Integrating the Healthcare Enterprise

IHE Radiology
Technical Framework Supplement

AI Results
(AIR)

Revision 1.1 – Trial Implementation

Date: July 16, 2020
Authors: Radiology Technical Committee
Email: rad@ihe.net

Please verify you have the most recent version of this document. See here for Trial Implementation and Final Text versions and here for Public Comment versions.

Foreword

This is a supplement to the IHE Radiology Technical Framework V18.0. Each supplement undergoes a process of public comment and trial implementation before being incorporated into the volumes of the Technical Frameworks.

This supplement is published on July 16, 2020 for trial implementation and may be available for testing at subsequent IHE Connectathons. The supplement may be amended based on the results of testing. Following successful testing it will be incorporated into the Radiology Technical Framework. Comments are invited and may be submitted at http://www.ihe.net/Radiology_Public_Comments.

This supplement describes changes to the existing technical framework documents.

"Boxed" instructions like the sample below indicate to the Volume Editor how to integrate the relevant section(s) into the relevant Technical Framework volume.

Amend Section X.X by the following:

Where the amendment adds text, make the added text **bold underline**. Where the amendment removes text, make the removed text **~~bold strikethrough~~**. When entire new sections are added, introduce with editor's instructions to "add new text" or similar, which for readability are not bolded or underlined.

General information about IHE can be found at www.ihe.net.

Information about the IHE Radiology domain can be found at ihe.net/IHE_Domains.

Information about the organization of IHE Technical Frameworks and Supplements and the process used to create them can be found at http://ihe.net/IHE_Process and http://ihe.net/Profiles.

The current version of the IHE Radiology Technical Framework can be found at http://ihe.net/Technical_Frameworks.

CONTENTS

How to read this Supplement 6
Introduction to this Supplement 7
Closed Issues 9
IHE Technical Frameworks General Introduction 17
9 Copyright Licenses 17
9.1 Copyright of Base Standards 17
9.1.1 DICOM (Digital Imaging and Communications in Medicine) 17
9.1.2 HL7 (Health Level Seven) 17
9.1.3 LOINC (Logical Observation Identifiers Names and Codes) 17
9.1.4 SNOMED CT (Systematized Nomenclature of Medicine -- Clinical Terms) 18
10 Trademark 18
IHE Technical Frameworks General Introduction Appendices 19
Appendix A – Actor Summary Definitions 19
Appendix B – Transaction Summary Definitions 19
Appendix D – Glossary 20
Volume 1 – Profiles 21
49 AI Results (AIR) Profile 21
49.1 AIR Actors, Transactions, and Content Modules 21
49.1.1 Actor Descriptions and Actor Profile Requirements 23
49.1.1.1 Evidence Creator 23
49.1.1.2 Image Manager / Image Archive 24
49.1.1.3 Image Display 24
49.1.1.4 Imaging Document Consumer 25
49.2 AIR Actor Options 26
49.3 AIR Required Actor Groupings 26
49.4 AIR Overview 26
49.4.1 Concepts 26
49.4.1.1 Result Primitives 26
49.4.1.2 Result Filtering and Navigation 27
49.4.1.3 Result Presentation 27
49.4.1.4 Codesets 28
49.4.1.5 Confidence 30
49.4.1.6 Negative and Partial Results 30
49.4.1.7 Analysis Results vs Workflow Status 31
49.4.1.8 Result Approval, Retention, & Feedback 32
49.4.1.9 AI Algorithm Deployment 33
49.4.2 Use Cases 34
49.4.2.1 Use Case #1: Store, Retrieve & Display 34
49.4.2.1.1 Store, Retrieve & Display Use Case Description 34
49.4.2.1.2 Store, Retrieve & Display Process Flow 34
49.4.2.2 Use Case #2: Auxiliary Usage 36
49.4.2.2.1 Auxiliary Usage Use Case Description 36

49.4.2.2.2 Auxiliary Usage Process Flow 36
49.5 AIR Security Considerations 37
49.5.1 Security Considerations for Actors 37
49.5.2 Security Considerations for Analysis Results 37
49.6 AIR Cross-Profile Considerations 37
Volume 1x - Appendices 40
Appendix L – Analysis Result Filtering and Navigation (Informative) 40
L.1 Hypothetical Behavior: Result "Display Protocols" 41
L.2 Hypothetical Behavior: Root Results, Layers of Detail 44
L.3 Hypothetical Behavior: Algorithm Rendering Intent 44
L.4 Hypothetical Behavior: Longitudinal Navigation 45
Volume 2 – Transactions 46
4.136 Display Analysis Result [RAD-136] 46
4.136.1 Scope 46
4.136.2 Actor Roles 46
4.136.3 Referenced Standards 46
4.136.4 Messages 47
4.136.4.1 Display Results 47
4.136.4.1.1 Trigger Events 47
4.136.4.1.2 Message Semantics 47
4.136.4.1.3 Expected Actions (i.e., Display Requirements) 47
4.136.4.1.3.1 General Result Display Requirements 48
4.136.4.1.3.2 Display of Qualitative Findings 49
4.136.4.1.3.3 Display of Measurements 49
4.136.4.1.3.4 Display of Locations 49
4.136.4.1.3.5 Display of Regions 50
4.136.4.1.3.6 Display of Parametric Maps 51
4.136.4.1.3.7 Display of Tracking Identifiers 52
4.136.4.1.3.8 Display of Image References 53
4.136.5 Protocol Requirements 53
4.136.6 Security Considerations 53
4.136.6.1 Security Audit Considerations 53
4.136.6.2 Display Specific Security Considerations 53
4.108 Store Instances Over the Web [RAD-108] 54
4.107 WADO-RS Retrieve [RAD-107] 54
4.129 QIDO-RS Query [RAD-129] 54
4.129.4.1.2 Message Semantics 55
4.129.4.1.3 Expected Actions 56
4.137 Query Analysis Results [RAD-137] 56
4.137.1 Scope 56
4.137.2 Actor Roles 57
4.137.3 Referenced Standards 57
4.137.4 Messages 57

4.137.4.1 Query Analysis Results 57
4.137.4.1.1 Trigger Events 58
4.137.4.1.2 Message Semantics 58
4.137.4.1.3 Expected Actions 59
4.137.4.2 Return Responses 59
4.137.4.2.1 Trigger Events 59
4.137.4.2.2 Message Semantics 60
4.137.5 Protocol Requirements 60
4.137.6 Security Considerations 60
4.137.6.1 Security Audit Considerations 60
4.14 Query Images [RAD-14] 60
4.14.1 Scope 60
4.14.2 Actor Roles 60
4.44 Query Evidence Documents [RAD-44] 61
4.44.1 Scope 61
4.44.2 Use Case Roles 61
4.44.3 Referenced Standards 62
4.44.4 Messages 62
4.44.4.1 Query Evidence Documents 62
4.44.4.1.1 Trigger Events 62
4.44.4.1.2 Message Semantics 63
4.44.4.1.3 Expected Actions 64
4.44.4.1.3.1 Mammography Image Profile 64
Volume 2x – Appendices 65
Volume 3 – Content Modules 66
5.1 ITI-20 Record Audit Event 66
6.5 Imaging Analysis Result Content 66
6.5.1 Scope 66
6.5.2 Referenced Standards 68
6.5.3 Analysis Result Encodings 69
6.5.3.1 General Result Encoding Requirements 69
6.5.3.2 Qualitative Findings 71
6.5.3.3 Measurements 73
6.5.3.4 Locations 74
6.5.3.5 Regions 75
6.5.3.6 Parametric Maps 77
6.5.3.7 Tracking Identifiers 78
6.5.3.8 Image References 78
Appendices to Volume 3 79
Appendix A – Example Analysis Result Encodings 79
A.1 Common Header Attributes 79
A.2 Qualitative Finding 80

How to read this Supplement

For an overview of the scope and goals, see Introduction to this Supplement (below), and Volume 3 Section 6.5.1.

For an overview of the participating actors and the transactions between them, see Volume 1 Section 49.1 and Section 49.4.2.

For key concepts that affect implementation and deployment and may be helpful in understanding the rest of the profile, see Volume 1 Section 49.4.1.

For the key technical requirements on encoding and displaying analysis results, see Volume 3 Section 6.5.3 and Volume 2 Section 4.136.4.1.3 respectively.

Introduction to this Supplement

This AI Results Profile addresses the capture, distribution, and display of medical imaging analysis results. The central use case involves results generated by artificial intelligence (AI Model) algorithms. Key considerations include:

- Interoperable Results:

  The results need to be presented to radiologists in their reading environment. This depends on interoperability between result generation products and radiology reading workstation systems/software.

- Study Integrated Results:

  Radiologists expect AI-generated results to be presented in the context of the study to which they apply and expect them to supplement (rather than replace) traditional image analysis results and thus a given study will be composed of acquired images, AI results, and traditional clinical data. The result objects defined in this profile are defined for the existing imaging data storage infrastructure.

- Effective Presentation:

  Effective use of results hinges on presenting them in conjunction with the associated images during the busy process of reading the study.

- Convergence of Result Encoding:

  Many AI results are results that a human could otherwise have produced, and those human results may be used as training data for the AI. Analysis results might also be used by other AIs, e.g., in a GAN (generative adversarial network). AI and non-AI results need to be handled together. Convergent encoding of results facilitates this, as well as data pooling and sharing between sites.

- Display Primitives:

  It is unrealistic to expect radiology displays to implement specific display capabilities for each of the myriad of algorithms being developed. To minimize implementation complexity for displays, and avoid needing different software for each new analysis result, compose analysis results from a reasonable set of primitives.

This profile will establish baseline data handling and presentation capabilities for an image display product to be "AI-Ready". Result generation products will be similarly motivated to support these data formats so that their results can be compatible with a variety of displays and site workflows.

There are many other radiology applications of AI that are not about processing images. This profile is image-centric (i.e., processing imaging inputs).

While the profile will often refer specifically to radiologists, the profile is applicable to a variety of imaging clinicians and a variety of imaging specialties.

Depending on their validity and relevance, as determined by the imaging clinician, results (in whole or in part) may be incorporated into the diagnostic report and/or into the Electronic Medical Record. This profile introduces no changes to current practices for encoding findings into reports and medical records.

AI Workflow for Imaging is addressed in a separate IHE Radiology profile.

Importantly, many questions and challenges that are interesting and worthwhile are out of scope, including, but not limited to:

- handling data flow from image acquisition to image analysis,
- scheduling and managing AI analysis,
- analyzing imaging data that is not images,
- analyzing non-imaging data,
- "interactive" use of AI
- capturing clinician confirmation/contradiction/adjustment feedback
- conveying clinician feedback into algorithm monitoring and/or training
- applying results to clinical processes other than imaging interpretation,
- prioritizing reading worklists,
- performing patient risk stratification,
- submitting analysis results to the EMR for use outside radiology,
- training and validating algorithms such as AI models,
- tracking/communicating the training/provenance/validation of a model instance,
- mapping primitives to/from FHIR Observation (https://www.hl7.org/fhir/observation.html)

Note: Out of scope does not mean products that conform to this profile are not permitted to address such needs, but rather that this profile does not address those questions or specify interoperable methods to meet those challenges.

This profile also does not address encoding results that lack machine-readable semantics (e.g., using Secondary Captures, or Softcopy Presentation States). Implementations that support such encodings as a fallback in addition to the methods required in this profile may refer to the IHE Consistent Presentation of Images Profile for some guidance.

Closed Issues

	Q. What is the scope of the profile? A: Results of image analysis that inform an imaging clinician during reading of the study This focus is a deliberate strategy to avoid scope creep. AI is a large space and there is much fun to be had.
	Q. How to capture confirmation/contradiction/adjustment of the result by a clinician? A: Out of Scope (but really important) Collecting such data may be critical to monitoring quality of results (particularly in the early stages) and to driving improvement of AI models through expert feedback. This was highlighted by both Luciano and Tessa. The mechanism would likely need to capture such accept/reject/(adjust) input during review and convey it to QA processes (similar model to dual-read discrepancy resolution?) and to the training environment for the AI model. The DICOM Verification Flag (0040,A493) is one possible recording mechanism. Consider addressing this in Phase II, and perhaps coordinating with AIW.
	Q. Should we include a Secondary Capture primitive? A. No The goal of the profile is raising the bar to be interoperable/scalable. Including it would undercut the interoperability and shut down use of AI results to drive workflow, provide clinical decision support, populate the patient records, be included in clinical databases, etc. Beside, anyone that wants to store/display Secondary Capture doesn't need a profile to help them do it. Counterpoint: Should we mention it as a "fall through" in the primitives section? Even if not part of the profile, a default "rendering" involves little/no integration cost for new algorithms, or for dumb displays generally. One might argue that algorithms will understand the nature of their results "best", although they won't be able to address the individual preferences of the radiologist and the "analog" result won't work with databases, clinical decision support, etc., etc., etc.

	Q. Is it unreasonable for a display to support the full set of primitives? A. No. It's a bit of work, but it's what is needed Without a full set of primitives, we don't get effective interoperability. You can't really know which primitives the AI Models will need/use so if "some of them" are unsupported it is disruptive. That said, there is an open issue to have two specific sets with a large set of basic primitives and a small set of advanced primitives.
	Q. Should Surface Segmentations and RT Structure Sets be included as Named Options? A. No. RT Structure Sets are fairly complex and are out of our primary radiology interpretation scope. Surface Segmentations (mesh-based) are also fairly complex and are not yet widely used in this space. Most segmentations we imagine Evidence Creators doing is contour or voxel oriented. When interest emerges, these could be added as Named Options.
	Q. Specify use of existing STOW to store existing JSON (or binary) of SR now? A. Not now. Wait until DICOM WG-27 makes the transaction. When creating the IHE transaction, consider whether it should be specific to AI Results or general for Evidence documents.
	Q. Specify WADO-RS request to retrieve an SR, rendered into Sup219 representation? A. Not now. Wait until DICOM WG-27 makes the transaction.
	Q. Include "Report snippets"? (text rendering of a result for inclusion in a report) A. No. Out of scope for this profile. Might consider as part of a reporting profile that looks at the reporting process in more detail.

	Q. Make a new storage transaction for non-image instances? A. No. We have storage transactions for Images, Creator Images, Evidence Documents (just SR), Presentation States, Key Image Notes, Reports which should suffice.
	Q. Is there anything significant to be borrowed from IHE Evidence Documents? A. Not really. We have already borrowed the Store/Q/R transactions. Otherwise, it describes use of SWF and PPWF but we will likely be using AIW. It does make a requirement that the product DCS list all the SOP Classes supported, but that doesn't really seem to be an issue.
	Q. Should SR requirements describe each of the Enhanced, Comprehensive, and Comprehensive 3D IODs A: No. Just Comprehensive 3D. Since "lesser" SR IODs are all valid instances of "higher" IODs, Evidence Creators can simply label all their output as Comprehensive 3D. Since the Image Displays and Consumers should be prepared to handle 3D coordinates, they shouldn't mind and it removes a source of variability.
	Q. Should existing DICOM CAD SR templates be described in this supplement? A. No, not normatively. Section 49.6 "AIR Cross Profile Considerations" highlights the existing Mammo CAD and Chest CAD Profiles (although Chest CAD saw no uptake) that specify requirements on the mammo-specific and chest CAD-specific SR templates and associated display requirements. Interested Evidence Creators, Image Displays, and Image managers can support and claim those profiles. One could theoretically encode much of the content of those CAD IODs in TID 1500 and use them in this profile however that is not being specifically called out or analyzed here. If this profile is successful the committee could explore whether the benefits of harmonization/convergence would be worth the re-implementation costs.

1	Q. Should Parametric Map support be an Option or required? A. Required Due to the prevalence of saliency maps and the likelihood of other parametric map-based results, it was felt that it should be part of the basic capabilities. Also, if it were only an option, the likelihood of displays supporting it might go down and evidence creators that needed to use it would be unable to test their systems.
2	Q. Should probabilistic segmentations be stored as Segmentations or Parametric Maps? A. Probabilistic Segmentations Probabilistic Segmentations: • are specifically designed to support such usage • include metadata to describe segmentation semantics • are well-suited to choosing a threshold then working with the segment • have values that reflect confidence/applicability of a common statement • can represent decimal probability values (as scaled integers) • are part of the basic AIR Profile Parametric Maps • have values that "have a broader meaning" • can represent floating point probability values • but require advanced support by the display

3	Q. Should we help displays find summary/"entry points" into the collection of results? A. Not now. The Public Comment draft included a section on Root Results. Work will continue to explore the use of some kind of Key Object Selection type of instance that summarizes/organizes/serves as an entry point for a set of results in the study. These objects should be readily filterable/findable by a simple display that can use the top level to communicate the summary finding to a user and then let the user choose to display additional lower levels within the "tree" that would likely contain supporting details. The expectation is that there could be multiple "root results" in a study. It would not make sense to mandate that a single result catalog object index **all** the results in the study, since that would require continually revising the catalog object each time one of many algorithms stored new results to the study. It would be challenging, not to mention handling competing updates by when multiple algorithms happen to complete at the same time.
4	Q. How should algorithms encode confidence in each result? A. Don't know. This is an area of ongoing research. Have highlighted the value in Appendix L and the underlying issue in 49.4.1.
11	Q. Should we define encoding/display behavior for failed processing in more detail? A. No. Radiologists, and other interested systems and people, will be interested in tracking and perhaps doing exception handling when analysis fails, however that is more appropriately addressed in workflow artifacts rather than encoding such details in the results themselves. If, for example, there is a need to create a persistent stored instance of completed UPS Workitems, that can be addressed separately. It is, however, useful to give guidance about encoding a finding that a sought feature, e.g., pneumothorax, was found to be absent when processing completed successfully. This is a valid clinical result and can also distinguish cases where nothing was found from cases where the algorithm was not run successfully.

12b	Q. Should we add informative/normative detail about preliminary/partial/verified/etc.? A. No. The states are mentioned briefly in Appendix L, but it is not clear what the best mechanism is for this. The Preliminary flag was introduced for a different purpose and there may be more nuance to what is needed here. Think through how image managers might expose/suppress unverified results, or how they might implement deletion/retention policies.
14	Q. Should the algorithm description inside results include whether it is approved or for what? A. No. It is conceivable that displays might have different behavior depending on whether the algorithm is approved or not, or what the "approved paradigm of use" is" E.g., a result from an algorithm that has been approved for use in a "second reader" paradigm should not initially be presented to the radiologist. However, it is not clear what the full model is that should be encoded, and more importantly, it is unlikely this should be coded into the results themselves, rather than simply identifying the algorithm and leaving it to other out-of-band mechanisms to capture/communicate the current status of the identified algorithm.
15	Q. Is "Rendering Intent" an appropriate flag to indicate results that are not intended for imaging clinician review? A: No. Not at this time. In principle, if an algorithm creates one SR that captures the actual finding(s) for the imaging clinician and another SR that is intended to drive worklist, it would be useful for the Evidence Creator to have a flag that could be set to allow the imaging clinician to suppress the worklist SR elements. This use case will be folded into future discussions about the organization of results, but will not be addressed in this first edition of the profile. Counterpoint: we should assume that each display implementation will analyze the entire content of the SR Tree for all the results and make its best guess if each is something the radiologist needs to see or not. The algorithm is probably unable to determine the purpose of the results it is creating.

