

A Distributed Workflow Management System for Automated Medical Image Analysis and Logistics

J.G. Snel^{a*} S.D. Olabariaga^b J. Alkemade^a H. Gratama van Andel^a A.J. Nederveen^a
C.B. Majoie^a G.J. den Heeten^a M. van Straten^c R.G. Belleman^b

a. Academic Medical Center, Meibergdreef 9, 1105 AZ Amsterdam, The Netherlands

b. Informatics Institute, University of Amsterdam, Kruislaan 403, 1098 SJ Amsterdam, The Netherlands

c. Institute of Medical Physics, Henkestr. 91, 91052 Erlangen, Germany

Abstract

Advances in medical image analysis have increased the need to integrate and deploy image analysis software in daily clinical routine and in epidemiological studies. We developed a Distributed Workflow Management System (DWMS) that supports a wide portfolio of image analyses in different CT and MRI application domains. The DWMS supports software components for image import/export, caching, processing and notification that are distributed on a heterogeneous grid of commodity computers. Communication between the components is performed by exchanging SOAP messages on request of standard compliant Web services. The workflows are executed fully automatically upon receipt of the medical images. After processing, the results are routed to a workstation for review and further analysis or to an image archive (PACS). A web-based monitor shows the status of running, pending and terminated workflows. The DWMS improves the interoperability between image acquisition devices, clinicians and researchers by making image analysis applications available in a transparent way, which accelerates the uptake of new research techniques. Through distributed computing, the workload is balanced and results can be obtained quicker. As the availability is guaranteed at a 24/7-hour basis, the system provides a reliable and completely automated solution for demanding image analysis tasks in a multi-vendor environment.

1. Introduction

Advances in medical image analysis (MIA) and new imaging technologies have increased the computational and storage demands. Next to vendor supplied visualization and

image processing software which runs on dedicated computer hardware, there is a growing interest in scientific MIA software developed for special purposes. Application domains include MRI, functional MRI (fMRI), perfusion MRI (pMRI), diffusion tensor imaging (DTI), and CT angiography (CTA).

MIA software originates from an experimental development phase and has the potential to pass through a multi-disciplinary scientific stage with possible spin-offs to the clinical practice. It evolves through phases that include development, testing, parameter optimization and tuning, (clinical) validation, and finally deployment in the clinical routine or in epidemiological studies [10]. From development to deployment, the software is adopted in setups that go from purely experimental to purely clinical, requiring an increasingly higher level of connectivity to varied and physically dispersed enterprise systems (different scanners or image archives). Moreover, the background of users shifts from technical to clinical, requiring the collaboration among researchers with expertise in a variety of different scientific areas (e.g., Radiology, Neurology, Psychology, Physics, Computing).

Deployment of MIA in daily clinical routine and epidemiological studies requires various steps to be carried out repetitively by the human operator. These steps include the necessary data logistics to transfer the medical images between scanning devices, image archives, DICOM¹ nodes, and MIA processing units which are possibly geographically dispersed over the Radiology enterprise. Once the images have arrived for analysis, the data and MIA methods should be selected and loaded before program execution. Finally the resulting data are transferred to a centralized storage facility (CSF) for later review or further analysis. Consequently, a great deal of expert knowledge and time is

¹Digital Imaging and Communications in Medicine, <http://medical.nema.org>

*Email: j.g.snel@amc.uva.nl

required to carry out these tasks successfully.

For effective deployment in a clinical setting, the MIA software is subordinate to a number of strict conditions regarding workflow integration, automation, availability, stability, and control. It needs to be integrated into the radiological environment such that interoperability among researchers, clinicians and imaging devices is guaranteed in a multi-vendor environment. No expert knowledge on data handling and program settings should be required, and the method should be invoked in a fully automatic fashion, independently from the operator. Moreover, the software must be available at all times, providing stable and reliable results. Finally, problems (failure) should be clearly indicated to the users that depend on the results, as well as to the technical support.

In this paper we present a distributed workflow management system (DWMS) for deployment of custom MIA applications in a clinical setting. The system aims at supporting the data logistics and analysis for various clinical tasks, such as diagnostics [7, 11], surgery planning [9], and epidemiological studies. We first present an overview of the major applications, illustrating how the system is currently used (Section 2). The system architecture is described in Section 3 and the implementation is outlined in Section 4. Section 5 presents a discussion and concluding remarks.

