

*Application of the ILO International Classification of Radiographs of
Pneumoconioses to Digital Chest Radiographic Images*

A NIOSH Scientific Workshop
March 12-13, 2008
Washington DC, USA

Standardizing file formats, security, and integration of digital chest
image files for pneumoconiosis classification

Author: David A. Clunie, CTO, RadPharm, Inc.

© Copyright 2008 David A. Clunie.

Permission is granted to NIOSH to reproduce this work in part or entirety in any form.
All other rights reserved.

Background, scope, and assumptions

This paper considers the classification of PA projection chest radiographs for the purpose of identification and scoring of pneumoconiosis, by trained and certified human readers (NIOSH B readers), using radiographic images acquired, distributed, and displayed in digital form. Images digitized from traditional film-screen radiographs will not be considered, except with respect to the provision of reference images.

It is expected that the B readers are not necessarily affiliated with or credentialed by the acquisition site, and hence may not be providing direct patient care but rather independent review. This means that they may need to use their own local reading equipment, manufactured by a different vendor than the equipment used for acquisition. Further, the readers may or may not have duty of care to patient, which means that they may or may not have legitimate access to patient's identity (such as SSN).

Types of digital data

As discussed by others in these proceedings, there are essentially two families of imaging technology used for the acquisition of chest radiographs, Computed Radiography (CR), which uses a cassette-based workflow that involves a cassette "reader", and more recently Digital Radiography (DR), which uses a sensor that is in the x-ray path and directly connected to the processing and storage equipment. The DR sensor may be of the direct or indirect type, the latter using a phosphor in addition to the detector. The type of acquisition device makes little difference in the context of file format standardization, except to note that there are both historical factors as well as processing differences that result in difficulties in handling and displaying images with a comparable appearance.

File format

There is, for medical imaging, only one file "format" in widespread use. That is the Digital Imaging and Communications in Medicine, or DICOM, standard. Over the last 15 years, DICOM has become ubiquitous, and is supported by all modern devices sold in all countries. DICOM is a global standard. The committee that manages the standard is an international one, and recently the International Standards Organization has adopted

DICOM, as ISO 12052. DICOM is the only non-proprietary inter-vendor standard in use between acquisition devices (modalities) and Picture Archiving and Communications Systems (PACS) and display workstations.

A format like DICOM is required for medical imaging, rather than consumer formats like TIFF and JPEG. There is a need to encode images with greater than an 8 bit depth, to preserve the full dynamic range of the types of sensors current utilized in acquisition devices. There is also a need to standardize and encode additional information in the image header, including patient identification and demographics, management information such as the date and time, as well as descriptors of radiograph acquisition techniques, such as kVP and exposure.

However, the DICOM standard defines many different capabilities for different applications, and also is evolving, adding new features to support new technologies. It is therefore important to define which “flavor” of DICOM format is used for the storage of digital chest radiographs for classification. In DICOM, each service is linked to a storage object and referred to as a service-object pair (SOP) class.

For projection radiographs there are two choices, the “old” CR SOP class, and the “new” DX SOP class. The CR SOP class was present in the first published DICOM standard in 1993. It was designed specifically for Computed Radiography, and predated the development of direct digital acquisition technology. It is very loosely constrained in terms of which attributes are required or optional, and does not define a consistent grayscale space. This makes it difficult to assure that images are displayed with a similar appearance on devices provided by different vendors. The more recent DX SOP class was added to DICOM in 1998 in order to support all forms of projection radiography, whether acquired using CR technology or DX technology. DICOM took the opportunity to incorporate the lessons learned from half a decade of experience with the CR SOP class, and defined a more robust and consistent object. To reiterate, the DX SOP class supports encoding of images regardless of the acquisition technology, and specifically was intended to replace the “old” CR SOP class for both CR and DX applications.

Unlike the CR SOP class, the DX SOP class clearly distinguishes two types of images, depending on the phase of processing:

- ***For Processing***
- ***For Presentation***

For Processing images are those that require further processing before they are suitable for viewing by a human, whereas ***For Presentation*** are those that are ready to view. The reason to make the distinction is to ensure that in a multi-vendor environment that there will be no confusion about which images a workstation should make available for display, and no ambiguity about which device is responsible for performing the processing. Furthermore, DICOM specifies that all devices are required to support ***For Presentation*** images, and that ***For Processing*** images are optional. This requirement is to prevent the possibility that one device produces only ***For Processing*** images and the other displays only ***For Presentation*** images, and hence are incompatible.

