

Five Levels of PACS Modularity: Integrating 3D and Other Advanced Visualization Tools

Kenneth C. Wang · Ross W. Filice · James F. Philbin ·
Eliot L. Siegel · Paul G. Nagy

Published online: 8 February 2011
© Society for Imaging Informatics in Medicine 2011

Abstract The current array of PACS products and 3D visualization tools presents a wide range of options for applying advanced visualization methods in clinical radiology. The emergence of server-based rendering techniques creates new opportunities for raising the level of clinical image review. However, best-of-breed implementations of core PACS technology, volumetric image navigation, and application-specific 3D packages will, in general, be supplied by different vendors. Integration issues should be carefully considered before deploying such systems. This work presents a classification scheme describing five tiers of PACS modularity and integration with advanced visualization tools, with the goals of characterizing current options for such integration, providing an approach for evaluating such systems, and discussing possible future architectures. These five levels of increasing PACS modularity begin with what was until recently the dominant model for integrating advanced visualization into the clinical radiologist's workflow, consisting of a dedicated stand-alone post-processing workstation in the reading

room. Introduction of context-sharing, thin clients using server-based rendering, archive integration, and user-level application hosting at successive levels of the hierarchy lead to a modularized imaging architecture, which promotes user interface integration, resource efficiency, system performance, supportability, and flexibility. These technical factors and system metrics are discussed in the context of the proposed five-level classification scheme.

Keywords PACS · 3D imaging (imaging, three-dimensional) · Computer systems · Advanced visualization · Server-based rendering · Application hosting

Background

Advanced Visualization

Three-dimensional imaging in radiology has been the subject of research and technology development for many years, spanning a period which approximately mirrors the emergence of digital image management through DICOM and PACS. The terms “3D” and “advanced visualization” have been used interchangeably in radiology to refer to a range of specific graphics processing techniques, including maximum intensity projection (MIP), minimum intensity projection (MinIP), multiplanar reconstruction (MPR), curved MPR, isosurface rendering (i.e., shaded surface display) and volume rendering. Key clinical areas for such techniques include vascular applications (e.g., CT and MR angiography of the head and neck [1, 2], of the trunk and extremities [3, 4], as well as of the coronary arteries [5, 6]), and CT colonography [7, 8]. Three-dimensional evaluation through volumetric post-processing has also been described in a wide range of other clinical settings, including

K. C. Wang (✉) · E. L. Siegel
Department of Radiology, Rm. 1B-112,
Baltimore VA Medical Center,
10 North Greene Street,
Baltimore, MD 21201, USA
e-mail: kenwang@stanfordalumni.org

R. W. Filice
Center for Drug Evaluation and Research,
Food and Drug Administration,
Silver Spring, MD, USA

J. F. Philbin · P. G. Nagy
School of Medicine, Russell H. Morgan Department of Radiology
and Radiological Science, Johns Hopkins University,
Baltimore, MD, USA

traumatic injury [9, 10], tumor mapping [11, 12], bronchial visualization [13, 14], and urinary tract imaging [15]. Advancements in thin-slice multidetector CT, and the advent of high-resolution isotropic MRI, have only increased the opportunities and need for volumetric image post-processing in radiology.

However, while PACS has become widespread over the past 20 years, adoption of 3D imaging tools has been much more limited. Although PACS viewing applications have begun to incorporate certain functions which were previously considered advanced techniques such as MIP and MPR, volumetric imaging (especially with regard to volume rendering) remains by and large outside the typical clinical workflow. Technological factors have been a primary contributor to this situation. Three-dimensional post-processing methods, especially volume rendering, have historically depended on specialized hardware at the user's desk. This typically meant a dedicated graphics workstation in the radiologist's reading area, designed solely for the purpose of 3D post-processing and display, separate from the primary PACS viewing system. These stand-alone systems were associated with idiosyncratic user interfaces, specialized image routing mechanisms, additional equipment costs, and additional support costs.

