

Interoperable clinical data integration for distributed medical AI: The Open Medical Inference Gateway

Interoperable Integration klinischer Daten für verteilte medizinische KI: Das Open Medical Inference Gateway

Abstract

Background: Integrating AI inference services into clinical workflows remains a persistent challenge, primarily because inference backends and healthcare data standards operate with fundamentally different assumptions. The EU AI Act further requires transparency and traceability for high-risk AI systems in healthcare.

Methods: This paper presents the Open Medical Inference (OMI) Gateway, an interoperability layer that mediates between FHIR- and DICOM-based clinical data environments and heterogeneous AI inference backends. The gateway was implemented in Python using FastAPI with asynchronous job processing via Celery and RabbitMQ. The architecture follows a service-oriented plugin model that separates shared interoperability concerns from service-specific transformation logic. Input and output data are validated against FHIR profiles and DICOM specifications at multiple pipeline stages, producing OperationOutcome resources as a machine-readable audit trail.

Results: End-to-end pipeline latency was measured across 100 sequential inference requests, yielding a median latency of 263 ms and a 100% success rate. The design was evaluated against interoperability requirements from German Medical Informatics Initiative (MII) Data Integration Centers. The system is deployable with a single command and available as open-source software.

Discussion: The results indicate that embedding validation and structured artifact handling directly into the inference execution path yields a more governable and auditable integration model than ad-hoc service adapters.

Keywords: medical informatics, Open Medical Inference, gateway, FHIR, DICOM, AI inference, interoperability, open source, MII, service-oriented architecture

Zusammenfassung

Hintergrund: Die Integration von KI-Inferenzdiensten in klinische Arbeitsabläufe bleibt eine anhaltende Herausforderung, vor allem weil Inferenz-Backends und Standards für Gesundheitsdaten von grundlegend unterschiedlichen Annahmen ausgehen. Der EU AI Act fordert zudem Transparenz und Nachvollziehbarkeit für Hochrisiko-KI-Systeme im Gesundheitswesen.

Methoden: Diese Arbeit stellt das Open Medical Inference (OMI) Gateway vor, eine Interoperabilitätsschicht, die zwischen FHIR- und DICOM-basierten klinischen Datenumgebungen und heterogenen KI-Inferenz-Backends vermittelt. Das Gateway wurde in Python mit FastAPI und asynchroner Auftragsverarbeitung über Celery und RabbitMQ implementiert. Die Architektur folgt einem serviceorientierten Plug-in-Modell, das gemeinsame Interoperabilitätsaufgaben von dienstspezifischer Transformationslogik trennt. Ein- und Ausgabedaten werden an mehreren Stellen der Pipeline gegen FHIR-Profil und DICOM-Spezifikationen va-

Obioma Pelka^{1,2}

Jan Eil¹

Kaushik Manjunatha¹

Nimesh Singh¹

Jazia Omeirat²

Stefan Sigle³

Georg Mathes³

Simon-Tobias

Schweizer⁴

Shura-Roman Stump⁵

Gerda Girdziunaite⁶

Felix Nensa¹

1 Institute for Artificial Intelligence in Medicine (IKIM), University Hospital Essen (UME), Essen, Germany

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

3 MOLIT Institute for Personalized Medicine, Heilbronn, Germany

4 GECKO Institute for Medicine, Informatics and Economics, Heilbronn University of Applied Sciences (HHN), Heilbronn, Germany

5 Department of Radiology, University Medical Center, Freiburg, Germany

6 Medical Data Integration Center (MeDIC), Charité – Universitätsmedizin Berlin, Berlin, Germany

lidiert und erzeugen OperationOutcome-Ressourcen als maschinenlesbaren Audit-Trail.

Ergebnisse: Die Ende-zu-Ende-Latenz wurde über 100 aufeinanderfolgende Inferenzanfragen gemessen und ergab eine mediane Latenz von 263 ms bei einer Erfolgsrate von 100%. Das Konzept wurde anhand von Interoperabilitätsanforderungen aus Datenintegrationszentren der deutschen Medizininformatik-Initiative (MII) evaluiert. Das System ist mit einem einzigen Befehl einsatzbereit und als Open-Source-Software verfügbar.

