Parameter Sweeps for Functional MRI Research in the "Virtual Laboratory for e-Science" Project

Sílvia D. Olabarriaga (1,2)
silvia@science.uva.nl

Aart J. Nederveen (1)
a.j.nederveen@amc.uva.nl

Breanndán Ó Nualláin (2)
bon@science.uva.nl

(1) Academic Medical Center, (2) Informatics Institute
University of Amsterdam, The Netherlands

Abstract

Image analysis is an important component of neuroscience research. The ICT infrastructure and technical knowledge needed to perform (large scale) neuroimaging studies, however, is often not available to the neuroscientists. The "Virtual Laboratory for e-Sciences" project provides an advanced (grid) infrastructure offering data and computing services to researchers from several application domains. In this paper we describe how this infrastructure is used in the context of functional Magnetic Resonance Imaging (MRI) research, which is devoted to the study of brain activity due to stimulation. Our experience in using a generic application (Nimrod, Monash University, Australia) to manage large parameter sweep experiments is presented. These experiments were performed to investigate the effect of one parameter (delay in the hemodynamic response function) in the analysis result, but also to evaluate the available infrastructure (grid resources and experiment management services). Initial results indicate that it is feasible and simple to perform large image analysis experiments using Nimrod once the environment has been properly prepared.

1. Introduction

Computerized image analysis has become an important component of neuroscience research. The variety, amount and size of available images create difficulties for neuroscientists, who typically do not have the necessary technical means to organize their workflow adequately. Firstly, setting up the proper image analysis steps might require at least some parameter tuning for the type of images or focus of the study. Secondly, running the image analysis requires large ICT capacity for computation and storage, which usually is not available at the organizations that perform neuroscience research. And finally, there is growing interest in combining data and metadata generated in different studies, requiring mechanisms for controlled access to distributed data. The "Virtual Laboratory for e-Sciences" project (VL-e) investigates means to reduce the difficulties faced by researchers in various application domains by offering grid-enabled facilities and general services. One of the subprograms focuses on medical imaging, providing the context of the work presented in this paper. Here we describe the current attempts to facilitate performing large scale experiments for the analysis of (brain) images on distributed computing resources. A generic tool, Nimrod [2], manages experiments involving parameter sweeping and optimization. The initial application is functional Magnetic Resonance Imaging (fMRI).

The paper introduces the VL-e project and the medical subprogram in section 2. Section 3 presents the context of the fMRI application, and section 4 briefly describes the infrastructure under development, coined "Virtual Lab for fMRI" (VL-fMRI). Section 5 describes the experimental set-up and section 6 presents initial results. We conclude the paper with a discussion of related work, an evaluation of our experiences and future perspectives.

2. VL-e and VL-e Med

VL-e (*www.vl-e.nl*) is a Dutch project that brings together developers and users of grid technology with the goal of "boosting e-Science by creating an e-Science environment and carrying out research on methodologies." This project is carried out by a consortium of 40 academic and industrial partners from the private and governmental sectors. VL-e is organized into subprograms that address different components of an e-Science environment: applications, virtual laboratory methodologies (e.g., workflow, parallel programming, databases, integration) and infrastructure (data and computing resources). A "proof-of-concept" environment (VL-e PoC) implements the physical infrastructure (storage, computing) and general ICT services shared by different applications. The VL-e PoC resources are currently hosted by the SARA Computing and Networking Services and the Dutch Institute for Nuclear Physics and High Energy Physics (NIKHEF).

The “Medical Diagnosis and Imaging” subprogram (VL-e Med) aims at developing a grid-enabled problem-solving environment for medical imaging. Several activities are developed by a multi-disciplinary team of scientific programmers, system (integration) experts, physicists and end-users, involving several institutes of the University of Amsterdam (UvA), the Free University (VU), Philips, SARA and NIKHEF. One of these activities focuses on a specific users community: neuroscientists who share the 3.0 Tesla Philips MRI (3T MRI) scanning facilities available the Academic Medical Centre (AMC) of the UvA. These facilities are used for (clinical) research and advanced clinical applications, in particular for functional MRI (50%).