Server-based Rendering and Client Thinness

As image resolution and data size scaled upward, a rendering architecture known as remote, or server-based rendering emerged [16, 17], whereby sophisticated rendering computation is carried out on a server instead of on a user's local machine, and rendered views instead of source data are transmitted to the client. This approach centralizes the specialized rendering hardware resources required for 3D post-processing, thereby making these techniques more widely accessible. At the same time, network bandwidth requirements are potentially reduced, since rendered views will generally tend to be smaller than the models or images used to generate them. Reduction in required bandwidth is replaced, however, by requirements for sustained low network latency, since real-time interactive visualization depends upon on-going transmission of display data. Application of server-based rendering (SBR) has been described in other graphics domains, such as molecular modeling [18] and gaming, as well as in radiology [19]. Over the past few years, radiology vendors have begun to offer commercial products using SBR for clinical use.

Image Archive Architecture: Storage and Communications

While a detailed examination of PACS archive architectures is beyond the scope of this work (see instead, for example,

[20–22]), selected aspects of this topic are quite relevant to the integration of advanced visualization tools with primary PACS viewers. In particular, communications protocols and storage resource allocation are two key factors considered here. While modern PACS archives use DICOM to receive images from modalities, the mechanisms by which an archive communicates with front-end viewing applications are in general non-standard and proprietary in nature. Deviation from DICOM at this level likely provides important performance benefits. However, this situation implies dedicated links between archives and applications. In certain circumstances where there are multiple image-display applications (e.g., a primary PACS viewer and a separate advanced visualization viewer), this typically leads to duplication of storage requirements.

Methods

The current market for PACS and 3D products consists of many disparate systems, patchy options for interoperability, non-uniform user interfaces, proprietary communications protocols, and workflow challenges. We propose that modularity in PACS technology may serve to address these issues while also promoting innovation. Conversely, lack of modularity may slow technology adoption. For example, while one PACS vendor may provide exceptional worklist management, support for computer-aided detection may be relatively rudimentary. Or while a particular 3D vendor may offer a compelling CT colonography package, its patient database may not be competitive. To the extent that vendors are able to focus on specific core technologies, and users are able to combine these as needed to suit a particular practice setting, this should promote technology adoption and innovation. It is not our intent to describe or evaluate any specific vendor or product. Instead, we present an analysis of system modularity based on architectural features. This analysis is based on three key questions and several system factors, culminating in five levels of modularity.

First, three key questions. (1) *Where does the image data reside?* Images in a PACS generally reside on a long-term archive, which may include multiple levels of cached storage, possibly utilizing several types of storage media. Applications may additionally cache data on client machines. Advanced visualization systems may utilize dedicated image storage mechanisms outside of the long-term archive. (2) *Where is image computation performed?* Requirements for image computation arise from any operation where the displayed image deviates from the original image received from the modality. In the case of two-dimensional image display, window-level manipulation, and image filtering represent post-processing oper-

ations. MIP and MPR require projection and interpolation calculations. Prior to the introduction of SBR, image manipulations generally depended on computation performed on the client machine (although selected operations could be performed on the modality console, with the resultant images sent to the client via the image archive). With the advent of SBR, image computation is off-loaded from the client to a centralized server. Clients are often characterized by their thickness or thinness [23], with SBR leading to thinner clients. Conversely, a stand-alone 3D workstation utilizing its own specialized graphics hardware for volume rendering would be considered a thick client. (3) *How do processes communicate?* Communications between components of an imaging system may be characterized by the protocols used, the performance of those protocols, and whether or not they are standardized. As discussed above, while image transfer to an archive is typically DICOM-based, communications between an archive and viewing application are typically proprietary.

These three questions, regarding data storage, image computation, and communications protocols, reflect system architecture. Any such system may then be evaluated in terms of several factors: user-level integration, resource efficiency, performance, functionality, supportability, interoperability, and cost. User-level integration refers to the user's experience in utilizing the tools within a clinical environment. Resource efficiency in this context may relate to the degree of storage duplication, or access to a given graphics processing engine. For example, while a dedicated 3D workstation may only be used by a single user in a specific location, a system using SBR generally allows multiple users to access a graphics processing engine from anywhere on a given network thereby achieving greater resource efficiency. Performance of an integrated PACS/3D system may be measured with a variety of timing metrics, such as time to open study, time to access prior studies, scroll speed, and speed of other specific operations (e.g., display oblique MPR; rotate volume rendered view; bone removal). Functionality encompasses imaging capabilities (e.g., 2D display and navigation; image fusion; 3D rendering; clinical 3D packages such as coronary artery analysis), worklist management, as well as support for additional image-based operations such as computer-aided detection and image annotation. Supportability refers to the effort required to maintain systems, and is affected by system complexity and how resources are distributed within an enterprise. As clients become thinner and as computing power becomes more centralized, the complexity of client machines decreases, and the difficulty of supporting these clients also decreases [23]. Interoperability of component technologies depends in large part on the storage strategies and communications protocols employed, and determines the degree of user's choice in mixing and matching