17	Q. Is the measurement group in TID 1500 sufficient for establishing relationships between related findings? A. For now, Yes. TID 1500 has good basic relationship linking mechanisms. To "organize" result sets is more complex and may need a separate mechanism like Root Results. Per Closed Issue 3, that has been factored out of this draft for more investigation.
18	Q. Should we constrain pre-coord/post-coord of measurements or require displays to handle all conceivable variations? A. Don't constrain but provide guidance. Based on experience with Echo SR, unconstrained structure complexity and varying degrees of pre-coordination/post-coordination were a severe challenge for receiving systems and resulted in low interoperability. 6.5.3.3 proposed use of either fully precoordinated codes with well-chosen code meaning text, or a "simple" post-coordinated structure. Neither is mandated, so evidence creators can handle greater complexity if needed. Some hope that highly complex cases will instead result in special purpose templates like has been done for other focused tasks such as TID 5300, or perhaps particular popular RadElement Sets. While an enumerated list of pre-coordinated codes may sound tedious, and require configuration to handle new codes, an advantage is that it is mechanical and unambiguous. The business logic is a fairly simple "if code=X, then semantics are Y" as opposed to trying to figure out how to parse all the different possible scenarios in the post-coordinated structure (including all the unexpected value combinations "wait, what does it mean to take that measurement in M-mode?"). People may reason it out but it's harder for software. Post-coordinated coding feels like it can avoid configuration by encoding the semantics of each post-code and combining, but the experience has not borne that out. Each new code value and each new post-coordinated concept has the potential to affect the meaning of the other post-coordinated concepts and values. Plan to monitor this issue during Trial Implementation and early Connectathons and potentially "iterate" the specification. If TIDs are created for specific RadElement Sets, should consider profiling 1500 so they are still valid instances of 1500 to minimize additional logic for displays. Might need collaboration between DICOM and IHE for such maintenance.

21	Q. Do we need to address in detail the use of Patient Orientation and Spatial Locations Preserved in location display (per MAMMO Profile)? A. No. Too specific to the mammography use case of matching CAD output computed from For Processing versus display on For Presentation images - remain silent on this issue. The CXCAD Profile chose to skip it.
22	Q. Should we permit other Linear/Area/Volume TIDs? A. No. TID 1500 and it's sub-templates are adequate. TIDs 1400 Linear Measurement, 1402 Volume Measurement, 1404 Numeric Meas, 1406 3D Linear Meas are not necessary.
23	Q. How does a display know whether to display a polyline as a closed polygon? A. Per DICOM, the last point must be encoded the same as the first, otherwise it is an open polyline.
24	Q. Is Contributing Equipment Sequence the best way to record algorithms? A. It's a good start. It has the advantage that it is present for all IODs and gives the display a uniform place to look for the information. Many deployment models involve the AI Models being separate from the Evidence Creator (which would appear in General Equipment) that packages their results. In terms of lower level TID elements, those can be encoded in SR, maintaining the correspondences described in Table 6.5.3.1-1
25	TO DO: Monitor pre-coord vs post-coord coding patterns during Trial Implementation, RSNA Demo, and early Connectathons and potentially "iterate" the specification.

IHE Technical Frameworks General Introduction

The IHE Technical Framework General Introduction is shared by all of the IHE domain technical frameworks. Each technical framework volume contains links to this document where appropriate.

9 Copyright Licenses

IHE International hereby grants to each Member Organization, and to any other user of these documents, an irrevocable, worldwide, perpetual, royalty-free, nontransferable, nonexclusive, non-sublicensable license under its copyrights in any IHE profiles and Technical Framework documents, as well as any additional copyrighted materials that will be owned by IHE International and will be made available for use by Member Organizations, to reproduce and distribute (in any and all print, electronic or other means of reproduction, storage or transmission) such IHE Technical Documents.

The licenses covered by this Copyright License are only to those copyrights owned or controlled by IHE International itself. If parts of the Technical Framework are included in products that also include materials owned or controlled by other parties, licenses to use those products are beyond the scope of this IHE document and would have to be obtained from that other party.

9.1 Copyright of Base Standards

IHE technical documents refer to and make use of a number of standards developed and published by several standards development organizations. All rights for their respective base standards are reserved by these organizations. This agreement does not supersede any copyright provisions applicable to such base standards. Copyright license information for frequently referenced base standards is provided below.

9.1.1 DICOM (Digital Imaging and Communications in Medicine)

DICOM® is the registered trademark of the National Electrical Manufacturers Association for its standards publications relating to digital communications of medical information.

9.1.2 HL7 (Health Level Seven)

HL7®, Health Level Seven®, CDA®, and FHIR® are registered trademarks of Health Level Seven International.

Health Level Seven, Inc. has granted permission to IHE to reproduce tables from the HL7 standard. The HL7 tables in this document are copyrighted by Health Level Seven, Inc. All rights reserved. Material drawn from these documents is credited where used.

9.1.3 LOINC (Logical Observation Identifiers Names and Codes)

LOINC® is registered United States trademarks of Regenstrief Institute, Inc.

9.1.4 SNOMED CT (Systematized Nomenclature of Medicine -- Clinical Terms)

Some IHE Profiles incorporate SNOMED® CT, which is used by permission of the International Health Terminology Standards Development Organisation. SNOMED CT© was originally created by the College of American Pathologists. SNOMED CT is a registered trademark of the International Health Terminology Standards Development Organisation, all rights reserved.

10 Trademark

IHE® and the IHE logo are trademarks of the Healthcare Information Management Systems Society in the United States and trademarks of IHE Europe in the European Community. They may only be used with the written consent of the IHE International Board Operations Committee, which may be given to a Member Organization in broad terms for any use that is consistent with the IHE mission and operating principles.

IHE Technical Frameworks General Introduction Appendices

The IHE Technical Framework General Introduction Appendices are components shared by all of the IHE domain technical frameworks. Each technical framework volume contains links to these documents where appropriate.

Appendix A – Actor Summary Definitions

*Add the following **new or modified** actors to the IHE Technical Frameworks General Introduction Appendix A:*

New (or modified) Actor Name	Definition
Image Display	A ~~**part of a**~~ system that ~~**can access imaging evidence objects (images, Presentation States, Key Image Notes, Evidence Documents) through network query/retrieve or reading interchange media and allow the user to view these objects**~~ **presents medical images and associated imaging data.**

The table below lists *existing* actors that are utilized in this profile.

Complete List of Existing Actors Utilized in this Profile

Existing Actor Name	Definition
Evidence Creator	A system that creates ~~**additional**~~ evidence **data**~~**objects**~~ such as images **or measurements**, **through a process other than data acquisition.**~~**presentation states, Key Image Notes, and/or Evidence Documents and transmits them to an Image Archive. It also makes requests for storage commitment to the Image Manager for the data previously submitted. It may also retrieve worklist entries for post-processing steps from the Post-Processing Manager and provide notification of completion of the step, allowing the enterprise to track the status of post-processing work.**~~
Image Manager / Image Archive	A system that ~~**provides functions related to safe data**~~ stor**es**~~**e**~~**age** and **manages** imag**eing** data ~~**handling**~~. ~~**It supplies image availability information to the Department System Scheduler. It is always grouped with an Image Archive to provide long-term storage of images, presentation states, Key Image Notes, and Evidence Documents.**~~
Imaging Document Consumer	**A system that makes use of**~~**Retrieves**~~ imaging data ~~**based on references in a retrieved imaging manifest**~~.

Appendix B – Transaction Summary Definitions

*Add the following **new or modified** transactions to the IHE Technical Frameworks General Introduction Appendix B:*

New (or modified) Transaction Name and Number	Definition
Display Analysis Result [RAD-136]	**Presents one or more results of imaging analysis to a user.**
Query Analysis Results [RAD-137]	**Query for a list of DICOM analysis results matching a filter.**

Appendix D – Glossary

*Add the following **new or updated glossary** terms to the IHE Technical Frameworks General Introduction Appendix D.*

New (or modified) Glossary Term	Definition
Imaging Analysis Result	**A result produced by analyzing imaging data. The term encompasses results from classical image processing, machine learning and artificial intelligence, including traditional CADe and CADx as well as deep convolutional neural network.**

Volume 1 – Profiles

Add new Section 49 for AIR Profile

49 AI Results (AIR) Profile

This AI Results Profile specifies how imaging analysis results can be reliably stored, retrieved, and displayed. The motivating use case involves results generated by artificial intelligence (AI Model) algorithms, although the profile applies equally to non-AI-based analysis. The profile will refer generally to "algorithms" and "analysis results" throughout.

This profile defines content for data encoding, transactions for moving that content around, and behaviors for basic handling of the content.

49.1 AIR Actors, Transactions, and Content Modules

This section defines the actors, transactions, and/or content modules in this profile. General definitions of actors are given in the Technical Frameworks General Introduction Appendix A. IHE Transactions can be found in the Technical Frameworks General Introduction Appendix B. Both appendices are located at http://ihe.net/Technical_Frameworks/#GenIntro

Figure 49.1-1 shows the actors directly involved in the AIR Profile and the relevant transactions between them. If needed for context, other actors that may be indirectly involved due to their participation in other related profiles are shown in dotted lines. Actors which have a required grouping are shown in conjoined boxes (see Section 49.3).

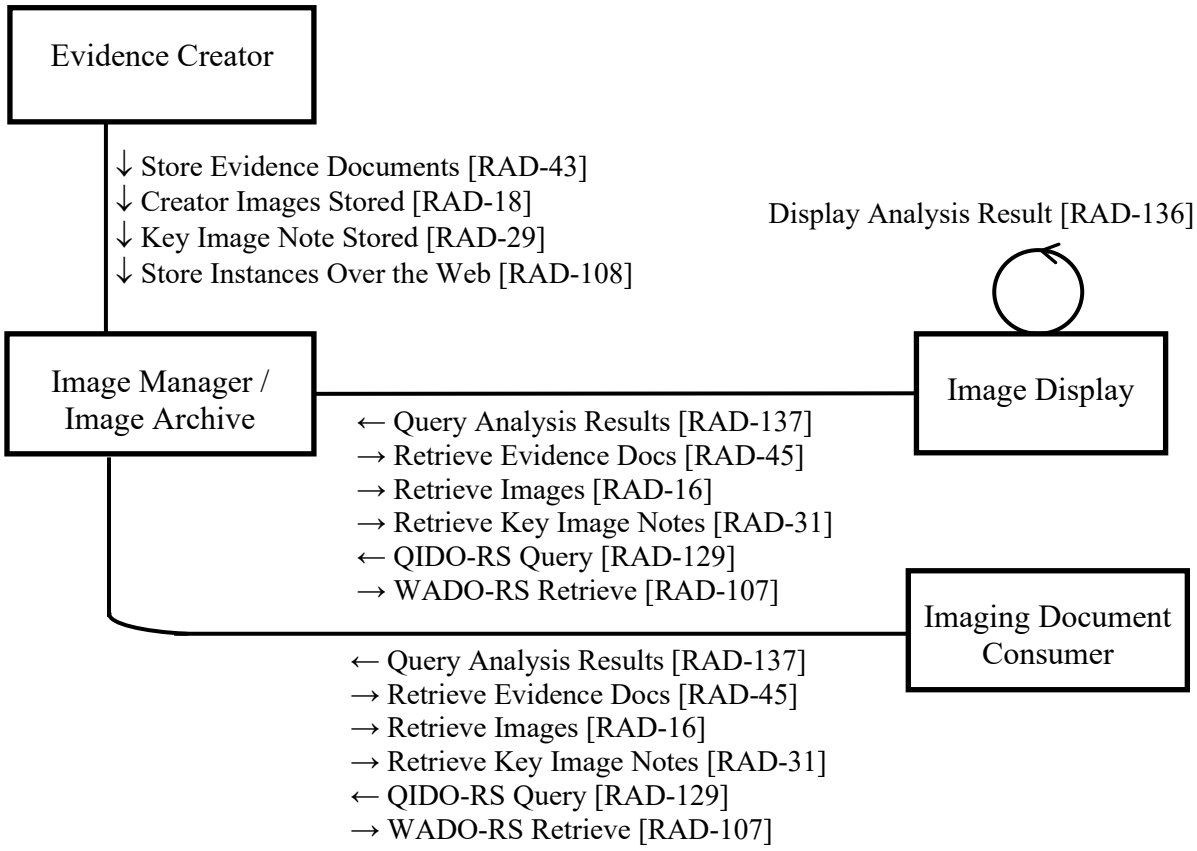

Figure 49.1-1: AIR Actor Diagram

345 Table 49.1-1 lists the transactions for each actor directly involved in the AIR Profile. To claim compliance with this profile, an actor shall support all required transactions (labeled "R") and may support the optional transactions (labeled "O").

Table 49.1-1: AIR Profile - Actors and Transactions

Actors	Transactions	Initiator or Responder	Optionality	Reference
Evidence Creator	Store Evidence Documents [RAD-43]	Initiator	O (See Sec. 49.1.1.1)	RAD TF-2: 4.43
	Creator Images Stored [RAD-18]	Initiator	O (See Sec. 49.1.1.1)	RAD TF-2: 4.18
	Key Image Note Stored [RAD-29]	Initiator	O (See Sec. 49.1.1.1)	RAD TF-2: 4.29
	Store Instances Over the Web [RAD-108]	Initiator	O (See Sec. 49.1.1.1)	RAD TF-2: 4.108
Image Manager/ Image Archive	Store Evidence Documents [RAD-43]	Responder	R	RAD TF-2: 4.43
	Query Analysis Results [RAD-137]	Responder	R	RAD TF-2: 4.137
	Retrieve Evidence Documents [RAD-45]	Responder	R	RAD TF-2: 4.45
	Creator Images Stored [RAD-18]	Responder	R	RAD TF-2: 4.18

Actors	Transactions	Initiator or Responder	Optionality	Reference
	Retrieve Images [RAD-16]	Responder	R	RAD TF-2: 4.16
	Key Image Note Stored [RAD-29]	Responder	R	RAD TF-2: 4.29
	Retrieve Key Image Notes [RAD-31]	Responder	R	RAD TF-2: 4.31
	QIDO-RS Query [RAD-129]	Responder	R	RAD TF-2: 4.129
	WADO-RS Retrieve [RAD-107]	Responder	R	RAD TF-2: 4.107
	Store Instances Over the Web [RAD-108]	Responder	R	RAD TF-2: 4.108
Image Display	Query Analysis Results [RAD-137]	Initiator	O (See Sec. 49.1.1.3)	RAD TF-2: 4.137
	Retrieve Evidence Documents [RAD-45]	Initiator	O (See Sec. 49.1.1.3)	RAD TF-2: 4.45
	Retrieve Images [RAD-16]	Initiator	O (See Sec. 49.1.1.3)	RAD TF-2: 4.16
	Retrieve Key Image Notes [RAD-31]	Initiator	O (See Sec. 49.1.1.3)	RAD TF-2: 4.31
	QIDO-RS Query [RAD-129]	Initiator	O (See Sec. 49.1.1.3)	RAD TF-2: 4.129
	WADO-RS Retrieve [RAD-107]	Initiator	O (See Sec. 49.1.1.3)	RAD TF-2: 4.107
	Display Analysis Result [RAD-136]	Initiator	R	RAD TF-2: 4.136
Imaging Document Consumer	Query Analysis Results [RAD-137]	Initiator	O (See Sec. 49.1.1.4)	RAD TF-2: 4.137
	Retrieve Evidence Documents [RAD-45]	Initiator	O (See Sec. 49.1.1.4)	RAD TF-2: 4.45
	Retrieve Images [RAD-16]	Initiator	O (See Sec. 49.1.1.4)	RAD TF-2: 4.16
	Retrieve Key Image Notes [RAD-31]	Initiator	O (See Sec. 49.1.1.3)	RAD TF-2: 4.31
	QIDO-RS Query [RAD-129]	Initiator	O (See Sec. 49.1.1.4)	RAD TF-2: 4.129
	WADO-RS Retrieve [RAD-107]	Initiator	O (See Sec. 49.1.1.4)	RAD TF-2: 4.107

49.1.1 Actor Descriptions and Actor Profile Requirements

Most requirements are documented in RAD TF-2 Transactions. This section documents any additional requirements on profile's actors.

49.1.1.1 Evidence Creator

Evidence Creators represent a source of results. The scope of this profile begins when the Evidence Creator has created a result to store. Methods to manage the Evidence Creator or allow it to obtain the inputs used in processing are not covered here. See the AI Workflow Profile for details on one approach.

The actor may be implemented by an analysis or AI software package itself, or it may be a proxy, gateway, or partner system that encodes and transmits results on behalf of such software. The profile does not distinguish between analysis executed locally, in the cloud, hosted in processing servers, or running in standalone workstations. For additional discussion, see Section 49.4.1.9 "AI Algorithm Deployment".

Evidence Creators shall encode their results using the Result Primitives described in Section 49.4.1.1.

Note: Evidence Creators are only required to support the Result Primitives needed to encode their results. This means they might not need to support all Result Primitives.

Evidence Creators shall support the SOP Class(es) corresponding to the Result Primitive(s) they support, as defined in RAD TF-3: 6.5 “Imaging Analysis Result Content”.

Evidence Creators shall support a corresponding transaction listed in Table 49.1-1 for each primitive they produce. Evidence Creators may choose to support both the RESTful version of a transaction and the conventional DIMSE version of that transaction, or they may choose to support one or the other.

Result encoding is converged on the Comprehensive 3D SR IOD and the TID 1500 SR template.

Notes: 1. Since the Comprehensive 3D SR IOD is a superset of both the Comprehensive SR IOD and the Enhanced SR IOD, any content that an Evidence Creator might otherwise have created as an instance of the Comprehensive SR SOP Class or the Enhanced SR SOP Class can be safely relabeled as an instance of the Comprehensive 3D SR SOP Class.

