

DICOM Implementations for Digital Radiography¹

Cost-effective deployment of digital radiographic (DR) facilities depends on efficient image management and soft-copy reading. Requirements for productivity and quality led the Digital Imaging and Communications in Medicine (DICOM) Standards Committee to extend the standard to include a family of image storage objects that specifically addressed the needs of DR. These digital x-ray (DX) objects incorporate lessons learned from experience with integrating computed radiography (CR), picture archiving and communications systems (PACS), and workstations. Factors addressed include the mandatory incorporation of information required for correct routing of images, appropriate organization of images for display (driving hanging protocols), and consistent appearance of gray-scale contrast in a distributed environment. The use of coded terminology to describe projections, techniques, and anatomic structures was emphasized.

This chapter describes the requirements and design of the DX family of image storage objects and the mechanisms by which the equipment can gather the necessary information efficiently. Advice is given to users in the form of purchasing guidelines, which emphasize key factors in the successful deployment of integrated DR equipment. The status and the adoption of the DX objects by vendors are reviewed, and the reasons for the relatively slow adoption of the objects are considered.

BACKGROUND

The DICOM standard defines protocols, objects, services, and conformance requirements. Support for the standard is a prerequisite for effective communication between radiologic imaging devices. The DICOM image objects include those designed for projection radiography. DICOM image objects are defined by "information object definitions," which include information encoded in attributes, together with the actual pixel data. These attributes describe identification, management, and acquisition technique information. The information object definitions are combined with services, such as the "storage service class," to produce "service-object pair classes," which are the unit of conformance in DICOM. A device may act as a "service class user" (eg, one that sends images) or a "service class provider" (eg, one that receives images).

Most of the benefits of flat-panel digital detector technology are unrelated to DICOM. The quality of the image and the potential for lower dose to the patient are a

Advances in Digital Radiography: RSNA Categorical Course in Diagnostic Radiology Physics 2003; pp 163–172.

¹From Princeton Radiology Pharmaceutical Research (RadPharm), 943 Heiden Rd, Bangor, PA 18013 (e-mail: dclunie@dclunie.com).



consequence of the sensor itself. The rapid room turn-around is largely a result of almost instant feedback (with the exception of cassette-based CR systems), with no wait for film cassette processing. Indeed, digital detectors have been deployed at some sites without integration into a PACS or soft-copy reading environment, with direct printing to film after acquisition. Realistically, however, most digital detector implementations are likely to involve centralized archiving of images, distributed soft-copy reading, and digital distribution of images to referring physicians.

DICOM is merely an interface tool for such integrated environments, providing standardization of communication across the boundaries between devices. As such, it is no panacea, nor is it the primary factor in cost-effective deployment of digital detector technology. In the absence of DICOM, however, one would be forced to depend on the completely proprietary equipment of a single vendor throughout. The DICOM standard includes services that (a) address work flow, both within an examination room and throughout the enterprise, (b) support efficient organization of images for display (hanging protocols), and (c) support distributed consistency of image appearance. Many of these services are generic, are not specific to digital flat-panel sensor technology, and are well described in the documentation of the Integrating the Healthcare Enterprise (IHE) initiative (1). IHE documentation does not, however, address the “payload” of image objects, whether it is specific to flat-panel digital sensors or any other acquisition modality. The content of an image object remains the province of the DICOM standard itself and is primarily what will be addressed here.

WORK FLOW

Early experience with integrating images for general single-frame projection radiography applications with PACS was largely based on the use of CR technology. Initially, the CR work flow mimicked the screen-film cassette work flow, often involved printing to film, and was labor intensive, with little attention paid to downstream integration by the vendors. Manual entry of patient and study identification information by technologists and radiographers led to frequent mistakes. Such mistakes were not terribly problematic in a print-to-film environment but resulted in incorrectly managed or missing studies and mismatched requests as PACS became more sophisticated.

Experience not just with CR but also with other modalities, such as computed tomography, magnetic resonance imaging, and ultrasound, resulted in the introduction of the DICOM “modality worklist” (MWL) service to address this problem. Nowadays, MWL support is universally regarded as fundamental to the efficient operation of modalities. This support allows an

Table 1
Categories of Diagnostic Radiology-related Image Storage Information Objects in DICOM

Category	DICOM Modality
Projection radiography	
Computed radiography	CR
Secondary capture (as used for screen-film radiography)	OT
X-ray angiography	XA
X-ray fluoroscopy	XRF
Digital x ray	DX
Digital mammography	MG
Intraoral radiography	IO
Cross-sectional imaging	
Computed tomography	CT
Magnetic resonance	MR
Ultrasound	US
Nuclear medicine	NM
Positron emission tomography	PET

operator to select a patient to be examined together with the request, rather than to manually enter the information. MWL support greatly reduces operator error and allows the DICOM image header to be populated not only with correct identifiers but also with other relevant information that does not need to be manually entered. The importance of support for the MWL in all modalities, including DR, cannot be over-emphasized and leads to *purchasing guideline 1: Do not buy a DR or mammography system or PACS without MWL*. MWL is the single greatest DICOM-related contributor to improved system productivity.

