

A standards-based infrastructure for distributed AI inference in medical imaging

Eine standardbasierte Infrastruktur für verteilte KI-Inferenz in der medizinischen Bildgebung

Abstract

Background: Integrating artificial intelligence (AI) into clinical workflows requires secure, interoperable infrastructures that connect heterogeneous clinical data systems while respecting governance and regulatory constraints. In medical imaging, PACS vendors implement the DICOMweb standard with significant variation, creating interoperability challenges for federated AI inference.

Methods: This work presents a standards-based DICOMweb adapter that mediates between heterogeneous clinical PACS implementations and federated AI services through protocol translation between DICOMweb and DIMSE, configuration-driven normalization of vendor-specific deviations, and explicit extension points for governance services such as de-identification. The approach was evaluated through clinical deployment at a university hospital and a reproducible multi-PACS simulation with two independent Orthanc backends, with a comparative benchmark against Orthanc's DICOMweb plugin for context.

Results: The adapter showed consistent performance across heterogeneous backends (median QIDO-RS latency 46–50 ms, 100% success rate). The benchmark showed that, in its default configuration, Orthanc does not proxy DICOMweb requests to remote DIMSE-based PACS and returns HTTP 404 for all WADO-RS retrievals requiring the multipart content type expected by downstream AI pipelines. The adapter served all requests successfully.

Discussion: Existing open-source tools are not designed for the mediation role required in federated AI workflows; a dedicated, standards-based mediation layer is necessary to bridge the gap between heterogeneous clinical PACS and distributed AI services.

Keywords: distributed artificial intelligence, medical imaging, DICOMweb, interoperability, PACS heterogeneity

Obioma Pelka^{1,2}

Nicolas A. Kamzol²

Kaushik Manjunatha¹

Nimesh Singh¹

Jazia Omeirat²

Felix Nensa¹

1 Institute for AI in Medicine (IKIM), University Hospital Essen, Germany

2 Data Integration Center (DIC), University Hospital Essen, Germany

Zusammenfassung

Hintergrund: Die Integration künstlicher Intelligenz (KI) in klinische Arbeitsabläufe erfordert sichere, interoperable Infrastrukturen, die heterogene klinische Datensysteme verbinden und zugleich Governance- und regulatorische Vorgaben berücksichtigen. In der medizinischen Bildgebung setzen PACS-Hersteller den DICOMweb-Standard sehr unterschiedlich um, was die Interoperabilität für föderierte KI-Inferenz erschwert.

Methoden: Diese Arbeit stellt einen standardbasierten DICOMweb-Adapter vor, der zwischen heterogenen klinischen PACS-Implementierungen und föderierten KI-Diensten vermittelt: durch Protokollübersetzung zwischen DICOMweb und DIMSE, konfigurationsgesteuerte Normalisierung herstellerspezifischer Abweichungen sowie explizite Erweiterungspunkte für Governance-Dienste wie die De-Identifikation. Der Ansatz wurde durch den klinischen Einsatz an einem Universitätsklinikum sowie eine reproduzierbare Multi-PACS-Simulation mit zwei unabhängigen

Orthanc-Backends evaluiert; zur Einordnung diene ein vergleichender Benchmark gegen das DICOMweb-Plugin von Orthanc.

Ergebnisse: Der Adapter zeigte über heterogene Backends hinweg eine konsistente Leistung (mediane QIDO-RS-Latenz 46–50 ms, Erfolgsrate 100%). Der Vergleich zeigte, dass Orthanc in der Standardkonfiguration DICOMweb-Anfragen nicht an entfernte DIMSE-basierte PACS weiterleitet und für alle WADO-RS-Abrufe mit dem von KI-Pipelines erwarteten Multipart-Content-Type HTTP 404 zurückgibt. Der Adapter beantwortete alle Anfragen erfolgreich.

