Incorporating novel image processing methods in a hospital-wide PACS

Erwin Bellon^{a,*}, Michel Feron^a, Paul Neyens^a, Klaas Peeters^a,
Matthias Sweertvaegher^a, Frederik Maes^b,
Paul Suetens^b, Bart Van den Bosch^a

^a*Department of Information Systems, University Hospitals Leuven, Belgium*

^b*Medical Image Computing (ESAT and Radiology), Katholieke Universiteit Leuven, Belgium*

Abstract. We claim that integration of new and even experimental methods for image processing and analysis into a routine PACS is not only needed to enhance the functionality of that PACS, but also to foster introduction of those image processing methods into clinical routine. From that latter perspective it must be possible to incorporate such methods without requiring specific developments from the side of the PACS provider. In this paper we discuss some of the possibilities and difficulties. We illustrate those using initial examples of integration of image processing, by a university hospital, into a commercial PACS in routine hospital-wide operation. © 2005 Published by Elsevier B.V.

Keywords: PACS (Picture Archiving and Communication Systems); Image processing; Image analysis

1. Introduction

The basic technological and organizational problems in PACS have been solved, and there is sufficient experience available to state that within a few years the digital working organization will be the standard. However, up to now a PACS merely improves on functionality that was already available with film. In contrast, computer-aided image processing and analysis provides functionality that is *fundamentally impossible to achieve using traditional film*. Examples are: automated registration of data sets obtained from the same patient using different imaging modalities, or of data sets obtained at different time

* Corresponding author.

E-mail address: erwin.bellon@uzleuven.be (E. Bellon).

points, or against an anatomical atlas [1]; (semi-)automated delineation of anatomical structures of interest, which in itself is a requirement for many 3D visualization and surgery planning applications [2]; and automated lesion detection or extraction of quantitative measurements in the context of computer-aided diagnosis [3].

Progress in medical image analysis has been steady but slow over the last decades [4]. It is fair to state that, at the present time, a PACS is not introduced in the hospital for its abilities to provide advanced image processing.

2. Purpose

Nevertheless, we are convinced that the PACS community should start exploring ways to more flexibly (and more rapidly) integrate methods for image processing and analysis.

A first motivation is that there is room now in PACS research and development to spend energy on new challenges, with image processing being an important opportunity.

A second motivation for such integration is to leverage the PACS infrastructure for image processing research, which is characterized by the need to validate each individual method on a large number of cases. The hospital's PACS can make it easier to collect clinically meaningful images. The PACS can also provide to applied image processing research professional user interfaces for interactive viewing and evaluation.

A third and most important motivation for integrating promising image processing methods into a routine PACS is to foster their transition into clinical routine. This requires understanding of the clinical usefulness (as opposed to technical correctness) of the methods, and experience about the situations in which they can best be applied. To obtain this experience those methods must be used extensively in a "routine" situation, be it probably in a university hospital setting at first and with considerable caution. This evaluation can only be performed on images that are relevant to the physicians at that very moment. Moreover, in this phase it is paramount, for practical reasons, to make the processing methods available on the physicians' PACS console instead of inviting physicians to the imaging laboratory.

Examples of such image processing methods that are being used clinically in our hospital are multimodality registration for oncology and radiotherapy planning (involving PET, CT and MRI) or statistical processing for functional MRI. Although those algorithms have already been found promising, their application is limited to selected cases because too much manual work is required to deploy them in full clinical production. But PACS vendors can only be expected to incorporate these methods in their products and thus make it straightforward to invoke them, after they have passed this stage of validation in a realistic setting on many cases.

It is to shorten this cycle that we are experimenting with principles to incorporate image processing and analysis methods into our hospital-wide PACS, without for every method requiring dedicated support by our commercial PACS provider.

3. The PACS environment

Our hospital-wide PACS (Impax, Agfa) is accessed from 35 radiological workstations and more than 1500 clinical workstations. It is currently used primarily for images from

Radiology and Nuclear Medicine (with a yearly production of 400,000 exams) but also stores dynamic cathlab and ultrasound exams from Cardiology.

