, 2012), subsequently integrated within the HERIe decision-support platform (Kupczak *et al.* 2018), which enabled dose-based risk estimates for a limited set of materials; however, its adoption within the conservation community has remained restricted

[1] https://lenticularfilmhitl.github.io/

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

by a narrow pigment coverage and the absence of systematic experimental validation against independent datasets. More accurate predictive approaches have since emerged at the research level: Ciortan *et al.* (2023a, b) demonstrated the application of tensor decomposition to hyperspectral time-series data, enabling both pigment characterisation and spatio-temporal simulation of colour change with higher spectral resolution, though the method remains computationally demanding and has not yet been translated into operational conservation tools. Similarly, Dang *et al.* (2024) proposed a photodamage prediction model explicitly accounting for the spectral composition of the light source through the use of spectrally tunable LED systems, demonstrating the importance of spectral modelling beyond broadband illuminance metrics; this approach, however, remains at the research stage and has not been validated across a broad range of pigment systems or integrated into accessible decision-support workflows.

Inherited LDE from PERCEIVE (TRL 5 Baseline)

The Light Damage Estimator (LDE) is a web-based decision-support tool for predicting and visualising cumulative light-induced colour change in photosensitive cultural heritage materials, developed within the PERCEIVE project (Ciortan et al. 2026) and currently assessed at TRL 5. The tool is implemented in Python as a modular API-based service within the PERCEIVE platform microservices infrastructure under a GPL-3.0 licence.

The LDE is structured around two distinct and sequential functional components. The first component addresses light source characterisation and exposure scenario definition. The tool implements the spectral classification framework (Siozos et al. 2025), which decomposes the spectral power distribution of light source into three bands. The user defines an exposure scenario by selecting a light source from the tool's built-in spectral power distribution library and specifying the illumination intensity in either illuminance (lx) or irradiance (μW/cm²) units, together with the exposure duration expressed as a combination of years, days per year, and hours per day. The light source library, spectral band definitions, and all model parameters are fixed within the tool; no user modification of the underlying spectral database or model coefficients is supported at the current development stage.

The second component addresses the prediction of colour change in each pigment present on the artwork surface. The tool accepts a standard sRGB digital image of the artwork as input. Pigment regions are identified through binary spatial masks, where each mask is associated with a single pigment and takes a yes/no format, indicating whether each pixel belongs to that pigment region or not. Each pigment

Funded by the European Union

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

is treated independently as a pure single-component system; mixture behaviour is not modelled at the current TRL 5 stage. For each masked region, the tool retrieves the corresponding photodegradation kinetics from a built-in accelerated ageing dataset, derived from controlled ageing experiments on single-pigment mock-ups. These kinetics, combined with the band-specific radiant exposures computed in the first component, drive the photochemical model to predict the cumulative colour change for each pigment region. All predictions are expressed in the CIELAB 1976 colour space, enabling direct comparison with conservation thresholds and published literature. The pigment database, ageing kinetics, and all associated model parameters are integral to the tool and are not modifiable by the user.

The tool has been validated on two case studies within the PERCEIVE project: Edvard Munch's The Scream (ca. 1910), representing the paintings and works on paper scenario, and a textile specimen. In both cases, predictions were generated for individual pure pigments identified within the respective artworks. The tool was presented to internal and external users at the Digital Heritage 2025 Conference, Siena, Italy, in September 2025. The current validated capability is confined to single-pigment systems, unvarnished paintings, sRGB image input, and binary pigment masks; The constraints are addressed by the planned developments under COLOURS.