2. MIA Applications

A large variety of scientific software is routinely adopted in our hospital (Academic Medical Center, Amsterdam, The Netherlands) including techniques for image segmentation and registration, fiber tracking in DTI, blood perfusion analysis, and fMRI analysis. These consist of newly developed algorithms and methods, as well as opensource library components or command line utilities developed by other research institutes.

The fMRI Brain Software Library (FSL²) of the Oxford Centre provides utilities for brain image analysis, such as segmentation and registration, fMRI (calculation of activation maps), and DTI analysis [6]. Mainly psychologists, psychiatrists, MRI specialists and radiology researchers in epidemiological studies use it in our hospital. Currently, we are also investigating fMRI as an augmented tool for pre-operative planning and post-operative evaluation with brain tumor patients. Here FSL is used to produce brain activation maps to identify the motor and sensory cortices [9].

Other image processing software (e.g. APP³) was developed in our hospital to produce perfusion maps from pMRI multi-volume data [3] and apparent diffusion coefficient (ADC) and fractional anisotropy (FA) maps from DTI data [1].

²<http://www.fMRIB.ox.ac.uk/fsl>

³AMC Postprocessing Package, <http://amcpostpack.sourceforge.net>

Finally, a preprocessing technique was developed to enhance the visualization of CTA scans of the head, which are typically inspected with maximum intensity projection (MIP) reconstructions. The matched masked bone elimination (MMBE) technique [7, 11] eliminates bone pixels from CTA source images by using an additional non-enhanced (blanco) CT scan. Different bone registration methods, respectively rigid and piecewise, are adopted for the brain and the neck. In this manner, the method creates bone-free MIP images that enable better inspection of the vascular structures in the head and neck, a tool that is regularly adopted in emergency care in our hospital.

The above MIA applications have been introduced in the clinical routine via the DWMS. After image acquisition, the operator simply has to send the images to a particular DICOM node connected to the system. The adequate image logistics and analysis steps (workflow) are automatically determined and executed. The user (e.g., radiologist) is notified via e-mail when and where the analysis results are stored. The technical support is notified when problems occur.

3. Distributed Workflow Management System

The DWMS is based on a service oriented architecture that takes care of the processing of the workflows in a coordinated and controlled way. It is composed of a set of loosely coupled *modules* (distributed software components) to perform specific tasks of the workflow.

3.1. Architecture

Each *module* supports the following generic (cross-platform) functionalities:

- Operating system background service (UNIX Daemon or MS Windows Service) supporting start / pause / resume / stop functions to activate the *module* without the requirement of a user to be logged in to run the service;
- SOAP⁴ core to offer and request standard compliant Web services;
- Multi-threading to handle multiple (queued) requests;
- SQL database connectivity for the management and control of (installed) modules, workflows, computer resources, and users.

The *modules* are installed on a heterogenous grid of commodity computers running MS Windows or Linux. Multiple instances of modules, especially *Workers*, can be installed to balance the load due to excessive computations in simultaneously processed workflows. The communication

⁴Simple Object Access Protocol, <http://www.w3.org/TR/soap>

between the modules (remote invocation of Web services) is based on exchanging SOAP messages. The messages pass through a SOAP router that forwards them to the module of the requested type with the lowest busy level. This level is based on the occupancy of the module's request queue.

The following modules are involved in performing the basic tasks within the workflow. The *DICOM receiver* serves as a DICOM node that is capable of storing medical images at the CSF. The received images come from different image sources such as scanning devices, PACS, and other (trusted) DICOM nodes using the DICOM network protocol. After receiving all required DICOM series (scans), a workflow is instantiated by the *Workflow Manager*, which is responsible for the coordinated execution of workflows. The *Workflow Manager* instructs the *Worker Manager* to select a *Worker* that meets the specific requirements on computers resources to perform the requested MIA. The *Worker* downloads the DICOM data from the CSF, executes the MIA utility and uploads the metadata to the CSF. Depending on the format of the MIA metadata, the results will be exported by the *Exporter* using the DICOM (pixeldata) or (secure) FTP (non-pixeldata) protocol to the location of destination, e.g. a diagnostic viewing station for further analysis or review. Imported DICOM series and MIA metadata are cached on the CSF. After expiration the *Cacher* trashes the data accordingly. Finally, the subscribers are notified about the progress (instant messaging) and termination (e-mail) of the requested workflows by the *Notifier*. The *Dicom Receiver* and *Worker* use the secure FTP network protocol for transferring files to/from the CSF; the *Cacher* uses this protocol for (remote) data removal on the CSF.

