Designing and Implementing a PACS-Aware DICOM Image Object for Digital X-ray, Mammography and Intra-oral Applications

David A. Clunie
Quintiles Intelligent Imaging
521 Plymouth Road, Suite 115, Plymouth Meeting PA 19462
mailto:dclunie@idt.net

ABSTRACT

The introduction of new digital detector technology for projection radiography created an opportunity to revisit the support for X-ray images in the DICOM standard. A new family of Digital X-ray (DX) objects has been developed with greater emphasis on the productivity and workflow requirements of Picture Archiving and Communications Systems (PACS) and soft copy reading on workstations. The use of these DX objects presents new design challenges for acquisition and display systems.

1. INTRODUCTION

The image objects and storage classes in the existing DICOM standard have served well to meet the needs of cross-sectional modalities as well as various forms of projection radiography including optically scanned conventional film-screen radiographs, Computed Radiography (CR), radio-fluoroscopy (RF) and X-ray angiography (XA). All of these objects have been widely used in various Picture Archiving and Communications Systems (PACS) and considerable experience has been gained. The limitations of the information stored in existing DICOM objects with respect to improving the productivity of PACS users are well understood.

The recent introduction of new acquisition technologies incorporating solid state flat panels and other types of sensors provided an opportunity to revisit the design of DICOM objects for projection radiography. DICOM Working Groups on Digital X-ray and Digital Mammography reviewed the ability of the existing DICOM objects to encode images from new detectors. For the clinical applications considered ¾ general and intra-oral radiography, diagnostic and screening mammography ¾ it was determined that the existing objects lacked features necessary to describe the acquisition. Thus it was decided to design a new family of DICOM Digital X-ray (DX) objects, defined in Supplement 32[1], which became part of the standard in September 1998.

This paper will review:

- the experience with existing DICOM image objects,
- the reasons for the development of a new family of DICOM objects,
- the features of the Digital X-ray (DX) family of image objects, including those for
  - general radiography,
  - digital mammography, and
  - intra-oral radiography,
- the implementation of DX objects from the perspective of the designer of
  - storage class providers (i.e. modalities or acquisition devices), and
  - storage class users (i.e. gateways, archives and workstations)

2. EXISTING DICOM OBJECTS

The existing DICOM objects for single frame projection radiography are the Computed Radiography (CR) object used for images from storage phosphor plates and the Secondary Capture (SC) object generally used for images obtained by optically scanning conventional X-ray films. The X-ray Angiography (XA) and X-ray Radio-flouroscopy (XRF) objects are used for multi-frame images (usually cine), but these will not be considered here. For cross-sectional images the Computed Tomography (CT), Magnetic Resonance (MR), Ultrasound (US), Nuclear Medicine (NM) and other objects are used. These also will not be considered here, except when relevant experience using particular features of these objects in PACS is applicable.

The existing DICOM CR object was considered as a candidate for storing digital X-ray images from the new detectors. However, the CR object doesn't describe specific characteristics of the new detectors well, and contains attributes that are not applicable to the new detectors. Attributes describing the detector characteristics are important for both quality control and sometimes interpretation of the image.

The definition of a CR Series prevents useful grouping of images by series. For example, a left and right image of the same body part cannot be contained in the same series, and a PA and lateral chest x-ray cannot be contained in the same series. This limitation essentially restricts CR images to have only one image per series, eliminating the series level from the DICOM Information Model.

The CR object also permits the common film-screen practice of allowing multiple exposures on the same plate (for example, by using a lead sheet to mask unexposed areas). Having two images of different body parts or orientations within the same image makes it difficult to automate organization of images in a meaningful manner on a display or within a PACS.

Furthermore, the anatomy, view, laterality and orientation of an image are poorly described by the CR object. Though defined terms for these attributes are present in the standard, they are optional, frequently not implemented, or vendors use their own defined terms or allow users to enter free text. The workstation or PACS cannot reliably use these attributes to make routing or display decisions. Laterality in particular, despite the definition in the standard that implies it should be sent unless unknown, is rarely, if ever, implemented. Thus the workstation does not know whether an image of a hand is the left hand or the right.

The grayscale of a CR image is not precisely defined, and in particular whether there is a positive or negative linear or non-linear (logarithmic) relationship of the stored pixel values to x-ray intensity is not defined. This makes it impossible to know whether or not a low pixel value is air or bone, even though it is possible to specify whether a low pixel value is displayed as black or white (Photometric Interpretation). It also makes it impossible for a display station or printer to choose the appropriate contrast curve to map the image to luminance or optical density in a consistent manner. Though a device or user may choose a recommended window center and width or non-linear lookup table (LUT) to be applied to the pixel data, there is no requirement that this contrast transformation must be applied.