In our setup the PACS is tightly integrated into the EPR (Electronic Patient Record), which is developed in-house. Many of the management functions are shifted to the overall system, having the PACS concentrate on the specifics of image handling [5,6]. For example, clinical users can from within their overall clinical workstation request image exams to be transmitted to dedicated workstations such as for surgery planning. Our commercial PACS partner assisted us with this integration by providing access to the internal PACS management layer and to internals of the clinical viewer.

4. Possible approaches to integrating image processing and analysis

4.1. Vendor proprietary integration

The PACS provider has the most options to efficiently integrate image processing methods. For example, geometric registration of two data sets does not necessarily have to result in the generation of new, reformatted images. Instead, the registration algorithm needs only to determine the geometrical transformation matrix to be applied to the data set. The actual transformation can be performed on the fly as part of displaying the data. This requires dedicated support in the viewing software; it also requires proprietary management of that transformation matrix.

Nevertheless, although we can only encourage PACS manufacturers to incorporate their own optimized algorithms, we also want to be able to integrate processing methods ourselves. No PACS provider will provide all the methods we require. Moreover, as was mentioned before, we also want to integrate new, experimental methods for validation, just to gather the clinical experience to demonstrate to our commercial partners that their integration efforts will pay off. For that purpose, however, our own integrations need not necessarily be thoroughly optimized.

4.2. Front-end “plug in” components

In this approach, processing is started interactively while the physician is studying the images. Processing can include interactive actions as well, e.g. for semi-automated measurements. On the other hand, processing must not take too long.

This amounts to extending the viewer of the PACS with our own image processing components (or more generally with 3rd party components). Fig. 1 illustrates a simple extension we made to the commercial PACS viewer to present a fused view of PET and CT. This extension is in routine use in our environment. This integration eliminates the need in clinical case discussions to switch to a dedicated application and load images that were manually put on the hard disk of the PC before.

An important aspect here is that this component leverages functionality already available in the encompassing viewer. The processing component need not provide mechanisms for selecting the study.¹ Quite importantly, the processing component is not

¹ In fact, we explicitly want this component to not provide such selection functionality, as that would not only hinder overall workflow support but also jeopardize our fine-grained access control.

Fig. 1. Our own component for presenting the PET information from a PET/CT scanner as color overlay on the corresponding CT slice (right) was “plugged in” into the commercial clinical viewer of our PACS (left). Our component focuses on facilities to interactively change color blending and scrolling through slices, while delegating such aspects as exam selection, image retrieval over the network, and interactive presentation of the original stacks to the overall PACS framework.

concerned with fetching the pixels from the archive, as it gets those from the PACS framework. The overall viewer also provides presentation of the individual stacks, so our component can focus on its specific task.

This kind of integration requires a “plug in” architecture, with the encompassing PACS viewer providing a public API (Application Programming Interface) for such actions as accessing pixels. Some, albeit still few, commercial PACS vendors provide such public APIs that 3rd parties could use to augment their viewer.

Although the basic technology for such component integration is available, it is far from clear whether we can come to a standard high-level API for use among different vendors. Our goal is not to propose such a general API, but rather to convince PACS vendors that they should design their products using such extensibility in mind.

4.3. Front-end application synchronization

In the previous approach we aimed at (1) leveraging the technological functionality already available in the PACS framework, and (2) providing an integrated interface to our medical users that makes it efficient for them to call the image processing functions as part of their routine workflow. Even if the former goal (making life easier for the engineers) cannot be met, the latter goal (making life easier for the clinical users) remains important in routine operation or when evaluating practical usefulness of image processing.

A potential approach in this situation is to have different applications share a central context: when the user switches to a new study in e.g. the viewer, the processing application is notified about that event so it can also switch to that study. The processing application (for example for 3D presentation) has to fetch the images from the archive, even if those images are already loaded in the overall application.