Special modules are required for logging facility (*Logger*) and fault recovery (*Checker*). An overview of the DWMS embedded in the Radiology network is depicted in Figure 1.

3.2. Workflows

Designing workflows. The workflows are defined in an XML⁵ structure which are called *workflow templates*. They are composed of *units* that correspond to previously described tasks regarding image import, export, cache, analysis, and notification. A specific unit can be handled by the accompanying module, e.g. a *Worker* accepts *Worker* units to perform a MIA procedure. We developed a Workflow Designer application to create and modify the *workflow templates*. To process the units in a specific – possibly parallel – order, dependencies between units can be indicated by the Workflow Designer. The *Worker* unit for example will be made dependent on the *DicomSeries* unit(s) to extract information about the location of the received DICOM series that are input(s) for the MIA procedure. Af-

⁵eXtensible Markup Language, <http://www.w3.org/XML>

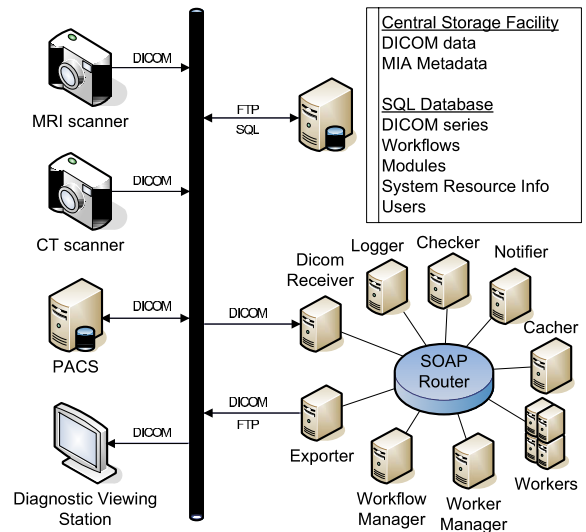


Figure 1. DWMS embedded in the Radiology network environment including imaging devices, image archives and diagnostic viewing stations.

ter creation the *workflow templates* are stored in the SQL database and can be activated for MIA. An example of a *workflow template* for the MMBE analysis (see Section 2) is indicated in the unit diagram of Figure 2. Each *workflow template* is registered on a specific TCP/IP network channel that is used by the *DICOM Receiver*. Ideally, for each *workflow template* another channel is chosen, such that only a single workflow will be carried out upon receiving data through that channel using virtual port (channel) addressing. However, it is important to limit the number of DICOM nodes (and thus the number of network channels) to be managed on scanners and the PACS. As such we register multiple *workflow templates* on the same network channel. The *Workflow Manager* uses a special mechanism based on regular and boolean expressions to match the DICOM tags of received series to instantiate the correct *workflow template*. Conflicting cases are solved by choosing other network channels. An example of a *DicomSeries* unit is depicted in Figure 3 where the DICOM tags regarding 'Acquisition Protocol', 'Manufacturer Model', 'Contrast Bolus Agent', and 'Body Part Examined' are used to match the non-enhanced series (blanco scan) in the MMBE workflow for the processing of CTA brain images (see Section 2). Generally, MIA procedures are applicable to a set of DICOM series possibly originating from various imaging studies. Following the hierarchy of the DICOM model, we use the DICOM level to group multiple import series in workflows. We distinguish SERIES, STUDY, PATIENT,

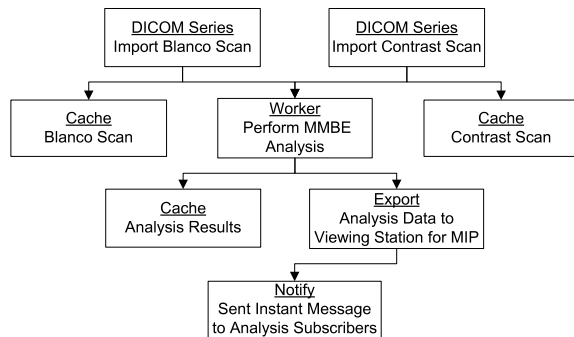


Figure 2. An overview of the units in the MMBE CTA workflow. Arrows indicate dependencies between units.

and HOSPITAL level. The respective IDs (e.g. series ID) are used in the *workflow instances* to couple (multiple) DICOM series according to the chosen DICOM level. If SERIES level is established only DICOM images that originate from the same DICOM series can be assigned to the workflow. In the example of the MMBE analysis STUDY level was chosen to combine a non-enhanced and an enhanced CT scan of the same study. Higher levels can be chosen for workflows regarding multi-modality analysis (PATIENT) or even multi-patient (HOSPITAL).