Related to the grayscale issues is the question of whether the image pixel data is "processed" or "unprocessed". Some implementers store only images (or images and a LUT) that are intended for immediate printing or display, but others send data in a "raw" form that requires choice and application of an appropriate LUT or more extensive processing (such as spatially adaptive contrast transformation or unsharp masking). This inconsistency makes it very difficult for systems to interoperate, and often leads to the need for arbitrary system wide decisions that restrict the choice of acquisition devices and workstations that can be incorporated satisfactorily.

The Secondary Capture (SC) object has many of the same problems as the CR object problems, but in addition is unconstrained in terms of whether it contains a grayscale or color object, defines no modality specific technique attributes, and contains no anatomic, projection, laterality or orientation information. Essentially the SC object circumvents the conformance mechanism inherent in modality specific DICOM objects, and requires a proprietary understanding of mutual capabilities for two devices to interoperate satisfactorily.

3. DIGITAL DETECTOR REQUIREMENTS

The established techniques for digital projection radiography include the use of storage phosphor plates (Computed Radiography), optical scanning of conventional film/screen images, as well as the use of selenium drums (Philips Thoravisionä), and the already widespread application of CCDs for imaging of small areas (such as dental intra-oral radiography and digital mammographic biopsy systems). The new technologies being introduced for general radiography and full field diagnostic mammography applications include large solid state flat panels (with or without a scintillator), as well as various forms of large area CCD systems, and spot, slot or slit scanning systems.

Though the specific characteristics of each technology varies, many characteristics are shared or can be described similarly in an image object. The previous DICOM objects do not contain sufficient attributes to describe many of the features of the new detectors and acquisition protocols. Optional attributes could have been added to existing objects, or the meanings of existing attributes broadened, or many of the new descriptions left to private attributes, but the users and vendors participating in the

working groups felt that it was both possible and desirable to define a new common set of descriptive attributes. Several factors contributed to this decision, but in particular it was agreed that:

- quality control (QC) would be a high priority in order to achieve and maintain image quality, and that common standard descriptive attributes would facilitate the use of distributed QC data gathering tools, and
- capture of technique and dose related information in standard attributes would facilitate compliance with increasingly strict regulatory requirements.

The possible characteristics of image data from the different technologies were also reviewed, and it was determined that a well defined model defining and describing stored pixel data, relationship to X-ray intensity, contrast transformations and image processing would be required to achieve interoperability. In particular it was felt necessary to distinguish between data that has been "detector corrected" (the first relatively technology neutral stage in processing) but is not yet ready for presentation (display or printing), since appropriate filtering or contrast transformation has not yet been applied, and data that is ready "for presentation". Images in both forms are useful, since later reprocessing may be necessary or advanced applications such as Computer Assisted Diagnosis (CADx) may prefer one form or another. More precise description of the form of the pixel data was also felt important in order to drive processing and presentation choices. For example images linear to X-ray intensity need to be transformed into a logarithmic space before subtraction from each other to prevent shading gradients, and users generally don't like the appearance of untransformed linear images since they don't "look like film". Whether an implementer chooses to store data in a linear space (as may be more convenient for some detector types) or not, if the image is declared to be 'for presentation", then the object must contain an appropriate transformation (non-linear LUT) into something that is suitable for the radiologist to interpret.

4. PACS REQUIREMENTS

Supporting new detector requirements required an entirely new family of objects. The creation of these objects presented an opportunity to incorporate support for the requirements of PACS and diagnostic workstation users.

Experience using existing DICOM objects in a PACS had highlighted a number of issues. The DICOM services (such as Storage, Query/Retrieve, Print, Modality Worklist, and Storage Commitment) are largely adequate to support the basic functions of a PACS. However, at the application level there are clear limitations, especially related to the productivity of softcopy reporting. These include the automated:

- routing of images (either to a reading worklist or particular reporting station),
- identification of image or study type,
- grouping of images for display (especially with the corresponding previous images),
- layout of images (so-called "hanging" or "default display" protocols), and
- determination of the grayscale appearance of images.

It is important to specifically address the requirements of PACS, since modality and PACS vendors and groups traditionally have separate management and independent goals. In this instance, since the cost effective deployment of digital detector technology may well depend on efficient image management and efficient soft copy reading, it is important to address these PACS requirements in the new DX objects. In this manner it is hoped that digital detector purchases can be justified not only by better image quality and new applications, but also by improved PACS usability and productivity.