IMAGE TRANSFER AND STORAGE

The most widely known DICOM services are those related to the transfer of images and related objects across the network, referred to in the standard as “storage.” Images are transferred (stored) in a modality-specific form, which includes common identifying information as well as information about the specific acquisition technique. Pixel data are also stored in a modality-specific form. Two broad categories of image storage information objects can be identified, those for projection radiography and those for cross-sectional imaging, as described in Table 1. Specific attention will be paid to the advantages of the DX family of image objects (DX, digital mammography, and intraoral radiography) over the older CR object.

In designing the DX objects, the responsible DICOM working group determined that the other DICOM services (related to study and work-flow management and query and retrieval) were sufficient to support DR applications. Attention was therefore directed only to designing new image objects.

CR OBJECT LIMITATIONS

The purpose here is not to compare CR with other digital technologies but rather to address the limitations of the early DICOM CR image storage object for storing projection images. It would have been (and still is) entirely possible to encode DR images as CR objects.

First, the CR object does not provide attributes to describe the digital flat-panel acquisition process. Apart from some rudimentary technique attributes related to x-ray exposure and projection geometry, there is little more than a few text fields that describe the body part examined, the view, and the cassette or plate. In practice, when additional information has been required, manufacturers have encoded it in private attributes.

The anatomic structures and view are defined by a small set of defined terms, which are text strings. Experience has shown that these text strings are insufficient for many common applications and are often supplanted by the manufacturer's own string values. In many cases, because the standard allows for these attributes to be left empty, vendors or users simply fail to send a value. Indeed, almost every attribute of the CR image may either be omitted or left empty "if unknown." A workstation receiving such an image cannot depend on (a) any of these attribute values being present or (b) the attribute fields containing useful or consistent information.

Also problematic is the lack of definition of what the image pixel values mean and how they are to be displayed. The gray-scale space is not formally defined, the relationship of the pixel values to x-ray intensity is not defined, and whether or not the image data have been processed into a form suitable for presentation is not specified. Because of variation in how vendors encode CR image pixel data and what transformations workstations subsequently apply (whether linear window operations or nonlinear lookup table operations), there is considerable inconsistency in image appearance. The consequence has been considerable dissatisfaction among users trying to integrate equipment from different vendors, which is alleviated only by extensive tuning of lookup tables until a tolerable (if not desirable) result is achieved.

DX DESIGN GOALS

A primary design goal was to support the new flat-panel digital sensor technologies, both direct and indirect (ie, with or without a scintillator), as well as various approaches that used charge-coupled devices or slit-scanning. In addition, it was expected that existing technologies, such as CR, selenium drums, and even optical scanning of film, could be supported to take advantage of the new features of the DX objects.²

The DX objects had to take into account the characteristics of the new technologies, both in terms of the image

pixel data and the description of the acquisition process. Support for nonlinear lookup table gray-scale contrast transformations of pixel values was required because linear windowing of x-ray data does not typically produce a satisfactory display appearance. The relationship of pixel intensity to x-ray intensity (positive or negative, linear or logarithmic) needed to be recorded. Of particular importance was the need to encode quality control information related to the acquisition, dose, detector behavior, and detector identification.

From the perspective of the requirements of PACS, it was suspected that cost-effective deployment of digital detector technology might well depend on efficient image management and efficient soft-copy reading. One goal, therefore, was to encourage the advantages of digital detectors by focusing on opportunities to improve PACS usability and productivity. Because the modality groups and PACS groups within manufacturers typically operate separately (and, sadly, often at cross purposes), it was deemed necessary to include many mandatory attributes in the DX objects to encourage modality vendors to take PACS needs into account.

IDENTIFYING PACS NEEDS

Two major categories of PACS needs were identified: those related to image management functions and those related to soft-copy reading functions. The image management functions include matching the images with the request and with images from previous studies, as well as routing of images to the appropriate soft-copy reading workstation. The soft-copy reading functions include presenting the images in the correct order and orientation and with the appropriate gray-scale contrast.

Failure to meet these needs has important implications for productivity and usability. Radiologists cannot read images without the request, the request without the images, current images without the prior images, or images that have not successfully been included in the worklist or made available to the workstation. Furthermore, radiologists will not read, or will read slowly, images that (a) are in the wrong order, (b) are upside down or the wrong way round, or (c) have the wrong contrast.