Diskussion: Bestehende Open-Source-Werkzeuge sind nicht für die in förderierten KI-Workflows erforderliche Vermittlerrolle ausgelegt; eine dedizierte, standardbasierte Vermittlungsschicht ist notwendig, um die Lücke zwischen heterogenen klinischen PACS und verteilten KI-Diensten zu schließen.

Schlüsselwörter: verteilte künstliche Intelligenz, medizinische Bildgebung, DICOMweb, Interoperabilität, Heterogenität von PACS

1 Introduction

The integration of artificial intelligence (AI) into routine clinical workflows is a central objective of biomedical informatics [1], [2]. While machine learning has produced models with expert-level diagnostic performance [3], translation into clinical practice remains limited, largely due to the complexity of integrating AI systems with heterogeneous clinical data infrastructures while maintaining governance and regulatory compliance [4], [5].

Distributed AI inference, the coordinated application of trained models across multiple institutions, makes AI capacity available where clinical data reside while preserving institutional data sovereignty [6], [7]. Unlike federated learning, where data never leave the institution, inference workflows transfer image data in pseudonymised form from the data-holding site to the site hosting the AI service, with results returned to the origin. Practical realization is challenging due to heterogeneous clinical information systems. Medical imaging infrastructures are built around Picture Archiving and Communication Systems (PACS) using the DICOM standard, with DICOMweb [8] providing RESTful access. Structured metadata exchange in the broader OMI platform [9] uses HL7 FHIR [10]; this paper focuses on the DICOM imaging path. In practice, PACS vendors implement these standards with significant variation, creating interoperability challenges insufficiently addressed by standards alone [11].

Prior work on distributed AI has focused primarily on algorithmic aspects such as federated learning protocols [6], [7], [12], with comparatively little attention to the infrastructural mediation layers required for clinical deployment. This paper presents a standards-based DICOMweb adapter that mediates between heterogeneous PACS implementations and distributed AI services, evaluated at the University Hospital Essen within the Open Medical Inference (OMI) platform [9] and validated through multi-PACS simulation and a comparative benchmark.

2 Background and related work

Federated learning (FL) has emerged as a dominant paradigm for privacy-preserving multi-institutional model training [6], [7], [12], with projects such as FeTS [13] and EXAM [14] demonstrating feasibility. However, most FL research focuses on model aggregation and training orchestration, while the data access infrastructure connecting local imaging repositories to distributed services receives little attention. Distributed inference shares these infrastructural requirements but additionally needs low-latency, on-demand data access integrated with clinical workflows.

Clinical imaging systems traditionally communicate through the DICOM Message Service Element (DIMSE) protocol, in which application entities (AETs) exchange images over dedicated network associations using operations such as C-FIND (query), C-MOVE and C-GET (retrieve), and C-STORE (store); DICOM Part 10 defines the file format of stored objects. DICOMweb [8] re-expresses these operations as RESTful web services and should in principle enable vendor-neutral access through QIDO-RS (query), WADO-RS (retrieve), and STOW-RS (store), which are more readily consumed by modern AI pipelines. In practice, however, PACS vendors implement the standard with significant variation: incomplete metadata, non-standard content-type headers, vendor-specific authentication, and proprietary query parameters are commonly observed [15]. Open-source implementations such as Orthanc [16] and dcm4chee [17] serve as research-oriented PACS but are typically deployed as standalone servers rather than mediation layers for existing clinical infrastructure. Commercial enterprise-imaging products and vendor-neutral archives also address imaging integration, but as proprietary, site-licensed platforms rather than open, standards-based mediation layers embeddable in a distributed research infrastructure.

IHE integration profiles such as AI Results (AIR) and AI Workflow for Imaging (AIW-I) [18] address how AI results




Figure 1: Distributed AI infrastructure overview. The DICOMweb adapter is the central per-institution mediation component between the local PACS and the distributed services. Image data leave the institutional boundary only after pseudonymisation, and inference is performed at the service site with results returned to the origin.

are structured and inference is orchestrated, but assume conformant DICOMweb interfaces. Inference frameworks such as MONAI Deploy [19] provide model containerization and execution but do not mediate between heterogeneous PACS. The present work is complementary: DICOM-RST provides the data access mediation layer these higher-level tools presuppose.