2. This profile does not require the Evidence Creator to support TIDs other than the TIDs described in RAD TF-3: 6.5, and their included sub-templates. It is expected that many applications for which a specialized template exists (e.g., TID 4000 Mammography CAD, or TID 5300 Simplified Echo Procedure Report) will continue to use those specialized templates rather than follow this profile for general analysis results. Such support is outside the scope of this profile. Some related profiles are discussed in Section 49.6.

The Evidence Creator is responsible for appropriately populating the Patient-level and Study-level attributes of the DICOM instances it creates. Since that information might come from the headers of the instances the Evidence Creator is processing or from the attributes in a processing workitem, as described in the AI Workflow for Imaging Profile, neither of those sources is specifically mandated in this profile.

49.1.1.2 Image Manager / Image Archive

Image Manager / Image Archive actors store analysis results for a study. Typically, the study will also contain modality and human-generated results (such as measurements or annotations).

Image Manager / Image Archives may include data caches or proxy systems in addition to the more typical PACS and VNA systems.

Image Manager / Image Archive actors shall support all the SOP Classes listed in RAD TF-3: 6.5.3 “Analysis Result Encodings”.

49.1.1.3 Image Display

Image Displays present results, typically together with the other associated images and data relevant to a study.

Image Displays shall support the retrieval of all IODs described in RAD TF-3: 6.5.3. This means Image Displays are required to support a set of corresponding transactions listed in Table 49.1-1; however, Image Displays may choose to support both the RESTful version of a transaction and the conventional DIMSE version of that transaction, or they may choose to support one or the other.

Image Displays are required by [RAD-136] to support specific display capabilities for the IODs described in RAD TF-3: 6.5.3.

Display capabilities for SR are specifically focused on support of the Comprehensive 3D SR IOD and the TID 1500 SR Template.

Notes: 1. The Comprehensive SR IOD and the Enhanced SR IOD are subsets of the Comprehensive 3D SR IOD, so an Image Display that has implemented support for the Comprehensive 3D SR IOD will have implemented all the capabilities to support the Comprehensive SR IOD and the Enhanced SR IOD, however this profile does not require the Image Display to do so.

2. This profile does not require the Image Display to support TIDs other than the TIDs described in RAD TF-3: 6.5, and their included sub-templates. It is expected that many applications for which a specialized template exists (e.g., TID 4000 Mammography CAD, or TID 5200 Echocardiography SR) will continue to use those specialized templates rather than follow this profile for general analysis results. Image Displays should consider supporting additional TIDs as appropriate. Some related profiles are also discussed in Section 49.6.

3. Image Displays may, but are not required to, support other features and display capabilities for the IODs listed in RAD TF-3: 6.5.3. Such other features are not necessary for conformance to this profile.

49.1.1.4 Imaging Document Consumer

Imaging Document Consumers make use of results in ways other than displaying them.

Imaging Document Consumers may include decision support systems, clinical databases, and report creators. An Evidence Creator might implement a grouped Imaging Document Consumer to access existing results as inputs.

Although this profile does not explicitly address adding results directly to an EMR, it is conceivable that a PACS or interface engine might implement an Imaging Document Consumer that selects, extracts, and transcodes results for insertion into an EMR using an unspecified mechanism.

Imaging Document Consumers shall support at least one of the Result Primitives described in RAD TF-3: 6.5.3.

Note: Imaging Document Consumers are only required to support Result Primitives needed to support the intended functionality of the Image Document Consumer system. This means they might not need to support all Result Primitives.

For each supported Result Primitive, Imaging Document Consumers shall support all the SOP Class(es) corresponding to the Result Primitive(s) they support, as defined in RAD TF-3: 6.5.3 “Analysis Result Encodings”.

Note: This means, for example, that an Imaging Document Consumer that consumes volumetric segmentations will need to handle both segmentations encoded in the Segmentation SOP Class, and segmentations encoded as a set of contours in the Comprehensive 3D SR SOP Class.

Imaging Document Consumers shall support corresponding query and retrieval transactions listed in Table 49.1-1 for each primitive they consume. Imaging Document Consumers may choose to support both the RESTful version of a transaction and the conventional DIMSE version of that transaction, or they may choose to support one or the other.

49.2 AIR Actor Options

Options that may be selected for each actor in this profile, if any, are listed in the Table 49.2-1. Dependencies between options, when applicable, are specified in notes.

Table 49.2-1: AI Results – Actors and Options

Actor	Option Name	Reference
Evidence Creator	No options defined	--
Image Manager / Image Archive	No options defined	--
Image Display	No options defined	--
Imaging Document Consumer	No options defined	--

49.3 AIR Required Actor Groupings

An actor from this profile (Column 1) shall implement all of the required transactions and/or content modules in this profile ***in addition to all*** of the requirements for the grouped actor (Column 2).

Section 49.5 describes some optional groupings that may be of interest for security considerations and Section 49.6 describes some optional groupings in other related profiles.

Table 49.3-1: AI Results - Required Actor Groupings

AIR Actor	Actor(s) to be grouped with	Reference	Content Bindings Reference
Evidence Creator	ITI CT / Time Client	ITI TF-1: 7	--
Image Manager / Image Archive	None	--	--
Image Display	None	--	--
Imaging Document Consumer	None	--	--

49.4 AIR Overview

49.4.1 Concepts

49.4.1.1 Result Primitives

It is unrealistic to expect radiology displays to implement specific display capabilities for each of the myriad of algorithms that have been developed and will continue to be developed. To make the implementation complexity for displays manageable, and to avoid needing different software for each new analysis result, this profile follows the direction set by DICOM WG-23 to determine a reasonable, finite, set of primitive elements from which analysis results can be composed.

The primitives include things like measurements, qualitative findings, locations, and segmented regions. For a more complete description see RAD TF-3: 6.5.1.

Evidence Creator actors are required to render their results using the defined set of primitives encoded in specific DICOM SOP Classes (see RAD TF-3: 6.5.3 "Analysis Result Encodings") and Image Display actors are required to support basic presentation of those primitives (see RAD TF-2: 4.136.4.1 "Display Results").

Note: IHE profiles do not typically prohibit additional behaviors or functionality beyond that specified in the profile, for example, Evidence Creators that conform to this profile may have ways of encoding results in addition to those described in the profile.

49.4.1.2 Result Filtering and Navigation

Any given set of images (e.g., a CT Chest Study) may be analyzed by multiple algorithms evaluating multiple conditions and hence multiple, perhaps many, results may be generated. A central use case of this profile involves making these study results easily consumable by radiologists, and other imaging clinicians, during the study interpretation process. Allowing radiologists to effectively filter and navigate results is likely to be a very important capability and challenge for Image Displays. It will be correspondingly important for Evidence Creators to make appropriate metadata available to the Image Displays (preferably using standard rather than private codes) and for Image Displays to help users to leverage that metadata.

This profile does require common metadata be populated by the Evidence Creators, does require support for common matching and return fields in the data query transactions, and does require certain baseline display capabilities in the Display Analysis Result [RAD-136] transaction, but otherwise leaves filtering and navigation capabilities up to the implementation.

For a further discussion of possible considerations, see RAD TF-1x: Appendix L, which describes a variety of filtering concepts that came up during profile development and discussions with radiologists.

49.4.1.3 Result Presentation

This profile places a number of basic presentation requirements on the Image Display in the Display Analysis Result [RAD-136] transaction.

Due to the variety of analysis results, details about how results and the associated images are presented is largely left to the implementation. Further, there is significant potential for display products to devise and provide display capabilities that present the results in ways that very effectively support the interpretation process. This profile is not intended to discourage such functions in any way.

Image Display implementations might consider some of the display capabilities described in the DICOM Volumetric Presentation States. In those objects, the Graphic Annotation Module includes Tracking UID (0062,0021) which can be used to reference an entity in a Segmentation or Structured Report instance to associate it with a Presentation State Graphic Annotation, and the Volumetric Graphic Annotation Module includes Referenced Structured Context Sequence (0070,1903) which can be used to reference a node in a Structured Report instance to associate it with an Presentation State Graphic Annotation.

49.4.1.4 Codesets

While this profile does not mandate the use of particular codesets for most coded observations in the results (beyond the structural codes required by the DICOM standard templates), agreeing within the local site or organization on common codesets will be a key prerequisite of any effective deployment of this profile. Conformance to Regional or National codesets would be a forward-looking step that might yield benefits in the future when evaluating patients who have been at multiple institutions, when applying national clinical guidelines, or when compiling training datasets from diverse sources, etc.

Consistent use of common codes can facilitate many automatic, or semi-automatic, functions such as:

- Correlating current observations with prior observations
- Inserting findings into reports in dictation systems
- Using findings to drive decision support tools

Even more simply, consistent use of codes will make it far more practical to configure Image Displays to behave in sensible ways.

Finding Site

When encoding values for the (363698007, SCT, “Finding Site") of an observation, a variety of recommended codesets may be found in DICOM PS3.16 Context Groups (CID Tables).

Anatomy codes are available that offer Creators a wide range of granularity, e.g., from fine-grained to coarse-grained, codes exist for (36371001, SCT, “Left Sinus of Valsalva"), (81128002, SCT, “Sinus of Valsalva"), (8128003, SCT, “Aortic Root"), (54247002, SCT, “Ascending Aorta"), and (113262008, SCT, “Thoracic Aorta"). Maps of the hierarchical relationships between such codes are available that may be helpful to receiving systems.

Anatomic Region and Primary Anatomic Structure

Many of the same DICOM Context Groups apply when populating the Anatomic Region Sequence (0008,2218) and Primary Anatomic Structure Sequence (0008,2228). These, and Finding Site, will be particularly helpful when organizing results anatomically.

Qualitative Finding Concepts and Values

For the coded Concept of a Qualitative Evaluation, and for corresponding coded Values, the DICOM PS3.16 Context Groups again provide a starting set of codes.

The Reporting And Data Systems (RADS) vocabulary curated by the American College of Radiology provides organized sets of finding codes for a variety of evaluations including: BI-RADS (Breast), C-RADS (CT Colonography), CAD-RADS (Coronary Artery Disease), LI-RADS (Liver), Lung-RADS (lung nodules), NI-RADS (Head & Neck), O-RADS (Ovarian-Adnexal), PI-RADS (Prostate), and TI-RADS (Thyroid).

The RadLex terms curated by the Radiological Society of North America (http://radlex.org) includes a variety of codes for clinical findings and imaging observations.

Another potential source of codes for measurements and qualitative findings are the Common Data Elements for Radiology (RDE), curated by RSNA and ACR (http://radelement.org). As an example, the code RDE436 represents a finding about the "Presence of pulmonary embolism" and is constrained to have a value of Absent, Present, or Indeterminate.

Some RADELEMENT data elements and values may have equivalent concepts in terminologies such as SNOMED CT or LOINC, in which case the use of the latter is preferred in DICOM, due to their more widespread use in the EMR and the availability of mappings in the UMLS.

In Segmentations, sometimes the property being segmented is the anatomy, and sometimes it is something else that also needs an anatomical description, in which case the anatomy is coded separately. For the property being segmented, both a category and a type are specified, and the category determines whether the property type is anatomy, or not.

For example:

- Segmented Property Category Code (0062,0003) = (123037004, SCT, "Anatomical Structure")
- Segmented Property Type Code (0062,000F) = (10200004, SCT, "Liver")

Compared to:

- Segmented Property Category Code (0062,0003) = (49755003, SCT, "Morphologically Abnormal Structure")
- Segmented Property Type Code (0062,000F) = (50960005, SCT, "Hemorrhage")
- Anatomic Region Sequence (0008,2218) = (10200004, SCT, "Liver")

For further information, refer to PS3.16 CID 7150 Segmentation Property Categories.

Code Sources and Maintenance

DICOM codesets are largely constructed from SNOMED CT codes (that are free for use worldwide in the context of DICOM), LOINC codes (that are free for use worldwide), and DICOM codes (that are free for use worldwide).

Many of these codes are subject to maintenance and hence change over time as the terminology improves and evolves as well as errors being found and corrected. Accordingly, Evidence Creators need to be aware that they should be using the most recent codes defined in the DICOM Standard, but also remain sensitive to the needs of the installed base of Image Displays and Image Consumers. Likewise, Image Displays and Image Consumers may want to support both old and new codes in order to function with new objects from new devices, old objects from old devices, and various combinations. The change from SNOMED RT style codes to SNOMED CT style numeric identifiers has been particularly challenging in this respect, but also has significant benefits for harmonization and integration with other clinical systems like EMRs.

49.4.1.5 Confidence

Estimating and encoding the “confidence” that an AI model has in the result it produces, particularly in the medical context of a “finding”, is an area of ongoing research.

As indicated in RAD TF-1x: Appendix L “Analysis Result Filtering and Navigation”, confidence estimates are of significant interest to interpreting radiologists and might be very useful for the filtering and organization rules that radiologists might configure on their displays; for example,confidence thresholds might be used to determine whether a result is highlighted or not displayed at all.

Some AI Models may be able to generate quantitative estimates of confidence based on solid statistical modeling, others may be able to estimate rough relative confidence, and other Models may not be able to provide any indication of confidence at all.

For present/absent-type findings or staging/classification findings, a certainty score may be appropriate. For quantitative results, a 95% confidence interval for the result value might be more appropriate.

Note that the AI Model may have internal thresholds that control when it outputs a result at all, and those thresholds might be fixed as part of the regulatory process. Once a result is output, sites might also have locally controlled user thresholds for when a result is displayed, reported on, etc.

49.4.1.6 Negative and Partial Results

Many encoding descriptions in this profile describe “positive results”, where an algorithm has detected or measured or located something. However, a “negative result”, where an algorithm has reached a determination that something is not present, is also important clinical information and is expected to be encoded in analysis results.

Methods of encoding negative findings is an area of ongoing investigation.

A common concept-value pattern for a positive result is:

(121071, DCM, “Finding") = (59282003, SCT, “Pulmonary embolism")

A corresponding negative result could be recorded as:

(444436002, SCT, “Clinical finding not suspected") = (59282003, SCT, “Pulmonary embolism")

Alternatively, a negative result pattern introduced by DICOM TID 1350 is to attach a subordinate modifier (which makes it critical that the parsing system not overlook the modifier which reverses the sense of the modified concept):

(121071, DCM, “Finding") = (59282003, SCT, “Pulmonary embolism")

> (121052, DCM, “Presence of property") = (272519000, SCT, “Absent")

While some implementers might consider using a finding code, like the above pulmonary embolism code, as a concept and assigning it a value of absent, SNOMED has indicated that is not an appropriate use of a finding code.

So a third alternative is use a code that *has* been defined as a concept that takes a value, such as:

(RDE436, RADELEMENT, "Presence of pulmonary embolism") = (272519000, SCT, “Absent")

In addition to communicating useful clinical information, the presence of a negative result helps the radiologist differentiate the situation where nothing was found from the situation where the analysis was not performed.

Cases where the algorithm generated partial results that are in some sense valid, but processing did not fully complete, may decide whether or not it is appropriate to record those analysis results. DICOM CID 6402 includes a code for Partial Success, but this profile does not provide specific guidance on its use.

49.4.1.7 Analysis Results vs Workflow Status

While this profile addresses the encoding of analysis results, the planning and management of analysis processing tasks that generate those results is addressed by the AI Workflow for Imaging Profile. Since analysis results are one outcome of an analysis processing task, the distinction between results and task status can become blurred. It is useful to be clear about what information is recorded where.

Note: The terms used here are to illustrate concepts, not to propose a formalized model of results and status.

The previous section describes two cases where an analysis result is expected to be created:

- Positive result, where analysis processing completes successfully, and a positive finding or measurement is generated and encoded in a result object.
- Negative result, where analysis processing completes successfully, and a negative finding is generated and encoded in a result object.

Users may also have an interest in knowing about:

- Failed Processing, where analysis processing was run on a given dataset but was not able to complete successfully and stored no clinical results.
- Unattempted Processing, where analysis processing was not run (or perhaps not even scheduled).
- Partial Processing, where analysis processing was only partially completed and some clinical results may have been created.

This profile considers Failed Processing or Unattempted Processing information to be workflow status not clinical results, and thus is more appropriately addressed in workflow artifacts rather than encoding such details in the results themselves. Similarly, as described in the previous section, partial results might be encoded using this profile, however details of the processing that

was attempted and what challenges the processing encountered are better encoded in workflow artifacts. The AIW-I Profile represents tasks, including their processing status, using DICOM UPS Workitem instances. UPS instances are not usually stored for very long beyond the completion of the task, (although conceivably a specification could be created to store them persistently), so capturing workitem status for analysis or communication to the radiologist or other user would depend on additional behaviors by systems such as a Watcher or Task Manager.

49.4.1.8 Result Approval, Retention, & Feedback

This profile is focused on the storage, retrieval, and display of analysis results in the image interpretation process. A full lifecycle management of results goes beyond the scope of this profile, but the associated questions and issues are important for implementations to consider.

Three specific aspects highlighted here are:

- Result approval, meaning policies and practices for the review and acceptance of results prior to them being made available for wider use. Policies might include requiring that results be approved before they are made available to other clinicians and/or incorporated into the electronic medical record.
- Result retention, meaning policies and practices for what results are and are not persistently retained in the imaging record, and perhaps for how long. Policies might include deciding what results are deleted shortly after the study is reported, and which are kept for some period of time. Decision factors might include whether the result was viewed and/or referenced in the reporting process, whether the result was part of a billable activity, whether the result was above or below a threshold for confidence or severity, whether the result was the primary finding or subordinate data, or whether the result is covered by local medicolegal retention practices.
- Result feedback, meaning policies and practices for human experts to affirm, contradict, or modify results, to validate algorithms and/or provide improved training data for this or other algorithms. Policies may relate to localization of algorithms, local validation of algorithms, or wider area programs

While technical mechanisms to support each of the above would likely be helpful, it will be necessary to first establish the types of policies and practices needed and related details. Such information will be sought during trial implementation of this profile and may be the subject of future IHE work.

Result feedback is expected to be particularly important during the early years of the adoption of AI-based algorithms. It is worth noting that all of the result encodings described in RAD TF-3: 6.5.3 would also be appropriate for capturing results that have been generated or edited by a human, such as the radiologist. For example, when editing a contour, or adjusting a measurement, or reversing a finding, a new SR instance could be stored with the new values. TID 1500 includes TID 1001 and TID 1002, where (121005, DCM, "Observer Type") can be set to (121006, DCM, "Person"), and details in TID 1003 can identify the specific person and their

role in the organization. The new result can reference the original analysis result that was edited in the Predecessor Document Sequence (0040,A360).