3. Functional MRI

Functional MRI is a modality that enables the observation of brain activity during physical or cognitive stimulation [11]. The activated brain regions receive oxygen-rich blood, and the changes in the oxygenation level can be detected with MRI. The acquired scan contains a series of 3D datasets (volumetric images) with measurements through time, some of which obtained during stimulation and some at “rest”. The scans, stimulus data and other signals (e.g., EEG) are analysed in a sequence of steps including temporal and spatial motion correction, normalisation, and statistical analysis. The goal is to generate maps indicating the location of the regions where brain activation has been detected. Activation maps are used to characterize brain function in individuals (e.g., for neurosurgery planning [13]) or in populations (e.g., for neuroscience research).

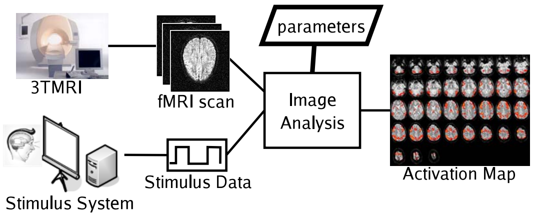

Figure 1. Simple workflow for fMRI data acquisition and analysis for an individual.

Figure 1 illustrates a simplified acquisition and analysis workflow for one individual – see more details in [17]. From the raw scan to activation maps, the image analysis process can take from minutes to hours, generating many intermediate results (500MB to 1GB) that need to be stored for each individual until the study is completed. For population studies, an additional statistical analysis step is performed based on the activation maps of all individuals. Here the individual scans also need to be aligned to a common reference coordinate system, e.g., an average brain.

The analysis of fMRI data is performed with sophisticated and complex methods that are still subject of research and development. Neuroscientists typically do not develop the methods themselves, but use existing software packages such as the fMRIB Software Library (FSL) [18] and Statistical Parametric Mapping (SPM) [10]. Although these packages hide much of the complexity of the image analysis process, the choice of parameters still plays an important role in fMRI analysis. Most researchers, however, adopt standard parameter values for spatial smoothing, hemodynamic response function and image registration. The differences between the final results of an fMRI study obtained with different parameter settings are often not investigated because of practical limitations regarding computing resources. We hypothesize that a practical implementation of parameter sweeps, supported by an adequate computing infrastructure could enhance fMRI data analysis.

4. Virtual Lab for fMRI

The Virtual Lab for fMRI (VL-fMRI) is an attempt to face logistics problems experienced by neuroscientists by providing a computational infrastructure to facilitate the storage, analysis and sharing of fMRI data. The motivation, challenges and structure of VL-fMRI are described in detail in [15]. For completeness, a summary of the main features relevant for parameter sweeps is presented below.

Figure 2 illustrates the main components of VL-fMRI: acquisition devices located at the AMC; computational resources (data storage and computing) located at the SARA and NIKHEF; low level services provided by standard grid middleware; generic virtual laboratory services; and application-specific services for data acquisition, storage, analysis and access.

Access to local and remote data resources is obtained with the graphical user interface of the Virtual Resource Browser (VBrowser, see [14]). All fMRI data (scans, signals, results) are stored in the Storage Resource Broker (SRB, *www.sdsc.edu/srb*), which provides a seamless interface to store and retrieve federated data and metadata across a wide area network [6].