Governance requirements – including de-identification, audit logging, and regulatory compliance [5], [20] – are critical for clinical AI deployment but not natively supported by DICOMweb. This work designs the data access pipeline with explicit extension points for governance services, treating governance as an infrastructural concern; the governance services themselves are under separate development within the OMI platform.

3 System architecture and DICOMweb adapter

The architecture is guided by five design principles:

1. **institutional autonomy by design;**
2. **standards-based interoperability;**
3. **explicit mediation** between specification and practice;
4. **governance as infrastructure;** and
5. **modularity and extensibility.**

The system consists of a local PACS, the DICOMweb adapter, client components, a gateway, and a service registry (Figure 1).

In a typical workflow, an inference request originates from a distributed AI service and reaches the institution through the gateway and service registry; the local adapter retrieves the requested study from the institutional PACS, the study is pseudonymised at the origin (by the

DE-IDicom service, see below) before it leaves the institutional boundary, inference is performed at the service site, and results are returned to the requesting clinician. The DICOMweb adapter, called DICOM-RST, is implemented in Rust and translates between the DICOMweb RESTful interface and the traditional DIMSE protocol used by clinical PACS. It implements a DIMSE backend that translates QIDO-RS, WADO-RS, and STOW-RS requests into C-FIND, C-MOVE, and C-STORE messages, providing DICOMweb access to any DIMSE-capable PACS regardless of its native web service support. An experimental S3 backend additionally supports cloud-based DICOM retrieval.

Each adapter instance maintains a site-specific YAML configuration encoding connection parameters, AET definitions, timeout settings, and protocol-level normalization rules for the respective PACS. It addresses known interoperability issues including legacy PACS that fail with compressed transfer syntaxes (configurable uncompressed mode) and systems that omit required message identifiers (sequential retrieval mode). The architecture provides extension points for governance services (de-identification, debranding) as DICOMweb-compatible proxies downstream of the adapter (Figure 2). These are not yet implemented; a dedicated de-identification service (DE-IDicom), providing DICOM tag de-identification and OCR-based debranding, is under development within the OMI platform, with defacing as a forthcoming plugin, and integrates with the adapter via DICOMweb. The source code is publicly available under the MIT license at <https://github.com/UMEssen/DICOM-RST>.




Figure 2: DICOMweb adapter processing pipeline. Steps 3 (de-identification and defacing) and 4 (logging and monitoring), shown with dashed borders, are governance services that are currently being implemented.

4 Implementation and evaluation

The adapter was evaluated in two complementary settings: a clinical deployment at the University Hospital Essen and a controlled multi-PACS simulation with two additional open-source PACS. A comparative benchmark against Orthanc's DICOMweb plugin contextualized the adapter's capabilities.

4.1 Deployment environment

Clinical site

The University Hospital Essen operates a clinical PACS (PACS-Clinical) that supports DICOMweb but exhibits vendor-specific deviations: incomplete QIDO-RS metadata, non-standard includefield handling, vendor-specific WADO-RS content-type headers, and proprietary authentication formats. It manages approximately 2,000–3,000 new studies daily, and the adapter was deployed as a containerized service within the hospital's internal network.

Simulation environment

A reproducible multi-PACS simulation was deployed using Docker Compose: two independent Orthanc 1.12 instances [16] (PACS-Orthanc-A, PACS-Orthanc-B) and one dcm4chee-arc-light 5.32 [17] (PACS-dcm4chee) with PostgreSQL. DICOM-RST was configured with three AET entries, each with independent connection pools, and a separate Orthanc instance (Orthanc-Proxy) served the comparative benchmark (Section 4.4). The complete configuration, benchmark scripts, and analysis tooling are in the repository (github.com/UMEssen/DICOM-RST). All containers ran on a single host (Apple M4, 10 cores, 32 GB RAM, macOS) over the local Docker network, so reported latencies characterise adapter processing under a single-host configuration rather than wide-area network transfer.