Planned New Features

Cloud deployment compliance:

The LDE is being migrated from the PERCEIVE microservices infrastructure to operate as a service within the COLOURS Vertical Application, in preparation for integration within the ECHOES-ECCCH cloud environment. This transition fundamentally changes the tool's data model: rather than relying exclusively on a fixed internal database, the LDE will support the contribution of user-owned data and bidirectional exchange with the Colour Knowledge Repository (CKR) as the central workflow hub:

- User-contributed data: sRGB images, HSI datasets, XRF maps, light source spectral power distribution (SPD) data, pigment reflectance spectra (artificial ageing data), stored securely under the user's own identity authenticated via the ECHOES AAI brokered by the COLOURS Platform and user data stored in the CKR/COLOURS storage, with full user ownership in line with ECHOES data governance and FAIR principles
- CKR to LDE: reference data (pigment sensitivity database, varnish absorption coefficients, mixture calibration data) pulled from the CKR at runtime

Funded by the European Union

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

- LDE to CKR: simulation outputs (predicted CIELAB trajectories, lighting risk assessments, cumulative dose estimates) pushed to the CKR as FAIR research outputs, accessible and reusable by the wider conservation community

Scientific developments

Three specific constraints of the TRL 5 baseline are addressed:

- Hyperspectral image (HSI) input: the sRGB image pathway of the TRL 5 baseline is retained. When HSI data are available, the tool additionally accepts hyperspectral imaging data as input. Per-pixel reflectance spectra are extracted from the HSI cube and converted to CIELAB via standard CIE colorimetry (CIE 1931 2° observer colour matching functions, D65 illuminant, XYZ to CIELAB), providing spectrally resolved colorimetric pipeline in place of the three-channel sRGB approximation
- Automatic pigment mask generation: performed directly within the LDE from multi-modal analytical imaging data, without reliance on external tools; spectral unmixing of the HSI cube identifies pigment regions from their spectral signatures; automatic segmentation based on elemental signatures from XRF mapping provides complementary spatial delineation; together these replace the manually drawn binary single-pigment masks of the TRL 5 baseline with automatically generated, analytically grounded pigment maps covering the full artwork surface
- Extension to colour and white pigment mixtures: the photodegradation model is extended from pure single-pigment systems to binary colour and white mixtures using the Kubelka-Munk framework; the extended reflectance relation is derived by substituting the initial reflectance of the mixture in place of that of the pure pigment, preserving the functional form of the single-pigment model; governing assumptions stated explicitly: additive mixing of absorption and scattering coefficients at each wavelength proportional to pigment concentration (Duncan additivity); white pigment treated as non-absorbing relative to the coloured pigment; key finding: whether the addition of a white pigment affects only optical scattering or also the underlying photochemical kinetics of the coloured pigment is colorant-pair-specific and cannot be assumed a priori, with direct implications for the tool's calibration requirements

Funded by the European Union

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

Evaluation and Validation

Software quality evaluation:

- Evaluation strategy consistent with the Cloud Assessment Framework (CAF) under development within the ECHOES project. Assessment follows the CAF self-assessment model in which the tool provider generates structured evidence using recommended instruments for review by the coordinating body; dimensions assessed at subsequent iterations via the same CAF self-assessment instruments: interaction capability, sustainability, and user adoption, once platform-level integration within ECHOES-ECCCH is finalised
- Core dimensions assessed: open-source licensing transparency (GPL-3.0); compatibility and interoperability with the COLOURS platform, the CKR, and the ECHOES-ECCCH infrastructure; adherence to FAIR principles for research software (FAIR4RS)

Scientific validation:

- Primary accuracy metric: CIELAB ΔE between model predictions and published measured colour change data for known pigments under known illumination conditions; primary target dataset is Samain et al. (2011) on Prussian blue fading
- HSI-to-CIELAB extraction pipeline validated against reference colorimetric measurements on known pigment samples
- Automatic pigment mask generation validated against manually annotated reference masks for known pigment distributions
- Full end-to-end validation against new pilot case studies with conservation professionals using their own including calibrated sRGB images or HSI datasets, and XRF maps, deferred to M24 second iteration