Beyond that, performing approval and feedback activities, or executing retention policies, are primarily workflow and higher-level data management functions that may be best addressed at the workflow layer above the basic result storage and display addressed by this profile. For comparison, image retention policies are generally not coded into the images themselves. Addressing these workflows, and perhaps generating audit messages associated with performing such tasks may be addressed in the AI Workflow of Imaging Profile.

49.4.1.9 AI Algorithm Deployment

Profile requirements about how results are encoded and submitted for storage are placed on the Evidence Creator, which represents the "hospital-facing" part of the system. The profile does not constrain how AI algorithms are integrated with the Evidence Creator or the analytic technology used by AI algorithms. Some products may choose to implement the AI algorithm and the Evidence Creator features in the same software package. It is also possible that the Evidence Creator will be implemented by a platform product (such as an AI marketplace) and one or more AI algorithms are then integrated with that platform to avoid having to individually implement the result coding functions. Some implementations may progressively assemble components of an analysis result, culminating in data that is conformant with this AIR Profile and is stored by the Evidence Creator as one or more instances. For an illustration of this concept, see Figure XXXX.1-1. "Example of Successive Refinement of JSON Payload to Complete SR" in DICOM Supplement 219 "JSON Representation of DICOM Structured Reports" available at: ftp://medical.nema.org/medical/dicom/supps/Frozen/sup219_fz_JSONSR.pdf

Or to put it differently, it is permitted, but not required, that the Evidence Creator perform the analysis. The Evidence Creator may receive results from other software and be responsible for ensuring the result encoding conforms to this profile and for sending the results to the Image Manager / Image Archive.

On a related topic, some implementations may involve "interactive" AI Models in the sense that a user might choose "on the fly" to invoke one analysis or another based on what they observe in the image, or a user might provide a needed seed point for a segmentation or select a particular region of an image for analysis, or perhaps there might be several iterative invocations of AI Models with user feedback in between. Regardless, it is expected that the generated results would be conformant with this profile and displayable by conformant Image Displays.

In such scenarios, an Evidence Creator would potentially be integrated with the Image Display. Alternatively, the Image Display might be grouped with a Task Requester in the AI Workflow for Imaging Profile to have the analysis performed by an external Task Performer. In either case, this profile could potentially still be applied to encoding, storing, and being able to display the results generated, although such workflows or iterative result updates have not been analyzed here.

Interactive AI raises a number of interesting issues. Due to the possibility of bottlenecks in processing and data transfer, some AI may be implemented similar to post-processing labs where

it is all done before reading starts. But as technology advances, more "on-the-fly" or interactive analysis may be possible, where the radiologist might request AI workup of a suspicious feature, perhaps to quantify its characteristics or to get a list of the most likely diagnostic findings.

Such issues (particularly relating to managing execution of AI Models) are out of scope of this profile. Implementers may also refer to Section 49.4.1.7 in the AI Workflow for Imaging Profile.

49.4.2 Use Cases

This profile is focused on handling Analysis Results that are generated from images and are applied in the context of "reading" the associated imaging study.

49.4.2.1 Use Case #1: Store, Retrieve & Display

49.4.2.1.1 Store, Retrieve & Display Use Case Description

Results are produced from the analysis of images by one or more Evidence Creators and are stored and later retrieved and displayed by an Image Display to support the interpretation of the associated imaging study.

As described in Section 49.4.1.9, the Evidence Creator is the "hospital-facing" facet of an algorithm.

An important characteristic of the Use Case is that the Image Display does not depend on specific knowledge of the algorithm or the details or internal structure of the result that the algorithm produces. This is intended to facilitate the use of a variety of image display products with a variety of products and architectures while minimizing the integration and customization overhead.

Although not shown, the use case expects there to be multiple Evidence Creators and a given study could easily have multiple results.

There is a possibility that some results may arrive after reading has been performed. Coordinating some kind of "Ready To Read" signal based on expected and received data and other aspects of reading workflow is outside the scope of this profile.

49.4.2.1.2 Store, Retrieve & Display Process Flow

Figure 49.4.2.1.2-1 shows storage and retrieval of SR instances using conventional DIMSE DICOM transactions. Alternatively, results could also, or instead, consist of image instances, and could be retrieved using RESTful DICOM transactions.

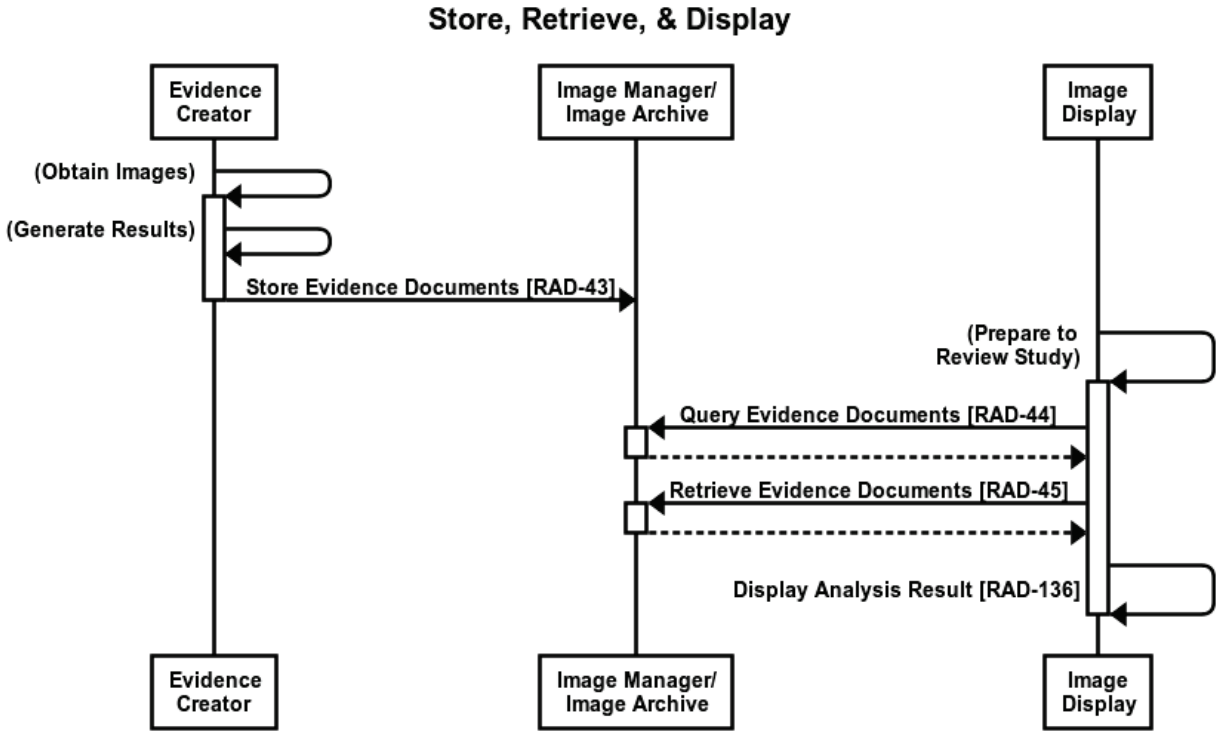

Figure 49.4.2.1.2-1: Store, Retrieve & Display Process Flow in AIR Profile

The text in Figure 49.4.2.1.2-2 was used to generate the diagram in Figure 49.4.2.1.2-1. Readers will generally find the diagram more informative. The text is included here to facilitate editing.

```
title Store, Retrieve, & Display

participant Evidence\nCreator as EC
participant Image Manager/\nImage Archive as IM
participant Image\nDisplay as ID

EC->+EC: (Obtain Images)
EC->EC: (Generate Results)
EC->-IM: Store Evidence Documents [RAD-43]
ID->+ID: (Prepare to\nReview Study)
ID->+IM: Query Evidence Documents [RAD-44]
IM-->-ID:
ID->+IM: Retrieve Evidence Documents [RAD-45]
IM-->-ID:
ID->-ID: \nDisplay Analysis Result [RAD-136]
```

Figure 49.4.2.1.2-2: Diagram Pseudocode for Store, Retrieve & Display Process Flow

49.4.2.2 Use Case #2: Auxiliary Usage

49.4.2.2.1 Auxiliary Usage Use Case Description

This use case highlights that non-display actors can make use of the same results used by the Image Displays in Use Case #1. Examples of Imaging Document Consumers include report creation systems that use result contents to populate fields in a draft imaging report, clinical databases that extract measurements and support functions like comparing values for a given patient over time, and clinical decision support systems that use result contents as input values to drive their decision logic.

Such Imaging Document Consumers will benefit from the standardized representations, but would not be required to support display behaviors.

This use case also demonstrates query and retrieval using DICOMweb (RESTful DICOM transactions – QIDO-RS and WADO-RS); it could alternatively have used conventional DIMSE DICOM.

49.4.2.2.2 Auxiliary Usage Process Flow

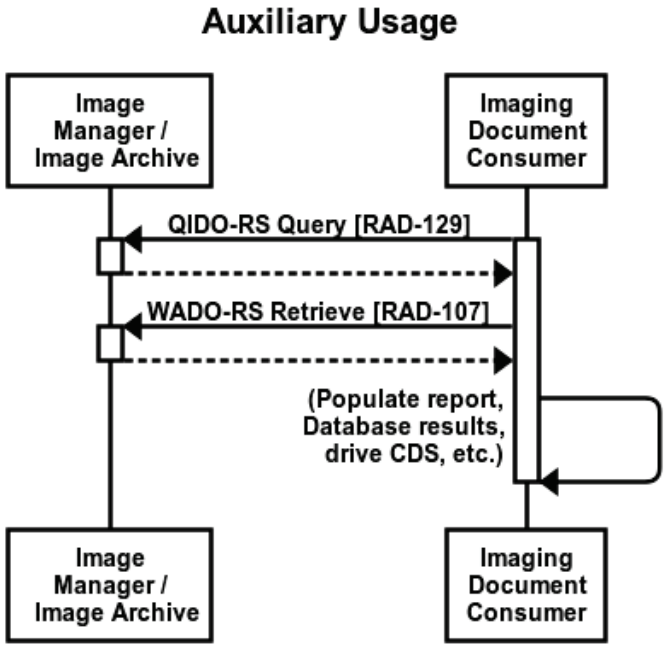

Figure 49.4.2.2.2-1: Auxiliary Usage Process Flow in AIR Profile

The text in Figure 49.4.2.2.2-2 was used to generate the diagram in Figure 49.4.2.2.2-1. Readers will generally find the diagram more informative. The text is included here to facilitate editing.


```
title Auxiliary Usage

participant Image\nManager /\n Image Archive as IM
participant Imaging\nDocument\nConsumer as IDC

IDC->+IM: QIDO-RS Query [RAD-129]
activate IDC
IM-->-IDC:
IDC->IM: WADO-RS Retrieve [RAD-107]
IM-->-IDC:
IDC->IDC: (Populate report,\nDatabase results,\ndrive CDS, etc.)
deactivate IDC
```

Figure 49.4.2.2.2-2: Diagram Pseudocode for Auxiliary Usage Process Flow

49.5 AIR Security Considerations

Refer to RAD TF-1x: Appendix F "Security Environment Considerations".

Protected Healthcare Information (PHI) is present in the DICOM instances being stored, retrieved, processed, and displayed.

49.5.1 Security Considerations for Actors

This profile strongly recommends all actors group with an ITI ATNA Secure Application or Secure Node Actor using the Radiology Audit Trail Option.

The ATNA Profile requires actors to implement:

- Record Audit Event [ITI-20] transaction which would record when and where analysis results are distributed and displayed.
- Authenticate Node [ITI-19] transaction to further ensure the integrity of transactions via node authentication and communication encryption.

This profile does not add security considerations beyond those already established for the transfer and storage of clinical data in other profiles.

49.5.2 Security Considerations for Analysis Results

Analysis Result instances as defined in this profile contain personal demographic information and clinical information. It is appropriate for products implementing the AI Results Profile to include appropriate PHI controls. Specifying such mechanisms and features is outside the scope of this profile.

49.6 AIR Cross-Profile Considerations

Since the data created and exchanged in the AI Results Profile are encoded using common DICOM instances, many other Radiology profiles that manage content could be used in conjunction with the content of the AI Results Profile:

- Cross-Enterprise Document Sharing for Imaging (XDS-I) and/or Cross-Community Document Access for Imaging (XCA-I) could be used to exchange analysis results within and between enterprises
- Import Reconciliation Workflow (IRWF.b) and/or Import and Display of External Priors (IDEP) could be used to localize and manage the import of analysis results
- Web-based Image Access (WIA) could be used to access analysis results using DICOMweb
- Invoke Image Display (IID) could be used to launch viewing of analysis results
- Portable Data for Imaging (PDI) could be used to distribute analysis results on portable media
- Audit Trail and Node Authentication (ATNA) (with the Radiology Option) is recommended to secure the communication of, and record audit trails for, analysis results.
- Imaging Object Change Management (IOCM) could be used to manage changes and quality control of analysis results including deprecating distributed instances that have been disapproved or superseded by revised instances.
- Standardized Operational Log of Events (SOLE) could be used to log processing events like the generation of image analysis results for operational analysis

Similarly, it would be sensible for Image Display systems to support the AI Results Profile together with other Profiles that include related display capabilities, such as:

- Scheduled Workflow (SWF.b) to display the analysis results together with the images from the study
- Mammography Image (MAMMO)
- Digital Breast Tomosynthesis (DBT)
- NM Image (NMI)

Both the AI Results Profile and the Evidence Documents (ED) Profile make use of the [RAD-43], [RAD-44], and [RAD-45] transactions to store, query, and retrieve SR instances.

Both the AI Results Profile and the Web Image Access (WIA) Profile make use of the [RAD-129] and [RAD-107] transactions to query and retrieve DICOM instances using RESTful DICOMweb. The WIA Profile also documents using the ITI XDS and MHD Profiles to support backend functions, but that is not really relevant to the AI Results Profile.

The Mammography Image (MAMMO) Profile introduces requirements on the creation, exchange, and display of mammography images and DICOM SR objects containing CAD results. While there are direct parallels between the latter and the results in this profile, the MAMMO Profile mandates the specialized SR template (TID 4000) and addresses many mammography-specific issues. Products that create and display mammography results are

advised to conform to the MAMMO Profile. The situation is very similar for the Chest X-Ray CAD (CXCAD) Profile.

The Results Distribution (RD) Profile addresses the communication of imaging results, but in that context, "results" refers to diagnostic reports, so there is no direct interaction between the RD Profile and the AIR Profile.

AIW-I – AI Workflow for Imaging Profile

A Task Performer in AIW-I might be grouped with an Evidence Creator to assign and track completion of the processing performed by the Evidence Creator (or its constituent algorithms) as well as provide the Evidence Creator with information about how to access its input data and allow the Evidence Creator to communicate the list of instances it has created and where they have been stored. The tasks also facilitate controlling the parameters fed to the algorithm, recording execution status (for example, billing may depend on whether a given algorithm was run even if it didn't generate any results), and allowing systems (like Reporting Worklist Managers) to subscribe to notifications about AI processing that has been scheduled and/or completed.

RRR-WF - Radiology Remote Reading Workflow Profile

A Task Performer in RRR-WF might be grouped with an Image Display in a reading workstation to access a defined reporting task that identifies the study/instances to be reported and where they may be obtained. Since those instances may include the Analysis Results, this is one of the ways that the Image Display can become aware of the Analysis Results that it Query/Retrieves in this profile.

Volume 1x - Appendices

Add new Appendix L.

Appendix L – Analysis Result Filtering and Navigation (Informative)

As AI adoption increases, a given study may be subjected to many analyses, potentially from different vendors, leading to many AI-generated results.

Consider a chest CT that is acquired and evaluated by algorithms that produce a collection of results, including findings or observations about:

- Cracked ribs,
- Pneumothorax,
- Pneumonia,
- Scoliosis,
- Aortic aneurysms,
- Coronary plaque, and
- Lung lesions

Further, some results may consist of multiple components; e.g.,

- A lung lesion result might include a finding that a lung nodule is present, a location for the nodule, a segmentation of its boundary, measurements of its dimensions and volume, an overall nodule classification as benign/suspicious/malignant/etc., a Key Object Selection indicating the image slice most representative of the nodule, etc.
- A cracked rib result might include a finding on recency; is it recently cracked, or an older injury.
- RadElement commonly groups multiple related elements into "element sets"

The software used to review imaging studies (studies which will now include all these analyses, both conventional and AI-based) will need to address the challenge of helping the radiologist be aware of the available information and review it as appropriate, while simultaneously avoiding the addition of extra time and tangible steps that are anathema in the imaging interpretation process.

Note: The display of individual result primitives is directly addressed by this profile in the form of baseline requirements in [RAD-136]. Specific result discovery or navigation behaviors by the Image Display are left to product design and not addressed by this profile; however, the profile indirectly addresses the related tasks of filtering and navigating the available results by trying to ensure sufficient relevant metadata is available in the result objects. A "zero-click interface" would certainly be appealing, but this profile does not attempt to design one.

The following "Hypothetical Behavior" sections are intended to facilitate discussion of possible use cases and mechanisms that would help imaging clinicians. If those discussions identify

certain key metadata needed to support such mechanisms, that may result in Change Proposals with requirements, constraints, or recommendations for Evidence Creators or Image Displays. These hypothetical behaviors are not requirements in this profile.

L.1 Hypothetical Behavior: Result "Display Protocols"

One could imagine an Image Display allowing users to configure rules and behaviors that controlled what results were initially displayed and how they were presented. This is analogous to conventional hanging protocols which use rules to control what images are presented and how they are organized on the available displays.

A presumption in this display protocol discussion is that none of the results in the study would be inaccessible to the imaging clinician, rather some results would be automatically prioritized, arranged and presented, while others might require interactions by the imaging clinician to access. Based on experience in mammography CAD, where the rate of false positives presented a significant challenge, some forms of result filtering will be necessary to make the display of results usable when more than a few are present in a study.