Instantiating Workflows. The *Dicom Receiver* activates all TCP/IP network channels that are defined in the *workflow templates* to communicate the DICOM data. A workflow can be instantiated by sending DICOM images from a (trusted) DICOM node to a network channel of the *Dicom Receiver*. Whenever a DICOM image of a new series has been received, the *Dicom Receiver* request the MATCH-SERIES Web service at the *Workflow Manager*. The *Workflow Manager* checks whether the incoming series should be included in an existing workflow instance or whether a new workflow instance should be created from a template.

Executing Workflows. The *Workflow Manager* is responsible for handling units in *workflow instances*. It starts with the units on top of the workflow (root units). Generally the root units are *DICOM Series* units that remain active until the last image of the DICOM series has been received. The *Workflow Manager* is notified about this by the SERIESFINISHED Web service that is requested by the *DICOM Receiver*. Next it processes the depending units by issuing Web services to modules of the requested type. A module receives a SOAP message encapsulating the unit from the SOAP router. Initially the SOAP request is placed in the module's request queue and will be handled by a module thread that marks the unit ACTIVE and starts processing. Finally, the unit is marked FINISHED after successful exe-

```

<unit xsi:type="xwf:DicomSeries" description="CTA Blanco">
  <dicomTagCondition>
    <tag tagId="0" description="Acquisition Protocol">
      <group>24</group><element>4144</element>
      <regEx>MMBE | CTA | CAROTIDEN | WILLIS</regEx>
    </tag>
    <tag tagId="1" description="Manufacturer Model">
      <group>8</group><element>4240</element>
      <regEx>MX8000</regEx>
    </tag>
    <tag tagId="2" description="Contrast Bolus Agent">
      <group>24</group><element>16</element>
      <regEx>CONTRAST</regEx>
    </tag>
    <tag tagId="3" description="Body Part Examined">
      <group>18</group><element>15</element>
      <regEx>BRAIN</regEx>
    </tag>
    <boolEx>tag(0) & tag(1) & !tag(2) & tag(3)</boolEx>
  </dicomTagCondition>
</unit>
  
```

Figure 3. An XML fragment of the Dicom Series unit to match dynamically the blanco scan (no contrast) for the CTA MMBE brain workflow requested by the Philips Mx8000 CT scanner using regular expressions (<regEx>) and a boolean expression (<boolEx>) based on DICOM tags indicated by group and element numbers.

cution, and the UNITFINISHED Web service is issued at the *Workflow Manager*. The *Workflow Manager* will continue handling the successive units until the whole workflow has been processed. As result, the workflow instance will be marked FINISHED in the SQL database.

Worker Manager and Workers. The most important task in the workflow is performing the MIA accomplished by the *Workers* that are managed by the *Worker Manager*. As such, we describe these modules in more detail.

A MIA procedure requires specific computational and network capacity to be carried out efficiently. The requirements regarding capacities on available CPU (GHz), internal memory (GB), hard disk space (GB), and LAN speed (MB/s) and are defined in the *Worker* unit. Each *Worker* stores information about the occupancy of these system resources in the SQL database and updates this information frequently by storing time averaged occupancies.

Worker units are initially directed by the *Workflow Manager* to the *Worker Manager*, which finds the least occupied *Worker* that meets the minimal system requirements. It behaves like a special router that interprets the SOAP message encapsulating the *Worker* unit. As a result, the *Worker Manager* forwards the SOAP message to the selected *Worker* that eventually performs the MIA on request.

In order to prevent *Workers* from being locked up by a MIA process, we use bounds on absorbed CPU and wall clock time that are indicated in the *Worker* unit. The *Worker*

frequently performs a self test to detect if it exceeds its maximum allowed computation time. It stops the analysis whenever this is the case and notifies the *Workflow Manager* by marking the unit FAILED using the UNITFAILED Web service.

Also the *Worker Manager* can interrupt a (low priority) analysis on a *Worker* to give way to a *Worker* unit of higher priority. The *Worker* unit is marked STOPPED in this case and returned to the *Workflow Manager* that will resume it in a later instance when computational resources become available again.