TRL Progression

- TRL 5 (current, M12): tool validated on two controlled case studies within PERCEIVE; limited to pure pigment systems, unvarnished paintings, sRGB image input, and manually defined pigment regions; operating within the PERCEIVE microservices infrastructure
- TRL 6 (target M24): Python implementation of all planned developments completed and demonstrated on real artwork data within the COLOURS environment; CKR integrated as central workflow hub; scientific validation against independent published dose-response data confirmed
- TRL 7 (target M36): full end-to-end validation with conservation professionals operating the tool within the COLOURS environment using their own artwork data, including calibrated sRGB images or HSI datasets, and XRF maps; expanded pigment and light source database; tool

Funded by the European Union

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

accessible to the wider conservation community via the COLOURS Applications

VarnishLab

The Varnish Removal/Revarnish Tool (VarnishLab) predicts the colour appearance of paintings following varnish removal or the application of a new varnish layer. It provides non-invasive previews to help conservators compare treatment options before physical intervention.

The simulation uses physics-based models of light interaction with varnish and paint layers, including absorption, scattering and reflection at interfaces. Initial models have been experimentally validated using hyperspectral captures of a nineteenth-century oil painting. Further refinement will draw on before-and-after treatment datasets and physical mock-ups to characterise different varnish materials and assess their effects on colour and gloss.

Within COLOURS, these models will be integrated into a tool supporting RGB and hyperspectral inputs, adjustable degrees of cleaning and new varnish thickness, and localised treatment simulations. Side-by-side visualisations will compare the untreated, virtually cleaned and re-varnished painting, complemented by spectral reflectance curves where available. Evaluation will compare simulated outcomes with physical treatments and assess robustness and usability. Simulation parameters and available material data will accompany exported results to document the assumptions underlying each prediction.

VarnishLab is a new tool introduced in COLOURS in response to needs expressed by painting conservator-restorers during the WP2 requirements analysis (D2.1). It was assessed at TRL 3 at project start, with a target of TRL 6 by the end of the project. A preliminary proof of concept was presented to conservator-restorers at the COLOURS–ECCCH workshop in Ljubljana (12 May 2026), where participants asked for adjustable degrees of varnish removal, an immediate before/after comparison, and operation on calibrated RGB images and point spectral

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

measurements. The tool addresses the first research question of Internal Pilot 1, the removal of non-original varnish.

Neg2Pos

Neg2Pos converts digitised colour negatives into positive images by simulating the photochemical printing process and subsequent viewing conditions. Its contribution to predictive analysis is to estimate how a negative would appear as a photographic print or projected film, supporting historically informed rendering of archival materials.

The approach uses physics-based models incorporating capture-system characteristics, negative-film properties and the response of the selected positive material. Simulation parameters include printing exposure and colour filtration, allowing users to explore their effects on colour balance, contrast and tonal reproduction. The theoretical framework and initial feasibility experiments provide the basis for developing these models into a practical tool.

Within COLOURS, development will integrate high-bit-depth image import, capture and film-stock specifications, and emulation of photographic papers and cinema release-print stocks. Evaluation will compare simulated outputs with reference prints or projected images where available, assessing colour and tonal fidelity, consistency and usability through internal and external pilots. Exported positive images will be accompanied by processing metadata documenting the selected materials, calibration information and simulation settings, supporting reproducibility and integration within ECCCH.

Neg2Pos is a new tool introduced in COLOURS; it was assessed at TRL 2 at project start, with a target of TRL 6 by the end of the project.

Degreening Tool

The Degreening Tool, inherited from the PERCEIVE project, already covers its intended core use case: the removal of greening defects from digitised

Funded by the European Union

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

Autochrome plates. Its deep-learning model detects affected areas reliably enough that only limited manual post-processing is generally required, substantially reducing the effort compared with fully manual correction using general-purpose image-editing software. Moreover, the WP2 activities conducted during the first project year did not identify additional functional requirements, partly because they focused primarily on paintings and polychrome sculptures rather than photographic materials. Consequently, no further algorithmic development is currently planned, and development resources are being prioritised for tools addressing needs that are only partially covered by existing solutions. Nevertheless, the Degreening Tool remains part of the COLOURS toolset. It will be integrated into the COLOURS Vertical Application as a platform-managed service, with outputs stored in the CKR, validated with conservation professionals during the pilots, and documented for use by photographic archives and conservation laboratories. These integration and validation activities are expected to support its progression from TRL 5 to TRL 7. If pilot validation reveals significant functional gaps, targeted enhancements will be considered.