Diskussion: Die direkte Einbettung von Validierung und strukturierter Artefaktverarbeitung in den Ausführungspfad der Inferenz ergibt ein besser steuer- und auditierbares Integrationsmodell als ad-hoc entwickelte Dienstadapter.

Schlüsselwörter: medizinische Informatik, Open Medical Inference, Gateway, FHIR, DICOM, KI-Inferenz, Interoperabilität, Open Source, MII, serviceorientierte Architektur

1 Introduction

Healthcare AI services are now routinely deployed as remote or containerized endpoints, but their operational integration into clinical information systems remains fragmented. Hospitals may operate FHIR servers, PACS archives, and Data Integration Centers (DICs), but connecting these to AI inference endpoints requires solving several distinct problems, including semantic alignment between healthcare standards and computation-oriented model protocols, lifecycle management for asynchronous jobs, validation of inputs and outputs, and consistent artifact storage. Without a dedicated integration layer, each new AI service introduces its own adapter logic, leading to systems that scale poorly and are difficult to govern or audit.

Several platforms address aspects of this challenge. Kaapana [1] provides workflow orchestration for medical data analysis. The Joint Imaging Platform (JIP) [2] enables federated analytics across the German Cancer Consortium. MONAI Deploy [3] offers containerized inference services. Inference platforms such as Triton [4] or KServe [5] optimize for throughput but leave semantic transformation to the caller. The IHE profiles AIW-I and AIR [6], [7] address result communication but not inference governance. None provide a complete, standards-compliant inference gateway with built-in governance mechanisms. The Open Medical Inference (OMI) protocol [8] defines standardized interfaces for AI service discovery, data exchange, and inference invocation based on HL7 FHIR (Release 4B, v4.3.0) [9] and DICOMweb [10]. The corresponding FHIR Implementation Guide is published on Simplifier [11]. The OMI framework was designed within the German Medical Informatics Initiative [12], [13], [14]. The OMI Gateway is the key execution boundary, providing an interoperability layer that validates, stages, and manages multimodal clinical data before and after model execution.

This work defines the architectural properties a gateway component needs to fulfill interoperability requirements in distributed medical AI, implements them in the OMI Gateway, and evaluates its operational functionality. Our contributions are threefold:

1. We describe the architecture of a production-capable inference gateway based on open standards.
2. We present a plugin architecture enabling AI developers to integrate services with minimal effort.
3. We evaluate the system's operational performance and validate the inference pipeline, including a multi-stage FHIR-based validation mechanism producing machine-readable audit trails.

2 Background

2.1 Clinical data standards in AI workflows

HL7 FHIR and DICOM occupy complementary roles in clinical AI workflows. HL7 FHIR provides structured representations of patient context, observations, and workflow parameters [9], while DICOM carries imaging data and modality-specific semantics [10]. In multimodal AI settings, neither standard alone is sufficient. As Lehne et al. [15] observe, interoperability is a prerequisite in digital medicine, not a peripheral concern.

Inference protocols such as KServe's Open Inference Protocol [5] assume inputs are already model-ready tensors. They do not define how FHIR resources or DICOM series should be validated or mapped. Placing this alignment inside each AI service leads to duplicated effort and inconsistent behavior.




Figure 1: Microservice architecture of the OMI Gateway connecting clinical data environments to distributed AI inference services

2.2 Data Integration Centers and the MII context

The German Medical Informatics Initiative (MII) provides a reference for federated clinical data integration [12], [13]. Data Integration Centers (DICs) consolidate institutional data and expose it through FHIR interfaces. Recent work has demonstrated scalable DICOM-to-FHIR pipelines [16], making AI-ready data provisioning feasible at scale. The OMI Gateway is positioned as the interface between DIC-level data structures and distributed AI services.

3 Methods

3.1 System architecture

The Gateway follows a microservice architecture implemented via Docker Compose with six services: the API server, an asynchronous worker, PostgreSQL, RabbitMQ, SeaweedFS (S3-compatible storage), and a FHIR validation service (Figure 1).