Ideally, the user-friendly aspects of conventional screen-film technology should be emulated successfully. The result achieved must be comparable with a correctly exposed, properly positioned set of films hung in the correct order with relevant prior examinations by an experienced technologist or file room clerk. The productivity of radiologists requires that their work flow not be interrupted by having to

² One disappointing aspect of the relatively slow adoption of the DX objects is that film-scanning software and CR devices have not begun to use the DX objects.

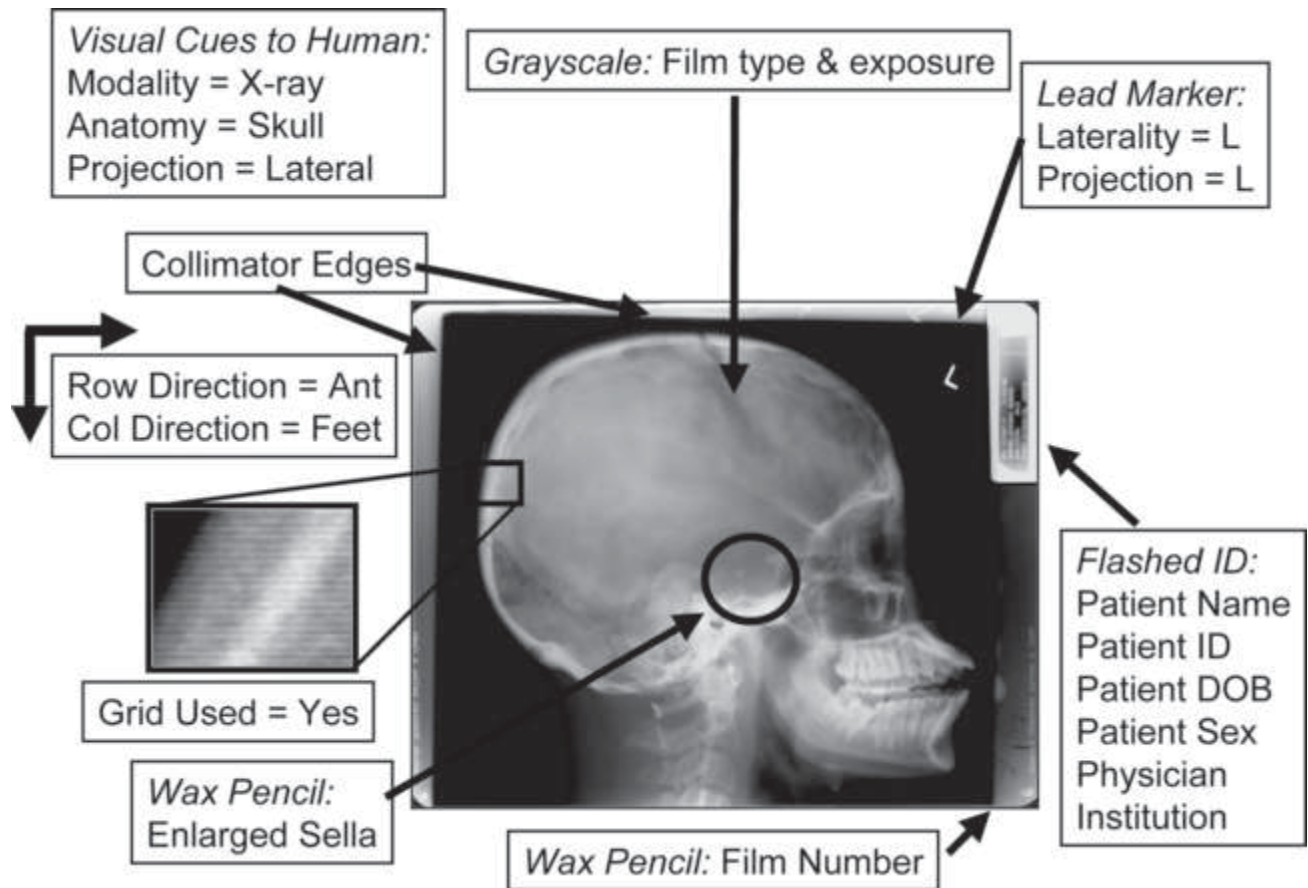


Figure 1. Features present on a film. From a film, a broad range of information is available, which either is explicitly “encoded” or may be implicitly derived by a knowledgeable human observer. *Ant* = anterior, *Col* = column, *DOB* = date of birth, *ID* = identification, *L* = left.

reorder, reposition, or adjust the gray-scale appearance of displayed images.

To emulate all of the functions of film successfully, one must (a) replicate the equivalent of the visual cues that the file clerk, technologist, or radiologist uses to hang a film; (b) replicate the function of the flashed identifier plate, any lead markers included in the image, or wax pencil marks used to annotate features; and (c) provide a well-defined and repeatable gray scale (Fig 1).

