RESETRAY INSIGHTS

What Will the Next Generation of Radiology AI Be Built On?

Are we investing too narrowly in models while underinvesting in the semantic infrastructure that future systems may need?

I-SCL

Imaging Semantic Continuity Layer

Semantic inheritance for the next generation of Radiology AI

Vladimir A. Novikov, MD

Radiologist | Head of Radiology Department | Co-founder, ResetRay

Perspective Article | Part II of II

Master website edition | Version 1.0 | July 2026

docs.resetray.com

Abstract

Radiology AI is moving beyond isolated task-specific algorithms toward foundation models, multimodal systems, large language models, and agentic workflows. Investment in model capability, larger datasets, computational power, validation, and clinical integration remains essential. Yet these advances do not automatically create continuity of information between models, vendors, platforms, and generations of technology.

This perspective argues that the next stage of radiology AI may depend not only on what new models can detect or generate, but also on what informational foundation future systems can inherit. It introduces the author-proposed concept of an Imaging Semantic Continuity Layer (I-SCL): a model-independent semantic infrastructure through which a durable, multidimensional digital imprint could be created from a specific DICOM imaging study.

I-SCL is not presented as an established standard, a completed universal architecture, or a replacement for DICOM, PACS, radiology reports, specialized AI models, or professional interpretation. It is a conceptual framework for preserving computationally reusable meaning beyond the lifetime of an individual model or product.

Using inheritance and variability as a deliberately limited biological analogy, the article proposes that different AI systems provide variability, while semantic continuity provides informational inheritance. The resulting investment question is therefore broader than what the next model can do. It also concerns what the next generation of systems will be able to understand, reuse, and build upon.

Keywords: *artificial intelligence; radiology; medical imaging; DICOM; semantic continuity; information infrastructure; foundation models; multimodal AI; AI agents; interoperability; quantitative imaging; reproducibility; semantic inheritance*

1. From clinical trust to informational inheritance

The first article in this two-part perspective examined the basis of clinical trust in contemporary medical AI. It considered observer variability, reference standards, external validation, provenance, interoperability, reproducibility, and lifecycle monitoring.

That discussion ended with a question:

What if part of the limitation lies not only in the model, but also in the form in which imaging information reaches it?

The present article begins from that point.

Radiology AI is evolving from narrow systems designed for individual tasks toward broader foundation models, multimodal architectures, and systems capable of coordinating tools and information across increasingly complex workflows [3-5].

This transition expands the capabilities of AI, but it also changes the infrastructure problem.

A narrow model may receive one study and generate one result. A future computational ecosystem may involve several models, an AI agent, imaging archives, reporting systems, prior examinations, and research or quality-assurance tools. Information may pass through many components before it contributes to a final action.

The central question is no longer only whether each component performs well.

It is also whether the informational value derived from an imaging study can persist when the components around it change.

2. We may be investing too narrowly

The current investment logic is understandable.

Better models are needed. Larger and more representative datasets are needed. Computational capacity, external validation, workflow integration, regulatory evaluation, and clinical oversight all require sustained investment.

The problem is not that these priorities are wrong.

The problem may be that they represent only part of the system.

Much of the visible competition in radiology AI concerns what the next model can detect, classify, segment, predict, retrieve, or generate. Foundation models extend this ambition by seeking reusable capabilities across multiple tasks and modalities [3-5].

Yet even a highly capable model does not automatically create a durable informational environment.

A model may interpret an imaging study, generate a result, and pass that result to another system. But when the model, interface, product, or vendor is replaced, some of the informational value may become difficult to retrieve, interpret, compare, or reuse.

The ecosystem therefore advances, but does not always accumulate.

New capabilities are added. New outputs are generated. New integrations are built. Yet the information created by one generation of technology may not become a stable foundation for the next.

We are not necessarily investing in the wrong technologies. We may be investing too narrowly in the most visible layer.

3. Progress does not automatically create continuity

Radiology already has a mature digital foundation.

DICOM is the international standard for the communication and management of medical imaging information and related data. It provides extensive mechanisms for representing, transmitting, and managing imaging objects and associated information [1,2].

PACS and enterprise imaging systems provide storage, access, distribution, and workflow support. Radiology reports preserve professional interpretation. Specialized applications and AI models produce task-specific results.

These layers are indispensable.

But the existence of digital infrastructure does not automatically mean that information produced by different computational systems will remain semantically continuous across time.