The image analysis workflow is run on the grid resources available at the VL-e PoC environment (currently two clusters, in total around 140 cores). The analysis is implemented by legacy applications activated via the command-line. Currently all computing nodes have the “VL-e PoC software installation”, which already includes FSL and SRB utilities. The module `feat` of FSL is used for fMRI analysis, encapsulating all steps into one script. A configuration file pro-

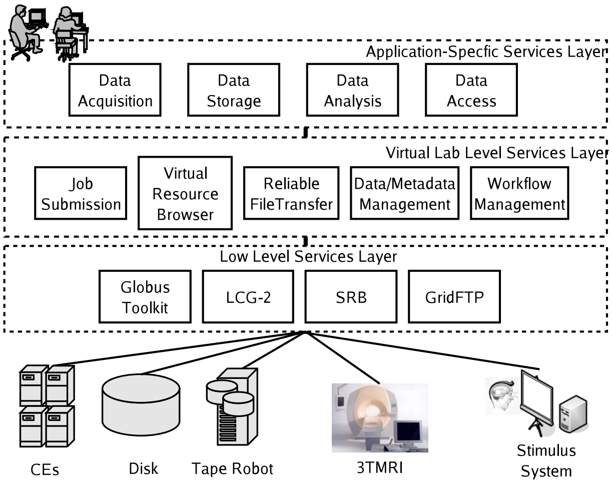

Figure 2. Components of VL-fMRI: acquisition devices, computing and storage resources, and services at the low, virtual laboratory and application-specific levels.

vides the parameters for the methods adopted in the analysis. The data is stored in the SRB, being staged in and out the computing node by a wrapper script. Also the `feat` configuration file is adapted automatically based on the local FSL installation and the staged data. When the analysis is completed, the results are uploaded into the SRB and can be interactively inspected by the user with the VBrowser. Computing jobs can be submitted using grid middleware of the EGEE project (Enabling Grids for E-science in Europe, *www.public.eu-egee.org*). Nimrod is used for large experiments, when image analysis is performed for varying parameter values (parameter sweep or optimizations) or for sets of images (job farming).

Nimrod is a system that facilitates the management of large experiments on distributed resources [2]. Nimrod/G [1] has tools to perform parameter sweeps (our focus here), while Nimrod/O [8] additionally has facilities to perform optimizations with a variety of methods. Nimrod can access computational resources through several widely-used interfaces for distributed computing (Globus, Condor, Legion). Nimrod dynamically discovers, acquires and monitors computational resources and adaptively schedules the tasks among them. It is fault-tolerant in the sense that a failure at a computing resource will not lead to failure of the entire computational experiment; Nimrod merely restarts the task on another computing resource. Further, additional resources can be added to a running experiment.

The computational experiment to be carried out is presented to Nimrod as a “plan file” using a declarative, parametric modelling language which has a straightforward syntax. The user provides a description of each parameter, its range of values, and the task to be carried out. Nimrod runs a task for each combination of parameter values, scheduling the resources automatically. It stores all task-related information in a database, which the user can query to monitor the experiment status or to extract performance statistics.

5. Experimental Set-up for Parameter Sweep

The main goal was to evaluate the usage of Nimrod available in the VL-e PoC infrastructure. A simple parameter sweep experiment was designed to vary the delay of the hemodynamic response function (HRF). The HRF is defined as the MR signal function related to a brief neuronal stimulus [5]. This signal reaches a peak after some time after the stimulus is applied, and this “delay” is an important parameter in HRF modelling. It is known that the form of the HRF varies among subjects. Furthermore some evidence exists that the HRF differs between brain regions [3]. Though attempts have been made to use basis sets of HRF functions in order to allow for variations between subjects [19], in general standard HRF functions are largely used in fMRI analysis, for example, as in FSL `feat` [18]. Here the approach is to sweep the delay parameter (from 4 s up to 8 s), while using only one HRF function (a mixture of two Gamma functions) per analysis.

Eleven healthy volunteers were scanned on a Philips 3.0 T Intera scanner during an event-related task paradigm (viewing of emotional pictures, IAPS). The task was performed twice for different echo times (22 datasets). The analysis was performed with `feat`, including the steps of motion correction, slice timing, Gaussian smoothing (FWHM 5 mm), and statistical analysis using the general linear model (GLM). All datasets were previously stored and prepared for analysis with jobs submitted via the EGEE middleware. Nimrod invoked the fMRI analysis and performed all resource management automatically based on the given plan file. One cluster was used to run jobs with a maximum duration of four hours (wall clock time).