The process of hanging films involves a series of steps that need to be emulated in the digital process, specifically (a) extracting the films from the patient’s folder, (b) sorting the films into old and new, (c) verifying the identification of the patient on each film, (d) arranging the films into the desired order, matching prior films with current films for the same anatomic structures and view, and (e) rotating or flipping the films to the correct orientation (eg, patient’s left on the right of the view box, feet on the bottom).

Similar steps need to be applied by a workstation “hanging” images for display: (a) reception or retrieval of all of the images for the current and prior studies from the reading worklist, (b) matching of the modality and anatomic structures and other attributes with the appropriate (radiologist-specific) hanging protocol, and (c) display of the images per the hanging protocol, including arrangement by prior and current images, arrangement by anatomic structures and view, rotating and flipping based on orientation, annotation as necessary, and selection from the available or protocol default gray-scale contrast choices.

Display workstations may support a range of different capabilities, from the simple display of images as encoded in their numeric order to user-customizable rule-based hanging protocols. Figure 2 provides an example of the relevant information present in a DICOM DX image.³ Note in particular the mechanism that is used to describe the orientation of an image with respect to the patient. The directions of the rows and columns are each specified as toward the patient’s anterior (A) or posterior (P), head (H) or feet (F), and left (L) or right (R). When the orientation is not parallel to one of these orthogonal axes but is instead oblique, multiple letters are used, such as in the example

³ Note that in the example in Figure 2, human-readable values for the attributes are shown, although in reality, these attributes would be coded as either string values (eg, L for left, MG for digital mammography) or as a triplet of values specifying a coding scheme, code value, and code meaning (eg, SNM3, T-04000, breast).




Figure 2. Information available in a DICOM DX image to help in hanging images with the correct order and orientation and matching them with relevant prior images. *A\FR* = anterior/foot right.

of a left mediolateral oblique mammogram, in which the column direction is toward the patient's feet and also toward the right (because it is a left mediolateral oblique view).

This mechanism of describing the orientation of an image is used consistently throughout the DICOM projection radiography image objects, but only in the DX family is it mandatory to supply this information. The rationale for omitting it from earlier image objects was that the information was difficult to obtain or that entering it required manual intervention by the operator, perhaps while reviewing the image on a quality control display. Workstations typically displayed the image with whatever orientation it happened to be encoded. Users have been forced to depend on lead markers burned in the pixel data to determine how an image is oriented with respect to the patient.⁴ For the DX objects, it was decided to force the issue and insist that the acquisition device or the operator supply the orientation information.

Similarly, designation of the laterality of an image (eg, whether it is the left or the right hand) is optional in other DICOM image objects but mandatory in the DX family of objects. Radiologists expect left and right or current and prior images to be matched and arranged appropriately, and without the laterality information, such arrangement would be impossible.

One side effect of these decisions is that only a single exposure may be encoded in one image. If multiple projections of a body part, for example, were included in the same image, it would not be possible to describe a consistent orientation and laterality. This matter was extensively discussed, and it was accepted that most of the flat-panel sensor technologies would not support such multiple exposures in a single "image." However, when screen-film or CR cassettes are used, it is not uncommon to include multiple exposures of small body parts (eg, hands and fingers), which makes it difficult to use the DX image object for such CR applications. One solution is to allow the user to designate the area of each exposure from the entire cassette image on a quality control workstation before the images are sent to the PACS.

This issue is typical of the trade-offs that are a consequence of attempts to improve usability and productivity in one area of the work flow, at the expense of increasing the effort required at another step. During the design of the DX objects, a conscious decision was made to improve downstream usability (eg, by radiologists and referring physicians), at the expense of a more complex technologist interface if necessary, on the premise that an image is acquired only once but may be used many times. In practice, creative implementation strategies in the acquisition device can minimize the burden on the operator. The operator must review all images before dispatch to the PACS to check for correct positioning, motion, and exposure. The in-room productivity gains of almost-immediate image display largely offset any minor inconvenience to the operator caused by requirements to orient and specify the laterality of an image.

Table 2 summarizes the fundamental differences between the CR and DX image objects with respect to the information required for arranging images for display. The key distinguishing features of the DX objects are that more critical attributes are required to be included with values by the acquisition device, and more critical attributes are coded with consistent values. The standard provides an extensive dictionary of coded values to describe the anatomic region and radiographic projection (including eponymous named projections).

⁴ Lead markers in the pixel data may interfere with automated image-processing algorithms; being so dense, they artificially skew the histogram of pixel values. However, it is probably unrealistic in the short term to expect users to forgo them completely and depend on DICOM header attributes. There would be a perceived, if not real, safety risk.