Semantification Tools

StyleShade3D: AI-Assisted Semantification of 3D Cultural Heritage Assets

StyleShade3D, originally developed within the PERCEIVE project, provides the semantic spatial representation required to connect predictive colour-change models with 3D cultural heritage (CH) assets. The tool was assessed at TRL 4 at the start of the project and has since advanced to TRL 5, reflecting the validation of its integrated prototype in a relevant environment using eight representative CH objects. It has not yet undergone a controlled end-user study or been demonstrated in routine professional production. The tool transforms a textured and UV-parameterised 3D object into a semantically segmented atlas in which relevant surface regions, such as materials, object components, painted areas,

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

previous interventions, or degradation zones, are represented by consistent labels.

This semantification step is necessary because colour degradation rarely affects an object uniformly. Different materials and surface regions may respond differently to light exposure, humidity, temperature, pollutants, ageing, or conservation treatments. Predictive models must therefore be able to associate their outputs with clearly defined areas of the digital object. StyleShade3D supplies this association and acts as an interface between the AI prediction services, the available scientific data, and the 3D visualisation environment.

StyleShade3D does not itself predict future colour change. Instead, it prepares the spatial and semantic structure on which predictions can operate. Its output allows predicted effects such as fading, darkening, discolouration, bleeding, spotting, or material loss to be assigned selectively to the affected regions of a 3D representation.

Technical approach

StyleShade3D follows a human-in-the-loop, atlas-based segmentation approach. Rather than relying exclusively on a 3D segmentation model, the tool renders a compact set of two-dimensional views of the object, applies AI-assisted segmentation to those views, and projects the resulting labels back onto the UV parameterisation of the original model. The technical approach and its evaluation on eight CH objects are described in Kühn et al. (2026).

The workflow consists of four principal stages:

1. **Viewpoint selection:** A set of candidate camera positions is generated around the object. Surface visibility is evaluated by ray casting, and a greedy set-cover strategy selects a compact subset of views that covers the sampled surface representation.
2. **AI-assisted semantic annotation:** The selected views are segmented using the SAM 2 foundation model integrated into an interactive annotation

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

environment based on Label Studio. Users can prompt, inspect, correct, and refine the proposed masks.

3. **Back-projection into texture space:** The validated masks are projected from image space onto the UV coordinates of the 3D model. The partial results obtained from the individual views are combined into a unified segmentation atlas.
4. **Semantic enrichment and downstream use:** Semantic identifiers and domain-specific attributes are associated with the atlas regions. These regions can subsequently receive material parameters, degradation predictions, colour transformations, uncertainty information, or other simulation results.

The use of a human-in-the-loop process is particularly important for CH applications. The definition of a meaningful region depends on the intended conservation question and may change when additional historical, material, or diagnostic information becomes available. For example, a conservator may require segmentation by visible object component, while a heritage scientist may require a distinction between pigments, substrates, restoration materials, or areas with different exposure histories. StyleShade3D therefore preserves expert control over both segmentation boundaries and their semantic interpretation.

In its current form, StyleShade3D does not provide diagnosis-driven segmentation: diagnostic or scientific information can inform an expert's manual definition and refinement of semantic regions, but it is not automatically processed to generate or adapt segmentation masks. The tool also does not currently support standalone 2D assets or an independent end-to-end 2D workflow. The rendered 2D views and the resulting 2D atlas are intermediate representations of a UV-parameterised 3D asset. Diagnosis-driven segmentation and standalone 2D support remain subjects for future development.

Input and output specifications

The initial implementation assumes that the input asset contains a valid UV parameterisation without overlapping UV islands. Where this is unavailable, UV

Funded by the European Union

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

generation must be performed as a preprocessing step. Overlapping regions introduce ambiguity because multiple surface locations may map to the same atlas texel.