4.2 Interoperability robustness

The deployment across three PACS backends revealed distinct deviation profiles (Table 1): the clinical PACS exhibited the most (5 categories), followed by dcm4chee (3) and Orthanc (2). The largely complementary profiles support the argument that PACS heterogeneity is a sys-

temic challenge requiring adaptive mediation. All deviations were addressed through configuration-driven normalization without changes to the core adapter code.

4.3 Operational performance

Clinical site

Over a two-week observation period, the adapter processed 4,218 QIDO-RS queries and 1,843 WADO-RS retrievals from three distributed AI inference workflows. QIDO-RS median response times were 45–80 ms (study-level, $n=2,671$) and 120–250 ms (series-level, $n=1,547$); WADO-RS ranged from 200 ms (single-frame) to 2.5 s (volumetric CT, 200+ slices). Normalization overhead, measured via a stratified random sample of 500 queries replayed with and without mediation, was median 12 ms (95% CI: 10.8–13.2 ms; IQR: 8–15 ms), with p95 overhead of 18 ms under up to 10 parallel requests. All 6,061 requests completed successfully (0.0% error rate).

Simulation environment

The standardized test protocol issued a fixed sequence of QIDO-RS (study, series, instance) and WADO-RS (study, instance) requests against PACS-Orthanc-A and PACS-Orthanc-B ($n=10$ iterations per endpoint per backend), recording per-request latency and HTTP status (Table 2). QIDO-RS latencies were consistent across both backends (median 46–50 ms), confirming that DIMSE translation introduces stable overhead independent of the PACS instance, and both backends returned identical DICOM tag sets (19 tags per study). The 100% success rate across all 80 requests validates the adapter's reliability under controlled conditions.

4.4 Comparative benchmark: DICOM-RST vs. Orthanc

To contextualize the adapter's contribution, a comparative benchmark used Orthanc-Proxy, an Orthanc 1.12 instance with the DICOMweb plugin, on the same Docker network. Orthanc's DICOMweb plugin was chosen as the baseline because it is the most widely used open-source DICOMweb implementation and is commonly proposed as a lightweight mediation option; it is examined here to test whether an off-the-shelf tool can fill the mediation role. Both systems accessed the same test data: DICOM-RST

Table 1: PACS deviations observed and adapter handling strategies

Category	Example	Handling	In
Incomplete metadata	Missing tags in QIDO-RS	Default value injection	C
Non-standard queries	Vendor includefield syntax	Query translation	C
Content-type deviations	Incorrect multipart boundaries	Header normalization	C, D
Vendor extensions	Proprietary API parameters	Extension filtering	C
Authentication variance	Non-standard token format	Pluggable auth module	C
Transfer syntax restrictions	Compressed C-STORE rejection	Uncompressed mode	O, D
Response ID omission	Missing C-STORE message ID	Sequential retrieval	O
Encoding variance	Non-UTF-8 tag values	Encoding normalization	D

C= PACS-Clinical, O= PACS-Orthanc, D= PACS-dcm4chee

Table 2: DICOM-RST latency across two Orthanc backends (n=10 per endpoint per backend). Values in ms.