The need to provide sufficient attributes to facilitate efficient and safe image arrangement for display leads to *purchasing guideline 2: Insist on DX support in acquisition devices for DR and CR, as well as PACS workstations*. To reiterate, it is entirely possible and appropriate to use DX objects to encode images acquired with CR technology, as long as (a) a user interface is provided to enter the necessary attributes, and (b) the potential for multiple exposures per plate is taken into account.

It is also important to emphasize that a PACS must not merely passively receive and regurgitate DX images. The workstations for radiologists and referring physicians⁵ must use the additional information in the DX object to order, orient, and annotate the displayed images appropriately. This behavior is outside the scope of the DICOM standard to specify and is also rarely described in conformance statements. This leads to *purchasing guideline 3: Insist on hanging protocols driven by DX attributes in PACS reading workstations*. Mandatory coded attributes from the modality yield no benefit if they are not used.

IMPLEMENTING DX ATTRIBUTES

As has been alluded to, the simplest approach to populating the mandatory DX attributes describing anatomic structures, view, orientation, and laterality is to provide an operator user interface. Such an interface can be incorporated into the normal image quality control step. It is a small matter to require the operator to confirm or correct these values.

Ideally, the anatomic structures and view can be determined automatically, or a small subset of possibilities can be identified, for example, from MWL information or information from the generator protocol used for the exposure. A multitude of data sources are potentially available, including (a) procedure codes and protocol codes from the worklist, (b) generator protocol selection, (c) collimator selection, (d) physical gantry configuration, (e) filtration selection, (f) grid selection, and (g) detector values and statistical analyses of detector performance.

Apart from providing appropriate patient and study identifying information, the scheduled procedure step selected by the operator from the worklist may well be sufficiently detailed to permit derivation of the anatomic structures, view, and perhaps laterality. The success of this approach depends on a consistent understanding by the PACS or information system and the modality of the codes exchanged between them. Progress in this regard has been considerably ham-

Table 2
Information for Arrangement of Images for Display: CR versus DX

Attribute	CR	DX
Modality	Nonspecific (CR)	More specific (DX, mammography; intraoral radiography)
Anatomy	Optional and text	Required and coded
Laterality	Optional	Required
View	Optional and text	Required and coded
Orientation	Optional	Required

pered by the lack of a standard set of radiology procedure codes for ordering; the coding schemes typically used for billing are often not appropriate for ordering and may be hampered by proprietary or licensing restrictions. Furthermore, modality vendors are wary of including dependencies on site-specific code tables in their products.

Integrated rather than add-on DR systems have a potential advantage in that the generator information may be available to the system building the DICOM image objects. Certainly, the information about exposure values and dose can be incorporated automatically, as well as used to drive choices of frequency and contrast processing algorithms. If the generator-user interface supports the concept of body part- and projection-specific protocols, this information can be reused in the DICOM object, as illustrated in Figure 3. Thus, the operator is not forced to enter the same information twice. Indeed, in a well-integrated system, the generator protocols can be made artificially more detailed to include such items as laterality. Thus, the operator could choose, for example, a left lateral knee rather than a right lateral knee and get the same exposure but more information with which to populate the image attributes.

Furthermore, for systems that have fixed geometry, such as an upright bucky for chest examinations, the orientation of an image with respect to the patient can be determined automatically. For example, even though the exposure may be the same for anteroposterior and posteroanterior chest images, distinguishing these two allows the DICOM subsystem not only to describe the view more precisely but also to determine the orientation to be "right/feet" rather than "left/feet" as appropriate. Indeed, in anticipation of the naive behavior of conventional workstations that do not make use of the orientation attributes, the system may choose to automatically flip the encoded image pixels to the most common left/feet orientation that users expect to see, regardless of whether the image was acquired as anteroposterior or posteroanterior. Otherwise, by default, the pixel data may simply be encoded as viewed from the tube side on every occasion.

⁵ Referring physicians are sometimes not considered in the ordering, orienting, and annotating of projection radiographic images. There is no reason, for example, why a Web browser interface to a PACS cannot use the DX attributes to render the study easier to review, rather than displaying the images in whatever arbitrary form they were received.




Figure 3. Generator protocol selection information can be re-used in DICOM attributes, to avoid the need to enter values manually.

Although in the design of DX objects, the efficiency of the downstream PACS work flow has been given priority over the convenience of the operator, there is still no reason to burden the operator unnecessarily. The more information that can be derived automatically, the better. Minimizing reentry of information and reuse of available information should be high priorities for the modality designer. This leads to *purchasing guideline 4: Choose a DR modality that fully populates attributes but has minimal effect on the productivity of the operator.* Many sources of information are automatically obtainable and reusable, and in-room productivity gains are too valuable to sacrifice unnecessarily.