components. Finally, initial system costs should be weighed against longer-term support costs.

Results

Based on these questions and factors, a classification scheme for PACS/3D integration architectures has been devised. This scheme consists of five levels of increasing modularity. While specific combinations of commercial products may or may not fall neatly into these categories, the scheme is intended to describe major degrees of modularity.

Level 1: Separate Turnkey Clients

Level 1 consists of a thick-client 3D application running on dedicated, specialized hardware, separate from the primary PACS viewer (see Fig. 1). Until recently, this was the predominant model for 3D post-processing in radiology. These 3D systems typically manage their own private, local cache of image data. This data may be retrieved from the primary archive, or received directly from a modality, via DICOM transfer. The 3D application typically requires the full image data set to be transferred across the network to the local cache. Access of non-cached studies, such as relevant priors, incurs additional network transfer time. Once the relevant images are received, graphics performance is generally very good as there is no further network dependence. The thick-client nature of the 3D system increases support costs [23]. In addition, since the 3D client

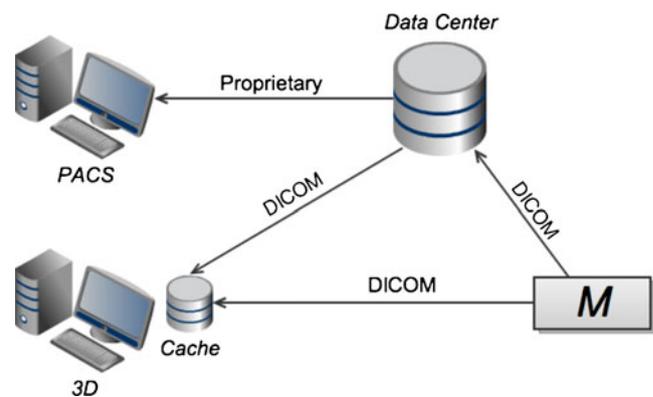


Fig. 1 At level 1 of the proposed classification scheme, the primary PACS viewing application and the advanced visualization/3D application are thick clients running on two separate workstations, with independent links to the imaging archive in the data center and/or direct connections to modalities (denoted *M*). Note that the 3D system in this case typically maintains its own local image archive, duplicating data, which also resides in the data center. Also note that image communication to the 3D workstation is typically achieved through DICOM transfer

runs on a dedicated machine, there is duplication of disk storage, computer memory, and display hardware. Lack of integration at the user-interface level is disruptive to workflow, and may discourage users from using the 3D application. However, in cases where the need for 3D imaging is limited, this type of arrangement may be advantageous due to its relative simplicity. In addition, independence of the PACS and 3D systems implies freedom of choice for users in selecting 3D vendors.

Level 2: Proprietary Integration Through Context-sharing

At this level, the PACS vendor provides the ability to launch advanced visualization tools, now running on the same physical machine as the primary viewer (see Fig. 2). This is typically achieved through dedicated, proprietary links for communicating imaging study metadata to specific third-party 3D applications. Image data, while still transferred in full to the client machine, is shared between the PACS viewer and the 3D application on the client machine. While this data sharing may be accomplished using relatively fast in-memory transfer, less efficient approaches, such as local network transfer within the client machine may also be utilized. In either case, there is no longer a separate, duplicated network transfer of image data from the data center to the client machine for the 3D application, as there was in level 1. Context sharing on a single physical machine eases the workflow barrier of level 1. That is, users no longer need to walk to a separate 3D system, but now have access to advanced visualization techniques on their primary workstation. Hardware redundancy of level 1 is reduced. However, specialized graphics hardware is still required at the user's desktop, the limitations of local

rendering remain (i.e., lack of widespread availability to users), and graphics performance may be adversely affected since local hardware resources are now shared between multiple client applications. In addition, users are limited to the proprietary integration options offered by a particular PACS vendor.