One model may produce a result for a defined task. Another model may later analyze the same source differently. An agent may combine several outputs. A future system may ask a question that was not anticipated when the original study was acquired or first processed.

If each new system must independently reconstruct the informational basis from the beginning, the ecosystem remains partly repetitive.

Continuity requires something more than storage and transmission.

It requires a way for information derived from a study to remain computationally intelligible and reusable without becoming permanently dependent on the first model that produced or interpreted it.

4. Inheritance and variability

Biological evolution depends on both inheritance and variability.

The analogy should not be understood literally. AI systems do not reproduce biologically, and computational information is not a genome.

Nevertheless, the distinction is useful.

Different AI models provide variability. They may analyze the same imaging study through different architectures, objectives, training strategies, and computational perspectives. Future systems may discover relationships or applications that are not available to current models.

But variability alone does not create cumulative development.

Something must also persist.

For radiology AI, this persistent element could be a stable semantic foundation that allows information derived from an imaging study to be carried forward across models and technological generations.

Models create variability. Semantic continuity creates inheritance.

Variability allows new interpretations to emerge.

Inheritance prevents every new system from beginning again with no structured informational continuity from what came before.

The purpose of inheritance is not to dictate future conclusions. It is to preserve a sufficiently stable informational foundation from which different conclusions can be developed.

5. Introducing the Imaging Semantic Continuity Layer

For the purposes of this perspective, I propose the term:

Imaging Semantic Continuity Layer - I-SCL

I-SCL is not an established international standard or consensus term. It is an author-proposed name for a conceptual semantic infrastructure whose purpose would be to preserve continuity of imaging-derived information across changing computational systems.

PROPOSED CONCEPTUAL DEFINITION

An Imaging Semantic Continuity Layer is a model-independent semantic infrastructure through which a durable, multidimensional digital imprint can be created from a specific DICOM imaging study for use by current and future computational systems.

The term layer is important.

I-SCL would not be the imaging study itself. It would not be the AI model analyzing the study. It would not be the radiology report or the clinical interpretation.

It would be an additional infrastructural level connecting the source study with a changing ecosystem of analytical systems.

The term continuity is equally important.

The objective is not simply to move information between two current applications. It is to allow relevant computational meaning to remain available when models, interfaces, platforms, and analytical methods change.

6. A multidimensional digital imprint

The proposed digital imprint should be understood carefully.

It would not be a diagnosis.

It would not be a clinical conclusion.

It would not be a complete representation of a patient.

It would not be a digital twin.

It would be a durable computational representation derived from one specific source study.

The original DICOM study would remain the primary imaging record and would remain available for direct analysis. The digital imprint would exist alongside it as an additional resource intended for computational reuse.

Its purpose would not be to determine how all future systems must interpret the study.

Its purpose would be to give future systems a richer and more persistent informational basis from which they could perform their own analysis.

Continuity of meaning at the informational layer does not require uniformity of interpretation at the model layer.

Different AI systems could still reach different conclusions.

They would simply not need to reconstruct the basic informational foundation independently every time.

7. DICOM and semantic infrastructure

The relationship should not be framed as:

DICOM or semantic infrastructure.

It should be framed as:

DICOM and semantic infrastructure.

DICOM already provides a sophisticated and extensible standard for medical imaging information [1,2]. Existing work also demonstrates the value of standardized encoding of annotations, measurements, metadata, and model outputs within imaging environments [8-15].

The proposed I-SCL concept does not argue that these mechanisms are inadequate or should be replaced.

It addresses a different level of the problem.

DICOM provides the essential imaging and communication foundation. I-SCL describes a possible additional semantic infrastructure for preserving a durable computational representation derived from a specific study across future analytical contexts.

The distinction is complementary rather than competitive.

DICOM preserves and communicates the imaging study.

I-SCL would support the continuity of a multidimensional computational imprint derived from that study.

Future models could retain access to both.

8. Machine-readable is not the same as semantically continuous

A file may be readable by software without its meaning being reusable across different systems.

An output may be transferred through an API while remaining dependent on a particular application. A value may be digitally stored but lose part of its interpretability when separated from the environment in which it was generated.

The FAIR principles emphasize that scientific data should be findable, accessible, interoperable, and reusable, with particular attention to machine actionability [6]. Work in digital medicine and medical imaging has similarly shown that technical connectivity alone is not sufficient; semantic and organizational interoperability are also required [7-10].