Category	Required or expected content
3D geometry	Polygonal mesh representing the CH object
Surface parameterisation	Non-overlapping UV map
Appearance data	Texture images and, where available, physically based material parameters
Camera information	Generated viewpoints, camera intrinsics, and camera extrinsics
Semantic definition	Region classes supplied or validated by CH experts
Primary output	Colour-coded or indexed semantic segmentation atlas
Additional outputs	Per-view masks, region identifiers, class labels, confidence information, and projection metadata
Integration output	Region-to-prediction mapping consumable by the simulation and 3D visualisation components
Optional contextual data	Material analyses, multispectral data, condition maps, exposure data, or conservation annotations that experts may use to guide manual semantic definition; these data are not currently processed automatically for diagnosis-driven segmentation

Each semantic region should be represented by a persistent identifier rather than solely by an RGB colour. The identifier can be associated with metadata including the semantic class, material type, confidence level, annotator, creation date, and provenance of the supporting evidence. This separation between visual encoding and semantic identity avoids ambiguities caused by colour-space conversion or texture compression.

Network design and human-in-the-loop operation

The AI component uses SAM 2 to generate segmentation masks from lightweight user prompts, such as positive or negative points and bounding boxes. The foundation model is applied to rendered views rather than directly to the 3D mesh. This makes it possible to exploit a general-purpose image segmentation model without requiring a large, domain-specific collection of manually segmented 3D CH assets.

Funded by the European Union

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

The model output is treated as a proposal rather than as authoritative ground truth. Domain experts remain responsible for deciding:

- which semantic entities are relevant to the prediction task;
- whether two visually similar regions should be represented as separate classes;
- how uncertain or partially visible boundaries should be handled;
- whether available diagnostic information justifies manually revising an existing segmentation; and
- whether the result is sufficiently accurate for simulation and visualisation.

Validated masks are propagated across views using consistent class identifiers. During back-projection, rendered face-index information establishes the correspondence between image pixels, mesh faces, and UV coordinates. If a surface area is visible in more than one view, the resulting labels must be consolidated. Conflicts can be resolved using expert validation, viewing-angle quality, model confidence, or an agreed priority rule.

Integration with predictive analysis

The segmentation atlas provides the spatial control layer for the predictive services. Each semantic region can be linked to a material or degradation profile and can receive a separate predicted temporal trajectory. The region identifier generated through StyleShade3D allows these outputs to be written into textures or auxiliary data layers. Depending on the prediction service, the results may be represented as:

- time-dependent colour or reflectance maps;
- masks indicating the predicted location and extent of degradation;
- scalar intensity maps for fading, darkening, or material loss;
- uncertainty maps;
- parameter layers controlling shaders or material models; or
- discrete scenarios representing alternative environmental assumptions.

Funded by the European Union

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

This structure supports both region-level and texel-level prediction. Region-level predictions are appropriate when only aggregate material or condition information is available. Texel-level predictions can be introduced where spatially resolved measurements and sufficiently dense training data exist.

Integration with the 3D visualisation system

The semantic atlas and the associated prediction layers will be transferred to the adapted MuLaX-based visualisation component. The visualisation system will use the atlas to identify which portions of the digital object should be modified for a selected simulation time or scenario.

The expected exchange sequence is as follows:

1. MuLaX loads the original geometry, UV coordinates, textures, and material definitions.
2. The semantic atlas assigns each surface texel to a persistent semantic region.
3. The prediction service stores time-dependent values for one or more regions in the CKR, from which MuLaX retrieves them..
4. The values are converted into colour-managed texture layers or shader parameters.
5. MuLaX renders the simulated state while preserving the original geometry and material context.
6. The user selects time points, exposure scenarios, or uncertainty levels and compares the predicted state with the documented baseline.

Colour transformations must be performed in a controlled colour space and should retain the metadata required to reproduce the result. Whenever possible, predicted changes should be applied to calibrated colour or reflectance data rather than directly to display-dependent RGB values. The visualisation should also distinguish measured information from model-generated changes and provide access to confidence or uncertainty indicators.