Level 3: Thin Clients with Standardized Context-sharing

At level 3, the 3D application and the primary PACS viewer consist of thin clients (see Fig. 3), reflecting the current trend toward client thinness. The 3D tool now uses SBR, and is integrated with the primary PACS viewer through a standardized context-sharing protocol. The Health Level 7 Clinical Context Object Workgroup (HL7-CCOW) standard is one such protocol [24]. SBR and thin PACS clients imply simpler client hardware, improving supportability. Standardized context-sharing between client applications facilitates interoperability and workflow integration, compared to levels 1 and 2. The thin clients remain largely separate applications, however, leading to non-uniform user interfaces. Also note that there is no server-side integration at this level of the classification scheme. The thin PACS client is served by the PACS archive, and the thin 3D application is served by the SBR server, each requiring its own copy of the image data in the data center, and each communicating with the user-level applications through proprietary protocols. Lack of server-side integration prevents unification of archive storage, and leads to duplication of disk storage on the server. Image data communication in the data center between the PACS archive and the 3D server may be achieved with DICOM transfers. The 3D server may also receive imaging data directly from a modality.

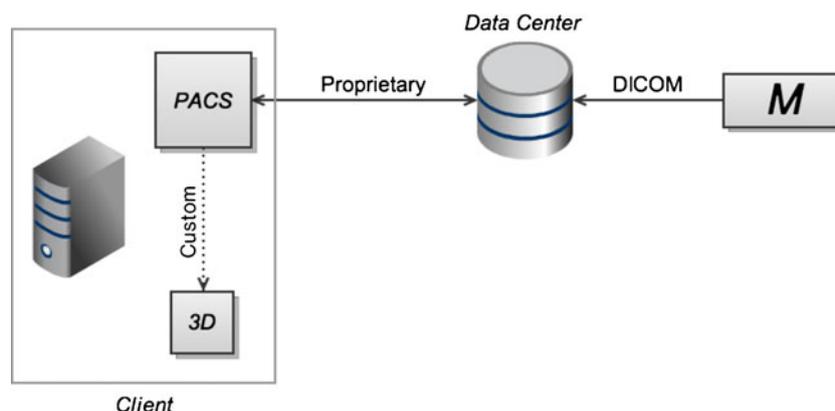


Fig. 2 At the second level of the scheme, both the primary PACS application and the 3D application consist of thick clients, running on a single client machine. In this scenario, a custom-built link (dotted line) between specific products (typically determined by vendor partnership agreements) allows the 3D application to be launched using context information from the PACS application. Users are

constrained by the integration options offered by specific vendors. Aside from context-based launching, the applications remain separate, with their own interfaces. The separate, duplicated network transfers of image data from the data center to the 3D application, which were required in level 1 (Fig. 1), are eliminated here

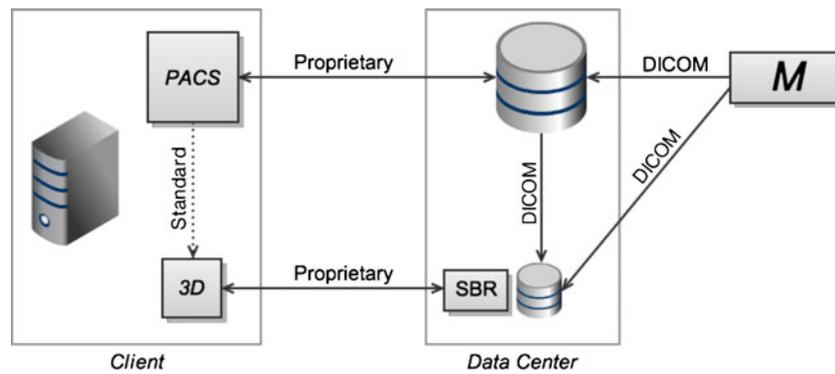


Fig. 3 Server-based rendering (SBR) is introduced at the third level, along with thin client applications. Thin clients have a smaller footprint on the client machine. SBR offloads graphics computation from the client machine to a server in the data center. At this level, the rendering server maintains its own private cache of image data, received from the primary archive or modality (*M*) via DICOM

transfer. While SBR provides several advantages including improved performance and supportability, duplicated storage and server-to-server DICOM transfers are inefficient. Context-sharing on the client machine is now achieved by a standardized protocol (*dotted line*), allowing for improved interoperability