The DX SOP class, like all “modern” DICOM image objects, also defines a standard grayscale output space for ***For Presentation*** images, in order to achieve consistency of appearance regardless of the display device. This is achieved by specifying an output space in “P-Values”, which have a defined meaning for a display device that is calibrated according to the DICOM PS 3.15 Grayscale Standard Display Function (GSDF). Note that the goal is only to assure that a single image will have consistent perceived contrast on different display devices. It is not to make different images appear consistent with each other. Nor does GSDF compliance necessarily result in a “better” displayed image than another choice of display function.

The DX SOP class also makes mandatory many attributes that were optional in the CR SOP class, provides standard codes for attributes that were previously text strings, to ensure that there will be sufficient information about such things as orientation and laterality to allow the images to be displayed correctly. The overall design objective of the new object was to enable all DX ***For Presentation*** images to be reliably and consistently displayed on any vendor’s equipment. By contrast, the CR SOP class images’ consistency depends on the type of image configured at acquisition. In particular

images may need further “processing” at time of display, but there is no standard way to convey whether or not this is the case.

Having the ability to further process an image whilst it is displayed is a desirable feature. At a single site, it is a nice feature to have in the PACS if one has the same PACS vendor as the CR or DX vendor. However, one cannot be sure of the presence of processing capability when the images are sent elsewhere to be read. This is why DICOM established the concept of *For Presentation* images as a mandatory baseline, with *For Processing* images as a desirable option. Theoretically, a “standard” could describe a “raw” space in which *For Processing* images could be encoded, after vendor and detector specific corrections had been performed. Subsequently, “standard” processing methods could be applied to the raw image, regardless of acquisition vendor. However, for the time being this remains an area for further research.

Accordingly, in order to achieve interoperability with the current state of the art, it is essential to insist that as a minimum “processed” images are sent, whether they be DX SOP class *For Presentation* images, or CR SOP class images that have been processed appropriately. This requires an *a priori* choice of processing algorithm and parameters, and it may be desirable for NIOSH to specify these in advance for each manufacturer and system, to achieve consistency of appearance. In reality, one must consider that it may be difficult to influence the acquisition sites to configure their systems to use a particular processing choice, particularly if they are using the system for other work.

DICOM – more than a file format

The DICOM standard also defines many services, including protocols for transferring images and objects across both local networks and the Internet, services for performing queries and retrieving lists of patients and studies, for supplying work lists, and for printing to film. Further, DICOM standardizes the use of interchange media, including the transfer of images using CD, DVD and USB. This allows for the establishment of a so-called “sneaker net”, which allows the transport or mailing of images. The purpose of these services is to allow automation of interoperability, and to avoid, for example, manual loading or dragging and dropping of image “files”. The relevant DICOM services

for workflow and image transfer need to be considered in the context of classification of chest images for pneumoconiosis.

First, the radiology facility acquires digital images. From that point, if the site has a PACS, a DICOM transfer to PACS will occur; if there is no PACS, then the modality will burn the image files to a DICOM-compatible CD. If the images are to be read locally, then they will either be displayed on a workstation built-in to the PACS, or on a 3rd party DICOM workstation attached to the PACS or the modality.

However, if the images are to be sent to an off-site B reader and classified elsewhere, they need to be exported from the modality or PACS, which can be done by:

- burning them to a DICOM-compatible CD,
- sending them via a network to the remote reading site, or
- making them available over a network for remote viewing

Any of these possibilities are feasible, but as a matter of expedience, in the absence of a secure network infrastructure, the use of DICOM-compatible CDs may be most practical.

Software compatibility issues

Standards such as DICOM are required to achieved interoperability, but are not entirely sufficient. There remains the potential for incompatibility between acquisition, transfer and display software, arising from areas not addressed by the standard, ambiguities in the standard, optional features of the standard that a vendor has chosen to not implement, and issues arising from software quality issues (bugs).

For an acquisition device, incompatibilities can arise because there is a choice of different SOP classes (CR, DX) as well as the difference between *For Processing* and *For Presentation* image types, as already discussed. Because the DICOM CR SOP class does not require use of GSDF P-values, there is also the problem of inconsistent grayscale contrast. For both CR and DX SOP classes, there may be problems caused by incorrect encoding of look-up tables applied to stored image pixels to derive display values; adherence to the standard in this respect is not universal among vendors, and display

software may need to account for this. Other configuration issues may arise as a consequence of image acquisition vendors' need to make their equipment highly configurable to work with the vagaries and limitations of a range of PACS, old and new, commercial and homegrown. Though the configuration may satisfy the local users, the configuration choices may have unintended and undesirable consequences when images are sent off site. Insistence upon compliance with the *DX For Presentation* SOP class mitigates this class of issues.