Failure to meet PACS requirements would impair the productivity of the users, particularly radiologists who are an expensive resource, often involved in purchasing decisions, and often very vocal about the inadequacies or inconveniences of a system. Radiologists can’t read images without the request, a request without the images, new images without the old images, and images that are not yet on their reading worklist or station. Radiologists refuse to read, or read more slowly, images that are in the wrong order or upside down, and images displayed with the wrong grayscale contrast.

5. FILM REPLACEMENT

Ideally, the new objects would provide a complete replacement for the traditional functions of radiographic film, including electronic "lead markers", manual routing, matching and hanging, and consistency of appearance. Additional human operators to interpret and make use of "visual cues" burned in the image thus become redundant.

These visual cues typically include

- a "flashed" identification panel containing the

    - patient's name,
    - patient's ID number,
    - patient's date of birth,
    - patient's sex
    - referring physician,
    - institution
- lead markers describing
    - laterality (e.g. left or right body part)
    - projection (e.g. left or right lateral or AP or PA)
    - technique modifiers (e.g. inspiration, post micturition)
- characteristics obvious to a trained observer
    - anatomy (e.g. skull, chest)
    - orientation (e.g. left side, feet)
    - projection (e.g. frontal or lateral)
    - modality (e.g. a projection radiography)
    - grid used (visible grid lines)
- quality control related features
    - visible collimator edges
    - correct exposure
- manual (wax pencil or sticky label) annotations
    - highlighting suspicious areas
    - noting artifacts
    - correcting incorrect burned in annotations (wrong side or name)

There are many places in the workflow where some or all of these visual cues are used. For example, the process of hanging a traditional set of films involves:

- extracting the films from patient's folder
- sorting them into old and new films
- verifying patient name and ID on each film
- arranging the films into desired hanging order given the available view box area
- matching old films with new films for same anatomy and view according to personal preference
- turning and flipping the films to the preferred orientation (e.g. left on right of view box, feet on bottom)
- turning on the view box, and optionally the use of a "bright light" to view dense areas

The process of displaying a set of images on a softcopy workstation is similar:

- select the appropriate images from a worklist or pre-fetch them
- match the modality and anatomy with the display protocol
- per the display protocol
    - arrange the old and new images
    - arrange the images by anatomy, laterality and view
    - rotate and/or flip the images based on preferred orientation
    - annotate the images according to preference
    - select the appropriate grayscale contrast choice

For a workstation to perform these tasks successfully it needs information that is accessible as attributes of the object, not as burned in visual cues. In principle, if the information in the objects is reliable, then many of the manual verification steps inherent in hanging films may be elided. In practice, until further experience is gained with film-less operation and the system is "trusted", users will likely make use of residual burned in information for verification. Burned in markers do interfere with automated image processing and grayscale transformation operations, however.

Accordingly, it was determined that the DX objects required attributes:

- that describe the projection and orientation of the image relative to the patient, allowing images to be "hung" automatically using display protocols, without the need for reorganization by the user;
- that contain standardized codes for technique, anatomy and projection to allow automated routing of images to reading worklists, and matching of what was performed with what was requested as well as matching with corresponding previous images;

- that define the relationship of pixels to X-ray intensity and define the "output" of the image object in terms of a device independent, perceptually linearized standard display function, the DICOM Grayscale Standard Display Function (GSDF).

Specifically, the following attributes are included in the DX objects, and in most cases they are both mandatory and use standardized codes:
- Image Laterality, which is mandatory and may be left, right, unpaired or both, and is instantiated at the image level (not the series level as in the CR object),
- Patient Orientation, which is mandatory and specifies in what direction (left/right or anterior/posterior or head/feet) the rows and columns of the image are orientated with respect to the patient,
- Anatomic Region, which is selected from a standard list of codes that is enumerated in the standard,
- View Code, which is also selected from a standard list of codes that is enumerated in the standard, and
- SOP Class UID, which defines what type of image is contained and further restricts the list from which codes such as anatomy and view may be drawn, and adds further mandatory requirements.

Other coded attributes further specify or qualify anatomy, view and radiographic technique (e.g. a code for a Towne's view of the skull).

These attributes should be sufficient to correctly match images with those of previous studies, as well as orient (hang) the image correctly for display. For example, a PA chest image is often obtained and stored in the image flipped from left to right relative to the way in which it is normally viewed, since the patient is facing away from the tube but the image is usually viewed as if it were an AP image. In this case, the row direction of the Patient Orientation will be towards the right rather than towards the left, and a hanging protocol should know to flip the image for display. Alternatively, if the acquisition system flips the image, then the row direction of the Patient Orientation will already be towards the left, and no flipping on the workstation is necessary. Either case can be transparent to the user.