Recovery of Failed Workflows. A special module called *Checker* serves as an autonomous agent to check the availability of modules. It frequently invokes (remote) system inquiries to test the responsiveness of all modules. In case a module remains unresponsive, it is deactivated by removing it from the SOAP router table.

The *Checker* is also notified about a possible unit failure by the *Worker Manager*. It looks up the respective module that produced the failure in the central database and de/activates it accordingly. The *Workflow Manager* retries an incorrectly processed unit (marked FAILED) for three more times after which the *workflow instance* will be marked FATAL FAILURE. A fall back workflow is issued instead to notify the requestors and technical support about this. This is also the case for *workflow instances* that were not completely instantiated due to incomplete DICOM series. They expire and are marked FAILED after pending for more than one day.

3.3. Control Center

The DWMS uses a centralized SQL database to store and retrieve varied information for the management of workflows, modules, (trusted) DICOM sources, users that subscribe to workflows and to keep track on software versions of modules and MIA utilities. The remote database is accessed via an SQL server. We developed a tool called *Control Center* that is used for the following purposes: installation and configuration of modules (optional), automatic updates to synchronize the DWMS without the need to reconfigure each PC separately and repeatedly, module management (adding/removing/(de)activating modules), system administration (users, DICOM sources, log), data management (downloading analysis results, adjustment of data expiration, trashing), combined workflow/DICOM series/analysis log view, instant messaging to notify subscribers about the progress of workflows. The tool is meant for system administrators, for PC users that offer their idle CPU time through *Workers*, for clinicians and researchers requesting workflows.

4. Implementation and Results

In the current DWMS the modules are distributed over seven PCs supporting both MS Windows and Linux OS. Five PCs were set up as *Workers*. The remaining modules were installed on the other two PCs, such that the incoming (*DICOM Receiver*) and outgoing (*Exporter*) network traffic was separated. We introduced an application repository from which *Workers* are configured to support the MIA applications described in Section 2. For each analysis a *workflow template* was designed by selecting the distinctive tags for the *DICOM Series* unit(s) and adding units to accomplish the desired MIA process. Additionally, DICOM nodes have been added to the MR and CT scanners to automatically invoke the workflows upon transmission of the medical images.

For the 3 Tesla MR scanner (Philips, Intera, Eindhoven), three workflows are directed to the processing of fMRI (brain activity maps), pMRI (blood perfusion maps), and DTI (ADC and FA maps). Varied maps are produced in a fully automated fashion and the requesting users (radiologists and neurosurgeons) are notified on completion. The maps are forwarded to the neuronavigation system at the neurosurgery department for pre-operative planning of patients with brain tumors. The maps are also cached on the CSF and accessed by radiologists for post-operative diagnostic evaluation.

The DICOM node has also been included into the angiography protocols (brain and neck) on all CT scanners (Philips, Mx8000, Brilliance 64, Eindhoven / Siemens, Sensation, Erlangen) of the Radiology department to activate the workflows for MMBE CTA analysis. After image acquisition, the scans are automatically sent to the *DICOM Receiver*, and the *Workflow Manager* triggers the proper MMBE CTA analysis. After processing, bone free enhanced scans are submitted to a DICOM viewing station for advanced visualization, and radiologists are notified that data have arrived for volume rendering and maximum intensity projection (MIP). As this procedure also involves patients that are suspected from brain aneurism, the highest priority has been assigned to the MMBE CTA workflows [8].

Since the introduction of the first prototype of the DWMS (mid 2004) more than 2000 scans have been automatically processed and the demand is increasing further.

5. Discussion and Concluding Remarks

The DWMS offers an effective solution for (bio)medical researchers and clinicians to perform custom MIA applications in an automated fashion for evaluation, epidemiological studies or in clinical routine. The system provides a transparent infrastructure to carry out data logis-

tics management and to control the computational load for high throughput. Heavy multi-volume image registration and statistical analysis are carried out immediately after acquisition of fMRI data to produce brain activation maps that are directly forwarded to a neuronavigation system. Bone-free head scans, obtained with the computationally heavy MMBE technique, are made available for visualization shortly after arrival of a patient at the emergency care.