The mechanics of transferring images using the DICOM protocol on a local network is rarely an issue, is widely tested, and is essentially a pre-requisite for PACS to work at all. However, DICOM CD compatibility is not universally reliable; some vendors default to proprietary CD formats, some vendors write incompatible DICOM CDs, while the use of data compression may raise issues. There is renewed emphasis by the industry on CD compatibility testing, but problems can arise with equipment that remains in the field. As a practical matter, given the limited number of sites acquiring and sending digital radiographic images for pneumoconiosis screening and research, each site can be prequalified and an appropriate transfer mechanism and procedure established and tested.

CDs are often burned with Windows-compatible digital image viewing software, nominally capable of displaying the images on the CD. Problems may arise with these viewers for a number of reasons. The issues relate to PC hardware and operating system versions, software installation, display speed (if the viewer is run from CD), and display compatibility, especially if a calibrated display is required. Security policy in place at the facility may impose restrictions, since executing code from outside local systems and networks creates a risk of propagating viruses, as well as a risk of interference with locally installed applications. There are also training and usability issues; readers may need to learn to use multiple different viewing software products, depending on the number of sites sending images. Furthermore, image viewing software is typically designed for review and not for primary interpretation, and may prominently display disclaimers to that effect. Such software may have limited functionality and in particular lack full grayscale pipeline support, so the displayed images may not demonstrate the intended grayscale contrast.

Accordingly, it is strongly recommended that a single dedicated viewing software product is used, and that the execution of the CD-based viewer be suppressed, either with an appropriate registry setting or holding the shift key down whilst inserting the CD.

Display software requirements

The display software product that is selected needs a number of specific features to minimize compatibility problems. The product must support the different DICOM SOP classes, specifically both CR and DX. It must utilize images that are *For Presentation* and “ready to view”, not *For Processing* or raw. The product must implement a full grayscale pixel processing pipeline, in order to produce a consistent grayscale contrast appearance, which includes the correct application of “lookup tables” in the image header, as well as support a pair of GSDF calibrated displays to enable side-by-side display of the selected ILO standard image and the subject image. The software should provide methods for managing the various LUT problems that have been recognized when using acquisition devices provided by some vendors.

A base set of features is essential, providing the user with the ability to zoom, pan, and window the image. Additional features can be considered, such as enhancement and image processing, and the ability to handle unprocessed images.

Image contrast adjustment is a particularly important feature. Despite the broad exposure latitude of digital acquisition, a single default presentation of image contrast is usually not sufficient. The user will often need to adjust the image to better evaluate light and dark areas. Traditionally, linear window center and width controls are available, but these work poorly for projection radiographic images beyond a very small range of adjustment. Non-linear contrast adjustment using a sigmoid shaped curve (analogous to the H and D curve of film response) may achieve more satisfactory results. This can be achieved using lookup tables (LUT) or continuous functions. Though the details are beyond the scope of this presentation, DICOM supports all three mechanisms. It is important that the display software product that is selected support them all, and allow user adjustments. Note also that the image (or saved presentation states that can be reapplied to the image) may

contain multiple choices of presentation, and the display software product should provide these choices to the user.

Compression

Compression was mentioned as a potential software compatibility issue. DICOM does support a range of ISO standard compression schemes, should they be necessary. Both lossless (reversible) compression (using JPEG, JPEG-LS or JPEG 2000) and lossy (irreversible) compression (using JPEG or JPEG 2000) are supported. The typical size of digital radiographic images is such that compression is rarely necessary on CD, though compression may increase the speed of network image transfer. However, given the uncertain impact on detection and characterization of subtle lung disease, it is recommended that lossy compression be avoided for pneumoconiosis classification functions.

Reference images

During the performance of pneumoconiosis classifications, the ILO requires that the subject radiograph be compared side-by-side to the appropriate comparison images from the ILO standard reference set of chest images. With digital image acquisition and soft copy display, there is an obvious need for the ILO standard reference chest images to also be digitally displayable alongside the patient images. Use of a separate light box on which to display reference hard copy films is impractical, since it degrades workflow and alters perception (by being too bright). There is therefore a need for a digital version of the reference sets, with comparable contrast and processing to digital acquisitions. The process of digitizing (previously copied) reference films may result in significant quality degradation and dramatically different contrast characteristics when compared with digitally-acquired images. Ideally, a new set of digitally-acquired (rather than digitized) reference films will become available. Even then, there remains the issue that different acquisition technology, vendor, and processing may alter the chest image appearance.