Some conceivable factors that might be incorporated into rules might include:

- Whether the result is "normal" or "abnormal" (e.g., might not initially display results that are normal or unremarkable)
  - Normal/abnormal may be treated as a continuum using a score or degree of severity
  - Value ranges for normal/abnormal will sometimes be locally defined
- The algorithm "confidence", perhaps in terms of the sensitivity/specificity of the algorithm, or the positive/negative predictive value, or in terms of some generated metric (e.g., might not display results that do not exceed a certain threshold, and/or might highlight results with a particularly high confidence). Note that "confidence" is easier said than estimated. See RAD TF-1: 49.4.1.5.
- The change or consistency of the current result when compared to prior results. (e.g., might highlight or display results that represent a change from the prior situation; a new lesion vs one that was previously present; a mass that has increased or decreased in size vs one that has not changed size)
- The criticality of an abnormal result (e.g., might display a result with high criticality, perhaps even if it was below the normal confidence threshold). The IHE Results Distribution Profile highlights the RADLEX codes for observation categories (from most to least severe, where the top 3 are differentiated by the rough time frame in which action should be taken):
  - RID49480^Category 1 Emergent Actionable Finding^RadLex
  - RID49481^Category 2 Urgent Actionable Finding^RadLex
  - RID49482^Category 3 Noncritical Actionable Finding^RadLex

925
    - RID50261^Non-actionable^RadLex
    - RID13173^Normal^RadLex (synonym “Unremarkable”)
- Whether the result has been reviewed/verified/approved by a human (e.g., might display results a resident has “approved/verified/confirmed”, perhaps even if it was below the normal confidence threshold).

930
- The relationship between the type or anatomical location of the results and the indications, reason for study or anatomical focus of the imaging procedure (e.g., might suppress cardiac results that are normal when viewing a lung study but displayed them in a cardiac study, or might provide an organ-based summary of findings)
- The algorithm make/model/version used (e.g., an imaging clinician might be evaluating a
935 particular algorithm and want to see the results of that algorithm for all studies during the period of evaluation; or an imaging clinician might not have confidence in a particular algorithm and not want those results “cluttering” the display)
- Whether the result is a “root result” (see the following Hypothetical Behavior: Root Results, Layers of Detail, Result Summarization) or a subordinate finding (e.g., might
940 display the root result that pneumothorax is present or a lesion was detected, but suppress the segmentation showing the lesion or pneumothorax location unless requested).
- Whether the algorithm completed successfully (e.g., algorithms that failed or only partially completed might still generate results useful for evaluating its performance, or simply knowing that it was run; results created with a failure status would not normally
945 be displayed). Based on experience in mammography CAD, the following states might be considered:
    - Successfully processed
    - Partially processed; incomplete
    - Processing failed; unsupported data (e.g., wrong body part or modality)
950
    - Processing failed; modality model not validated
    - Processing failed; low confidence
- The relative date/time of multiple results (e.g., current results might be presented next to prior results, and/or differentiated from a current result from re-processing a prior image)
    - Note that the date/time of the results reflects when they were created, as distinct from
955 the date/time of the images from which the results were generated
    - Multiple findings might also be present in the study for the same Concept when different algorithms make a determination, say, on whether pneumonia is present, or the same algorithm is run with different parameters.
- The approval status of the algorithm (e.g., during normal clinical reading might not
960 initially display the results of a research algorithm that is not yet FDA approved). The

approval status of the algorithm might not be encoded directly in the result but rather the Algorithm Identification coded into the result might be used to consult a database or configuration table that records the status for each algorithm locally in use.

- The approved paradigm of the algorithm (e.g., might not initially display the results of an algorithm approved for use as a "second reader" since that would invalidate its intended mode of use, or).
- The type of reader (e.g., might not initially display some types of results to Emergency Room readers, or residents)
- The nature of the current review (e.g., might present the "worst" case of several conflicting results when performing triage, while normal reading workflow might prioritize differently or might present both results).
- Known details about the patient (e.g., might not highlight a pneumothorax result for a patient known to have chest tube in place unless the pneumothorax size exceeds a certain threshold)
- Relationships between results (e.g., for a primary finding that a tumor is present, the related results that show the segmented surface of the tumor, and the numerical measurements of the dimensions and volume might be displayed together, or the primary finding might be displayed with an indicator that there are secondary/supporting results that can also be displayed. In addition to direct references between results, relationships can also be inferred by results sharing the same Tracking Identifier, the same Finding Site, or the same Frame of Reference)
- Results that correspond to fields in the current report template (e.g., findings that will automatically "flow through" to populate elements of the report based on the current configuration would be displayed or highlighted to make sure the radiologist is fully aware. To a certain extent, the report template encodes a certain focus of interest. More specifically, a report template that is tuned to the needs of a particular specialist (cardiologist, pulmonologist, neurosurgeon, family doctor, etc.) may focus on findings relevant to their interests. A study of coronary artery calcification might not highlight an abdominal aortic aneurysm of less than 3cm. As more analysis is performed, the ratio of observations that do not appear in the report might exceed the observations that do.)

Rules might also control how results are grouped or formatted on the display, for example grouping abnormal findings together, or grouping cardiac results separately from lung results, or grouping results that came as a set from a particular algorithm.

An additional complexity is that a study might contain conflicting results. Consider a general chest X-Ray algorithm that evaluates 6 conditions, which determined that cracked ribs were absent, and another special purpose algorithm (which might be considered to be more sophisticated) which determined that two cracked ribs were present. What if the findings were reversed, with the more specialized model indicating no fractures and the less specialized model indicating fractures were present? Perhaps the result with the higher confidence would be displayed with an indication that conflicting results are present.

SR instances from different Evidence Creators will be in different series which may provide additional clues for the Image Display.

L.2 Hypothetical Behavior: Root Results, Layers of Detail

It seems likely that displays might leverage component hierarchy by first presenting a summary "root" result or key value to an imaging clinician and offer the ability to explore additional layers of detail as needed, for example allowing the imaging clinician to expose the segmentation that underlies a volume measurement. This exploration might be done to gain greater confidence in the root result, or to comprehend more details and nuances of the finding(s).

In another example, a lung screening algorithm might record a LungRADS™ score as a root result, supported by secondary results consisting of multiple detected nodule locations and assessments of the size, solidity, and margin of each detected nodule. Another algorithm might record an SR finding of "pneumonia present" as a root result with a reference to a separate saliency map instance. An Image Display might then initially present the two root results (LungRADS = Category 3) and (Pneumonia present) and offer "drill down", rather than initially presenting 43 components consisting of location, size, solidity, margin, and LungRADS for each of 8 nodules, an overall LungRADS score, the pneumonia finding, and the pneumonia saliency map reference.

Root results would likely capture the hierarchy within a set of results that were generated together. If multiple applications were run on a study, each would likely generate its own root result. More advanced logic or analysis software might prioritize all results for a given study.

An effective root result mechanism should allow a display to use a simple query filter to get the first-order set of "summary findings". The references in each of those root results provide a logical next layer of detail. Ideally, some displays will develop much more sophisticated analysis and logic, or more advanced configurations, and more advanced navigation and display, while the root results would provide a first simple step up from the flat list of findings.

L.3 Hypothetical Behavior: Algorithm Rendering Intent

Mammography CAD allows the analysis algorithm to communicate recommendations about the behavior of the Image Display by flagging specific findings as having a Rendering Intent of either Presentation Required, or Presentation Optional. Evidence Creators could also encode an indication of the criticality of an abnormal result (see above) as a way to influence display behaviors.

Such an approach could make some display behaviors consistent for the same result across different displays. This supports some "centralization" of some display logic since it is configured/determined at the algorithm, rather than at each individual display. On the other hand, when there are many results, the problem remains of prioritizing/sequencing multiple algorithms all indicating presentation required.

In practical terms, the rendering intent of the algorithm is information that could be used by display logic, but would be unlikely to override the instructions of the user, and there is no

obligation for displays to anything at all with it, so it is currently not a reliable mechanism for Evidence Creators.

L.4 Hypothetical Behavior: Longitudinal Navigation

Longitudinal results (the same measurement or evaluation performed at several points in time) present another axis of result relationships commonly of interest to imaging clinicians. For example, the change in size of a given tumor or the stability of a stenosis grade for a given vessel may be of interest.

Tracking Identifiers provide a mechanism for indicating that the subject of two or more measurements or evaluations is the same entity. Whether such longitudinal results are listed, graphed, scrolled through, or shown side-by-side, are design choices left to Image Display implementations.

RAD TF-2: 4.136.4.1.3.7 contains associated requirements for the display of Tracking Identifiers. The value of Finding Site may also be useful to correlate findings or results from different timepoints that may be related due to being at the same finding site. Note that finding sites may describe anatomy, devices (stents, pins, biopsy clips, etc.), or, when associated with Finding, a pathology.

Volume 2 – Transactions

Add Section 4.136

4.136 Display Analysis Result [RAD-136]

4.136.1 Scope

This transaction is used to present image analysis results to someone, such as a radiologist interpreting a study.

This transaction is not a typical IHE transaction between two devices; the primary focus is on the required behavior of the display rather than messaging between two actors. This can be thought of as an "informational transaction" between a display device and a user.

The specification is organized around a defined set of result data structures ("primitives") that may be encountered by the display, and the baseline display behaviors required for each. As with many IHE specifications, the display may have behaviors in addition to those required by this transaction.

Methods for selecting the data to be displayed are outside the scope of this transaction.

4.136.2 Actor Roles

The roles in this transaction are defined in the following table and may be played by the actors shown here:

Table 4.136.2-1: Actor Roles

Role:	Display: Presents results visually to a user, such as a radiologist.
Actor(s):	The following actors may play the role of Display: Image Display

Transaction text specifies behavior for each role. The behavior of specific actors may also be specified when it goes beyond that of the general role.

4.136.3 Referenced Standards

- DICOM PS3.3: A.35.4 Key Object Selection Document IOD
- DICOM PS3.3: A.35.13 Comprehensive 3D SR IOD
- DICOM PS3.16: TID 1500 Measurement Report
- DICOM PS3.3: A.51 Segmentation IOD

- DICOM PS3.3: A.75 Parametric Map IOD

4.136.4 Messages

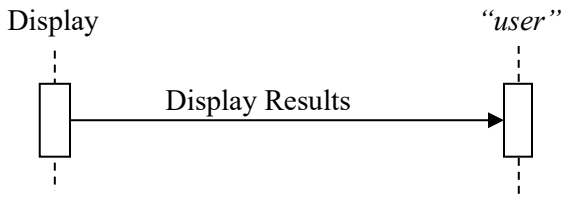

Figure 4.136.4-1: Interaction Diagram

4.136.4.1 Display Results

The Display presents the analysis results to the user.

The results being presented may or may not have been previously verified by a human and may or may not be accurate.

4.136.4.1.1 Trigger Events

A user or an automated function determines that one or more results should be presented.

4.136.4.1.2 Message Semantics

The results are encoded as described in RAD TF-3: 6.5.3.

This transaction does not depend on how the instances that contain the encoded results are represented (DICOM binary, DICOM XML, DICOM JSON), or the messaging protocol by which the instances containing the encoded results were transferred to the Display. If the Display receives results by a profiled mechanism such as DICOM C-STORE, or DICOMweb WADO-RS, the messaging protocol is specified in that corresponding transaction. If results are accessed by being grouped with another actor such as an Image Manager / Image Archive or Evidence Creator, there is no messaging protocol involved.

4.136.4.1.3 Expected Actions (i.e., Display Requirements)

The behaviors in this section are specified as baseline capabilities. Displays may have additional or alternate capabilities that may be invoked or configured.

Displays shall support the capabilities described in this section for result primitives (See RAD TF-3: 6.5.3) encoded in instances of the following:

- Key Object Selection Document IOD
- Comprehensive 3D SR IOD
- Segmentation IOD
- Parametric Map IOD

4.136.4.1.3.1 General Result Display Requirements

The Display:

- shall make the user aware that results are available for display
- shall support multiple results collected together in one instance
  - e.g., multiple qualitative findings, measurements, locations, and/or regions in a single SR instance, multiple region segments in a single Segmentation instance, etc.
- shall support sets of results that are spread across multiple instances in one or more series in current and prior studies being reviewed by the user
- shall be able to display each result primitive (as defined in RAD TF-3: 6.5) in a complete, unambiguous, clinically useful manner without excessive user interaction
- shall make clear the specific image/frame to which the result primitive applies
  - this may involve toggling a GSPS-like overlay on/off
- shall be able to superimpose each primitive on the appropriate underlying images in a manner appropriate to the primitive
- shall allow the user to control which result primitives are displayed on an individual basis (see RAD TF-1: 49.4.1.2 for some approaches and considerations)
  - that control shall include being able to select which, if any, locations, regions, measurements, and qualitative findings are displayed superimposed on images
  - that control shall include being able to resolve overlapping results (i.e., results where one could obscure the other when rendered together). The method is left to the implementation (e.g., allowing the user to cycle through them, choose between them, etc.).
- shall make available for display the unique identifier, version, name, and manufacturer of the algorithm that generated each result
  - RAD TF-3: 6.5.3.1 describes the location(s) where this information is encoded
- shall make available for display the following information about each result
  - Content Date (0008,0023) and Content Time (0008,0033) of the primitive instance
- shall, if an image reference with concept of (121200, DCM, “Illustration of ROI") is associated with a primitive, make the user aware if it (unobtrusively) and be able to display that image.

If the Display is unable to present a result, it is up to the Display whether to notify the user, make a record in a log file, or fail silently.

Displays should note that instances of Segmentations and Parametric Maps both belong to the family of Enhanced Multiframe objects which make use of the mechanism that splits certain details between the Shared Functional Group Sequence (5200,9229) and the Per-frame

Functional Groups Sequence (5200,9230) depending on whether the values apply to specific frames or to all frames.

4.136.4.1.3.2 Display of Qualitative Findings

The Display shall be able to display qualitative findings.

For qualitative findings that are associated with an image, location, or region, the Display:

- shall be able to display qualitative findings in the context of the image with which they are associated
- shall be able to display qualitative findings in the context of the locations or regions with which they are associated, when the latter is activated for superimposed display, regardless of whether the location or region encodings are coordinate, contour, or pixel/voxel-based
- shall be able to display not only the concept name of the qualitative findings but also any associated modifiers for the finding or enclosing Measurement Group including but not limited to Finding Site and Finding

For qualitative findings that are not associated with an image, location, or region, the Display:

- shall provide a method for discovering and displaying the qualitative findings, however the specific presentation details are left to the discretion of the implementation of the Display.

4.136.4.1.3.3 Display of Measurements

The Display:

- shall be able to display measurements in the context of the image with which they are associated
- shall be able to display measurements in the context of the locations or regions with which they are associated, when the latter is activated for superimposed display, regardless of whether the location or region encodings are coordinate, contour, or pixel/voxel-based
- shall be able to display not only the concept name of the qualitative findings but also any associated modifiers for the finding or enclosing Measurement Group including but not limited to Finding Site and Finding

Displays should be prepared for a degree of variety in the measurement codes they receive and the patterns of pre-coordination and post-coordination used. See RAD TF-3: 6.5.3.3.

4.136.4.1.3.4 Display of Locations

The Display:

- shall be able to display point locations as a marker on the referenced image

- shall be able to display line locations as graphics on the referenced image
- shall be able to display the concept name of the location and any associated modifiers for the finding or enclosing Measurement Group including but not limited to Finding Site and Finding (e.g., "Lesion, Center")

The following capabilities are not required for conformance, but were described by domain users as potentially useful and may be worth consideration by implementers.

- may be able to display point locations as markers on images other than the referenced image which are in the same Frame of Reference
- may be able to present an MPR (Multi-Planar Reformat) view centered on a selected point location

The form in which the location marks are displayed may influence user performance, and hence it may be necessary to display them in a manner prescribed by the vendor that generated the results, which is not encoded in the DICOM object. The form of the location mark rendering is not specified by this transaction.

4.136.4.1.3.5 Display of Regions

The Display:

- shall be able to display each region in a way that the user can see the boundary of the planar or volumetric region
- shall be able to display each region as either:
  - a shaded or wire frame volume if the Display supports 3D visualization, or
  - a shaded region superimposed on the appropriate images, e.g., with the region color-coded, or
  - region contours superimposed on the appropriate images
- shall be able, when superimposing regions that are encoded as Segmentation instances, to handle
  - different in-plane or cross-plane extent than the underlying images
  - different in-plane or cross-plane sampling rate (Pixel Spacing or Spacing Between Slices) than the underlying images
- shall be able to display Segmentations that are either BINARY or FRACTIONAL (including fractional types of OCCUPANCY or PROBABILITY)
- shall be able to superimpose multiple selected regions
- shall be able to display the associated region metadata, if available, including:
  - if encoded in an SR instance:

- the concept name of the location and any associated modifiers for the finding or enclosing Measurement Group including, but not limited to,
  - (363698007, SCT, " Finding Site")
  - (121071, DCM, "Finding")
- if encoded in a Segmentation instance:
  - Tracking ID (0062,0020)
  - Segment Number (0062,0004)
  - Segment Label (0062,0005)
  - Segment Description (0062,0006)
  - Segment Algorithm Type (0062,0008)
  - Segmented Property Category Code Sequence (0062,0003)
  - Segmented Property Type Code Sequence (0062,000F)
  - Anatomic Region Sequence (0008,2218)
  - Anatomic Region Modifier Sequence (0008,2220)

- shall allow the user to toggle the display of each region on/off.

The following capabilities are not required for conformance, but were described by domain users as potentially useful and may be worth consideration by implementers.

- may be able to display a group of related regions in the same view
- may be able to generate a volumetric view of the region(s) embedded in the images to which they apply
- may be able to superimpose a region with a different orientation (e.g., Image Orientation (Patient)) than the underlying images (note that Creators may be configurable to produce segmentations that are aligned with the underlying data)
- may be able to apply a Volumetric Presentation State, if present, to the rendering of the region.

Note: An ROI that is encoded in an SR with a reference to a Segmentation may have information duplicated in the SR and Segmentation metadata, e.g., the same value for Finding Site in the SR as in Anatomic Region Sequence, or the same value in the Segmented Property Type Code Sequence as the Finding. In such cases the Display may omit the separate display of duplicate information.

4.136.4.1.3.6 Display of Parametric Maps

The Display:

- shall be able to display a parametric map alone as a color image

- shall be able to identify corresponding images from which a parametric map was derived, as identified in the Source Image Sequence (0008,2112) of the Derivation Image Sequence (0008,9124)
- shall be able to display the meanings of the name-value pairs of the items of Quantity Definition Sequence (0040,9220), including but not limited to the value of (246205007, SCT, "Quantity")
- shall be able to display the meaning of the units from Measurement Units Code Sequence (0040,08EA)
- shall be able to display the values of LUT Label (0040,9210) and LUT Explanation (0028,3003) if Quantity Definition Sequence is absent or empty
- shall be able to display a parametric map superimposed on appropriate underlying images in a manner that is functionally equivalent to the pipeline defined in DICOM PS3.4 Section N.2.4
- shall allow the user to control blending parameters of the display
- shall support resampling of one or both datasets as part of the superimposition process (i.e., to handle cases where the extent, Pixel Spacing, and/or Spacing Between Slices differs)

The following capabilities are not required for conformance, but were described by domain users as potentially useful and may be worth consideration by implementers.