The architecture rests on four design decisions:

1. **enforced conformance** – validating all inputs and outputs against FHIR/DICOM profiles;
2. **separation of concerns** – the platform manages authentication, job lifecycles, and storage, while plugins handle service-specific logic;
3. **asynchronous execution** – decoupling processing from HTTP handling;
4. **artifacts as first-class objects** – storing inputs and outputs as addressable entities for auditability [17].

The API server, implemented using FastAPI [18], creates a job record on each submission, uploads FHIR resources and DICOM files to S3-compatible storage encrypted at rest, and dispatches the task via RabbitMQ. An HTTP 202 response with the job identifier is returned immediately. The worker retrieves data from storage, executes the pipeline, and uploads results back, allowing independent scaling of servers and workers.

3.2 Processing pipeline

Each request passes through six phases (Figure 2), with the inter-component message flow shown in Figure 3:

1. **Pre-validation:** FHIR resources are validated against the service's input StructureDefinition using an external FHIR validator. DICOM files are validated within the worker process for required tags (e.g., Patient ID, Modality), correct VR types, and VM ranges [19]. Both produce OperationOutcome resources.
2. **Service-specific validation:** The plugin performs domain-specific checks, e.g., verifying CT modality.
3. **Preprocessing:** Clinical inputs are transformed into model-ready representations.
4. **Inference:** The plugin executes locally or delegates to a remote backend.
5. **Postprocessing:** Predictions are translated into FHIR resources or DICOM segmentation objects.
6. **Post-validation:** Outputs are validated against the declared output StructureDefinition.

Each validation stage produces OperationOutcome resources stored alongside job artifacts. If a plugin raises an unhandled exception, the worker catches the error, sets the job to *failed*, and generates an OperationOutcome with severity *fatal*.




Figure 2: Six-stage pipeline. Validation stages (blue) produce OperationOutcome resources as an audit trail. Processing stages (green) handle transformation and inference.




Figure 3: End-to-end job execution sequence. The client receives HTTP 202 immediately; processing continues asynchronously. OperationOutcome resources are persisted at each validation stage. If validation fails, the job is marked as failed and the OperationOutcome is returned via the job endpoint.

Three failure modes illustrate this:

1. a DICOM file missing the *Modality* tag triggers pre-validation termination;
2. a FHIR resource with an undefined element produces a *warning* that may allow continuation depending on severity thresholds;
3. missing DICOM files trigger a fatal OperationOutcome during service validation.

3.3 Plugin architecture

AI services are integrated through Python entry points [20]. Developers implement the abstract *BaseService* class (Table 1), packaged as an independent Python package. On startup, the gateway discovers and loads all installed service packages. The gateway handles data transport, encryption at rest, validation, and job lifecycle; developers only implement the service-specific methods.

Table 1: BaseService interface methods

Method	Description
input_parameters	FHIR StructureDefinition for input validation
output_parameters	FHIR StructureDefinition for output validation
is_ready() / is_live()	Service readiness / liveness checks
validate() *	Domain-specific input validation
preprocess()	Input data transcoding
infer()	AI model execution
postprocess()	Output transcoding to FHIR and DICOM

* Default no-op implementation

Table 2: Pipeline durations across 100 requests (dummy service, single worker)

Pipeline stage	Mean (ms)	SD (ms)	Median (ms)	IQR (ms)	P95 (ms)
Upload (API server)	29	4.1	30	26–31	36
Processing (worker)	232	21.4	234	232–237	240
End-to-end	261	22.0	263	259–268	276

SD= standard deviation; IQR= interquartile range; P95= 95th percentile

3.4 API design and security

The gateway exposes RESTful endpoints for service listing (*GET /services*), readiness and liveness checks, input validation (*POST /services/{id}/validate*), inference submission (*POST /services/{id}/infer*), and job retrieval (*GET /jobs/{id}*). Input data is submitted as multipart form data with HL7 FHIR conformant JSON and DICOM files. Authentication uses API keys (*X-API-Key* header) stored as SHA-256 hashes with role-based permissions. *OperationOutcome* resources at each stage, together with timestamped job status logs, provide a complete audit trail aligned with the EU AI Act [21], and AMIA's AI principles [17]. Where an integrated AI service qualifies as a medical device, the EU Medical Device Regulation (MDR) [22] additionally applies; the gateway's persistent, machine-readable audit trail can support the traceability and post-market surveillance obligations arising from the MDR.