The use of the DICOM protocol allows for a seamless integration of the DWMS into the daily routine of the radiology department. All MIA applications in the DWMS portfolio can be invoked fully automatically at a 24/7-hour basis, independently from the operator. As such no expert knowledge or special skills are required from the operator, which is especially important for the MMBE CTA service. The focus can be put on the patient and the operation of the scanner, while the request for the particular MIA is generated on the fly and the radiologist is notified automatically.

The system records all *workflow instances*, keeping track of the provenance of the performed analyses. This information can be used to discard computations that were performed earlier in the context of other workflows that cached their (intermediate) results in the CSF. Moreover, it could be used as feedback to the development team, since, by executing MIA algorithms on different heterogeneous systems, software deficiencies that not were observed on the development platform might become apparent. As a result, sharing of MIA algorithms and provenance data could improve research and development and lead to good clinical practice.

Currently the workflows encompass units that denote high-level autonomous processes. Our next goal is to enrich workflows with increased granularity and enable user intervention in multi-step MIA protocols. Multi-step MIA algorithms could give rise to multiple *Worker* units that can be processed in parallel on different *Workers*. User intervention could be used to control the application settings before performing the image analysis, to approve analysis data before continuing a complex workflow at a critical stage, or whenever data are pending for final archival. User intervention could be accomplished by introducing a *User Agent* unit that pauses the workflow execution and notify the user accordingly. In this manner, the participants with different roles and skills would handle the varied interactive workflow tasks, allowing for synergy among researchers and clinicians in a virtual collaborative environment.

The DWMS builds on existing Web technologies to provide a scalable, configurable, and extensible cross-platform solution for MIA computing. It enables the reallocation and virtualization of CPU power of commodity computers, which are abundant in hospital environments. Other systems that have been recently developed for workflow management are based on Grid middleware (Globus toolkit⁶)

⁶www.globus.org

which provides tools for sophisticated authorization, access control and job scheduling [4, 5]. They are directed to archive and process medical images on large scale distributed systems. Grid technology offers an enormous potential by exposing massive computational power and provides added value through the integration of application metadata facilities into the middleware. The workflow management systems based on Grid middleware, however, are still in a maturing phase as they suffer from various limitations regarding the support for non-UNIX OSs, data management, state of service, and security [2]. We will consider integrating these Grid facilities into DWMS without sacrificing its current availability, reliability and security in the clinical production process.

References

- [1] E. Akkerman. Efficient measurement and calculation of MR diffusion anisotropy images using the platonic variance method. *Magnetic Resonance in Medicine*, 49:599–604, 2003.
- [2] T. Glatard, J. Montagnat, and X. Pennec. Grid-enabled workflows for data intensive applications. In *Proc. 18th IEEE Symp. on Computer Based Medical Systems (CBMS'05), Dublin, Ireland, June 23-24*. IEEE, 2005.
- [3] F. Calamante et al. Measuring cerebral blood flow using magnetic resonance imaging techniques. *Journal of Cerebral Blood Flow & Metabolism*, 19:701–735, 1999.
- [4] J. Montagnat et al. Medical images simulation, storage, and processing on the European DataGrid testbed. *Journal of Grid Computing*, 2(4):378–400, 2004.
- [5] A. Rowland et al. Information extraction from images (IXI): Image processing workflows using a grid enabled image database. In *Proc. of the DiDaMIC Workshop (MIC-CAI'2004)*, 2004.
- [6] S. Smith et al. Advances in functional and structural MR image analysis and implementation as FSL. *NeuroImage*, 23:208–219, 2004.
- [7] M. van Straten et al. Removal of bone in CT angiography of the cervical arteries by piecewise matched mask bone elimination. *Med Phys*, 31(10):2924–33, October 2004. Erratum in: *Med Phys*. 2005 Jan;32(1):303.
- [8] M. van Straten et al. Automatic bone removal in CT angiography. *Medica Mundi*, 49(1):4–8, May 2005.
- [9] E. Vlieger et al. Functional magnetic resonance imaging for neurosurgical planning in neuro-oncology. *Eur Radiol*, 14:1143–1153, 2004.
- [10] S. Olabariaga, J. Snel, C. Botha, and R. Belleman. Integrated support for medical image analysis methods: from development to clinical application. *Accepted for publication in IEEE Trans. Information Technology in Biomedicine*, 2006.
- [11] H. Venema, F. Hulsmans, and G. den Heeten. CT angiography of the circle of Willis and intracranial internal carotid arteries: maximum intensity projection with matched mask bone elimination-feasibility study. *Radiology*, 218(3):893–8, March 2001.