Semantic continuity extends this logic over time.

Interoperability asks whether systems can exchange information.

Semantic continuity asks whether the information can retain sufficiently stable computational meaning when the receiving system, analytical purpose, or technological generation changes.

This becomes increasingly important for AI agents.

An agent may retrieve information from several systems, call specialized models, compare results, and initiate additional operations. Its effectiveness will depend not only on the quality of its reasoning, but also on the stability of the informational environment it receives.

A more capable model cannot fully compensate for information whose meaning must repeatedly be guessed or reconstructed.

Ambiguity should be reduced when information is created, not repeatedly delegated to every downstream model.

9. From isolated outputs to cumulative infrastructure

An isolated AI output may be valuable for the task for which it was generated.

But it does not necessarily become a durable asset for future systems.

When information remains tied to one product, interface, or model generation, technological replacement may force later systems to begin again from the original study.

That may be acceptable for some tasks. It is less satisfactory as the scale and complexity of radiology AI increase.

A semantic continuity layer could change the development pattern.

A model would still be free to produce a new interpretation.

A future system would still be able to return to the source images.

But information derived previously from the study would not necessarily disappear with the system that first used it.

It could remain available as part of a persistent computational foundation.

- new systems would not be limited to repeating earlier extraction;
- earlier information could remain available for new analytical questions;
- different models could build on a shared informational basis without becoming identical;
- the value of an imaging study could extend beyond the technological generation in which it was acquired.

The objective is not to freeze interpretation.

It is to prevent avoidable informational amnesia.

10. A broader investment framework

The investment question should therefore not be reduced to a choice between models and infrastructure.

Both are necessary.

Radiology will continue to require investment in:

- specialized and generalist imaging models;
- multimodal systems;
- AI agents;
- representative datasets;
- computational capacity;
- clinical validation;
- safety and lifecycle monitoring;
- workflow integration;
- imaging interoperability.

But the framework should also include the long-term semantic infrastructure through which information can remain reusable when individual technologies change.

This work may be less visible than a new model.

It may not produce the same immediate demonstration value.

Its benefit may become most apparent only when systems are replaced, information must be reused, or a future analytical method requires access to what was derived years earlier.

Infrastructure is rarely the most visible part of innovation.

It is often the part that determines whether innovation becomes cumulative.

The key investment question is therefore not only:

What will the next model be able to do?

It is also:

What informational foundation will the next generation of systems inherit?

11. Boundaries and unresolved questions

I-SCL is proposed here as a conceptual direction, not as a completed standard or validated clinical system.

Substantial questions would need to be addressed before such an infrastructure could be implemented broadly:

- alignment with existing imaging and health-information standards;

- governance across institutions and jurisdictions;
- security and access control;
- long-term maintenance;
- validation of computational consistency;
- versioning and technological evolution;
- responsibility for interpretation and clinical use.

The existence of a semantic layer would not make every downstream conclusion correct.

It would not eliminate the need for validation, quality control, human oversight, or direct access to the original images.

It would solve a narrower but foundational problem:

How can information derived from a specific imaging study remain computationally reusable across systems and time without being permanently tied to one model?

This perspective does not claim to provide the final architecture.

It argues that the question deserves a place alongside model performance, data scale, and workflow integration in the future strategy of radiology AI.

Conclusion

The next generation of radiology AI will require better models, stronger evidence, representative datasets, safer deployment, and more effective clinical integration.

But those advances may not be sufficient on their own.

Radiology AI may also require a durable semantic infrastructure through which a multidimensional digital imprint can be created from each specific DICOM imaging study and remain available to current and future computational systems.

For the purposes of this perspective, that concept is called the Imaging Semantic Continuity Layer - I-SCL.

I-SCL would not replace DICOM, PACS, radiology reports, specialized AI models, or professional interpretation.

It would complement them.

Models would continue to provide variability.

Semantic continuity would provide informational inheritance.

Together, they could allow radiology AI to evolve cumulatively rather than repeatedly reconstructing its informational foundation from the beginning.

We may not be investing in the wrong technologies.

We may simply be investing too narrowly in the most visible part of the system.

The future of radiology AI may be shaped not only by what the next model can see, but also by what the next generation of systems will be able to understand, reuse, and build upon.

Disclosure

The author is a Co-founder of ResetRay. This article presents an independent conceptual perspective. It does not describe an established technical standard, a validated clinical product, an autonomous diagnostic system, or a replacement for professional radiological interpretation.