4 Results

We now evaluate the Gateway's operational behavior and its fulfillment of the interoperability requirements introduced above. The gateway was validated using dummy inference service implementing the complete BaseService interface. Testing confirmed the successful processing of multi-file DICOM submissions alongside FHIR metadata through programmatic access and Swagger UI. The asynchronous job pattern correctly transitions between pending, executing, and completed.

End-to-end latency was measured across 100 sequential requests using two DICOM files (*CT_small.dcm*, 38 KB and *MR_small.dcm*, 10 KB), chosen to isolate gateway overhead. Table 2 shows all 100 requests completed

successfully. The median 263 ms demonstrates low overhead; in production, inference time is additive.

4.1 Interoperability evaluation

The plugin interface allows routing requests to either a central remote inference service or a local compute node, such as a RACOON-connected GPU server [23], without changing the client-facing API. Consider a segmentation service accepting a CT study with a FHIR parameter set: the gateway validates both inputs, dispatches processing asynchronously, and reconstructs the segmentation as a DICOM object with an optional FHIR diagnostic report. The gateway realizes the following architectural properties derived from MII-aligned interoperability requirements: FHIR profile conformance through profile-aware input/output validation producing *OperationOutcome* resources; DICOM handling through structural validation against specification expectations; asynchronous execution via background workers decoupled from HTTP handling; service discoverability through a versioned plugin registry with queryable FHIR profiles; multimodal data staging with separate artifact directories per modality; audit and traceability through job lifecycle records combined with *OperationOutcome* resources; and backend flexibility through plugin routing that transparently directs to remote services or local GPU nodes. Iterative testing revealed five integration issues across storage, validation, and deployment. The system is deployable with *docker compose up*. The source code will be made available upon publication.

Table 3: Comparison with related platforms

Feature	OMI GW	Kaapana	JIP	MONAI
OMI protocol compliance	Yes	No	No	No
FHIR validation pipeline	Yes	No	No	No
OperationOutcome audit trail	Yes	No	No	No
Plugin-based extensibility	Yes	Yes	Yes	Yes
DICOMweb support	Yes	Yes	No	Yes
Open source	Yes	Yes	Yes	Yes
Federated learning	No	Yes	Yes	Yes

5 Discussion

The OMI Gateway's principal contribution lies in relocating interoperability enforcement from individual service adapters into a shared integration component. When validation is handled centrally, all AI services benefit from consistent conformance checking, and standards-conformant outputs can be consumed downstream without additional transformation. This positions the gateway as a genuine integration component rather than a thin proxy [17].

Compared to existing platforms (Table 3), the gateway occupies a distinct position: Kaapana, JIP, and MONAI Deploy provide federated model training and workflow orchestration but lack a FHIR-based validation pipeline with structured audit trails. The OMI Gateway focuses on standardized inference invocation with built-in governance. The FHIR-based validation pipeline creates a complete, machine-readable audit trail supporting EU AI Act [21] transparency requirements. Together with the OMI DICOMweb adapter [24], it provides the core standards-based infrastructure for clinical AI inference.

5.1 Limitations and future work

The current implementation uploads files synchronously, which may introduce latency for large DICOM datasets. Worker concurrency is limited to a single process; horizontal scaling is supported via additional Celery workers. Federation features including registry synchronization and heartbeat exchange are extension targets. Consent management is expected to rely on surrounding DIC components.

Future work includes integration with the OMI Service Registry [25] through a FHIR-based heartbeat mechanism, federation with RACOON [23], multi-step pipeline orchestration, and governance-aligned monitoring [17]. Validation with production AI models at clinical sites is planned.