For our purpose of fostering applied research, that is not all that bad. There are many situations in which adhering to the PACS framework may be difficult. Fig. 2 illustrates an application to evaluate the clinical relevance of automated extraction of the cardiothoracic ratio from chest radiographs. We do not want the hospital-wide PACS to be “polluted” at this

Fig. 2. To assess the clinical relevance of a method for automated extraction of the cardiothoracic ratio, we included assessment thereof into the routine workflow of some of our radiologists. Image processing and result management are performed outside the PACS, but that is a detail at this stage. For this application we found it easier to call up processing results from within the RIS, as that subsystem has most management information available about the type and purpose of the exam (left). The radiologist is suggested to have a quick look at image analysis results for this specific routine case and to fill in an evaluation form (right). Illustration courtesy of Dieter Seghers, Katholieke Universiteit Leuven.

stage with these analysis results. On the other hand, we want our radiologist's opinion about the usefulness of this automated measurement tool in routine application. For that purpose it is perfectly reasonable to have the processing done completely outside the PACS, as long as our radiologists do not loose time to assess the results and give feedback on relevance.

4.4. Back-end image processing chains

Processing that does not require user interaction can be performed in the “back office”. An advantage is that processing duration is less critical. The results must be stored to be

Fig. 3. In this application for functional MRI, a few images with statistical information were generated automatically out of the huge data set of primary images, and included into the PACS. Other applications may generate information that is more difficult to manage than images.

presented later. A typical result is a set of images, which can be conveniently stored in the PACS (Fig. 3). Other results such as contours on images may be managed in the PACS as DICOM objects as well, but there are other results that cannot readily be stored in this standard format. Even if the PACS enables “foreign” data to be archived, it cannot itself interpret or present those data. Some results may more naturally be stored into the EPR.

In practice there are some problems to overcome. Firstly, one must decide which processing to apply on what data. For small-scale experimental validation this could amount to manually send specific studies to a processing server, but for large-scale routine application this will require automation. The information for this decision may be available in the EPR rather than in the PACS. Secondly, the processing chain may require a step for interactive quality assessment (QA) by an expert. It remains paramount that all steps of the processing chain, including notification that results are waiting for QA, are managed by an automated system instead of requiring manual actions. Thirdly, it must not be overly difficult to incorporate new methods (that often will come from research groups) into the processing chain. Our current approach is to provide a set of tools that make it easier for the image processing researchers to provide “official” DICOM results and to comply with a general PACS integration API.

5. Conclusion

We anticipate a steep increase in the use of image processing and analysis in clinical practice, in part due to cross fertilization with PACS. On the one hand, now that we have PACS, we no longer have any excuse to not invest in further developing and integrating promising processing methods that have been waiting in the lab. On the other hand, new image processing and analysis will become a driving force for PACS developments.

More experience is required to gain a better understanding of the tools and APIs needed to more flexibly integrate image processing and analysis into overall organization and into overall PACS operations. Considerations to take into account are not only technological, but also are about organization, access control and a clear separation of responsibilities.

References

- [1] F. Maes, D. Vandermeulen, P. Suetens, Medical image registration using mutual information, *Proceedings of the IEEE—Special Issue on Emerging Medical Imaging Technology*, vol. 91(10), 2003, pp. 1699–1722.
- [2] J. Van Cleynenbreugel, et al., A semiautomatic three-dimensional segmentation method for disarticulation of bone structures on spiral computed tomography images, *Journal of Digital Imaging* 8 (4) (1995) 156–161.
- [3] G. Kiss, et al., Computer-aided diagnosis in virtual colonography via combination of surface normal and sphere fitting methods, *European Radiology* 12 (1) (2002) 77–81.
- [4] J.S. Duncan, N. Ayache, Medical image analysis: progress over two decades and the challenges ahead, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22 (1) (2000) 85–106.
- [5] E. Bellon, et al., PACS: a hospital-wide perspective, *Imaging Management* 4 (3) (2004) 16.
- [6] E. Bellon, et al., Integrating PACS into overall hospital informatics, *Proceedings of EuroPACS-MIR*, ISBN: 88-8303-150-4, 2004, pp. 75–78.