Six experiments were performed to investigate different aspects of the provided facilities. In the first five experiments the complete workload (22 scans $\times$ 17 parameters per scan, in total 374 jobs) was split for each group of volunteers. Nimrod was configured to use a maximum of 8 computing nodes to run the jobs. The sixth experiment covered half of the workload (11 scans $\times$ 17 parameters per scan, in total 187 jobs) and Nimrod was configured to use a maximum of 16 computing nodes simultaneously. The following measurements were performed: total elapsed time to perform the experiment (from submission to completion of all jobs, provided by Nimrod); time to start the experiment due to queuing for resources (from submission until the first job is started); computing time of each job, including data staging and data analysis (wall clock time, command `date`). The standard output and standard error were inspected manually for errors.

6. Results

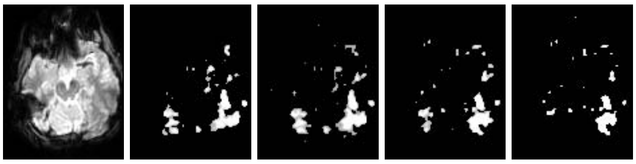

Figure 3. Slice of a structural scan and activation maps for increasing HRF delay (4,5,7 and 8 s).

Figure 3 presents an impression of the obtained results, showing a single slice of activation maps generated with different HRF delays for the same subject. Note that the size of the activated area is different for each value.

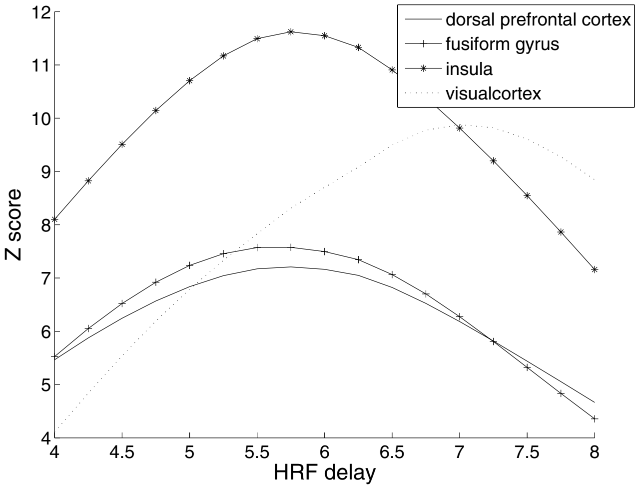

Figure 4. Activation (Z-scores) versus HRF delay for four different brain areas.

In Figure 4 the statistical brain activation (Z-score) obtained for each parameter is plotted for four different areas (dorsal prefrontal cortex, fusiform gyrus, insula, visual cortex). Note that the optimal delay (to obtain maximal activation) differs per area. This result is consistent with our present knowledge of the HRF variability throughout the brain, and indicates the usefulness of performing such experiments to better understand and choose parameter values for the analysis of fMRI data.

The complete experiments consisted of 561 fMRI analyses that were performed in 63.8 hours, generated around 90 GB of data and consumed a total of 416.7 h of computing time. Besides monitoring the status and checking the results for errors, no manual intervention was required to perform the experiments. This was a significant improvement with respect to the previous workflow management, which used EGEE middleware directly from custom scripts for parameter crude experiment management.

A summary of the measurements is presented in Table 1. From the 561 jobs, six failed due to an expired grid proxy and two failed due to temporary connectivity problems between the node and the SRB server. The large variation in queuing time (from 3 to 55 min) reflects the load and share of resources available for the VL-e Med virtual organization in Nov-Dec, 2006. Also note the large variation in the time to complete each job (from 30 to 148 min), which cannot be due to differences in the image analysis: all jobs perform essentially the same task, the images have approximately the same size, and the variations in the parameter value do not have consequences for the computation time. Since the cluster used in these experiments is composed of computers with homogeneous architecture, this difference can only be explained by some overhead introduced by experiment management or by bottlenecks in the local and remote resources, such as memory and disk[1] and the SRB server.