The ILO digital reference image set should be encoded as DICOM DX *For Presentation* SOP class, to maximize consistency of display on different workstations and to allow images to be stored within the PACS for comparison. Digital encoding also provides the

opportunity to distribute the standard images at negligible cost on the Internet, and hence increase their availability to the public, for training and research.

Displaying the reference images requires them to be made available as either a reference set “built in” to the display software, or as a “pseudo-patient” with dummy identifiers, allowing them to be used with ordinary software. In general, however, neither existing PACS workstations nor off-the-shelf (OTS) 3rd party DICOM generic workstations have explicit mechanisms to display “reference” images. For safety reasons, many systems prevent the showing of multiple “patients” simultaneously.

A customized workstation may need to be developed, dedicated to the classification task performed by B readers. Such a workstation could perform a DICOM query to the PACS to retrieve a patient’s images, and/or read them from CD, and have the ILO standard reference images installed and available locally. There are sufficient customizable open source and commercial toolkits for DICOM support and image display available to make this a feasible option.

Display equipment

As discussed elsewhere in these proceedings, achieving digital display quality comparable to film typically requires dedicated high intensity grayscale display devices, currently provided by flat panel digital displays rather than CRTs. These can be expensive, so the number of displays should be carefully considered. Traditional PACS workstations use a pair of portrait three-megapixel (3MP) displays, side-by-side, with the intent of displaying a current and prior or PA and lateral chest image pair with one image filling the area of each display.

The current film-based ILO viewing guidelines require simultaneous display of two films, the subject plus a single reference, though three are recommended to allow the subject film to be hung between references. One can simulate such a comparison with two digital displays, particularly if the software allows the user to rapidly toggle between one reference and another on a single comparison display; this presupposes that the software can order the reference images by increasing profusion and size. Accordingly, two displays are recommended.

Given the cost, some B readers may opt to re-use their existing infrastructure, particularly if classifications are for patients “outside” their hospital or office. This is analogous to importing patients’ images for “consultation”, a common practice. Some PACS already have explicit support for importation and reconciliation of foreign identifiers (see, for example, the IHE Import Reconciliation Workflow (IRWF) profile). Alternatively, readers may be able to view images from a CD inserted locally. When not possible, whether by policy or technology limitations, one can install a separate computer but share the expensive high resolution grayscale PACS displays by using KVM (keyboard-video-monitor) switches; devices that can support 3MP displays are available.

Remote reading

Rather than sending the images to the reader, it is possible to read “remotely”. Images can be provided on a central server, and accessed via a secure Internet connection. Typically, the display software is remotely accessed, managed and maintained, and the reader’s local machine provides Internet access and high quality monitors.

The mechanism of implementation varies, and may involve the use of a browser applet or plug-in, or an installable local client (ActiveX or Java Web Start). Acceptable performance and a satisfactory user experience is largely a function of the speed of the connection. Additionally, the patient and reference images can be pre-loaded locally in anticipation of their being needed (e.g., by looking ahead at work list tasks). Lossy compression is sometimes used to accelerate responsiveness, but this is unlikely to be acceptable for pneumoconiosis classification applications.

In the future, remote reading could be enabled by NIOSH providing a central archive server and a remote access mechanism, and supplying the same client software for all readers.

Cross-enterprise sharing

The problem of sharing patient related images and documents between loosely coupled organizations is not new. The Integrating the Healthcare Enterprise (IHE) effort initially focused on using existing standards (DICOM and HL7) to support workflow within an

enterprise, but is now expanding to cross-enterprise document sharing (XDS), including images (XDS-I), and is starting to address cross-enterprise user authentication (XUA) and patient identifier cross-referencing (PIX).

The National Cancer Institute (NCI) Cancer Biomedical Informatics Grid (caBIG) also provides a mechanism for secure access to shared resources & services, including support for DICOM images.

In the future, it may be feasible and routine for authorized B readers to obtain remote access to images stored at the acquisition site, using an extension of infrastructure in place for routine patient care.

Security and privacy

Films, digital images, reports and forms may contain individually identifiable health information (IIHI), which is “protected” (PHI) under the HIPAA Privacy Rule.

Unauthorized access to IIHI or PHI is a bad thing. It is necessary to either protect or remove such information.