Funded by the European Union

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

Initial evaluation and observed limitations

The underlying atlas-based workflow has been demonstrated on eight CH objects with diverse geometries and surface appearances, including statues, busts, furniture, weapons, a shield, and a manuscript. The experiments showed that usable segmentation atlases could be generated for all tested objects. For the seven objects with recorded annotation times, the reported effort ranged from approximately 15 to 35 minutes, with an average of about 21 minutes.

The evaluation also identified recurring challenges:

- small or narrow structures often require manual correction;
- visually similar neighbouring regions can be merged by the segmentation model;
- low-texture objects depend strongly on suitable rendering and lighting;
- cavities and strongly occluded surfaces may be observed from unfavourable angles;
- fine semantic details can be difficult to edit efficiently in the annotation interface;
- coverage of sampled points does not necessarily guarantee an optimal projection angle for every point on the high-resolution mesh; and
- the absence of ground-truth 3D segmentations currently limits quantitative accuracy assessment.

These observations confirm the importance of expert validation. They also indicate that segmentation quality must be evaluated not only geometrically, but in relation to the downstream prediction task. A minor boundary error may be acceptable for broad visualisation, while the same error may be unacceptable when two adjacent regions contain materials with substantially different degradation behaviour.

Funded by the European Union

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

Contribution to technology readiness

StyleShade3D raises the technology readiness of the predictive toolchain by converting experimental segmentation methods into a controllable workflow with explicit inputs, outputs, correction mechanisms, and integration points. Its principal contribution is to bridge the gap between AI predictions expressed as numerical or image-based outputs and the semantically meaningful regions required by CH professionals.

In the first implementation phase, the focus is on robust atlas generation, expert correction, and data exchange. Subsequent phases will address tighter integration with prediction services, colour-managed temporal visualisation, task-specific model adaptation, and validation in pilot workflows. The objective is not to remove expert judgement, but to reduce repetitive manual processing and provide a transparent, reusable semantic basis for simulating colour change on complex CH objects.

Inspection/Visualization Tools

ATON

ATON is an open-source Web3D/WebXR framework developed by CNR ISPC to present and interact with complex/multi-dimensional data on the web [Fanini et al. 2021]. It enables the rapid creation and deployment of interactive Web3D/WebXR applications accessible through common web browsers. ATON provides a modular, cross-device platform for presenting and interacting with 3D models, scenes, and associated metadata online using mobile devices, desktop workstation, museum kiosks, and up to immersive VR/AR devices.

Feature list:

- Adaptive 3D presentation across mobile, desktop, and XR devices
- Immersive AR/VR experiences through WebXR support

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

- Semantic annotations and volumetric shapes
- Multi-resolution streaming of meshes, point-clouds and 3DGS via Cesium 3D Tiles (OGC standard)
- Support for open formats (e.g. glTF - ISO standard)
- Advanced rendering (PBR materials) and support for Khronos extensions
- Real-time multi-user collaboration features
- Spatial UI components for XR sessions

The base framework also provides a built-in front-end ("Hathor") with several features directly accessible via UI [Amadei et al. 2025]. Besides a presentation mode, the front-end also includes an editor mode, allowing authenticated users to create and organize a 3D scene.

The framework offers a plug&play model to accelerate the creation and deployment of vertical solutions built on top of ATON features [Fanini et al. 2025]. ATON is thus a foundational framework for tools like MuLaX and CSE (Collaborative Simulation Environment), traceable within COLOURS' objectives to develop interoperable, web-based services/tools targeting cultural heritage research and dissemination.

Integration with predictive analysis

ATON provides a wide range of mechanisms and routines to be employed in advanced Web3D/WebXR applications developed on top of it, offering the possibility to create and design advanced prediction tools for interactive and visual inspection. Specifically, MuLaX and CSE (built on top of ATON) will access output data of different tools/services, including prediction data.