CONSISTENCY OF APPEARANCE

Distributed consistency of image appearance is something of a holy grail—much desired, poorly understood, and rarely attained for digital projection radiographs. The primary objective of most users is to achieve filmlike appearance and consistency. No matter what the other benefits of digital radiography may be, if the images do not “look right” everywhere, the result is extreme dissatisfaction. Incorrect image contrast is a source of inefficiency, fatigue, distraction, and potential diagnostic error. Radiologists cannot realistically be expected to manipulate every image to achieve a suitable image appearance. Appropriate defaults are essential.

The builders of CR and DR systems take considerable effort to develop appropriate image-processing algorithms to optimize the appearance of their images. These algorithms may involve the use of simple or conventional gray-scale contrast techniques, tuned for specific body parts or exposure ranges, and may provide for frequency-selective contrast enhancement. Regardless, all such algorithms are designed and tested with a particular output device in mind. Tuning the algorithms and their parameters is a remarkably empirical process, often involving a cycle of repeated expert evaluations until the desired effect is achieved.

Evaluations are typically performed by printing the processed images to film and then distributing the films among the experts. When these algorithms are deployed in the field, unless the display or print device is exactly the same as was used in the tests, including the same lighting conditions, the result may diverge from what is expected.

The solution to this problem is to make use of an idealized gray-scale “space” that is chosen to allow images to be rendered on a variety of output devices, such that a human observer perceives approximately the same equivalent contrast.⁶ In principle, at least for gray-scale images, this approach allows an image to “look the same” regardless of whether it is printed on film and hung on a light box or displayed on a low- or high-luminance monitor. The number of individually detectable gray-level differences⁷ may vary between output devices, but the overall contrast appearance is similar enough to satisfy most users. Remarkably, despite the vast difference in luminance range of soft-copy display devices and printed films hung on a light box, reasonable consistency can be achieved.

The DICOM standard specifies such a gray-scale space in the form of the gray-scale standard display function (GSDF) (Fig 4). It is widely adopted by PACS workstation vendors and supported by the manufacturers of both cathode-ray tube and flat-panel gray-scale displays intended for medical applications.

The CR image object in DICOM has no defined gray-scale output space. It is entirely up to the manufacturer how to encode the image pixel data and any window or lookup table information, and there is no consistency among vendors as to how the result is displayed. Essentially, the appearance on film or on the screen is arbitrary and dependent on such factors as the gamma of the device and any corrections that may or may not be applied by the display application, the operating system, or the display driver.

Furthermore, even if one were to assume that display devices have similar gray-scale contrast characteristics, there is no requirement for the display device to apply any window values or lookup tables that may be present in the CR image object. Indeed many display workstations ignore the lookup tables entirely. Worse, some CR vendors have encoded their lookup tables incorrectly, which results in black images on those workstations that do apply the lookup table.

A different approach is used for the DX family of objects. Every DX object that is intended for presentation to the user is targeted at the DICOM GSDF, and all window values or lookup tables are required to be

⁶ This approach is referred to as “perceptual linearization.”

⁷ Formally referred to as “just-noticeable differences.”

