

## Deliverable D4.3

### Updated web front-end

Lead beneficiary	SAI
Author(s)	Gustav Mueller-Franzes - Bulat Maksudov – Architecture, development, testing, user training Maria Londono – Project Management
Reviewer(s)	UZH
Dissemination level	PU
Type	DEC
Delivery date	22 December 2025

*ODELIA is funded by the European Union's Horizon Europe Framework  
under Grant Agreement 101057091*



**Funded by  
the European Union**

## TABLE OF CONTENTS

Executive summary.....	3
Introduction .....	3
Feedback Interface.....	4
FEEDBACK DASHBOARD .....	5
USER AUTHENTICATION.....	6
REAL MODEL INTEGRATION .....	7
Model initialization:.....	7
AI Result representation: .....	7
Viewer Documentation .....	9
Conclusion .....	9
Appendix A: MEDIA STORAGE CONFORMANCE STATEMENT .....	10
INTRODUCTION .....	10
SUPPORTED SOP CLASSES .....	10

## DISCLAIMER

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

## EXECUTIVE SUMMARY

We report the successful completion of Deliverable 4.3. This deliverable extends the initial web-based front-end developed in D4.1, transforming it into a comprehensive platform for evaluating and refining AI models within the ODELIA consortium. This milestone augments the viewer with critical capabilities that strengthen human–AI interaction, model assessment, and system robustness.

The updated viewer now empowers consortium members to:

- **Provide structured feedback** on AI-generated diagnostic results to support continuous model refinement and validation.
- **Visualize aggregated feedback** through a dedicated dashboard, enabling systematic monitoring of model performance and data quality.
- **Authenticate users securely**, ensuring traceability, accountability, and role-based access control within the evaluation process.
- **Interact with fully integrated "Real" AI models**, specifically the baseline model from the ODELIA MICCAI 2024 Challenge, supported by enhanced workflows.
- **Benefit from technical maturity**, including streamlined deployment documentation and optimized inference handling.

These enhancements build upon the foundational functionalities delivered in D4.1—manual annotation, submission of MRI studies to AI, and visualization of AI results—while significantly advancing the platform’s readiness. The updated viewer now provides a rigorous, interactive environment that supports the consortium’s efforts to validate and improve the decentralized AI models developed within the ODELIA project.

## INTRODUCTION

The objective of Work Package 4 (WP4) is to develop a robust, user-friendly front-end that supports both the evaluation of AI models and their eventual integration into clinical-research workflows. Following the initial delivery of a functional web viewer in D4.1, which introduced core capabilities for study visualization, manual annotation, and AI inference routing, the present deliverable advances the system toward a mature, validation-ready platform.

The enhancements implemented in this phase focus on enabling structured user interaction, strengthening model evaluation workflows, and ensuring system security. These developments mark the project’s transition from prototype functionality to a platform capable of supporting iterative model improvement and realistic clinical testing.

This document outlines the design principles guiding the extended viewer, details the new functionalities introduced (ODV-005 to ODV-009), and describes the technical improvements implemented since the previous deliverable. It further describes the workflow implications of these additions, particularly regarding the handling of heterogeneous MRI sequences, and provides an overview of the updated user experience.

For Deliverable D4.1, we addressed the implementation of features ODV-001 through ODV-004. In the current Deliverable D4.3, we address features ODV-005 through ODV-009.

ID	Feature	Functionality
ODV-001	Launch ODELIA web viewer	Launch a web viewer, select particular MRI examination from a list and of studies, review it in main view, do basic manipulations with the images.
ODV-002	Provide manual annotation tools	Tools for annotating lesions within the MRI for AI model training.
ODV-003	Send to AI	Select AI model from list of available AI models, add new models, send an MRI examination for processing by AI, enable user to select the correct sequences for AI inference.
ODV-004	View AI results	Select AI results produced from different algorithms, display classification results and confidence, and explainability maps (if provided by the model).
ODV-005	Feedback Interface	Feedback interface for the standardized collection of feedback on AI model performance.
ODV-006	Feedback Dashboard	Customizable dashboard to visualize feedback results.
ODV-007	User Authentication	Secure user authentication to link reviewer identity to the feedback provided and control access.
ODV-008	Real Model Integration	Improved AI model integration, including the capability to automatically download and use the baseline model from the ODELIA MICCAI 2025 Challenge.
ODV-009	Viewer Documentation	Comprehensive documentation for viewer setup and deployment.