- may be able to perform spatial transformations to handle cases where the Parametric Map is not aligned to the axes of the underlying image
- may be able to use instances of DICOM Blending Softcopy Presentation State and/or DICOM Advanced Blending Presentation State and support more advanced blending capabilities such as references to the fused sets of images as well as compositing factors and pseudo-color palettes.

Related functionality is also described in the Image Fusion (FUS) Profile.

4.136.4.1.3.7 Display of Tracking Identifiers

The Display:

- shall be able to show the Tracking Identifier for any given result
- shall be able to identify results that relate to the same entity, as indicated by sharing the same value for (112040, DCM, “Tracking Unique Identifier"), or within a single instance sharing the same value for (112039, DCM, “Tracking Identifier")
- shall be able to associate results that are related to the same entity, e.g., by allowing the user to sequence through them, showing prior and current side-by-side, or listing them in chronological order

The following capabilities are not required for conformance, but were described by domain users as potentially useful and may be worth consideration by implementers.

- may be able, when the same type of result is identified as relating to the same entity at two different timepoints, to calculate changes, for example the % increase, or the volume doubling time for a tumor

Note: If such derived values are simply present as additional result primitives, having been pre-calculated by an Evidence Creator, then the Image Display will be able to display them, per the requirements above to display basic findings. The same is true for longitudinal assessments.

4.136.4.1.3.8 Display of Image References

The Display:

- shall be able to display the image(s) referenced by any result unless the SOP Class of the image is unsupported by the Display
- shall be able to display the image(s) referenced by a Key Object Selection and shall be able to convey the Document Title and the (113012, DCM, "Key Object Description") content item, if present

4.136.5 Protocol Requirements

N/A.

4.136.6 Security Considerations

This transaction involves presenting DICOM objects that typically constitute personal health information (PHI) to human observers who are typically clinicians. Typical access controls and audit trails in accordance with local policies would be appropriate.

4.136.6.1 Security Audit Considerations

The Radiology Audit Trail Option in the ITI Audit Trail and Node Authentication Profile (ITI TF-1:9) defines audit requirements for IHE Radiology transactions. See RAD TF-3: 5.1.

4.136.6.2 Display Specific Security Considerations

Since this transaction involves the display of PHI, it may be reasonable for Image Displays to implement typical access controls for patient records, such as logins for users and role-based access policies.

Since this transaction involves parsing datasets generated by other systems, it may be reasonable for Image Displays to implement basic digital hygiene, such as sanitizing datasets to avoid malicious executable scripts that might be executed by a browser-based viewer.

In RAD TF Volume 2: RAD-16 & 45 require no changes (Already has Image Display, IM/IA, Imaging Document Consumer)

RAD-18 & RAD-43 require no changes (already have Image Display, IM/IA)

In the WIC TI Supplement, modify Table 4.108.2-1 as shown.

4.108 Store Instances Over the Web [RAD-108]

…

Table 4.108.2-1: Actor Roles

Role:	**Sender**: Creates and sends well-formed DICOM composite objects
Actor(s):	The following actors may play the role of Sender: Image Capturer Lightweight Modality **Evidence Creator**
Role:	**Receiver**: Receives objects from the Sender
Actor(s):	The following actors may play the role of Receiver: Image Manager/Archive

Delete the bubble-box Use Case Diagram in both RAD-107 & RAD-129 (in WIA)

4.107 WADO-RS Retrieve [RAD-107]

4.129 QIDO-RS Query [RAD-129]

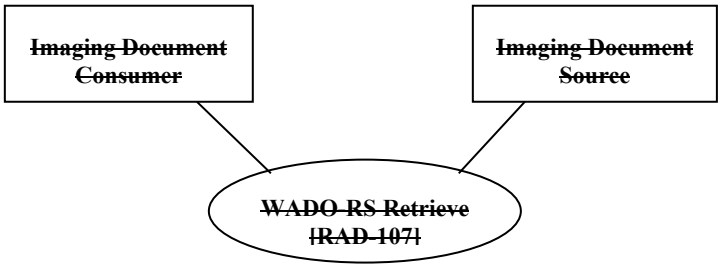

~~Figure 4.*(107 and 129)*.2-1: Use Case Diagram~~

In the WIA TI Supplement, modify the Actor Roles table in both RAD-107 & RAD-129 as shown.

Table 4.(*107 and 129)*.2-1: Actor Roles

Role:	Requester: Submit retrieve DICOM object requests
Actor(s):	The following actors may play the role of Requester: **Image Display** Imaging Document Consumer
Role:	Responder: Returns the requested DICOM object
Actor(s):	The following actors may play the role of Responder: **Image Manager / Image Archive** Imaging Document Source

In the WIA TI Supplement, add SCU keys to Table 4.129.4.1.2-4 as shown.

4.129.4.1.2 Message Semantics

…

The Requester shall support all keys required for the SCU as defined in RAD TF-2: Table 4.14-1.

In addition, the Requester shall be capable of using the following attributes as matching key and return key:

- Issuer of Patient ID (0010,0021)

Note: Issuer of Patient ID is used in conjunction with Patient ID (0010,0010).

- Issuer of Accession Number Sequence (0008,0051)

  Note: Issuer of Accession Number Sequence is used in conjunction with Accession Number (0008,0050). Issuer of Accession Number Sequence is useful when the Requester is dealing with multiple domains.

  Note: For XDS-I Backend Option related requirements on accession number, see Table 4.129.4.1.3.1.

The Requester may implement one or more of the following sets of matching or return keys for the Query SCU:

Table 4.129.4.1.2-4: Additional SCU Query Keys

Query Key Specifications	Type of objects
RAD TF-2: Table 4.15-1	Presentation State Objects
RAD TF-2: Table 4.26-1	DICOM Structured Report Objects
RAD TF-2: Table 4.30-1	Key Image Notes
RAD TF-2: Table 4.44-1	**Evidence Document Objects**
RAD TF-2: Table 4.137-1	**Analysis Result Objects**

In the WIA TI Supplement, add SCP keys to Table 4.129.4.1.3-1 as shown.

4.129.4.1.3 Expected Actions

…

The Responder shall support all keys required for the SCP as defined in the sections referenced in Table 4.129.4.1.3-1:

Table 4.129.4.1.3-1: SCP Query Keys

Query Key Specifications	Type of objects
RAD TF-2: Table 4.14-1	Image Objects
RAD TF-2: Table 4.15-1	Presentation State Objects
RAD TF-2: Table 4.26-1	DICOM Structured Report Objects
RAD TF-2: Table 4.30-1	Key Image Notes
RAD TF-2: Table 4.44-1	**Evidence Document Objects**
RAD TF-2: Table 4.137-1	**Analysis Result Objects**

Add Section 4.137

4.137 Query Analysis Results [RAD-137]

4.137.1 Scope

This transaction is used to query for image analysis results.

4.137.2 Actor Roles

The roles in this transaction are defined in the following table and may be played by the actors shown here:

Table 4.137.2-1: Actor Roles

Role:	Initiator: Queries for analysis result objects.
Actor(s):	The following actors may play the role of Initiator: Image Display Image Document Consumer
Role:	Responder: Returns analysis result entries matching the request.
Actor(s):	The following actors may play the role of Initiator: Image Manager / Image Archive

Transaction text specifies behavior for each role. The behavior of specific actors may also be specified when it goes beyond that of the general role.

4.137.3 Referenced Standards

DICOM PS3.4: Query/Retrieve Service Class

4.137.4 Messages

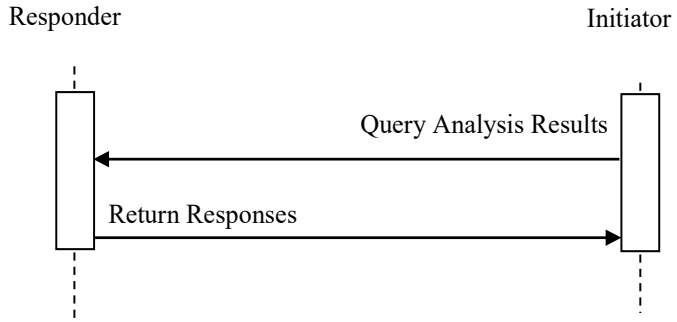

Figure 4.137.4-1: Interaction Diagram

4.137.4.1 Query Analysis Results

The Initiator provides a matching filter in a request for matching analysis results that are available on the Requester.

The Responder shall support handling such messages from more than one Initiator. The Initiator shall support making requests to more than one Responder.

4.137.4.1.1 Trigger Events

A user or an automated function on the Initiator needs information about analysis results available on the Responder.

Typically, the Initiator intends to subsequently retrieve appropriate matching instances.

4.137.4.1.2 Message Semantics

The message is a DICOM C-FIND request of the Query/Retrieve SOP Classes (Study Root – FIND and optionally Patient Root – FIND). The Initiator is the SCU, and the Responder is the SCP.

The Initiator shall send a C-FIND Request from the Study Root Query/Retrieve Information Model – FIND SOP Class or the Patient Root Query/Retrieve Information Model – FIND SOP Class to the Responder.

The Initiator uses one or more matching keys as filter criteria to obtain the list of matching entries in the Responder at the selected level (Patient & Study/Series/Instance).

In addition to the required and unique keys defined by the DICOM Standard, the IHE Radiology Technical Framework has defined matching and return keys to be supported by query SCUs and SCPs. The conventions for key usage are defined in RAD TF-2: 2.2.

The Initiator (SCU) and the Responder (SCP) shall support the following:

- keys identified as required in Table 4.14-1 (Query Images)
- keys identified as required in Table 4.44-1 (Query Evidence Documents)
- keys identified as required in Table 4.137-1 (Query Analysis Results)

Notes: 1. Some of the attributes listed below are indicated as being specific to particular types of instances and the SCP is not required to return those attributes for other types of instances. E.g., attributes under "Segmentation Specific – Instance Level" are only required to be returned for Segmentation instances.

2. The attributes in RAD TF-2: Table 4.44-1 specifically target SR instances, including the Template Identifier (0040,DB00) which will be useful for finding instances using TID 1500. Per DICOM, the value Template Identifier is a string of digits, without leading zeroes, and does not include the string "TID", so for TID 1500 it will be "1500".

Table 4.137-1: Analysis Result Instance Specific Query Matching and Return Keys

Attribute Name	Tag	Query Keys Matching		Query Keys Return	
		SCU	SCP	SCU	SCP
General Analysis Result Specific - Instance Level					
Anatomic Region Sequence	(0008,2218)	O	O	R+	R+
>Anatomic Region Modifier Sequence	(0008,2220)	O	O	R+	R+
Primary Anatomic Structure Sequence	(0008,2228)	O	O	R+	R+

Attribute Name	Tag	Query Keys Matching		Query Keys Return	
		SCU	SCP	SCU	SCP
>Primary Anatomic Structure Modifier Sequence	(0008,2230)	O	O	R+	R+
Image Laterality	(0020,0062)	O	O	R+	R+
Image Type	(0008,0008)	O	O	O	R+
Series Description Code Sequence	(0008,103F)	O	O	O	R+
Contributing Equipment Sequence	(0018,A001)	O	O	O	R+
>Manufacturer	(0008,0070)	O	O	O	R+
>Manufacturer's Model Name	(0008,1090)	O	O	O	R+
>Software Versions	(0018,1020)	O	O	O	R+
>Device UID	(0018,1002)	O	O	O	R+
Segmentation Specific - Instance Level					
Segment Sequence	(0062,0002)	O	O	O	R+
>Segment Number	(0062,0004)	O	O	O	R+
>Segment Label	(0062,0005)	O	O	O	R+
>Segment Description	(0062,0006)	O	O	O	R+
>Segment Algorithm Type	(0062,0008)	O	O	O	R+
>Segmented Property Category Code Sequence	(0062,0003)	O	O	R+	R+
>Segmented Property Type Code Sequence	(0062,000F)	O	O	R+	R+
>Tracking ID	(0062,0020)	O	O	R+	R+
>Tracking UID	(0062,0021)	O	O	R+*	R+
Parametric Map Specific - Instance Level					
Quantity Definition Sequence	(0040,9220)	O	O	O	R+

4.137.4.1.3 Expected Actions

The Responder shall accept and process the request. This involves parsing the matching key values provided by the Initiator, using those to determine matching analysis result instances, and composing response entries, containing the requested return keys, for return to the Initiator in the Return Responses message.

4.137.4.2 Return Responses

The Responder sends matching entries back to the Initiator.

4.137.4.2.1 Trigger Events

The Responder receives a Query Analysis Results message.

4.137.4.2.2 Message Semantics

The message is a DICOM C-FIND response of the Query/Retrieve SOP Classes (Study Root – FIND and optionally Patient Root – FIND). The Initiator is the SCU, and the Responder is the SCP.

4.137.5 Protocol Requirements

N/A.

4.137.6 Security Considerations

The patient demographics and clinical record details returned in the response, and potentially matching details contained in the query, typically constitute personal health information.

4.137.6.1 Security Audit Considerations

This transaction is associated with a Query Information ATNA Trigger Event.

The Radiology Audit Trail Option in the ITI Audit Trail and Node Authentication Profile (ITI TF-1:9) defines audit requirements for IHE Radiology transactions. See RAD TF-3: 5.1.

In RAD TF Volume 2: Modify RAD-14 as shown (Add Image Document Consumer; already has Image Display, IM/IA)

In 11 places in 4.14.4, replace Image Display with Initiator

In 9 places in 4.14.4, replace Image Archive or Image Manager or Image Manager/Archive with Responder

4.14 Query Images [RAD-14]

4.14.1 Scope

The **Initiator**~~**Image Display**~~ queries the **Responder**~~**Image Archive**~~ for study, series and image instances for retrieval.

4.14.2 Actor Roles

~~**Actor: Image Archive**~~

~~**Role: Responds to queries for Studies, Series, and Images.**~~

~~**Actor: Image Display**~~

~~**Role: Issues Queries for Studies, Series, Images**~~

The roles in this transaction are defined in the following table and may be played by the actors shown here:

Table 4.14.2-1: Actor Roles

Role:	**Initiator:** **Issues queries for Studies, Series, Images**
Actor(s):	**The following actors may play the role of Initiator:** **Image Display** **Imaging Document Consumer**
Role:	**Responder:** **Responds to queries for Studies, Series, and Images**
Actor(s):	**The following actors may play the role of Responder:** **Image Manager / Image Archive**

Transaction text specifies behavior for each role. The behavior of specific actors may also be specified when it goes beyond that of the general role.

In RAD TF Volume 2: Modify RAD-44 as shown (Add Image Document Consumer; already has Image Display, IM/IA)

Since RAD-137 points the Image Document Consumer to RAD-44, it is better if IDC is formally mentioned.

4.44 Query Evidence Documents [RAD-44]

4.44.1 Scope

This **~~section describes the sequence of~~** transaction**~~s required for the Image Display to~~** queri**es~~y the Image Archive~~** for instances of Evidence Documents.

4.44.2 Use Case Roles

~~Actor: Image Display~~

~~Role: Query for Evidence Documents objects (generally in order to retrieve them).~~

~~Actor: Image Archive~~

~~Role: Respond to queries from the Image Display for Evidence Documents objects.~~

The roles in this transaction are defined in the following table and may be played by the actors shown here:

Table 4.44.2-1: Actor Roles

Role:	**Initiator:**

	Queries for Evidence Documents objects (generally in order to retrieve them)
Actor(s):	**The following actors may play the role of Initiator:** **Image Display** **Imaging Document Consumer**
Role:	**Responder:** **Returns Evidence Document entries matching the request**
Actor(s):	**The following actors may play the role of Responder:** **Image Manager / Image Archive**

Transaction text specifies behavior for each role. The behavior of specific actors may also be specified when it goes beyond that of the general role.

4.44.3 Referenced Standards

DICOM PS3.4: Query/Retrieve Service Class

4.44.4 Messages

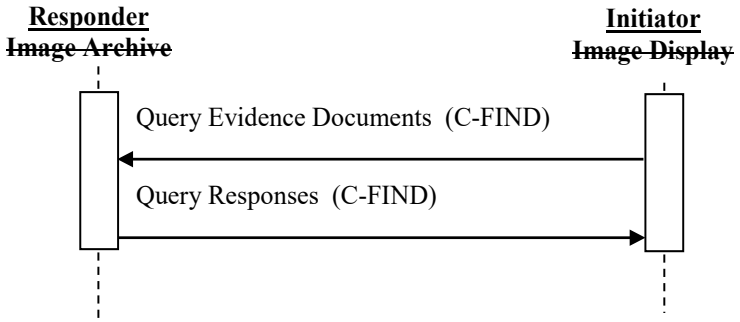

Figure 4.44.4-1: Interaction Diagram

4.44.4.1 Query Evidence Documents

The Query (Study Root – FIND and optionally Patient Root – FIND) SOP Classes shall be supported. Refer to DICOM PS3.4: Query/Retrieve Service Class for detailed descriptive semantics.

4.44.4.1.1 Trigger Events

~~**Image Display**~~**The Requester** needs to obtain information about Evidence Documents.

4.44.4.1.2 Message Semantics

The message semantics are defined by the DICOM Query/Retrieve SOP Classes.

The Initiator shall send a~~A~~ C-FIND Request from the DICOM Study Root Query/Retrieve Information Model – FIND SOP Class or the DICOM Patient Root Query/Retrieve Information Model – FIND SOP Class ~~**shall be sent from the Image Display**~~ to the **Responder~~Image Archive~~**.

The **Initiator~~Image Display~~** uses one or more matching keys as filter criteria to obtain the list of matching entries in the **Responder~~Image Archive~~** at the selected level (Patient & Study/Series/Instance).

In addition to the required and unique keys defined by the DICOM Standard, the IHE Radiology Technical Framework has defined matching and return keys to be supported by query SCUs and SCPs. The keys are defined in RAD TF-2: 4.14.4.1.2 and Table 4.14-1. The conventions for key usage are defined in RAD TF-2: 2.2. For the **Initiator~~Image Display~~** (SCU) and the **Responder~~Image Archive~~** (SCP) the additional Evidence Document Instances specific keys are defined in Table 4.44-1.