Digital data is at risk when in physical form, for example when exchanged on CD, in the same manner as such information on film and paper. On-line digital information is also at risk, both locally, to access by unauthorized staff, and remotely, to access by unauthorized individuals when in transit.

Protection of PHI in transit requires encryption on the network, for example by using SSL, as is used in electronic commerce on the Internet. Also required are authentication, such as by login with username and password, access control such as by configuring access rights that are constrained based on identity, and maintenance of an audit trail, a record of who saw and did what, when and where.

DICOM provides network security services built on top of existing security mechanisms, using virtual private networks (VPN) or SSL connections for privacy. The user identity can be conveyed in a DICOM connection. The PACS can constrain access and maintain an audit trail. There are also standard mechanisms defined for web-based access to DICOM images (WADO), which use normal browser security mechanisms, and can also

use VPN and SSL support. In such cases, the web server handles authentication, access, and audit trails.

An alternative approach is to remove PHI in digital images before transfer; if the information isn't there, then it does not need to be protected. Whenever readers do not need the patient's true identity, one can replace the patient's name, SSN and other identifiers with pseudonymous numbers. The association of the pseudonym and true identity can be securely maintained centrally. This process is used in clinical trials, in which independent readers are "blinded" to a subject's identity, both to reduce bias and protect privacy.

In DICOM images, the identity is normally stored in the digital image file headers, not burned in to the image pixels, which makes it relatively easy to automate such de-identification.

Integrating results with images

The DICOM standard addresses not only the encoding of images, but also image-related structured information, through the Structured Reporting (SR) mechanism. DICOM SR allows for the storage of codes, text, measurements, coordinates, references to images, locations and outlines, in a hierarchical organization. The structure can be pre-defined by "templates" for specific applications. Examples of such templates in the standard include the basic radiology report, computer-aided detection (CAD) results for mammography and chest images, radiation dose reports, and measurements for echocardiography, obstetrics, and cardiovascular CT.

The DICOM SR approach shares the same header structure as images, and in particular uses the same patient, study and series model. DICOM SR is widely used by modalities to encode measurements. They are easy to store in, and retrieve from, the PACS. They convert easily into other forms, and can, for example, be transformed into HL7 CDA XML, or extracted and rendered as plain text or PDF. The contents can be searched and specific structures and codes extracted. SR can also be considered as a DICOM "form", and can reference images and locations within images by coordinates.

Potentially, DICOM SR templates could be defined for ILO classification data, and could encode both the NIOSH Roentgenographic Interpretation and Miner Identification forms. Such templates would include standard codes for each concept, a reference to the unique identifier (UID) of the image being read, and possibly references to prior images used for comparison. Additionally, one could save image “annotations”, including pointers to locations of representative abnormalities in the image. An SR can also be digitally signed, as can any images referenced, using a standard cryptographic public key-based mechanism.

Conclusion and recommendations

An entire infrastructure already exists to support the clinical use of digital projection radiographs, based on the use of the DICOM standard between modalities, PACS, and workstations, with networks and CDs. Many sites now have considerable experience with exporting and providing outside access to digital images, including *For Presentation* digital radiographs. Correct choice, or construction, of an appropriate image viewer should allow consistent display and reliable review of images, side-by-side with ILO standards or equivalent reference images. Expensive displays already installed can easily be utilized.

If advantageous, the results of pneumoconiosis classifications can be stored as a DICOM Structured Report, and the DICOM organization can help with adding templates and codes as requested. Matters of security and privacy can and should be addressed through conventional means.

In terms of specific recommendations for the NIOSH B reader program, both CR and DX DICOM images should be permitted, due to the large installed base of devices and the continuing availability of equipment that does not support the DX object. Processed (*For Presentation*) images should be required, to avoid the B reader being dependent on proprietary processing in display workstations. However, it is desirable to also obtain unprocessed (*For Processing*) images if possible, to allow for future research into CAD, as well as to permit development of standardized mechanisms for image processing.

Display workstations that support B readers should be qualified and certified; they must be fully compatible with test images from different vendors and software, must support all variations of encoding and grayscale pipeline, and must be able to display reference images side-by-side with the patient's image.

Images should be de-identified before sending for reading if possible, to minimize the risk of leakage of IIII/PHI.

Ideally a digital, rather than digitized film, reference set should be created and released, which is comparable in contrast and resolution to CR and DX images.

Initially, it may be expedient to distribute images to B readers on DICOM CDs, but NIOSH should explore the creation of a managed distributed or centralized infrastructure, to allow for remote reading from a central PACS, or an XDS or caBIG network.