Endpoint	Backend	Mean	Median	P95	Success
QIDO-RS studies	Orthanc-A /	45.6 /	45.6 /	48.0 /	100%
	Orthanc-B	46.8	46.9	48.6	
QIDO-RS series	Orthanc-A /	47.4 /	47.6 /	48.7 /	100%
	Orthanc-B	47.6	47.3	51.4	
QIDO-RS instances	Orthanc-A /	49.8 /	49.5 /	54.9 /	100%
	Orthanc-B	46.4	45.6	49.3	
WADO-RS study	Orthanc-A /	59.1 /	55.5 /	95.1 /	100%
	Orthanc-B	57.2	53.8	69.7	

P95= 95th percentile**Table 3: Comparative benchmark: DICOM-RST (DIMSE translation) vs. Orthanc DICOMweb (local database), n=20 per endpoint. Latencies in ms.**

Endpoint	DICOM-RST			Orthanc DICOMweb		
	Mean	Median	P95	Mean	Median	P95
QIDO-RS studies	46.0	46.8	49.1	3.6	3.8	4.9
QIDO-RS series	47.7	48.2	48.9	2.0	1.9	4.6
QIDO-RS instances	50.0	49.9	52.8	4.5	3.8	10.6
WADO-RS study	56.4	55.2	61.9	HTTP 404 (0% success)		
WADO-RS instance	57.0	56.9	65.5	HTTP 404 (0% success)		

P95= 95th percentile

via DIMSE translation to PACS-Orthanc-A, and Orthanc-Proxy via its native DICOMweb plugin serving locally stored data. Identical queries (n=20 per endpoint) were executed through both paths.

An important architectural distinction emerged: **Orthanc's DICOMweb plugin serves only locally stored data; it cannot translate DICOMweb requests into DIMSE operations for remote PACS systems.** The benchmark therefore evaluates Orthanc's best case (local database) against DICOM-RST's DIMSE translation path (Table 3).

For QIDO-RS, Orthanc's local database was 10–25× faster (2–5 ms vs. 46–50 ms) – expected, since DICOM-RST translates each query into a DIMSE C-FIND while Orthanc queries its local SQLite directly. The ~45 ms translation overhead is the inherent cost of enabling access to any DIMSE-capable PACS without data pre-loading.

Critically, Orthanc's DICOMweb plugin returned **HTTP 404 for all WADO-RS requests** when Accept: multipart/related; type="application/dicom" was specified (n=40, 0% suc-

cess). This content type is required by the DICOMweb standard for DICOM Part 10 objects and expected by downstream AI pipelines. DICOM-RST served all 40 requests successfully (100%, median 55–57 ms, 330 KB per response), demonstrating two limitations of Orthanc as a mediation layer: it cannot proxy DICOMweb to remote DIMSE-based PACS, and it does not support the multipart DICOM content type required for AI pipeline consumption.

5 Discussion

Infrastructural mediation layers are critical for enabling distributed AI workflows in clinical imaging. Deviation profiles differed substantially across backends, with largely complementary categories of non-conformance, reinforcing the need for a configuration-driven approach. The controlled benchmark confirmed consistent, reproducible performance across independent PACS instances

(median QIDO-RS 46–50 ms, 100% success across 80 requests), validating the clinical observations.

The comparative benchmark yielded two findings of practical significance: in its default configuration, Orthanc does not proxy DICOMweb requests to remote DIMSE-based PACS, and its DICOMweb plugin returned HTTP 404 for all WADO-RS requests requiring multipart/related DICOM Part 10 objects. Existing open-source tools are thus not designed for the mediation role required in distributed AI workflows; DICOM-RST addresses both gaps through DIMSE translation and standards-compliant response generation.

These results map onto the five design principles of Section 3: standards-based interoperability and explicit mediation via protocol translation and configuration-driven normalization; modularity and extensibility via handling new deviation profiles through configuration alone and pluggable governance proxies; governance as infrastructure via the DE-IDicom extension point; and institutional autonomy via the adapter's role as the per-institution access point within the OMI network.

Reproducibility

DICOM-RST is available as open-source software (MIT license) at <https://github.com/UMEssen/DICOM-RST>. The repository includes the multi-PACS Docker Compose configuration, automated benchmark scripts, and a Python analysis tool generating summary statistics from raw timing data.