In an attempt to identify bottleneck conditions, we used job start/end time to visualize how the jobs were distributed along time for each experiment. Figure 5 shows plots for experiments 4 and 5. In some cases resources were progressively obtained by Nimrod to run jobs (Figure 5-top), while in others all resources were obtained at once and remained available to Nimrod during the whole experiment (Figure 5-bottom). These results show that, once it gets started, Nimrod keeps a good occupancy of resources and the experiment progresses rapidly.

Experiment 6 was designed as a large-scale test, repeating a subset of tasks performed in the previous five and collecting more monitoring data. Contrary to expectation, a larger load on the cluster and the SRB server (up to 16 jobs were run simultaneously) did not cause a drop in performance (see Figure 6 and Table 1). Moreover, very long computation times were not observed in this case.

7. Final Remarks

Approaching image analysis in neuroimaging research as a parameter sweep problem is not new. In [16] the

---

[1] The cluster used has dual-core architecture.

Exp.	# scans	# Max. Nodes	# Jobs	Elapsed	Queuing	Computing	Job Computing Time (min)	
							Mean±Std.dev.	[Min, Max]
1	2	8	34	6.7 h	55.3 min	24.5 h	43.3 ± 7.2	[34.1, 62.9]
2	6	8	87	13.7 h	12.7 min	72.5 h	49.9 ± 18.7	[30.4, 148.4]
3	8	8	136	15.2 h	29.9 min	101.5 h	44.8 ± 12.0	[32.4, 113.4]
4	2	8	34(1)	4.8 h	7.2 min	25.5 h	46.3 ± 9.7	[32.4, 62.2]
5	4	8	68(6)	8.6 h	3.1 min	62.8h	60.8 ± 18.9	[35.8, 138.9]
6	11	16	187(1)	14.8 h	8 min	133 h	42.92 ± 2.77	[39.1, 55.2]
Total	33	-	561(8)	63.8 h	116.2 min	416.7 h	-	[30.4, 148.4]

Table 1. Results: number of processed scans, maximum number of computing nodes, number of jobs (failed jobs indicated in parenthesis), total elapsed time of experiment, queuing time, total computing time, and computing time per job (mean, standard deviation, minimum and maximum).

Pipeline system is used to manage the workflow of large scale image analysis experiments with custom middleware. In [12] the Gridbus resource broker is used to deploy neuroimaging MRI applications on Grids. The FSL package is the basis for the image services, and parameter sweep applications are composed in the XML-based Parametric Modelling Language. Nimrod/G has been used to analyse brain activity data acquired with magnetoencephalography [9], as well in related areas such as molecular modelling for drug design [7].

Our efforts in VL-e Med are concentrated on applying existing software (e.g., FSL, Nimrod) to build (e-Science) applications to be used by end users, neuroscientists in this case. In addition to running parameter sweeps, in this paper we also evaluated the environment for performing large scale image analysis experiments. We observed that Nimrod can facilitate enormously the management of experiments, in particular when compared with other options that handle individual jobs (e.g. EGEE job submission middleware). However, it took a significant amount of time and expertise to properly set-up the environment to successfully run the experiments described here. The preparations include: parameterise the task (for Nimrod); adapt the programs to run on grid nodes (with remote data resources); manual discovery and configuration of grid resources; and error handling. We are currently preparing templates to accelerate the development of new experiments for fMRI analysis with `feat`. It is unlikely, however, that in the short term the "real end users" will use this infrastructure without expert assistance. An important barrier is the user interface, since we currently only dispose of command-line utilities in the VL-e PoC installation. The installation of a Nimrod portal is under consideration to provide a friendlier user interface. In spite of the difficulties, Nimrod remains a promising tool not only for parameter sweeps, but also for optimizations. We are investigating Nimrod/O to automatically search for the optimal HRF delay for a given brain region as an aid for neuroscientists to tune image analysis parameters for specific studies. In the future, incorporation of Nimrod for multi-stage parameter sweeps and/or optimisation will be desirable. We believe that work currently in progress to incorporate Nimrod and SRB access into Kepler [4] will prove useful. This would enable the user to build complex scientific workflows which contain parameter sweeps and optimisations as components, as well as facilitate collaboration and provenance tracking.