Notes

Authors' ORCIDs

- Obioma Pelka: 0000-0001-5156-4429
- Jan Eil: 0009-0005-9879-5722
- Kaushik Manjunatha: 0009-0006-0993-6173
- Stefan Sigle: 0000-0001-7475-4583
- Georg Mathes: 0009-0000-2060-5558
- Simon-Tobias Schweizer: 0000-0002-0888-4301
- Shura-Roman Stump: 0009-0009-8542-6556
- Gerda Girdziunaite: 0009-0000-5562-453X
- Felix Nensa: 0000-0002-5811-7100

Author contributions

Obioma Pelka and Jan Eil contributed equally to this work.

Competing interests

The authors declare that they have no competing interests.

References

1. Akünal Ü, Bujotzek M, Denner S, Hamm B, Kades K, Schader P, Scherer J, Nolden M, Neher P, Floca R, Maier-Heine K. Kaapana: a comprehensive open-source platform for integrating AI in medical imaging research environments [Preprint]. arXiv. 2025. arXiv:2512.09644. DOI: 10.48550/arXiv.2512.09644
2. Scherer J, Nolden M, Kleesiek J, Metzger J, Kades K, Schneider V, Bach M, Sedlaczek O, Bucher AM, Vogl TJ, Grünwald F, Kühn JP, Hoffmann RT, Kotzerke J, Bethge O, Schimmöller L, Antoch G, Müller HW, Daul A, Nikolaou K, la Fougère C, Kunz WG, Ingrisch M, Schachtner B, Ricke J, Bartenstein P, Nensa F, Radbruch A, Umutlu L, Forsting M, Seifert R, Herrmann K, Mayer P, Kauczor HU, Penzkofer T, Hamm B, Brenner W, Kloeckner R, Düber C, Schreckenberger M, Braren R, Kaissis G, Makowski M, Eiber M, Gafita A, Trager R, Weber WA, Neubauer J, Reiser M, Bock M, Bamberg F, Hennig J, Meyer PT, Ruf J, Haberkorn U, Schoenberg SO, Kuder T, Neher P, Floca R, Schlemmer HP, Maier-Hein K. Joint Imaging Platform for Federated Clinical Data Analytics. JCO Clin Cancer Inform. 2020 Nov;4:1027-1038. DOI: 10.1200/CCI.20.00045
3. Cardoso MJ, Li W, Brown R, Ma N, Kerfoot E, Wang Y, Murrey B, Myronenko A, Zhao C, Yang D, et al. MONAI: an open-source framework for deep learning in healthcare [Preprint]. arXiv. 2022. arXiv:2211.02701. DOI: 10.48550/arXiv.2211.02701