Level 4: Thin Clients with Unified Archive

This leads to level 4, where multiple user-level applications are served by a single, unified archive (see Fig. 4). This implies a standardized archive communications protocol or application programming interface (API), capable of supporting the primary PACS viewer, the server-based rendering tool, and potentially, other applications such as computer-aided detection and computer-aided diagnosis packages. The duplication of server disk storage and server-to-server DICOM transfers required in level 3 is eliminated. Front-end (*i.e.*, user-level) integration continues to consist of loosely coupled clients with standards-based context sharing. Each client continues to provide its own user interface, and communicates separately with the server through the unified server API.

data center. Image computation is performed entirely on the imaging server. Communication between client and server becomes essentially a video stream of renderings computed on the server. The key feature of level 5 is a client-side API, which manages context data and communication with the server, allowing viewer applications to plug into an application framework (see Fig. 5), an arrangement also referred to as application hosting. This allows for the possibility of unified user interfaces. For example, such a system might provide 2D imaging in one viewport, 3D rendering in another viewport, and computer-aided detection in a third viewport, all sharing a common image manipulation interface.

Discussion

Level 5: Thin Client Platform with Unified Archive

Level 5 represents a vision of a completely modularized architecture. Image data resides on a unified archive in the

The market for PACS and 3D applications is currently in flux, with evolving options for stand-alone systems and varying levels of system integration. Architectural characteristics, especially archive structure, client thinness, and



Fig. 4 Level 4 introduces a unified server application programming interface (API), capable of serving multiple client applications. The duplicated storage and redundant DICOM image transfers required in

the third level (Fig. 3) are eliminated here. Integration on the client machine continues to consist of context-sharing through a standardized protocol (*dotted line*)

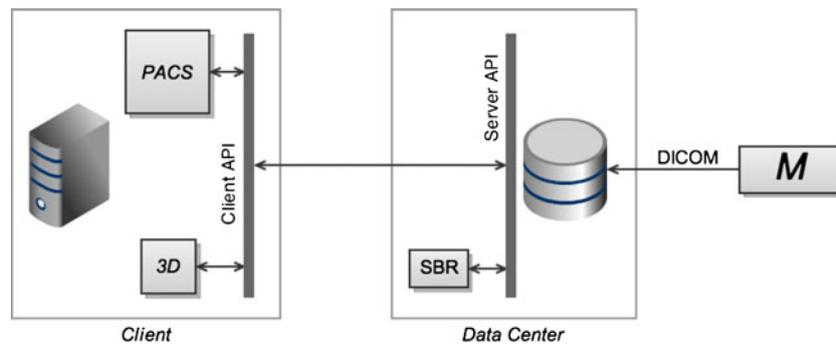


Fig. 5 The fifth level of the proposed classification scheme represents a fully modular framework, where client applications now plug into a client API. Client integration is now achieved through the shared application framework defined by this API, rather than through direct context-based links between applications. This allows for user

interfaces to be unified across multiple imaging modules. A standard for this type of architecture has been created by DICOM Working Group 23 [25], and a reference implementation has been developed by the National Cancer Institute's Cancer Bioinformatics Grid program [25, 26]

use of standardized protocols, determine the user experience and clinical utility of these tools. Increasing modularity in PACS and advanced visualization technologies should promote innovation by allowing vendors to concentrate on core strengths, as opposed to the current climate in which 3D vendors may need to develop certain fundamental PACS functions and vice versa. In evaluating advanced visualization options, departments and institutions should consider user-level integration, resource efficiency, performance, functionality, supportability, and interoperability in the context of the specific clinical applications of interest. These must be weighed against initial investment costs and longer-term maintenance costs. In general, thinner clients and SBR should lead to improved supportability and decreased longer-term costs. Greater modularity implies improved interoperability between products and vendors, with a consequent increase in freedom of choice for users. A summary of the relative strengths and weaknesses of the five levels with regard to selected system factors is shown in Table 1. Note that functionality and cost have been excluded from this table, since these may be more a function of specific products rather than the system architecture.