Threats to validity

The clinical evaluation covers one site with one commercial PACS, and the open-source systems were tested in a controlled environment that may not exhibit the full range of clinical deviations. Performance was assessed under moderate concurrent load (up to 10 parallel requests); high-concurrency behavior remains unassessed. The governance extension points are architectural, with empirical evaluation to follow completion of the DE-IDicom service. The comparative benchmark evaluates Orthanc in its default configuration; custom Lua scripting could address some metadata gaps but would not resolve the WADO-RS limitation or enable remote PACS access.

Future work

Planned work includes:

1. multi-site clinical deployment at the University Hospital Bonn for cross-vendor validation;
2. integration of de-identification and defacing services;
3. formal stress testing under high concurrency;
4. structured compliance-grade audit logging; and
5. a direct comparative benchmark against dcm4chee.

6 Conclusion

This work presented a standards-based DICOMweb adapter for distributed AI inference in medical imaging that mediates between heterogeneous clinical PACS and distributed AI services through DIMSE protocol translation, configuration-driven normalization, and an extensible governance architecture. Evaluation across three heterogeneous PACS backends showed that the adapter handles distinct deviation profiles through declarative configuration alone, with controlled benchmarks confirming consistent performance (median 46–50 ms QIDO-RS, 100% success). A comparative benchmark against Orthanc's DICOMweb plugin revealed that existing tools lack both remote DICOMweb-to-DIMSE proxying and the WADO-RS multipart support required by AI inference pipelines, confirming the necessity of a dedicated mediation layer.

Notes

Authors' ORCIDs

- Obioma Pelka: 0000-0001-5156-4429
- Kaushik Manjunatha: 0009-0006-0993-6173
- Felix Nensa: 0000-0002-5811-7100

Competing interests

The authors declare that they have no competing interests.

Acknowledgements

This work was funded by the German Federal Ministry of Research, Technology and Space (Bundesministerium für Forschung, Technologie und Raumfahrt, BMFTR) as part of the Medical Informatics Initiative, grant number 01ZZ2315A-P (OMI Consortium of the German Medical Informatics Initiative, <https://omi.ikim.nrw>).

The authors thank the clinical IT team at the University Hospital Essen for their support during deployment and evaluation.

References

1. Topol EJ. High-performance medicine: the convergence of human and artificial intelligence. *Nat Med*. 2019 Jan;25(1):44–56. DOI: 10.1038/s41591-018-0300-7
2. Esteva A, Kuprel B, Novoa RA, Ko J, Swetter SM, Blau HM, Thrun S. Dermatologist-level classification of skin cancer with deep neural networks. *Nature*. 2017 Feb;542(7639):115–118. DOI: 10.1038/nature21056
3. Litjens G, Kooi T, Bejnordi BE, Setio AAA, Ciompi F, Ghafoorian M, van der Laak JAWM, van Ginneken B, Sánchez CI. A survey on deep learning in medical image analysis. *Med Image Anal*. 2017 Dec;42:60–88. DOI: 10.1016/j.media.2017.07.005