Acknowledgements

Many thanks to colleagues who set-up and maintain the VL-e PoC: B. Heupers, M. Bouwhuis, J.J. Keijser, D. van Dok, and many others behind *grid.support@sara.nl* and *grid.support@nikhef.nl*. We are also grateful to Dr. M. Ruijter and Prof. Dr. D. Veltman for collaboration in the neuroscience study. Finally, we thank the support of the VL-e director, Prof. Dr. L.O. Hertzberger, and the VL-e Med coordinators at the AMC, Prof. Dr. C.A. Grimbergen and Prof. Dr. G.J. den Heeten. The VL-e project is supported by a BSIK grant from the Dutch Ministry of Education, Culture and Science (OC&W) and is part of the ICT innovation program of the Ministry of Economic Affairs (EZ).

References

[1] D. Abramson, J. Giddy, and L. Kotler. High performance parametric modeling with Nimrod/G: Killer application for the global grid? In *Intl. Parallel and Distributed Processing Symp. (IPDPS), Cancun*, pages 520–528, May 2000.

[2] D. Abramson, R. Sosic, J. Giddy, and B. Hall. Nimrod: A tool for performing parametised simulations using distributed workstations. In *The 4th IEEE Symposium on High Performance Distributed Computing*, 1995.

[3] G. Aguirre et al. The variability of human bold hemodynamic responses. *NeuroImage*, 8:360–369, 1998.

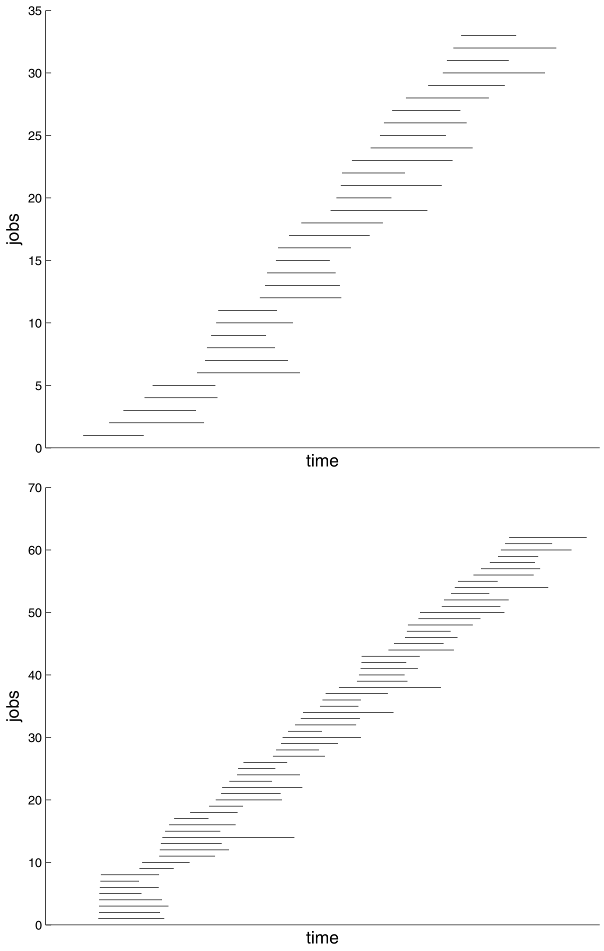

Figure 5. Time span of individual jobs for experiments 4 (top) and 5 (bottom).

[4] I. Altintas et al. A framework for the design and reuse of grid workflows. In *Intl. Workshop on Scientific Applications on Grid Computing (SAG'04)*, volume LNCS 3458, 2005.