As noted above, any particular deployment of systems may or may not fall easily into one of the five levels of this classification hierarchy. In particular, it is common for a relatively thick PACS client to be integrated with a thin 3D client using proprietary context-sharing, which may correspond to a hybrid between levels 2 and 3 of the hierarchy. However, it is our intent to use this five-level decomposition to highlight key system characteristics (e.g., level 2 for context-sharing and level 3 for thin clients), rather than to describe specific equipment installations.

Work on modularized imaging application architecture has already been undertaken. Formed in 2004, DICOM Working Group 23 (WG-23) has been working to define standards to allow client machines to host plug-in imaging applications [25]. In addition, the National Cancer Institute's Cancer Bioinformatics Grid program has developed an implementation of the WG-23 standard, known as the Extensible Imaging Platform (XIP) [25, 26]. A related open source framework has also been developed by ClearCanvas Inc. [27]. Wider industry adoption of application hosting standards would facilitate migration toward modular commercial systems, improving the usefulness of advanced

Table 1 For each of the five levels of the proposed classification scheme, a qualitative score is listed with respect to user-level integration, resource efficiency, performance, supportability and freedom of choice

	User-level integration	Resource efficiency	Performance	Supportability	Freedom of choice
Level 1	↓	↓	–	↓	↑
Level 2	–	↓	↓	↓	↓
Level 3	–	–	↑	–	↑
Level 4	–	↑	↑	↑	↑
Level 5	↑	↑	↑	↑	↑

Each score may have one of three possible values: relative strength (↑), neutral (–), and relative weakness (↓). Client thinness, server-based rendering, standards-based protocols, archive unification, and application hosting frameworks, reflected by increasing levels in the hierarchy, tend to promote these factors

visualization in routine clinical imaging and enabling customers to assemble best-of-breed technology solutions.

Conclusion

We propose a five-level classification scheme for PACS and advanced visualization integration. A spectrum ranging from stand-alone turnkey solutions to modularized application architectures is discussed. Departments and institutions evaluating advanced visualization options should consider several factors in selecting vendors and products.

Acknowledgments KCW gratefully acknowledges the support of RSNA Research and Education Foundation Fellowship Training Grant #FT0904, as well as that of the Walter and Mary Ciceric Research Award.