Table 4.44-1: Evidence Document Instance Specific Query Matching and Return Keys

Attribute Name	Tag	Query Keys Matching		Query Keys Return	
		SCU	SCP	SCU	SCP
Evidence Document Instance Specific Level					
Content Date	(0008,0023)	O	O	O	R+
Content Time	(0008,0033)	O	O	O	R+
Referenced Request Sequence	(0040,A370)				
>Study Instance UID	(0020,000D)	O	O	R+*	R+
>Accession Number	(0008,0050)	O	O	R+	R+
>Requested Procedure ID	(0040,1001)	O	O	R+	R+
>Requested Procedure Code Sequence	(0032,1064)				
>>Code Value	(0008,0100)	O	O	O	R+
>>Coding Scheme Designator	(0008,0102)	O	O	O	R+
>>Coding Scheme Version	(0008,0103)	O	O	O	R+
>>Code Meaning	(0008,0104)	O	O	O	R+
Content Template Sequence	(0040,A504)				
>Template Identifier	(0040,DB00)	O	O	R+	R+
Concept Name Code Sequence	(0040,A043)				
>Code Value	(0008,0100)	O	O	R+*	R+
>Coding Scheme Designator	(0008,0102)	O	O	R+*	R+
>Coding Scheme Version	(0008,0103)	O	O	O	R+
>Code Meaning	(0008,0104)	O	O	R+	R+

4.44.4.1.3 Expected Actions

The **Responder~~Image Archive~~** receives the C-FIND request, performs the matching on the provided keys and sends the list of matching records back to the **Initiator~~Image Display~~** via C-FIND responses.

The **Initiator~~Image Display~~** is expected to use the Template ID to select Evidence Documents for retrieval that it supports.

4.44.4.1.3.1 Mammography Image Profile

Initiator~~Image Display~~ and **Responder~~Image Manager/Image Archive~~** actors supporting the Mammography Image Profile shall support the Mammography CAD SR SOP Class.

Volume 2x – Appendices

No new Volume 2x appendices

Volume 3 – Content Modules

Add the following rows to RAD TF-3: Table 5.1-2

5.1 ITI-20 Record Audit Event

Table 5.1-2: IHE Radiology transactions and resulting ATNA trigger events

IHE Radiology Transaction	ATNA Trigger Event(s)	Actor(s) that shall be able to record audit event
Patient Registration [RAD-1]	Patient-record-event	…
…		
Display Analysis Result [RAD-136]	**Study-used**	**Image Display**
Query Analysis Results [RAD-137]	**Query Information**	**Image Display** **Image Document Consumer**

Add the following new sections to RAD TF-3 Section 6 IHE Radiology Content Specifications.

6.5 Imaging Analysis Result Content

6.5.1 Scope

This IHE Radiology Content Specification defines standard encodings for the results of medical imaging data analysis. It is specifically intended to cover the output of artificial intelligence algorithms but can also address the results of more conventional analysis software and can also be used for human-generated results.

The specification abstracts imaging analysis results as being composed from a set of primitives. Creators of such results are required, in Section 6.5.3, to be able to encode each primitive using specified DICOM IODs.

The set of primitives is comprised of the following:

Qualitative Findings

This primitive is a coded concept name paired with a coded or text value. This structure is common in SR instances. A qualitative finding may or may not reference a specific location or region.

Algorithms might use this for results like flagging presence/absence of a particular finding, categorizing structures in an image, staging pathologies, etc.

Examples of SR Concept-Value encoding include:

- (RDE436, RADELEMENT, "Presence of pulmonary embolism") = (272519000, SCT, "Absent")

- (129715009, SCT, “Breast composition") = (129719003, SCT, "Extremely Dense")

See RAD TF-1: 49.4.1.4 for a discussion of codesets for concepts and values.

Notes: 1. Whether “Score” findings are qualitative or measurements depends on the semantics. Scores that are rankings using discrete categorizations with specific criteria, like a BI-RADS Score, would be considered qualitative findings and use coded values, while scores that are numerical values on a continuum, like a pain score from 1-10, would be considered numerical measurement values, sometimes with no units.

2. Methods of encoding negative findings is an area of ongoing investigation. Additional guidance will be provided when it becomes available. See also RAD TF-1: 49.4.1.6.

Measurements

This primitive is a structure containing a coded concept paired with one or more numerical values and associated details like measurement units, etc. This structure is common in SR instances. A measurement may or may not be based on a specific location or region.

Algorithms might use this for results like measured tumor sizes, blood flow rates, etc.

Examples include:

- (81827009, SCT, "Diameter") = "26.43","mm"; where
  (363698007, SCT, “Finding Site") = (54247002, SCT, "Ascending Aorta")
- (410668003, SCT, “Length") = "52","mm"; where
  (121071, DCM, “Finding") = (27925004, SCT, "Nodule") and
  (363698007, SCT, “Finding Site") = (23451007, SCT, "Adrenal Gland")
- (285285000, SCT, “Cobb Angle") = "18.4","deg"; where
  (363698007, SCT, “Finding Site") = (122495006, SCT, "Thoracic spine")

Note: Other details about the Cobb Angle measurement, like the specific vertebra selected to derive the Cobb Angle and whether the top or bottom surface was chosen, might be encoded as additional subordinate results.

Locations

This primitive is a spatial location, in the form of coordinates for a point or a line. The coordinates are typically expressed in an imaging coordinate system in the frame of reference of an imaging dataset. The coordinates may be within a single “slice” (2D) or within an imaged volume (3D).

Algorithms might use this for results like the location of a tumor or calcification, a spatial fiducial point, or the tip of a needle or tube.

Regions

The primitive is a planar or volumetric region of an imaging dataset.

Algorithms might use this for results like classifying the type of tissue in different pixels/voxels, or defining the boundary of a particular organ, or encoding a bounding box showing the rough location or extent of an intracranial bleed, or encoding an arbitrary region within which a metric like average density has been calculated. Semantically, a single region can consist of multiple discontiguous parts.

Parametric Maps

This primitive is a set of 2D (or 3D) pixels (or voxels) that are image-like but the values encode a "parameter" such as a physical quantity or quantitative model parameter.

Algorithms might use this for results such as visual explanations (e.g., a “saliency map”, showing which parts of a source image influenced the algorithms decision), or quantitative analysis (showing calculated values like brain perfusion, apparent diffusion coefficients (ADC), or the parameters of a dynamic contrast model.

Note: Tissue classifications may be more appropriately handled as segmentations. Similarly, a probabilistic classification (showing for each pixel the probability it contains malignant cells, or the probability the voxel is part of a hemorrhage) is best encoded as a segmentation, using the FRACTIONAL Segmentation Type.

Tracking Identifiers

This primitive is an identification of a persistent entity in the form of a unique identifier.

Algorithms might use unique identifiers to designate the same lesion or tumor in a specific patient in each of a longitudinal set of studies over time. A display might use the corresponding human-readable identifier to communicate this to a user. Longitudinal analysis and evaluation of change are facilitated when the same Tracking Identifier is associated with other primitives like location coordinates, measurements, or qualitative findings (like tumor stage or therapeutic response).

Image References

This primitive is a reference to one or more image instances, in the form of an instance UID. When there is a need to be more specific for a multi-frame instance, the reference may include frame information.

Algorithms are also expected to commonly reference the “input” image that was analyzed, however that is expected to be part of the result metadata, rather than the result “payload”.

6.5.2 Referenced Standards

The DICOM Standard PS3 is available at: http://dicomstandard.org/current

- DICOM PS3.3: A.35.4 Key Object Selection Document IOD
- DICOM PS3.3: A.35.13 Comprehensive 3D SR IOD
- DICOM PS3.16: TID 1500 Measurement Report
- DICOM PS3.3: A.51 Segmentation IOD
- DICOM PS3.3: A.75 Parametric Map IOD

6.5.3 Analysis Result Encodings

Creators of imaging analysis results shall be capable of encoding them as described in this section. The general requirements in Section 6.5.3.1 apply to the encoding of all of the primitives described in the subsequent sections.

The encodings make use of the following SOP Classes:

- Comprehensive 3D SR Storage (1.2.840.10008.5.1.4.1.1.88.34)
- Segmentation Storage (1.2.840.10008.5.1.4.1.1.66.4)
- Parametric Map Storage (1.2.840.10008.5.1.4.1.1.30)
- Key Object Selection Document Storage (1.2.840.10008.5.1.4.1.1.88.59)

The correspondence between Primitives and the IODs used for encoding is summarized in Table 6.5.3-1.

Table 6.5.3-1: Primitives & Encoding IODs

	SR IOD (TID 1500)	Segmentation IOD	Parametric Map IOD	Key Object Selection Document IOD
Qualitative Findings	X			
Measurements	X			
Locations	X			
Regions	X	X		
Tracking Identifiers	X	X		
Parametric Maps			X	
Image References	X	X	X	X

Notes: 1. For extensive specific guidance on topics such as encoding planar measurements and finding-related coordinates in SR, it is strongly recommended that readers refer to: Clunie, D., DICOM SR for communicating planar annotations, An Imaging Data Commons (IDC) White Paper, 2019 https://docs.google.com/document/d/1bR6m7foTCzofoZKeIRN5YreBrkjgMcBfNA7r9wXEGR4/edit?usp=sharing

2. The use of the word “image” in this section is intended to mean both a single-frame image IOD and a specific frame of a multi-frame image IOD.

6.5.3.1 General Result Encoding Requirements

The Creator is expected to populate much of the contextual metadata (e.g., patient demographics, patient identifiers and issuers, study accession number, etc.) in the result instances based on those values in the medical imaging data being processed, and/or the UPS Workitem (if the Creator is being driven by the AI Workflow for Imaging Profile).

Per DICOM, the attributes of the General Equipment Module describe the overall equipment that created the result instance (i.e., the Creator as a whole, not individual subcomponents or algorithms).

The Creator shall describe each algorithm that was used to generate the results in the Contributing Equipment Sequence (0018,A001). Multiple items may be included. The Creator shall encode the following details in the Contributing Equipment Sequence:

- Purpose of Reference Code Sequence (0040,A170) shall be (Newcode1, 99IHE, "Processing Algorithm")
- Manufacturer (0008,0070)
- Manufacturer's Model Name (0008,1090)
- Software Versions (0018,1020)
- Device UID (0018,1002)

Note: For algorithms that include a human operator, the identity of the user can be encoded in the Operators' Name (0008,1070) attribute or the Operator Identification Sequence (0008,1072).

Each time an AI Model is modified, for example by training, it would be appropriate to update the Device UID.

In SR instances, when an algorithm is identified in the SR Tree (typically in TID 4019 and/or TID 1004) it will also be identified in the Contributing Equipment Sequence (0018,A001). The content items, if present, shall contain the same value as the corresponding attribute shown in Table 6.5.3.1-1.

Table 6.5.3.1-1: Corresponding Algorithm Identification Attributes and Content Items

Contributing Equipment Sequence Attribute	Device Observer Content Item (TID 1004)	Algorithm Identification Content Item (TID 4019)
Manufacturer (0008,0070)	(121014, DCM, "Device Observer Manufacturer")	
Manufacturer's Model Name (0008,1090)	(121015, DCM, "Device Observer Model Name")	(111001, DCM, "Algorithm Name")
Software Versions (0018,1020)		(111003, DCM, "Algorithm Version")
Device UID (0018,1002)	(121012, DCM, "Device Observer UID")	

The Author Observer Sequence (0040,A078) is not usually included in machine generated reports in which that information is the same as in the General Equipment Module (i.e., the authoring device is also the recording device).

DICOM SR Guidance

Content Date (0008,0023) and Content Time (0008,0033) are defined to be the date and time that the document content creation started. In the context of analysis results, these may be considered

to be the date and time that the analysis that generated the result(s) started executing. They are distinct from the Instance Creation Date (0008,0012) and Instance Creation Time (0008,0013), if present which may be some time later.

As an inherent part of the SR Content Tree encoding, each Content Item may have an Observation UID (0040,A171). Observation UIDs may be used to reference specific findings from another instance, but otherwise they are usually absent.

Notes: 1. A Tracking Unique Identifier uniquely identifies the entity that is the subject of an observation (e.g., this patient's aortic valve, or a specific lesion), while an Observation UID uniquely identifies the observation itself (e.g., a diameter measurement of that entity taken today at 2pm). So, a number of observations of the same entity will all have different Observation UIDs but the same Tracking Unique Identifier.

2. DICOM has published Supplement 219 (JSON Representation of DICOM Structured Reports) as a Frozen Draft for Trial Use and Comment. (ftp://medical.nema.org/medical/dicom/supps/Frozen/sup219_fz_JSONSR.pdf) It specifies a simplified JSON representation of the content of a DICOM Structured Reporting instance. The intention is to allow developers to encode image-derived results, in particular, measurements and annotations that might be generated by artificial intelligence (AI), machine learning (ML) and quantitative imaging (QI). The approach is applicable to both export of AI/ML results and also to encoding of truth data for AI/ML training, testing, and validation. Implementers are encouraged to experiment with the encoding and provide feedback.

Frame of Reference (FoR) Guidance

DICOM instances that encode spatial information typically include a Frame of Reference UID (0020,0052) that identifies the origin and coordinate system of coordinates in the instance (without specifically defining that origin and coordinate system). When two instances share the same Frame of Reference UID, coordinates in the two instances are nominally in the same coordinate space.

Creators are strongly encouraged to use shared Frame of Reference UIDs whenever appropriate to avoid placing the burden on downstream systems of performing otherwise unnecessary registrations and transformations. Since the coordinate systems are typically patient-relative, any change of the patient position or use of a different device would result in a new Frame of Reference UID.

Note: Some DICOM Image IODs (like CR images or Ultrasound images) do not necessarily contain a Frame of Reference UID. If DICOM Segmentation IODs are created for such images, the segmentation pixels are required by DICOM to correspond directly to the pixels of the segmented image. In that situation, the segmentation and the image are implicitly in the same Frame of Reference, and if the DICOM Segmentation IOD declares a Frame of Reference UID, the image may then be assumed to be in the same Frame of Reference.

6.5.3.2 Qualitative Findings

Qualitative findings shall be encoded in an instance of the DICOM Comprehensive 3D SR SOP Class using TID 1500 (Measurement Report) as the root template.

- http://dicom.nema.org/medical/dicom/current/output/chtml/part16/chapter_A.html#sect_TID_1500

The value for Procedure Reported (121058, DCM, "Procedure reported") shall describe the imaging procedure analyzed, not the algorithm used.

Depending on whether the qualitative findings refer to a planar region of an image, a volume, or does not refer to any imaging, the root template will contain, respectively, TID 1410 (Planar ROI

Measurement and Qualitative Evaluations), TID 1411 (Volumetric ROI Measurement and Qualitative Evaluations), or TID 1501 (Measurement and Qualitative Evaluations). The CONTAINER for (C0034375, UMLS, "Qualitative Evaluations") in TID 1500 shall not be used, as those can be encoded in TID 1501.

- http://dicom.nema.org/medical/dicom/current/output/chtml/part16/chapter_A.html#sect_TID_1410
- http://dicom.nema.org/medical/dicom/current/output/chtml/part16/chapter_A.html#sect_TID_1411
- http://dicom.nema.org/medical/dicom/current/output/chtml/part16/chapter_A.html#sect_TID_1501

Notes: 1. A qualitative finding that refers to an entire image may be encoded using TID 1501 with an IMAGE reference without spatial coordinates. See also DICOM CP 1999.

2. TID 1411 supports encoding a simple ellipsoid Volume Surface, or a set of isocontours, that may be used to define a bounding region around an entity such as a tumor without precisely delineating its surface. In such cases a (130400, DCM, "Geometric purpose of region") = (75958009, SCT, "Bounded by") modifier is appropriate.

3. Details about the patient condition that represent the context of other observations rather than being observations themselves are best encoded as Observation Context rather than analysis results.

Image references for qualitative evaluations in TID 1410 (Row 6 or 8) and 1411 (Row 6, 11, or 12) are mandated by DICOM for the ROIs that are the basis of the qualitative finding(s).

Image references, if any, for qualitative evaluations in TID 1501 are encoded in Rows 10b-10e. See DICOM CP 1999.

If multiple qualitative findings apply to the same entity (see Section 6.5.3.7 “Tracking Identifiers”), those findings may be encoded in a single invocation of TID 1501, TID 1410, or TID 1411.

A single SR instance may contain multiple invocations of multiple sub-templates (TID 1501, TID 1410, and TID 1411).

Implementations are encouraged to consider encoding anatomical locations for qualitative findings using the Finding Site concept. This can be useful to help organize and database the findings. For example:

(121071, DCM, “Finding") = (55584005, SCT, "Embolism")

(363698007, SCT, “Finding Site") = (39607008, SCT, "Lung")

While it is also possible to use

(121071, DCM, “Finding") = (59282003, SCT, "Pulmonary embolism"),

the Finding Site gives greater flexibility, such as,

(121071, DCM, “Finding") = (55584005, SCT, "Embolism")

(363698007, SCT, “Finding Site") = (45653009, SCT, “Upper lobe of lung")

>(272741003, SCT, “Laterality") = (24028007, SCT, “Right")

or for an intracranial hemorrhage,

(121071, DCM, “Finding") = (50960005, SCT, “Hemorrhage")

(363698007, SCT, “Finding Site") = (78277001, SCT, “Temporal Lobe")

>(272741003, SCT, “Laterality") = (7771000, SCT, “Left")

See also the related discussion at the end of the following Section 6.5.3.3 “Measurements”.

6.5.3.3 Measurements

Measurements shall be encoded in an instance of the DICOM Comprehensive 3D SR Storage SOP Class using TID 1500 (Measurement Report) as the root template.

Depending on whether the measurements reflect a planar region of an image, a volume, or are measurements that are not tied to a planar region or volume, the root template will contain, respectively, TID 1410 (Planar ROI Measurement and Qualitative Evaluations), TID 1411 (Volumetric ROI Measurement and Qualitative Evaluations), or TID 1501 (Measurement and Qualitative Evaluations).

Notes: 1. A measurement of an entire image may be encoded using TID 1501 with an IMAGE reference without spatial coordinates. See also DICOM CP 1999.

2. TID 1411 supports encoding a simple ellipsoid Volume Surface, which is often used to define a bounding region around an entity such as a tumor without precisely delineating its surface. In such cases a (130400, DCM, "Geometric purpose of region") = (75958009, SCT, "Bounded by") modifier is appropriate. An associated volume measurement would be expected to reflect the volume of the ellipsoid, not the volume of the tumor.

Image references for measurements in TID 1410 (in Row 6 or 8) and 1411 (in Row 6, 11, or 12) are mandated by DICOM for the ROIs that are the basis of the measurements.

Image references, if any, for measurements in TID 1501 are encoded in the invocation of TID 320 “Image or Spatial Coordinates” inside the invocation(s) of TID 300 “Measurement” used to encode the measurements inside TID 1501.