The base framework offers for such scenario:

- **Cloud integration**: mechanisms to access shared data layers with other COLOURS tools operating on the same artwork/item, thus the possibility to access data produced by other tools (e.g. StyleShade3D) or CKR data layer
- **Scene layers**: they can be used for multidimensional visualization, including prediction data processed by other COLOURS services/tools
- **Semantic annotations**: semantic shapes can be used to enrich digital replica of a given 2D or 3D artwork in virtual space, at runtime

Funded by the European Union

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

- **REST API** can be employed to enable bi-directional communication with CKR and other tools/services involved in computation or processing of predictive data

MuLaX

MuLaX (Multi-Layer eXtensible eXplorer) [Fanini et al. 2025] is an interactive Web3D/WebXR tool designed and developed under the PERCEIVE European project that will be evolved in COLOURS. It enables the online exploration of multi-layered 2D/3D collections enriched with diagnostic and analytical data. MuLaX is built on top of ATON framework and integrates high-resolution 3D models with multi-band imaging datasets (e.g., VIS, UVL, VIL) and punctual analyses (e.g., XRF, microscope, FORS), allowing users to visually compare and inspect heterogeneous scientific results within a single virtual environment. Due to its ATON foundations, the data-flow will inherit the same pipelines adopted in the framework, thus accessing directly CKR via REST API and its cloud space (for data derivatives).

Feature list:

- Interactive multi-layer discovery (lens and split modes) for analytical imaging (UVL, VIL, VIS)
- Integration of punctual analyses (XRF, microscope, FORS) as spatially referenced POIs
- Texture-based semantic masks and dynamic image-based masking
- Interactive Semantic Painting (ISP) to create expert-driven masks directly online Cloud integration for remote access to collections and analytical layers
- Responsive Web3D interface compatible with desktop and mobile devices
- WebXR immersive inspection in VR and AR modes
- Modular architecture supporting extension to paintings, textiles, and other domains.

Integration with predictive analysis

MuLaX offers a subset of advanced functionalities to conservation scientists, archaeologists, art historians, and diagnostic laboratories working with multi-band

Funded by the European Union

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

imaging and material analyses, that can be employed in predictive analysis, in particular:

- Multi-layer discovery, to interact with multiple predictive layers on the same item/artwork, through lens or split-mode discovery.
- Texture-based semantic masks and dynamic image-based masking
- Interactive Semantic Painting (ISP), to interactively create expert-driven masks on 2D/3D items through a web browser, and store results into a shared data layer accessible by other COLOURS tools/services

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

Conclusions

This deliverable has specified eight of the ten digital restoration tools that make up the COLOURS predictive and simulation toolset: doLCE 2.0, the Light Damage Estimator, VarnishLab, Neg2Pos, the Degreening Tool, StyleShade3D, ATON and MuLaX. For each of these tools, the document has described its current state of development — including, where applicable, new features and enhancements, the planned implementation approach, and the criteria by which its success will be assessed.

Across the specified tools, a common integration strategy is established: the Colour Knowledge Repository (CKR) acts as the central workflow hub through which tools exchange reference data, calibration information and simulation outputs, and the toolset as a whole is positioned for deployment as a single Vertical Application within the ECHOES-ECCCH platform, with authentication brokered by the COLOURS Platform through the ECHOES AAI (D7.1).

In line with the Grant Agreement, a second iteration of this deliverable is foreseen at Month 24, by which point the consortium expects the specifications presented here to be consolidated against validated end-user feedback gathered through the project's ongoing pilot and workshop activities.