## FEEDBACK INTERFACE

A critical requirement for validating AI models in a clinical setting is the ability for domain experts (radiologists) to validate model outputs. In D4.1, we identified the need for a "feedback functionality that will enable the structured collection of feedback". In D4.3, we have implemented a dedicated Feedback Interface panel.

When an AI result is displayed (as defined in ODV-004), the radiologist is now presented with a structured form to evaluate the diagnosis. This interface allows:

1. **Model Metadata & Traceability:** The panel clearly displays the Active Study UID, the specific Model Name (e.g., ODELIA-AI), Model Version, and the precise Result Timestamp. This ensures that all feedback is strictly coupled to a specific inference run, facilitating version-controlled performance tracking.
2. **Granular Prediction & Confidence:** The AI prediction is broken down by anatomy (Left Breast vs. Right Breast), displaying both the classification (e.g., Benign) and the specific confidence score (e.g., 40.6%).
3. **Structured Validation:** For each prediction, the radiologist can efficiently provide feedback using a simplified radio-button interface with three options: "Agree," "Unsure," or "Disagree." This design prioritizes speed and ease of use to minimize disruption to the clinical reading workflow.
4. **User Attribution:** The footer of the panel confirms the active user context ("Signed in as..."), ensuring that the feedback submitted is cryptographically linked to the authenticated user ID (as defined in ODV-007).
5. **Multi-Reviewer Functionality:** A single case can be reviewed independently by **multiple radiologists**, allowing the collection of inter-reader agreement statistics or consensus evaluation. Aggregated and comparative data can be accessed **post-hoc** through the Feedback Dashboard, which is intended for administrative or research analysis.



Figure 1: Feedback Interface

Once the "Submit Feedback" button is pressed, these structured responses are stored in the backend, feeding directly into the Feedback Dashboard (ODV-006) for aggregate analysis.

## FEEDBACK DASHBOARD

To make the feedback collected actionable, we have implemented a comprehensive Feedback Dashboard (ODV-006). While the Interface (ODV-005) focuses on single-case data entry, the Dashboard provides a centralized, aggregated view of validation efforts across the consortium.

Built on an open-source analytics platform (Grafana), the dashboard offers real-time insights into the validation process. As shown in the implemented design, the dashboard provides:

1. **Dynamic Filtering:** Top-level controls allow administrators to filter data by Model (e.g., comparing "ODELIA-AI" vs. "Cancer classification model") and by User. This capability is essential for isolating specific model versions or monitoring the activity of individual annotators.
2. **High-Level Metrics:** The dashboard presents key performance indicators at a glance, distinguishing between "Total submissions" and "Unique submissions." This logic accounts for scenarios where a user might edit their initial feedback, ensuring that the final statistics reflect the most current expert consensus.
3. **Verdict Distribution:** Visualizations break down feedback by anatomy (Left side verdicts vs. Right side verdicts). These bar charts provide an immediate visual check on model performance, displaying the count of "Agree," "Unsure," and "Disagree" responses side-by-side.
4. **Detailed Audit Trail:** A tabular view ("All submissions (history)") logs every interaction with the system. It tracks the submission kind (e.g., "initial" vs. "edit"), ensuring full traceability and data integrity for the validation study.

This dashboard enables project leads to track validation progress and identify model drift or data quality issues in real-time.