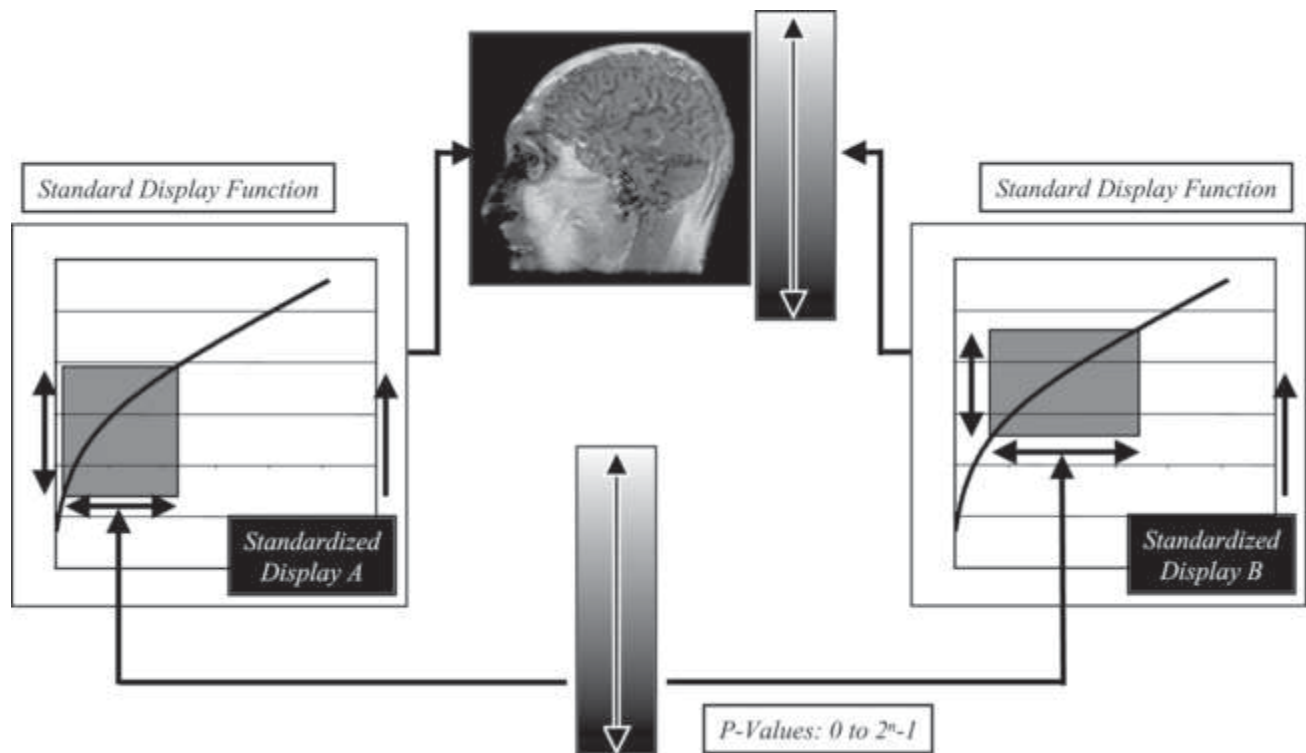


Figure 4. Use of the DICOM GSDF. Different display systems (shown here as A and B) are standardized such that the full range of possible input values (horizontal axis) is mapped to the full range of possible luminance values of which the display is capable (vertical axis), according to a standard function. The net effect is that a human observer perceives approximately the same gray-scale contrast on either device. Image pixel values that are used to drive standardized displays and span the full input range of the device are referred to in DICOM as “P values.”

applied.⁸ When a manufacturer tunes its image-processing algorithms, the parameters can be chosen to create a result appropriate for display with a GSDF-standardized device. When an image created by one vendor is displayed on another vendor’s workstation, because both are required to use the GSDF, the image will appear as expected. Most important, any linear window values or lookup tables chosen by the operator or radiologist will result in a consistent transformation of the image and a consistent appearance, regardless of the output device.

The GSDF is no panacea, and other factors, such as the luminance range, frequency characteristics, and spatial resolution of the output device, will affect the appearance. However, the gross variations in output appearance that have often been a feature of multi-vendor CR implementations can be greatly alleviated.

GSDF IMPLEMENTATION

Implementing the GSDF in DR images requires the modality manufacturers to allow the operator or

device to choose contrast transformations (window or lookup table) targeted at a GSDF-standardized and GSDF-calibrated display, rather than a specific film, camera, or monitor choice. Some vendors and users may balk at the initial and ongoing expense of a standardized and calibrated monitor on the acquisition device, but operators need this if they are to manipulate the contrast of an image before sending it to the PACS.⁹

A workstation implementation of the DX objects requires that window values and lookup tables be supported and that the display be both standardized to the GSDF and calibrated. Indeed, a regular quality control process needs to be in place to ensure that once displays are calibrated, they remain calibrated.

What are the consequences when a display device ignores the fact that DX images are intended for display according to the GSDF? The situation will certainly be no worse than it was before the GSDF, in that the results will be equally arbitrary and unpredictable and will likely not make the best use of all of the just-noticeable differences possible. As it happens, images intended for display according to the GSDF do not look particularly bad on a typical cathode-ray tube display that has not been standardized, calibrated, or gamma corrected.

⁸ More precisely, the transformations must be made available to the user for application to the image because there may be more than one. Regardless, they cannot just be ignored.

⁹ Also important is appropriate location of the operator’s display, with subdued and consistent ambient lighting.

This leads to *purchasing guideline 5: Insist on GSDF standardization and calibration and full DX image support in both modality and PACS workstations*. Distributed consistency of image appearance is impossible unless both ends are calibrated to similar expectations, and the DICOM DX and GSDF combination is the standard way to achieve this goal.

FOR PROCESSING OR PRESENTATION

There are two types of DX objects in DICOM, those “for presentation” and those “for processing.” The latter type is present to allow the storage of a “raw” form of image in a PACS archive that may be subsequently retrieved for reprocessing. It was recognized that this approach would create a risk that some vendors might use one form and some the other, leading to a lack of interoperability. Accordingly, both sending and receiving devices are required to support the for-presentation type and may optionally support the for-processing type.