Automating the correct selection of grayscale contrast is supported by:
- requiring images that are intended "for presentation" to contain a full definition of the appropriate transformation (i.e. the mandatory presence of a window center and width, non-linear LUT or statement that no transformation is necessary),
- requiring that the relationship to X-ray intensity be specified (so that display of air or bone as black or white can be chosen as preferred), and
- specifying the "output" of the object in the P-Values of the GSDF to achieve consistent appearance on different devices.

6. IMPLEMENTATION

Implementation of the DX objects presents challenges to both image producers (storage class users) and consumers (storage class providers).

DX image producers include both integrated and add-on digital detector acquisition systems, as well as CR and scanned film devices that may be upgraded to generate DX objects in order to support PACS productivity enhancement. The key challenge is to find a way to acquire the additional information such as anatomy, laterality and orientation that needs to be available before a DX image can be created.

This information must be entered manually in add-on systems, for example at the quality control stage, where operator inconvenience may be traded for improved PACS and reader efficiency. Even in an add-on system, if the Modality Worklist service is well supported, some of the required information can be selected from the worklist and matched, rather than manually entered. For example, if a worklist request for a PA and Lateral Chest X-ray is received then as part of the operator's QC step. the exposures can be matched not only with the correct patient and request, but also with the correct projection. Successful widespread deployment of such a strategy will probably depend on a standard set of procedure codes from the information system, something that is unfortunately not yet part of any standard, and largely remains site specific.

In an integrated system, effective coupling of generator and gantry controls can obviate the need for manual entry by providing information that is:
- already entered for other reasons, such as specifying radiographic technique protocol (e.g. "PA Chest"), and
- known from the geometry of the physical components (e.g. use of an upright bucky).

This potentially reduces error and avoids slowing down or annoying the operator by having them enter the same information multiple times. The prevalence of protocol based radiographic technique selection (as opposed to the old fashioned method of memorized techniques and guesswork, or tables taped to the wall beside the generator console) makes this approach feasible. Even if generator protocol data is too "coarse" (e.g. the same selection for similar exposures for lateral and oblique chests) it is preferable to refine the generator protocol interface to be more specific (e.g. separate the selections for lateral and oblique chests even if the same technique is delivered), than to force the operator to enter the selection again or to not send it at all. This approach can be refined further to add laterality (e.g. left hand) and projection (e.g. left lateral) to the protocol selection, saving another step.

When integrated with known characteristics of the physical setup (such as the use of an upright wall mounted bucky), this approach also makes it possible to automatically determine and describe (and if desired correct) the orientation of the rows and columns of the stored image. For example, the system can automatically determine the orientation of a PA image if the operator selects PA Chest rather than AP Chest from the generator protocol and the upright bucky is used. Similarly if a digital mammographic system "knows" which breast it is examining, then the column angulation allows it to automatically determine whether the view is cranio-caudal, caudo-cranial (from below), a medio-lateral oblique or a latero-medial oblique, as well as to detect from the separation between the sensor and the breast support the use of technique modifiers such as magnification.

Ideally the integration will extend to the transfer of descriptive attributes useful for interpretation, quality control and dose monitoring through connections with not only the generator and the gantry but also the collimator, the filtration on the tube, the grid and the detector controller.

Additional motivation for this level of integration, beyond support for the PACS, arises from the benefit to image processing of the raw detector image data if it is selected or parameterized according to technique, body part and orientation. This information also allows for more appropriate selection of appropriate grayscale contrast transformations for presentation.

The acquisition device implementer also needs to take into account the grayscale transformations defined in the DX objects, and provide stored image pixel data in the appropriate form, with the appropriate descriptive attributes (such as relationship to X-ray intensity) as well as a linear, non-linear or identity VOI transformation, with the understanding that the output will be interpreted as P-Values of the GSDF. Traditionally, CR and scanned film devices have applied proprietary transformations or lookup tables without specifying the intended display space, and usually the appearance of images on different stations and printers varies considerably. For any given configuration of display or printer, transformation of stored values into luminance or optical density, and viewing conditions, once a "desirable" appearance has been obtained empirically, the transformation used can be reformulated in terms of P-Values to achieve a device independent output.