Suggested citation

Novikov VA. What Will the Next Generation of Radiology AI Be Built On? ResetRay Insights. Perspective Article, Part II of II. Version 1.0. July 2026. docs.resetray.com

References

1. National Electrical Manufacturers Association. Digital Imaging and Communications in Medicine (DICOM) Standard, PS3.1: Introduction and Overview. Current Edition 2026c. Official standard; no DOI. Official standard
2. Kahn CE Jr, Carrino JA, Flynn MJ, Peck DJ, Horii SC. DICOM and radiology: past, present, and future. J Am Coll Radiol. 2007;4(9):652-657. DOI: 10.1016/j.jacr.2007.06.004
3. Moor M, Banerjee O, Abad ZSH, et al. Foundation models for generalist medical artificial intelligence. Nature. 2023;616:259-265. DOI: 10.1038/s41586-023-05881-4
4. de Almeida JG, Cerdá Alberich L, Tsakou G, et al. Foundation models for radiology - the position of the AI for Health Imaging network. Insights Imaging. 2025;16:168. DOI: 10.1186/s13244-025-02056-9
5. Wu C, Zhang X, Zhang Y, et al. Towards generalist foundation model for radiology by leveraging web-scale 2D and 3D medical data. Nat Commun. 2025;16:7866. DOI: 10.1038/s41467-025-62385-7
6. Wilkinson MD, Dumontier M, Aalbersberg IJ, et al. The FAIR Guiding Principles for scientific data management and stewardship. Sci Data. 2016;3:160018. DOI: 10.1038/sdata.2016.18
7. 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
8. Kondylakis H, Ciarrocchi E, Cerdá-Alberich L, et al. Position of the AI for Health Imaging network on metadata models for imaging biobanks. Eur Radiol Exp. 2022;6:29. DOI: 10.1186/s41747-022-00281-1
9. Kondylakis H, Kalokyri V, Sfakianakis S, et al. Data infrastructures for AI in medical imaging: a report on the experiences of five EU projects. Eur Radiol Exp. 2023;7:20. DOI: 10.1186/s41747-023-00336-x
10. Kalendralis P, Shi Z, Traverso A, et al. FAIR-compliant clinical, radiomics and DICOM metadata of RIDER, interobserver, Lung1 and Head-Neck1 TCIA collections. Med Phys. 2020;47(11):5931-5940. DOI: 10.1002/mp.14322
11. Zwanenburg A, Vallières M, Abdalah MA, et al. The Image Biomarker Standardization Initiative: standardized quantitative radiomics for high-throughput image-based phenotyping. Radiology. 2020;295(2):328-338. DOI: 10.1148/radiol.2020191145
12. Sullivan DC, Obuchowski NA, Kessler LG, et al. Metrology standards for quantitative imaging biomarkers. Radiology. 2015;277(3):813-825. DOI: 10.1148/radiol.2015142202
13. Kessler LG, Barnhart HX, Buckler AJ, et al. The emerging science of quantitative imaging biomarkers: terminology and definitions for scientific studies and regulatory submissions. Stat Methods Med Res. 2015;24(1):9-26. DOI: 10.1177/0962280214537333
14. Bridge CP, Gorman C, Pieper S, et al. Highdicom: a Python library for standardized encoding of image annotations and machine learning model outputs in pathology and radiology. J Digit Imaging. 2022;35(6):1719-1737. DOI: 10.1007/s10278-022-00683-y
15. Swinburne NC, Mendelson D, Rubin DL. Advancing semantic interoperability of image annotations: automated conversion of non-standard image annotations in a commercial PACS to the Annotation and Image Markup standard. J Digit Imaging. 2020;33(1):49-53. DOI: 10.1007/s10278-019-00191-6
16. Lekadir K, Frangi AF, Porras AR, et al.; FUTURE-AI Consortium. FUTURE-AI: international consensus guideline for trustworthy and deployable artificial intelligence in healthcare. BMJ. 2025;388:e081554. DOI: 10.1136/bmj-2024-081554
17. Tejani AS, Klontzas ME, Gatti AA, et al.; CLAIM 2024 Update Panel. Checklist for Artificial Intelligence in Medical Imaging: 2024 Update. Radiol Artif Intell. 2024;6(4):e240300. DOI: 10.1148/ryai.240300

— END OF MASTER EDITION —