[5] E. Amaro and G. Barker. Study design in fMRI: Basic principles. *Brain and Cognition*, 60:220–232, 2006.

[6] C. Baru, R. Moore, A. Rajasekar, and M. Wan. The SDSC storage resource broker. In *CASCON'98 Conference*, 1998.

[7] R. Buyya, K. Bramson, J. Giddy, and D. Abramson. The Virtual Laboratory: a toolset to enable distributed molecular modeling for drug design on the World-Wide Grid. *Concurrency and computation: practice and experience*, 2003.

[8] R. Buyya et al. Economic models for management of resources in peer-to-peer and grid computing. In *SPIE Intl. Symp. on The Convergence of Information Technologies and Communications (ITCom 2001), Denver, Colorado, USA*, August 2001.

[9] R. Buyya et al. Neuroscience instrumentation and distributed analysis of brain activity data: a case for eScience

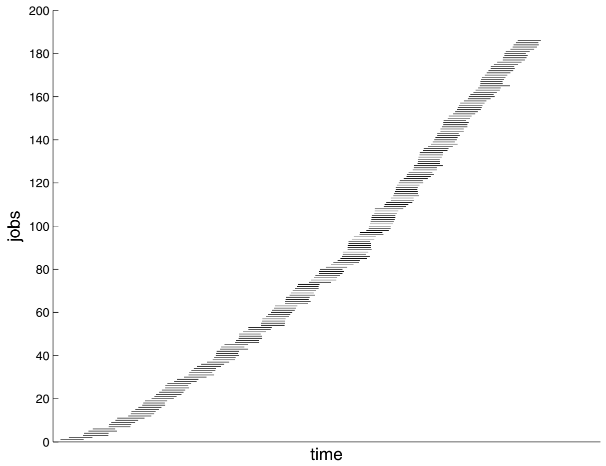

Figure 6. Time span of individual jobs (exp. 6).

on global Grids. *Concurrency and computation: practice and experience*, 2005.

[10] K. Friston. Statistical parametric mapping and other analysis of functional imaging data. In *Brain Mapping: The Methods*, pages 363 – 385. Academic Press, 1996.

[11] P. Jezzard, P. M. Matthews, and S. Smith. *Functional MRI: An Introduction to Methods the Art of Medical Image Registration: Correspondence, Homology and Quality*. Oxford University Press, 2003.

[12] S. Kolbe et al. A Platform for Distributed Analysis of Neuroimaging Data on Global Grids. In *Proc. of First IEEE Intl. Conf. on e-Science and Grid Computing*, 2005.

[13] E. Vlieger et al. Functional magnetic resonance imaging for neurosurgical planning in neuro-oncology. *Eur Radiol*, 14:1143–1153, 2004.

[14] S. Olabarriaga, P. de Boer, A. Belloum, K. Maheshwari, J. Snel, A. Nederveen, and M. Bouwhuis. Virtual Lab for fMRI: Bridging the Usability Gap. In *Proc. of Second IEEE Intl. Conf. on e-Science and Grid Computing*, 2006.

[15] S. D. Olabarriaga, A. J. Nederveen, J. G. Snel, and R. G. Belleman. Towards a Virtual Laboratory for fMRI Data Management and Analysis. In *Proceedings of HealthGrid*, pages 43–54. IOS Press, 2006.

[16] M. Pan and A. W. Toga. A grid enabled workflow management system for managing parameter sweep applications in neuroimaging research. In *Proc. of Sixth IEEE Intl. Symp. on Cluster Computing and the Grid Workshops*, 2006.

[17] S. Smith. Overview of fMRI analysis. *The British Journal of Radiology*, 77:S167–S175, 2004.

[18] S. Smith et al. Advances in functional and structural MR image analysis and implementation as FSL. *NeuroImage*, 23:208–219, 2004.

[19] W. Woolrich, T. Behrens, and S. Smith. The variability of human bold hemodynamic responses. *NeuroImage*, 21:1748–1761, 2004.