The rationale behind this choice was that all systems needed to be able to produce images that were usable everywhere, without special or vendor-specific image-processing capabilities being required on every workstation. Needless to say, the choice was controversial. It is vitally important for all users to guard against vendors who violate the standard and (a) send or display only the for-processing type of image, (b) claim to send images of the for-presentation type that are not ready for presentation (ie, do not contain appropriate window or lookup table information or are not targeted at the GSDF), or (c) claim to display images of the for-presentation type but ignore window values or lookup tables or do not display images according to the GSDF. This leads to *purchasing guideline 6: Insist on for-presentation support in both modalities and workstations, as explicitly required by the standard*. Interoperability goals could potentially be thwarted by vendors who fail to send or properly display for-presentation DX images.

ADOPTION OF DX

After the virtues of the DX objects have been so loudly proclaimed, it is only fair to say that the status of adoption in products, at least for general radiography, is somewhat disappointing. The mammography image object is an exception because almost all full-field mammographic modality, workstation, and computer-assisted detection device vendors use it.

Although vendors are not terribly forthcoming on these matters, a review of DICOM conformance statements available toward the end of 2002 revealed that three DR modalities supported the DX object. Of 13 PACS reviewed, nine supported the DX object, and four did not. In particular, some notable PACS vendors seemed content to ignore DX. However, although

a PACS is described as able to receive and store DX images, the quality and completeness of its workstations’ DX implementations are difficult to assess from a typical conformance statement. This level of detail is not normally addressed in such statements, particularly with respect to which attributes are required to drive hanging protocols, whether there is support for the GSDF and calibrated displays, and whether or not lookup tables are supported properly.

What are some possible reasons for the delayed adoption? As might be expected, to some extent this is a “chicken or the egg” problem. From the perspective of a modality vendor, there is a risk that the PACS of the user may not accept DX images. This risk is easily mitigated by implementing a fallback to encoding images as CR images if negotiation with the PACS indicates that DX cannot be used. However, this process may require changes in how the pixel data or lookup tables are encoded and hence may not be as trivial as it sounds.¹⁰ Furthermore, in the past, those who specify modality requirements have shown little awareness of the benefits of supporting features that enhance PACS work flow and productivity.

From the perspective of a PACS vendor, the vendor may see too few DX-capable modalities in the field to justify the expense of adding the support. This is particularly true of a PACS vendor that is not also a DR vendor or may be content with its CR implementation. In designing hanging protocol support in workstations, it may be risky to depend on the presence of DX attributes and may be tempting to depend only on the “lowest common denominator” of what might be present in CR images.

Furthermore, many users may not be terribly demanding, with respect either to hanging protocols or to distributed consistency of image appearance. Users will often tolerate extensive site-specific tweaking and work-arounds until things more or less work, experiencing pain not when the system is accepted but rather later, when new equipment is added.

Finally, it may be that the design assumptions in the DX object are incorrect. For example, there may be a flaw in the premise that a little inconvenience for radiographers and technologists is acceptable in return for downstream productivity gains for radiologists and referring physicians.

FUTURE STRATEGIES

One primary task is education. Users must be educated as to what is possible, so that they know what

¹⁰ Specifically, it makes sense to separate out nonlinear gray-scale transformations into a lookup table for a DX object because lookup table support is required, and the user can subsequently adjust these transformations. However, because lookup table support is patchy at best for CR objects, it may be safest to burn the transformation into the pixel data for CR fallback objects.

to ask for. Vendors must be educated about what users need to improve their productivity and quality of operation.

The IHE effort, which to date has focused more on work-flow and management issues, could and perhaps should address the “payload” of transactions, specifically which image objects are appropriate for specific applications. Furthermore, in the period since the DX objects were added to the standard several years ago, some weaknesses have been identified, and the DICOM standard may need to be corrected accordingly.

A potential contributor to progress is a new service being developed for encoding hanging protocols in DICOM. This proposed service is intended to allow hanging protocols to be archived and interchanged between different vendors’ workstations. Research into how to encode hanging protocols may lead to a better understanding of what information in the image objects is a prerequisite.

SUMMARY OF GUIDELINES

1. Do not buy a DR or mammography system or PACS without MWL.
2. Insist on DX support in acquisition devices for DR and CR, as well as PACS workstations.
3. Insist on hanging protocols driven by DX attributes in PACS reading workstations.
4. Choose a DR modality that fully populates attributes but has minimal effect on the productivity of the operator.
5. Insist on GSDF standardization and calibration and full DX image support in both modality and PACS workstations.
6. Insist on for-presentation support in both modalities and workstations, as explicitly required by the standard.

Reference

1. Integrating the Healthcare Enterprise: IHE technical framework. Radiological Society of North America Web site. Available at: www.rsna.org/IHE/tf/ihe_tf_index.shtml. Accessed September 19, 2003.