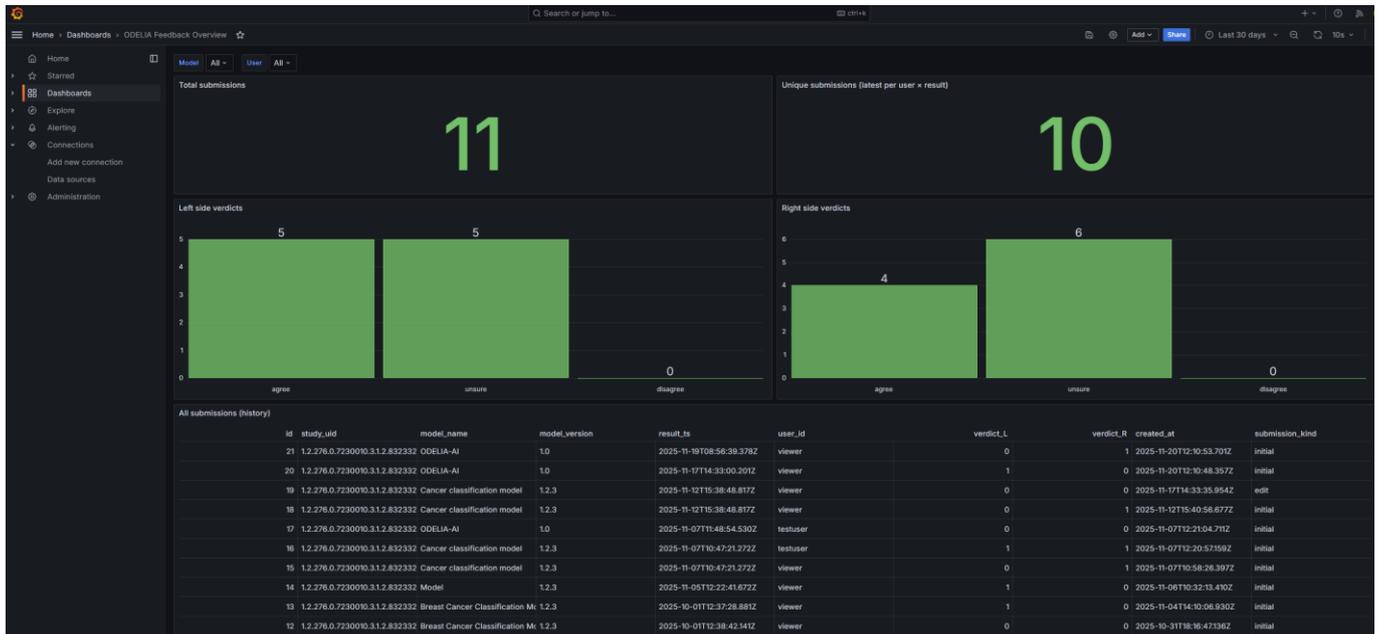


Figure 2: Feedback Dashboard

## USER AUTHENTICATION

As the viewer moves toward broader adoption and validation, tracking *who* provides annotations and feedback is essential for data integrity. We have implemented a secure User Authentication system.

- **Implementation:** We utilize an OpenID Connect (OIDC) compliant authentication flow (integrated with Keycloak or similar identity providers). This aligns with standard healthcare IT security practices.
- **Traceability:** Every feedback submission (ODV-005) is now stamped with the UUID of the authenticated user.
- **Access Control:** The system can now restrict access to specific features (e.g., only authorized radiologists can submit feedback; only administrators can access the Feedback Dashboard).
- **Password Management:** User credentials are managed by Keycloak, which stores passwords in a PostgreSQL database using secure hashing mechanisms. Passwords are not accessible in plain text. If users forget their password, an administrator can issue a password reset via the Keycloak administrative console.

This functionality addresses the requirement mentioned in D4.1 regarding the need to "link reviewer name to the feedback provided".