Support of the "for presentation" DX SOP Classes is mandatory for both SCU and SCP, but either SCU or SCP may also choose to support the optional "for processing" SOP classes for additional functionality. The choice of which form to use for storing the pixel data (e.g. pixel values linear to x-ray intensity with a VOI LUT that transforms pixels to a logarithmic space defined in P-Values, or pixel values in a logarithmic space with a VOI LUT that transforms pixels to P-Values, or pixel values already in P-values with an identity VOI LUT, or some other permutation) is affected by many factors. For example, if the "value of interest" operation is already "burned in", then the display user will not be able to adjust the contrast of the image in a useful manner. If a display SCP fails (in violation of the standard) to apply a VOI LUT that is present, then the quality of the displayed image will be low, through no fault of the SCU. For many implementers, the best choice will be to encode the pixel data in a logarithmic space, with a VOI LUT used to enhance body part specific characteristics of the image, such that a display SCP user may select an alternative VOI LUT to view other features.

The adaptation to DX of the design of image consumers such as PACS and workstations is less challenging, since many of the features in DX are already supported in proprietary or ad hoc ways. For example, the use of default display or hanging protocols is already widespread, but often depends on specific coded operator entries in free text fields (e.g. Study Description) or information that is known to be stored in a vendor's private attributes. Adding support for the standard coded attributes in the DX objects is relatively straightforward.

Probably the greatest challenge to the workstation designer is to effectively support the grayscale and presentation related features of the DX objects. These objects are the first to require the use of the Grayscale Standard Display Function (GSDF) and the first to specify their output in P-Values. Though an un-calibrated display station may choose to ignore this, and display the image using its default characteristics, the images will not have a consistent appearance with what was intended,

with what is printed, nor with what will be seen on other, calibrated, displays. Since achievement of high quality soft copy image presentation is generally considered to require monitor calibration and quality control on a regular basis anyway, it is likely that most of the support for the features of the GSDF are already present.

In addition to defining a grayscale output space, the DX objects also require an image consumer to support the non-linear value of interest lookup tables (VOI LUT), since these are defined as a "for presentation" requirement. A DX image storage SCP that fails to support the application of a VOI LUT that is present in the object when displaying a "for presentation" image is in violation of the standard. This requirement is in place since selection of the appropriate grayscale is crucial to the correct interpretation of an image, and the untransformed image stored in the pixel data attribute may be totally unsuited for consumption by a human observer. Many existing workstations ignore these VOI LUTs when they are present as options in the existing CR objects, and so incorporating support for the DX object may be more difficult.

Even though the intent of the "for presentation" object is to be immediately useful for display on a calibrated system, workstation implementers may add value by providing mechanisms for substituting alternative VOI operations. These include providing alternative VOI LUTs to enhance different features of the image (e.g. "bone" vs. "lung" vs. "soft tissue" on a chest X-ray), interpolating and parameterizing supplied VOI LUTs in the object to allow the user to adjust control points (and simulate the effect of window center/width operations), and support for the optional "for processing" SOP Classes to receive data than can be effectively "reprocessed" with spatially or frequency selective algorithms. Normally, any such processing necessary (depending on the detector type) would already have been performed by the acquisition device.

Another challenge for workstations displaying DX images is mandatory support for bit mapped overlays. Since many applications of projection radiography involve annotation of images (analogous to wax pencil marks on film) either by a user or an algorithm (such as CADx), support for selected types of bitmapped overlays is mandatory. A DX image storage SCP that fails to support the application of a bitmapped overlay containing annotation that is present in the object when displaying a "for presentation" image is in violation of the standard. Failure to comply with this requirement many not bring an immediate response from users, but may potentially increase the implementer's legal liability in the event that annotation crucial to a medical decision was not displayed correctly or at all. The burden on the implementer is somewhat reduced by requiring only a subset of overlay types, those stored in a separate attribute from the pixel data, but given the relatively poor support for overlays in existing workstations, this feature remains a challenge.

7. CONCLUSIONS

The deployment of new digital detector technologies for projection radiography, and the requirement for new DICOM image objects to support them, provided an opportunity to revisit the support for PACS productivity enhancement in DICOM objects. Use of the new Digital X-ray (DX) object family is intended to enhance PACS workflow, particularly with respect to support for display hanging protocols, image routing and reading worklists, and more consistent grayscale appearance. In many cases, the new objects can be created by acquisition implementations that reuse available data rather than requiring additional operator entry.

8. ACKNOWLEDGEMENTS

I would like to thank my many colleagues from the vendor and user community who participated in the DICOM digital X-ray, digital mammography and base standard working groups and worked hard to define the DX objects in record time. In addition, I gratefully acknowledge my former co-workers at GE Medical Systems in the PACS, workstation, platform, digital X-ray and digital mammography groups, whose enthusiasm and imagination in creating a new family of products helped bring this standard to life.

9. REFERENCES

1. DICOM, "Digital Imaging and Communications in Medicine (DICOM), Supplement 32: Digital X-Ray Supplement. NEMA, 1998.