Funded by the European Union

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

References:

- Siozos P, Monico L, Romani A, Miliani C, Doherty B, Sandu ICA, et al. Light source classification and colour change modelling for understanding and predicting pigments discolouration. J Phys Photonics [Internet]. 2025 [cited 2025 Aug 20]; 7:035020. https://doi.org/10.1088/2515-7647/ade8c2
- Ciortan, IM. *et al.* (2026). Tools and Services for Colour Change Analysis and Sharing of Knowledge. In: Ciortan, I.M., Ferdani, D., Hardeberg, J.Y., Sotiropoulou, S., Veggi, M. (eds) Chromatic Visions. Lecture Notes in Computer Science, vol 16450. Springer, Cham. https://doi.org/10.1007/978-3-032-17700-1_2
- Amadei C.F., Massidda M., Menconero S., Gosti G., Fanini B. (2025). Early Stage of Community-Driven UI/UX Refactoring of the Open-Source ATON Framework. Digital Heritage 2025. https://doi.org/10.2312/dh.20253115
- Fanini B., Massidda M., Ferdani D., Bonifazi F., Magrini D., Iannacone R., Barandoni C. (2025). Interactive and immersive discovery of diagnostic processes on multi-layered 3D collections on the web: the MuLaX tool. Digital Heritage 2025. https://doi.org/10.2312/dh.20253102
- Fanini B., Gosti G., Amadei C.F., Menconero S. (2025). Shaping a Web3D framework for different scientific communities: ATON as a service in H2IOSC. Digital Heritage 2025. https://doi.org/10.2312/dh.20253055
- B. Fanini, D. Ferdani, E. Demetrescu, S. Berto, E. d'Annibale (2021). ATON: An Open-Source Framework for Creating Immersive, Collaborative and Liquid Web-Apps for Cultural Heritage. Applied Sciences, 11(22), 11062. https://www.mdpi.com/2076-3417/11/22/11062
- International Organisation for Standardisation. ISO 105-B01:2014: Textiles — Tests for colour fastness — Part B01: Colour fastness to light: daylight [Internet]. Geneva: ISO; 2014 [cited 2026 Sep 28]. https://www.iso.org/standard/65210.html
- International Organization for Standardization. ISO 105-B02:2014: Textiles — Tests for colour fastness — Part B02: Colour fastness to artificial light: Xenon arc fading lamp test. Geneva: ISO; 2014. https://www.iso.org/standard/65209.html
- Michalski S. Light damage calculator [Internet]. Ottawa: Canadian Conservation Institute; 2012 [cited 2026 Sep 28]. https://app.pch.gc.ca/application/cdl-ldc/description-about.app?lang=en
- Kupczak A, Jedrychowski M, Strojecki M, Krzemien L, Bratasz L, Lukomski M, et al. HERIe: a web-based decision-supporting tool for assessing risk of physical damage using various failure criteria. Stud Conserv [Internet]. 2018 [cited 2026 Sep 28];63(sup1):151–5. https://doi.org/10.1080/00393630.2018.1504447
- Ciortan I-M, Poulsson TG, George S, Hardeberg JY. Tensor decomposition for painting analysis. Part 1: pigment characterization. npj Herit Sci [Internet]. 2023a [cited 2026 Sep 28];11:76. https://doi.org/10.1186/s40494-023-00910-x
- Ciortan I-M, Poulsson TG, George S, Hardeberg JY. Tensor decomposition for painting analysis. Part 2: spatio-temporal simulation. Herit Sci. 2023b; 11:84. https://doi.org/10.1186/s40494-023-00913-8
- Dang R, Hu Z, Cai X, Tang M, Tai L, Wang B, et al. Photodamage prediction model and optimal lighting system for polychrome artworks in museum. J Cult Herit [Internet]. 2024 [cited 2026 Sep 28];70:1–12. https://doi.org/10.1016/j.culher.2024.09.009
- Kühn, P. J., Sinha, S. N., Hansen, J., & Horst, R. (2026). Human-in-the-Loop Atlas-Based 3D Asset Segmentation for Interactive Content Workflows. arXiv:2606.17824v2. https://doi.org/10.48550/arXiv.2606.17824

Funded by the European Union

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

- Samain L, Silversmit G, Sanyova J, Vekemans B, Salomon H, Gilbert B, et al. Fading of modern Prussian blue pigments in linseed oil medium. J Anal At Spectrom. 2011;26(5):930–41. https://doi.org/10.1039/c0ja00234h

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.