4. Goodman KW. Ethics, medicine, and information technology: intelligent machines and the transformation of health care. Cambridge: Cambridge University Press; 2016.
5. European Commission. Ethics guidelines for trustworthy AI. Brussels: European Commission; 2019.
6. Rieke N, Hancox J, Li W, Milletari F, Roth HR, Albarqouni S, Bakas S, Galtier MN, Landman BA, Maier-Hein K, Ourselin S, Sheller M, Summers RM, Trask A, Xu D, Baust M, Cardoso MJ. The future of digital health with federated learning. NPJ Digit Med. 2020;3:119. DOI: 10.1038/s41746-020-00323-1
7. Sheller MJ, Reina GA, Edwards B, Martin J, Bakas S. Multi-Institutional Deep Learning Modeling Without Sharing Patient Data: A Feasibility Study on Brain Tumor Segmentation. Brainlesion. 2019;11383:92-104. DOI: 10.1007/978-3-030-11723-8_9
8. DICOM Standards Committee. DICOMweb. Rosslyn (VA): NEMA; 2024.
9. Pelka O, Sigle S, Werner P, Schweizer ST, Iancu A, Scherer L, Kamzol NA, Eil JH, Apfelbacher T, Seletkov D, Susetzky T, May MS, Bucher AM, Fegeler C, Boeker M, Braren R, Prokosch HU, Nensa F. Demokratisierung von KI im Gesundheitswesen mit Open Medical Inference (OMI): Protokolle, Datenaustausch und KI-Integration Democratizing AI in Healthcare with Open Medical Inference (OMI): Protocols, Data Exchange, and AI Integration. Rofo. 2026 Feb;198(2):173-184. DOI: 10.1055/a-2651-6653
10. HL7 International. HL7 FHIR. Ann Arbor (MI): HL7 International; 2024.
11. Lehne M, Sass J, Essenwanger A, Schepers J, Thun S. Why digital medicine depends on interoperability. NPJ Digit Med. 2019;2:79. DOI: 10.1038/s41746-019-0158-1
12. Kairouz P, McMahan HB, et al. Advances and open problems in federated learning. Found Trends Mach Learn. 2021 Jun 23;14(1-2):1-210. DOI: 10.1561/22000000083
13. Pati S, Baid U, Edwards B, Sheller M, Wang SH, Reina GA, Foley P, Gruzdev A, Karkada D, Davatzikos C, et al. Federated learning enables big data for rare cancer boundary detection. Nat Commun. 2022 Dec;13(1):7346. DOI: 10.1038/s41467-022-33407-5
14. Dayan I, Roth HR, Zhong A, Harouni A, Gentili A, Abidin AZ, Liu A, Costa AB, Wood BJ, Tsai CS, et al. Federated learning for predicting clinical outcomes in patients with COVID-19. Nat Med. 2021 Oct;27(10):1735-1743. DOI: 10.1038/s41591-021-01506-3
15. Genereaux BW, Dennison DK, Ho K, Horn R, Silver EL, O'Donnell K, Kahn CE Jr. DICOMweb™: Background and Application of the Web Standard for Medical Imaging. J Digit Imaging. 2018 Jun;31(3):321-326. DOI: 10.1007/s10278-018-0073-z
16. Jodogne S. The Orthanc Ecosystem for Medical Imaging. J Digit Imaging. 2018 Jun;31(3):341-352. DOI: 10.1007/s10278-018-0082-y
17. Warnock MJ, Toland C, Evans D, Wallace B, Nagy P. Benefits of using the DCM4CHE DICOM archive. J Digit Imaging. 2007 Nov;20 Suppl 1(Suppl 1):125-129. DOI: 10.1007/s10278-007-9064-1
18. Integrating the Healthcare Enterprise (IHE). IHE radiology technical framework supplement: AI results (AIR) and AI workflow for imaging (AIW-I). Oak Brook (IL): IHE International; 2023.
19. MONAI Consortium. MONAI Deploy App SDK: a framework for developing, packaging, and deploying AI applications in healthcare. 2023.
20. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations. Off J Eur Union. 2024 Jul 12. Available from: <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>

Corresponding author:

Obioma Pelka
Institute for Artificial Intelligence in Medicine (IKIM),
University Hospital Essen (UME), Girardetstr. 2, 45131
Essen, Germany
obioma.pelka@googlemail.com

Please cite as

Pelka O, Kamzol NA, Manjunatha K, Singh N, Omeirat J, Nensa F. A standards-based infrastructure for distributed AI inference in medical imaging. GMS Med Inform Biom Epidemiol. 2026;22:Doc12. DOI: 10.3205/mibe000310, URN: urn:nbn:de:0183-mibe0003100

This article is freely available from

<https://doi.org/10.3205/mibe000310>

Published: 2026-09-22

Copyright

©2026 Pelka et al. This is an Open Access article distributed under the terms of the Creative Commons Attribution 4.0 License. See license information at <http://creativecommons.org/licenses/by/4.0/>.