4. NVIDIA Corporation. Triton Inference Server. Santa Clara (CA): NVIDIA; 2026 [cited 2026 Jun 23]. Available from: <https://developer.nvidia.com/triton-inference-server>
5. KServe. Open Inference Protocol (V2). KServe; 2026 [cited 2026 Jun 23]. Available from: <https://kserve.github.io/website/docs/concepts/architecture/data-plane/v2-protocol>
6. IHE Radiology Technical Committee. IHE Radiology Technical Framework Supplement: AI workflow for imaging (AIW-I). Rev. 1.1. Oak Brook (IL): IHE International; 2020 Aug 6 [cited 2026 Jun 23]. Available from: https://www.ihe.net/uploadedFiles/Documents/Radiology/IHE_RAD_Suppl_AIW-I.pdf
7. IHE Radiology Technical Committee. IHE Radiology Technical Framework Supplement: AI results (AIR). Rev. 1.3. Oak Brook (IL): IHE International; 2025 Aug 8 [cited 2026 Jun 23]. Available from: https://www.ihe.net/uploadedFiles/Documents/Radiology/IHE_RAD_Suppl_AIR.pdf
8. 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
9. Health Level Seven International. HL7 FHIR. Release 4B (v4.3.0). Ann Arbor (MI): HL7; 2022 [cited 2026 Jun 23]. Available from: <https://hl7.org/fhir/R4B/>
10. 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
11. OMI Consortium. OMI Protocol Implementation Guide. Simplifier.net; 2024 [cited 2026 Jun 23]. Available from: <https://simplifier.net/guide/OMI-Protocol-IG/Open-Medical-Inference-Protocol>
12. Medical Informatics Initiative (MII). Data integration centres. Berlin: MII Coordination Office; 2024 [cited 2026 Jun 23]. Available from: <https://www.medizininformatik-initiative.de/en/consortia/data-integration-centres>
13. Medical Informatics Initiative (MII). Core data set. Berlin: MII Coordination Office; 2024 [cited 2026 Jun 23]. Available from: <https://www.medizininformatik-initiative.de/en/medical-informatics-initiatives-core-data-set>
14. Sigle S, Werner P, Schweizer S, Caldeira L, Hosch R, Dyrba M, Fegeler C. Bridging the gap between (AI-) services and their application in research and clinical settings through interoperability: the OMI-protocol. Hannover: Technische Informationsbibliothek (TIB); 2024. DOI :10.34657/13458
15. 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
16. Iancu A, Bauer J, May MS, Prokosch HU, Dörfler A, Uder M, Kapsner LA. Large-Scale Integration of DICOM Metadata into HL7-FHIR for Medical Research. *Methods Inf Med*. 2024 Sep;63(3-04):77-84. DOI: 10.1055/a-2521-4250
17. Solomonides AE, Koski E, Atabaki SM, Weinberg S, McGreevey JD, Kannry JL, Petersen C, Lehmann CU. Defining AMIA's artificial intelligence principles. *J Am Med Inform Assoc*. 2022 Mar;29(4):585-591. DOI: 10.1093/jamia/ocac006
18. Ramírez S. FastAPI. 2026 [cited 2026 Jun 23]. Available from: <https://fastapi.tiangolo.com/>
19. Pydicom. dicom-validator [software]. Version 0.8.2. PyPI; 2026 [cited 2026 Jun 23]. Available from: <https://pypi.org/project/dicom-validator/>
20. Python Packaging Authority. Entry points specification. PyPA; 2026 [cited 2026 Jun 23]. Available from: <https://packaging.python.org/en/latest/specifications/entry-points/>
21. 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 certain Union legislative acts (Artificial Intelligence Act). Off J Eur Union. 2024 Jul 12; L 2024/1689. Available from: <http://data.europa.eu/eli/reg/2024/1689/oj>
22. Regulation (EU) 2017/745 of the European Parliament and of the Council of 5 April 2017 on medical devices, amending Directive 2001/83/EC, Regulation (EC) No 178/2002 and Regulation (EC) No 1223/2009 and repealing Council Directives 90/385/EEC and 93/42/EEC. Off J Eur Union. 2017 May 5;60(L117):1-175.
23. Network University Medicine (NUM). RACoon: Radiological Cooperative Network. Berlin: NUM; 2024 [cited 2026 Jun 23]. Available from: <https://racoon.network/>
24. Pelka O, Kamzol NA, Manjunatha K, Singh N, Omeirat J, Nensa F. A standards-based infrastructure for federated AI inference in medical imaging. *GMS Med Inform Biom Epidemiol*. 2026;22:Doc12. DOI: 10.3205/mibe000310
25. Open Medical Inference (OMI). OMI service registry . GitLab; 2024 [cited 2026 Jun 23]. Available from: <https://gitlab.com/omi-partner/wp6/omi-service-registry>

Corresponding author:

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

Please cite as

Pelka O, Eil J, Manjunatha K, Singh N, Omeirat J, Sigle S, Mathes G, Schweizer ST, Stump SR, Girdziunaite G, Nensa F. Interoperable clinical data integration for distributed medical AI: The Open Medical Inference Gateway. *GMS Med Inform Biom Epidemiol*. 2026;22:Doc11. DOI: 10.3205/mibe000309, URN: urn:nbn:de:0183-mibe0003098

This article is freely available from

<https://doi.org/10.3205/mibe000309>

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/>.