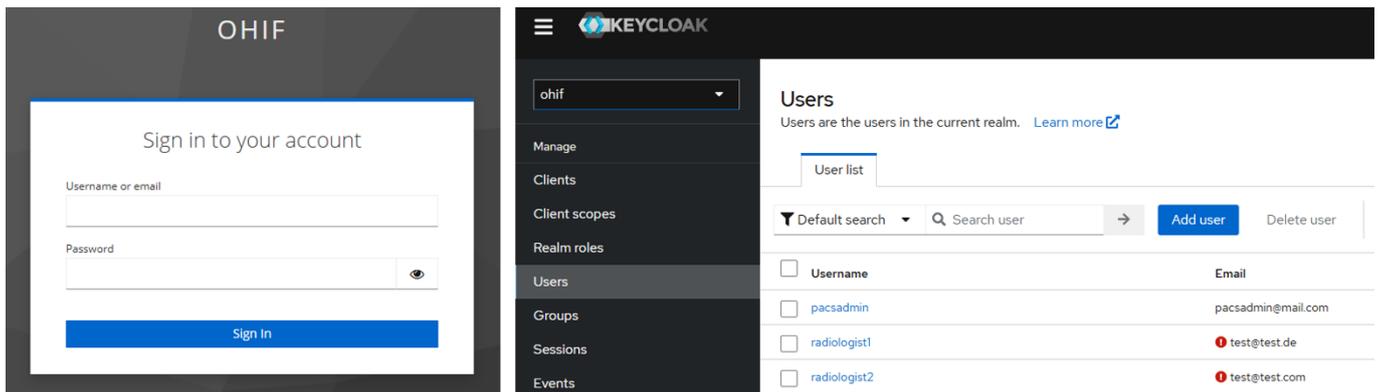


Figure 3: User Login (left). User Administration (right)

## REAL MODEL INTEGRATION

In D4.1, the viewer demonstrated capability using prototype back-end pipelines. For D4.3, we have integrated a "Real" deployment-ready model: the baseline model developed for the ODELIA MICCAI 2025 Challenge.

### Model initialization:

Model integration within the viewer is managed through a dedicated wrapper component that provides a standardized API interface, performing all model-specific preprocessing and postprocessing steps. The wrapper automatically downloads model weights from the Hugging Face repository, initializes the model, and exposes it as a self-contained inference service.

Each model is packaged as a Docker container built around this wrapper and deployed as part of the provided Docker Compose. This setup enables seamless installation and execution without manual configuration. Once deployed, the inference service is automatically recognized by the viewer, and its outputs are stored as standardized DICOM Structured Report (SR) and Secondary Capture (SC) objects to ensure interoperability within the ODELIA ecosystem.

We envision that this standardized interface will also facilitate the integration of customized or externally developed AI models into the ODELIA Viewer, allowing researchers and partners to extend the system with minimal adaptation. To our knowledge, this represents the first open-source implementation of the IHE AIW-I workflow, providing a transparent and reproducible reference for decentralized medical AI integration.

### AI Result representation:

AI inference results are saved as DICOM Structured Reports (SR) for classification outcomes and DICOM Secondary Capture (SC) images for visual annotations such as attention heatmaps. They follow the DICOM Standard (see Appendix A) and are stored in the Orthanc archive. These results can be exported as regular DICOM files and opened in any DICOM-compatible viewer or analysis tool.

Using standard medical coding schemes ensures that the exported data remain interoperable and can be easily reused for further review or research.

A significant challenge identified during the D4.1 discovery phase was the "variation in scanning sequences" and the "lack of standardization" in DICOM labels across consortium members. As anticipated in our previous report, we determined that an automated extraction solution was not yet feasible due to data heterogeneity.

To address this in ODV-008, we implemented a Manual Sequence Selection workflow:

1. **User Prompt:** When the user initiates "Send to AI," the viewer analyses the available series in the study.
2. **Selection Interface:** A modal window prompts the radiologist to map the specific sequences required by the model.
3. **Heatmap:** The integration now supports the retrieval and display of true attention maps generated by the MICCAI baseline model.

This approach ensures the MICCAI baseline model receives the correct input data regardless of the naming conventions used by the local MRI scanner, effectively solving the interoperability challenge identified in D4.1.