References

- Kato Y, Katada K, Hayakawa M, Nakane M, Ogura Y, Sano K, Kanno T: Can 3D-CTA surpass DSA in diagnosis of cerebral aneurysm? *Acta Neurochir* 143:245–250, 2001
- Matsumoto M, Kodama N, Endo Y, Sakuma J, Suzuki KY, Sasaki T, Murakami K, Suzuki KE, Katakura T, Shishido: Dynamic 3D-CT angiography. *Am J Neuroradiol* 28:299–304, 2007
- Rubin GD, Dake MD, Napel SA, McDonnell CH, Jeffrey RB: Three-dimensional spiral CT angiography of the abdomen: initial clinical experience. *Radiology* 186:147–152, 1993
- Willmann JK, Wildermuth S: Multidetector-row CT angiography of upper- and lower-extremity peripheral arteries. *Eur Radiol* 15: D3–D9, 2005
- Desjardins B, Kazerooni EA: ECG-gated cardiac CT. *Am J Roentgenol* 182:993–1010, 2004
- Lawler LP, Pannu HK, Fishman EK: MDCT evaluation of the coronary arteries, 2004: How we do it—data acquisition, postprocessing, display, and interpretation. *Am J Roentgenol* 184:1402–1412, 2005
- Macari M, Bini EJ, Jacobs SL, Lange N, Lui YW: Filling defects at CT colonography: pseudo- and diminutive lesions (the good), polyps (the bad), flat lesions, masses, and carcinomas (the ugly). *Radiographics* 23:1073–1091, 2003
- Shi R, Schraedley-Desmond P, Napel S, Olcott EW, Jeffrey RB, Yee J, Zalis ME, Margolis D, Paik DS, Sherbondy AJ, Sundaram P, Beaulieu CF: CT colonography: influence of 3D viewing and polyp candidate features on interpretation with computer-aided detection. *Radiology* 239:768–776, 2006
- Soto JA, Lucey BC, Stuhlfaut JW, Varghese JC: Use of 3D imaging in CT of the acute trauma patient: impact of a PACS-based software package. *Emergency Radiology* 11:173–176, 2005
- Rodt T, Bartling SO, Zajaczek JE, Vafa MA, Kapapa T, Majdani O, Krauss JK, Zumkeller M, Matthies H, Becker H, Kaminsky J: Evaluation of surface and volume rendering in 3D-CT of facial fractures. *Dentomaxillofacial Radiol* 35:227–231, 2006
- Horton KM, Fishman EK: Multidetector CT angiography of pancreatic carcinoma: part I, evaluation of arterial involvement. *Am J Roentgenol* 178:827–831, 2002
- Takeshita K, Kutomi K, Takada K, Kohtake H, Furui S: 3D pancreatic arteriography with MDCT during intraarterial infusion of contrast material in the detection and localization of insulinomas. *Am J Roentgenol* 184:852–854, 2005
- Silverman PM, Zeiberg AS, Sessions RB, Troost TR, Davros WJ, Zeman RK: Helical CT of the upper airway: normal and abnormal findings on three-dimensional reconstructed images. *Am J Roentgenol* 165:541–546, 1995
- Horton KM, Horton MR, Fishman EK: Advanced visualization of airways with 64-MDCT: 3D mapping and virtual bronchoscopy. *Am J Roentgenol* 189:1387–1396, 2007
- Croitoru S, Gross M, Barmeir E: Duplicated ectopic ureter with vaginal insertion: 3D CT urography with IV and percutaneous contrast administration. *AJR* 189:W272–W274, 2007
- Levoy M. Polygon-assisted JPEG and MPEG compression of synthetic images. *Proceedings Special Interest Group on Computer Graphics and Interactive Techniques*, 21–28, 1995.
- Yoon I, Neumann U (2000) Web-based remote rendering with IBRAC (image-based rendering acceleration and compression). *Eurographics* 19:321–330
- Bohne-Lang A, Groch W-D, Ranzinger R: AISMIG – An interactive server-side molecule image generator. *Nucleic Acids Res* 33:W705–W709, 2005
- Poliakov AV, Albright E, Corina D et al. (2001) Server-based approach to web visualization of integrated 3D medical image data. *Proc Am Med Inform Assoc Symp* 533–537
- Erickson BJ, Persons KR, Hangiandreou NJ, James EM, Hanna CJ, Gehring DG: Requirements for an enterprise digital image archive. *Journal of Digit Imaging* 14:72–82, 2001
- Liu BJ, Cao F, Zhou MZ, Mogel G, Documet L: Trends in PACS image storage and archive. *Computerized Medical Imaging Graphics* 27:165–174, 2003
- Huang HK: Enterprise PACS and image distribution. *Computerized Medical Imaging and Graphics* 27:241–253, 2003
- Toland C, Meenan C, Toland M, Safdar N, Vandermeer P, Nagy P: A suggested classification guide for PACS client applications: the five degrees of thickness. *Journal of Digital Imaging* 19:78–83, 2006
- Health Level Seven International. Available at <http://www.hl7.org/implementation/standards/ccow.cfm>. Accessed 13 July 2010
- Paladini G, Azar FS: An extensible imaging platform for optical imaging applications. *Proc SPIE* 7171:717108, 2009
- Prior FW, Erickson BJ, Tarbox L: Open source software projects of the caBIG in vivo imaging workspace software special interest group. *Journal of Digital Imaging* 20:94–100, 2007
- ClearCanvas Inc. Available at <http://www.clearcanvas.ca>. Accessed 13 July 2010