If multiple measurements apply to the same entity (see Section 6.5.3.7 “Tracking Identifiers”), they may be encoded in a single invocation of TID 1501, TID 1410, or TID 1411, i.e., in a single Measurement Group.

The anatomy, if relevant, is encoded in (363698007, SCT, "Finding Site"). See RAD TF-1: 49.4.1.4 for a discussion of codesets for Finding Site.

Note: TID 1501 (in TID 300), and TIDs 1410 and 1411 (in TID 1419), allow the (363698007, SCT, “Finding Site") to be encoded for each measurement, or if common, factored out to the Measurement Group level.

The finding to which a measurement pertains may be encoded in (121071, DCM, “Finding"), e.g., (86049000, SCT, “Neoplasm Primary") or (65818007, SCT, “Stent").

Implementations are encouraged to follow the pattern of encoding the anatomical location in Finding Site, the “pathology” or device, if appropriate, in Finding, and the rest of the semantics into the Concept Name of the measurement. Some examples are shown in Table 6.5.3.3-1. This is particularly useful when multiple measurements or observations are made of the same structure.

Table 6.5.3.3-1: Measurement Encoding Examples

Finding Site	Finding	Measurement Concept
(54247002, SCT, “Ascending Aorta")	n/a	(81827009, SCT, “Diameter")
(122495006, SCT, “Thoracic spine")	n/a	(285285000, SCT, "Cobb Angle")
(23451007, SCT, “Adrenal Gland")	(27925004, SCT, “Nodule")	(410668003, SCT, “Length")
(64033007, SCT, “Kidney")	(86049000, SCT, “Neoplasm Primary")	(118565006, SCT, “Volume")
(52433000, SCT, “Proximal Circumflex Coronary Artery")	(65818007, SCT, “Stent")	(410668003, SCT, “Length")

Applications may elect to use pre-coordinated codes, if available for specific findings, such as a measurement with a Concept Name of (18043-0,LN,"Left Ventricular Ejection Fraction by US"). Creators should be aware that receiving systems use codes to identify corresponding measurements to populate report fields, populate databases, or use values to drive clinical decision support. The difficulty for receiving systems increases when various Creators use different patterns of pre-coordination and post-coordination that mean the same thing, e.g.:

- Finding Site = Adrenal Gland, Finding = Nodule, Measurement = Length
- Finding Site = Adrenal Gland, Measurement = Nodule Length
- Finding = Adrenal Gland Nodule, Measurement = Length
- Measurement = Adrenal Gland Nodule Length

The encoding pattern in Table 6.5.3.3-1 is just a recommendation. It is beyond the scope of this Content Specification to constrain all possible encodings. In future, templates that specialize TID 1500 and constrain the choice of codes, may be defined for specific use cases. DICOM PS3.16 already contains many application-specific templates.

6.5.3.4 Locations

Point locations and line locations shall be encoded in an instance of the DICOM Comprehensive 3D SR Storage SOP Class using TID 1500 (Measurement Report) as the root template.

If the point location is a 2D point on a single image, the instance shall contain an SCOORD (111030, DCM, "Image Region") of Graphic Type POINT in TID 1501 (Measurement and Qualitative Evaluations).

If the point location is a 3D point, which may or may not be in the plane of a given image, the instance shall contain an SCOORD3D (111030, DCM, "Image Region") of Graphic Type POINT in TID 1501 (Measurement and Qualitative Evaluations). See DICOM CP 1999.

If the line location is a sequence of 2D points on a single image, the instance shall contain an SCOORD (111030, DCM, "Image Region") of Graphic Type POLYLINE in TID 1501 (Measurement and Qualitative Evaluations).

If the line location is a sequence of 3D points in an image volume, the instance shall contain an SCOORD3D (111030, DCM, "Image Region") of Graphic Type POLYLINE in TID 1501 (Measurement and Qualitative Evaluations). See DICOM CP 1999.

Image references for 2D point or line locations are mandated by DICOM in TID 1501 (in Row 10d).

Image references, if any, for 3D point or line locations in TID 1501 can be encoded in the Image Library of TID 1500 that is the parent of the TID 1501.

As described in DICOM, the Frame of Reference for the SCOORD location values is that of the mandatory referenced image (SELECTED FROM). The Frame of Reference for the SCOORD3D location values is encoded within the SCOORD3D content item data structure, and is generally expected to match the Frame of Reference of associated volume data. See the Frame of Reference Guidance in Section 6.5.3.1 for further details.

Locations may or may not be the basis for associated measurements or qualitative findings. If they are, the location will be encoded in conjunction with the associated measurement or qualitative finding. See Sections 6.5.3.2 and 6.5.3.3.

In addition to, or instead of, spatial coordinates, a location may be expressed in anatomical terms in the (363698007, SCT, "Finding Site") concept as described in Sections 6.5.3.2 and 6.5.3.3.

6.5.3.5 Regions

Both planar regions and volumetric regions may use either contour-based encodings, or pixel/voxel-based encodings. Points and lines are encoded as locations, see Section 6.5.3.4.

Regions may or may not be the basis for associated measurements or qualitative findings. When associated with measurements, contour-based regions will be encoded in conjunction with the associated measurements or qualitative findings in a sub-template of TID 1500, and pixel/voxel-based regions will be referenced from the associated measurement or qualitative finding. See Sections 6.5.3.2 and 6.5.3.3.

It is not necessary to create a region for measurements or qualitative findings of an entire image or series, since the contents of TID 1500 permit referencing the images directly.

Regions associated with an entity may define the surface of the entity, the entire internal volume of the entity, or may be a bounding box or ellipsoid that roughly encompasses the entity. Such semantics are encoded in (130400, DCM, "Geometric purpose of region") in TID 1410 or 1411.

Contour-based Regions

Contour-based regions shall be encoded as an instance of the DICOM Comprehensive 3D SR Storage SOP Class using TID 1500 (Measurement Report) as the root template. Planar regions are encoded in TID 1410 (Planar ROI Measurement and Qualitative Evaluations). Volumetric regions are encoded in TID 1411 (Volumetric ROI Measurement and Qualitative Evaluations).

If the contour-based planar or volumetric region is comprised of image-relative (2D) contour(s), the instance shall contain one or more SCOORD (111030, DCM, "Image Region") of Graphic Type POLYLINE, CIRCLE, or ELLIPSE.

Note: As defined in DICOM, to represent a closed planar polygon using image-relative coordinates, a POLYLINE is encoded with the last vertex the same as the first vertex.

If the contour-based planar or volumetric region is comprised of volume-relative (3D) contour(s), the instance shall contain one or more SCOORD3D (111030, DCM, "Image Region") of Graphic Type POLYGON, ELLIPSE, or ELLIPSOID.

Image references for contour-based regions are mandated by DICOM in TID 1410 (in Row 6 or 8) and TID 1411 (in Row 6, 11, or 12).

Pixel/Voxel-based Regions

Pixel/voxel-based regions shall be encoded as instances of the DICOM Segmentation Storage SOP Class. One or more segments (multiple regions, or multiple segments of a single region) may be encoded in a single Segmentation instance. A single region shall not span multiple Segmentation instances.

Notes: 1. The DICOM Surface Segmentation Storage SOP Class is not used in this profile.

2. Segmentation instances permit encodings that are binary (i.e., a single bit value defines whether a pixel/voxel is part of the region or not) or fractional (i.e., a scaled integer indicates either the percent probability the pixel/voxel is part of the region, or the fraction of the pixel/voxel that is part of the region).

3. Neither DICOM nor this profile supports label maps, in which a single encoded pixel represents an index into a map of labels for different segments. Instead it is expected that such information will be encoded with a separate bit plane per segment in a Segmentation instance.

4. Segmentation instances may be referenced in the relevant row of the corresponding template when qualitative findings or measurements are based on the pixel/voxel-based region. See Sections 6.5.3.2 and 6.5.3.3.

The Segmentation instances need not cover the same spatial extent nor be sampled at the same rate within plane (i.e., pixel spacing) or between planes (i.e., spacing between slices) as the images being segmented. For example, a Segmentation:

- might only encode frames that span the region of interest (e.g., a particular tumor) and not the entire volume of the images underlying the segmentation (e.g., the chest CT series)
- might cover a smaller region within each frame, sufficient only to include the region of interest
- might have a Pixel Spacing or a Spacing Between Slices that is smaller (super-sampled) or larger (sub-sampled) compared to the original images

The extent and sampling of the Segmentation are encoded within the attributes of the instance.

The Creator that encodes Segmentations shall be capable of encoding them with the same orientation as the associated image.

Note: This is required because image displays might not be capable of retrieving registration objects and/or perform such transformations. The Creator is not prohibited from encoding Segmentations in different orientations, as long as it can be configured to encode them with the same orientation.

Image references to images underlying Segmentations are recorded in the Source Image Sequence (0008,2112) in the Derivation Image Sequence (0008,9124) with a Purpose of Reference Code Sequence (0040,A170) value of (121322, DCM, "Source image for image processing operation") and a Derivation Code Sequence (0008,9215) value of (113076, DCM, "Segmentation").

6.5.3.6 Parametric Maps

Parametric Maps shall be encoded in an instance of the DICOM Parametric Map Storage SOP Class. For a given parametric map, the pixel values are either all floating point, all double, or all integer.

For Parametric Maps that represent saliency images for a finding (e.g., pneumothorax):

- the instance containing the finding shall reference both the Parametric Map instance, and the associated Image instance to which the finding pertains
  - the use of the (130401, DCM, "Visual explanation") concept in TIDs 1410, 1411, and 1501 for the reference to the Parametric Map is encouraged
- the Parametric Map instance shall use a value of (130404, DCM, "Saliency") for the Concept Name of (246205007, SCT, “Quantity") in the Quantity Definition Sequence (0040,9220)

The Parametric Map instances need not cover the same spatial extent nor be sampled at the same rate within plane (i.e., pixel spacing) or between planes (i.e., spacing between slices) as the associated “underlying” images. For example, a Parametric Map:

- might only encode frames corresponding to some of the associated images (or frames)
- might cover a region that is smaller or larger than the associated images (or frames)
- might have a Pixel Spacing or a Spacing Between Slices that is smaller (super-sampled) or larger (sub-sampled) compared to the associated images

The extent and sampling of the Parametric Map are encoded within the attributes of the instance.

The Creator shall be capable of encoding the Parametric Map with the same orientation as the underlying image.

For Parametric Maps that are intended to be available to overlay on an underlying image (e.g., the associated image from which the parametric map was derived), the Parametric Map shall use the Frame of Reference of the image on which it is intended to be overlaid. While the Creator is required to be able to encode the Parametric Map with the same orientation as the underlying image, it is permitted to have different pixel spacing, so the Image Display will not have to retrieve registration objects and perform transformations but may have to perform resampling. These requirements are intended to facilitate fused display by the Image Display without requiring sophisticated registration capabilities.

Image references to “underlying” images from which a parametric map was derived are encoded in the Source Image Sequence (0008,2112) of the Derivation Image Sequence (0008,9124) with a Purpose of Reference Code Sequence (0040,A170) value of (121322, DCM, "Source image for image processing operation"). If the parametric map is a saliency map, a Derivation Code Sequence (0008,9215) value of (130404, DCM, "Saliency") would be appropriate.

Blending Presentation State objects and Advanced Blending Presentation State objects may be created to encode fusion overlay parameters preferences, but the Creator should not assume that the Image Display supports those.

6.5.3.7 Tracking Identifiers

Tracking Identifiers for persistent entities shall, if encoded in instances of DICOM SR Storage SOP Classes, use concept (112039, DCM, "Tracking Identifier") for human readable Tracking Identifiers and concept (112040, DCM, "Tracking Unique Identifier") for globally unique Tracking Identifiers (UIDs). Such Identifiers allow correlating multiple different observations of the same entity (such as an organ or lesion) when an algorithm is capable of making that determination.

If a measured entity is a known entity (e.g., it is a particular tracked lesion), TID 4108 (Tracking Identifier) may be included in the invocation of TID 300 for the measurement and the Tracking Identifier(s) of the measured entity encoded therein, or the tracking identifier concepts in TIDs 1501, 1411, and 1410 may be used to identify the entity that is the subject of the entire Measurement Group.

Notes: 1. If there is a need to track specific observations (the diameter measurement recorded at 2:59pm on Aug 21 of lesion 45837459), each Content Item in the SR tree may have an Observation UID (and an Observation Datetime) associated with it as described in DICOM PS3.3 Table C.17-6.

2. TID is an abbreviation of Template ID in DICOM, not Tracking ID.

3. How Tracking Identifiers are assigned and correlated is not specified here. In some cases, an algorithm may assign them and may be able to inspect prior observations and determine the current entity is the same as the prior and thus it is appropriate to reuse the prior Tracking Identifier. In other cases, Tracking Identifier values may be provided to the algorithm by another system, e.g., as parameters in a workitem, or from an entity correlation service, or by reference to some external resource.

6.5.3.8 Image References

Image references that are subordinate to other results (e.g., references to images on which a measurement was taken, or references to the images that were processed to create a segmentation or parametric map) are discussed above in several sub-sections of Section 6.5.3.

Image references that comprise the result itself (for example identifying a relevant prior, or identifying an image that is most representative of an already identified finding) shall be encoded as an instance of the DICOM Key Object Selection Document Storage SOP Class.

Key Object Selection Documents communicate the purpose of the selection in the Document Title (i.e., the Concept Name of the root Container in TID 2010), with additional semantics optionally encoded in one or more (113011, DCM, "Document Title Modifier") values.

Appendices to Volume 3

Add new Appendix A to Volume 3.

Appendix A – Example Analysis Result Encodings

These examples were prepared in support of the AI Results (AIR) Profile. See Section 6.5.3 Imaging Analysis Result Content for the encoding specifications.

Due to the maintenance overhead for IHE Profile documents, appendices like this one do not typically contain extensive examples. Additional examples may be available as part of the implementation materials associated with this profile on the IHE website or in IHE Connectathons.

A.1 Common Header Attributes

The IODs used for Imaging Analysis Result Content (see Section 6.5.3) all include a number of common DICOM modules describing the Patient, Study, Series, etc. with which the analysis result is associated. Table A.1-1 provides an example of values for the attributes of such modules. Note that, depending on the specific IOD used, values such as SOP Class UID (0008,0016) and Modality (0008,0060) would differ and additional required attributes would be expected.

Table A.1-1: Common Imaging Analysis Result Header Example

Attribute	Tag	Value
SOP Class UID	(0008,0016)	1.2.840.10008.5.1.4.1.1.88.34
SOP Instance UID	(0008,0018)	1.2.3.4.5.6.7.300
Study Date	(0008,0020)	20191029
Content Date	(0008,0023)	20191029
Study Time	(0008,0030)	154500
Content Time	(0008,0033)	155210
Accession Number	(0008,0050)	123456
Modality	(0008,0060)	SR
Manufacturer	(0008,0070)	Aperture Labs
Referring Physician's Name	(0008,0090)	Johnson^Cave^^Dr.^M.D.
Manufacturer's Model Name	(0008,1090)	AI Portal Platform
Patient's Name	(0010,0010)	Person^Chell^^^
Patient's ID	(0010,0020)	000001498
Issuer of Patient ID	(0010,0021)	Mercy Hospital

Attribute	Tag	Value
Patient's Birth Date	(0010,0030)	19991109
Patient's Sex	(0010,0040)	F
Contributing Equipment Sequence	(0018,A001)	
>Purpose of Reference Code Sequence	(0040,A170)	(Newcode1, 99IHE, "Processing Algorithm")
>Manufacturer	(0008,0070)	Acme Algorithmics
>Manufacturer's Model Name	(0008,1090)	PneumoScan
>Software Versions	(0018,1020)	3.1.202006
>Device UID	(0018,1002)	1.2.8.7.6.5.78765234
Study Instance UID	(0020,000D)	1.2.3.4.5.6.7.100
Series Instance UID	(0020,000E)	1.2.3.4.5.6.7.2001
Study ID	(0020,0010)	345678
Series Number	(0020,0011)	3
…		

A.2 Qualitative Finding

As described in Section 6.5.3 Imaging Analysis Result Content, qualitative findings are encoded in DICOM SR objects using TID 1500 as the root template.

The example in Table A.2-1 represents the SR Content Tree inside the DICOM SR object. The header of the DICOM SR object would look like the example in Section A.1.

Notes: 1. The Image Library container is left empty since there is only a single source image which is considered to be reasonably well described by the imaging procedure.

2. The algorithm has also generated a saliency map image for the embolism finding. The saliency image is referenced in the Visual Explanation concept.

Table A.2-1: Qualitative Finding SR Content Example

Node	Code Meaning of Concept Name	Value	TID
1	Imaging Measurement Report		TID 1500
1.1	Language of Content Item and Descendants	(en-US, RFC5646, "American English")	TID 1204
1.2	Observer Type	(121007, DCM, "Device")	TID 1001 TID 1002
1.3	Device Observer UID	1.2.8.7.6.5.78765234	TID 1004
1.4	Device Observer Manufacturer	Acme Algorithmics	TID 1004
1.5	Device Observer Model Name	PneumoScan	TID 1004
1.3	Procedure Reported	(43778-0, LN, "Portable XR Chest AP supine")	TID 1500

Node	Code Meaning of Concept Name	Value	TID
1.4	Image Library		TID 1500
1.5	Imaging Measurements		TID 1500
1.5.1	Algorithm Name	PneumoScan	TID 1419
1.5.2	Algorithm Version	3.1.202006	TID 1419
1.5.3	Measurement Group		TID 1501
1.5.3.1	Tracking Identifier	Pneumo2394958	TID 1501
1.5.3.2	Tracking Unique Identifier	1.2.276.0.7230010.384756	TID 1501
1.5.3.3	Finding	(55584005, SCT, "Embolism")	TID 1501
1.5.3.4	Finding Site	(45653009, SCT, "Upper lobe of lung")	TID 1501
1.5.3.4.1	Laterality	(24028007, SCT, "Right")	TID 1501
1.5.3.5	Source Image	1.2.840.10008.5.1.4.1.1.1.1.1 *(Digital X-Ray Image Storage - For Processing SOP Class)* 1.2.3.4.5.6.7.8001 *(SOP Instance of XR Chest AP Image)*	TID 1501
1.5.3.6	Visual Explanation	1.2.840.10008.5.1.4.1.1.30 *(Parametric Map Storage SOP Class)* 1.2.3.4.5.6.7.9001 *(SOP Instance of corresponding saliency map)*	TID 1501