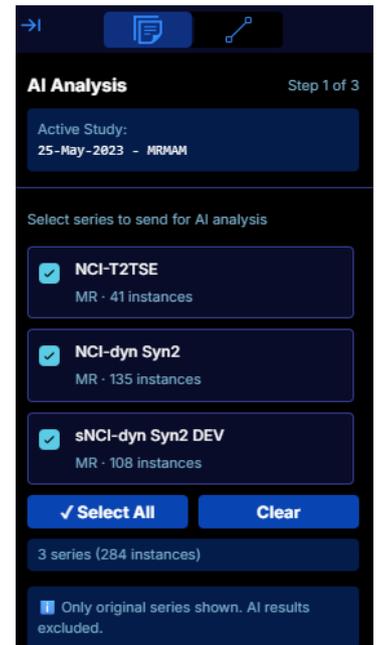


Figure 4: Series selection

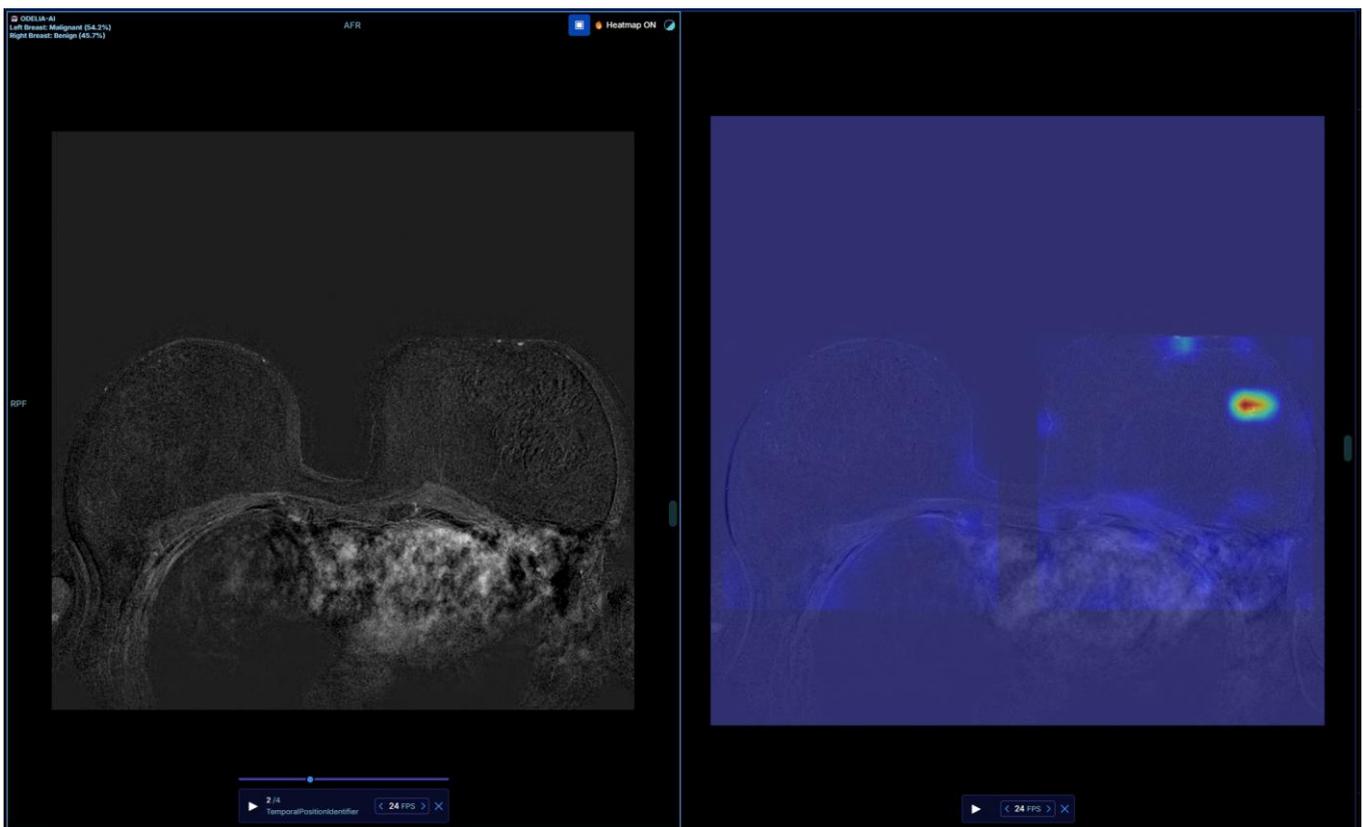


Figure 5: Heatmap visualization

## VIEWER DOCUMENTATION

To support the uptake of the ODELIA viewer by partners outside the core development team, we have significantly overhauled the documentation. Building on the basic installation instructions provided in D4.2, the new documentation includes:

- **Deployment Guides:** Detailed steps for deploying the viewer, the Orthanc backend, and the AI inference containers using Docker Compose.
- **Configuration Reference:** Explanations of environment variables required to connect the Viewer to the Authentication service (ODV-007) and the Model Registry.
- **User Manual:** A guide for radiologists explaining the new Sequence Selection workflow (ODV-008) and Feedback mechanisms (ODV-005).

This version of the ODELIA web-based MRI viewer has been published in Zenodo and is openly accessible under the DOI [10.5281/zenodo.17910662](https://doi.org/10.5281/zenodo.17910662)

## CONCLUSION

StratifAI has successfully delivered Deliverable 4.3, significantly maturing the ODELIA MRI viewer from a basic visualization tool into a platform capable of supporting rigorous AI model validation. By implementing structured feedback loops, secure authentication, and robust handling for real-world model integration, we have addressed the key technical and workflow challenges identified in D4.1. The introduction of the manual sequence selection workflow effectively mitigates the issue of heterogeneous MRI data structures, ensuring that consortium members can utilize the AI models regardless of their local scanner configurations. The system is now fully prepared for the pilot testing phase, where consortium members will use the viewer to validate the ODELIA Swarm Learning models.

With the core feature set for validation now complete, the focus for the remainder of WP4 will shift toward performance optimization and user experience refinement. Key priorities for the next phase include:

- **Improve Model Inference Speed:** We aim to optimize the communication overhead between the Orthanc backend and the AI containers to reduce the waiting time for radiologists.
- **UX/UI Refinement:** We will incorporate the qualitative feedback gathered via the new mechanisms in this release to refine the viewer's handling and design, ensuring high acceptance rates in clinical routine.

## APPENDIX A: MEDIA STORAGE CONFORMANCE STATEMENT

### INTRODUCTION

This section of the DICOM Conformance Statement specifies the Odelia Viewer compliance to DICOM requirements for the creation and storage of AI-derived Structured Reports (SR) and Secondary Capture (SC) images.

It details the DICOM Information Objects, Attributes, and roles which are supported by this software module.

The ODELIA Viewer is an OHIF extensions designed to display and evaluate AI results in form of DICOM SC/SR integrated into the ODELIA Viewer platform.

The ODELIA Viewer itself does not implement network communication (DICOM C-STORE SCP/SCU) or Media Interchange functions directly; those are inherited from the hosting DICOM Service (e.g., Orthanc Server or PACS gateway).

For full Media Interchange details, refer to the DICOM Conformance Statement of the host system.

The format of this section follows DICOM Standard PS 3.2 (Conformance). Readers are encouraged to consult that part of the standard for detailed definitions of Application Profiles, roles, and the semantics of “Used” and “Not Used” module entries.

### SUPPORTED SOP CLASSES

SOP Class Name	SOP Class UID	Role	Usage
Comprehensive SR Storage	1.2.840.10008.5.1.4.1.1.88.33	SCU	Creation of AI classification Structured Reports
Secondary Capture Image Storage	1.2.840.10008.5.1.4.1.1.7	SCU	Creation of RGB heatmap and text-overlay visualization images

No other SOP Classes are generated or interpreted by the Reporting Tool.

All generated DICOM objects are encoded using Explicit VR Little Endian (1.2.840.10008.1.2.1) Transfer Syntax.