



HHS Public Access

Author manuscript

Neuroinformatics. Author manuscript; available in PMC 2017 July 01.

Published in final edited form as:

Neuroinformatics. 2016 July ; 14(3): 305–317. doi:10.1007/s12021-016-9296-7.

Heterogeneous Optimization Framework: Reproducible Preprocessing of Multi-Spectral Clinical MRI for Neuro-Oncology Imaging Research

Mikhail Milchenko^a, Abraham Z. Snyder^a, Pamela LaMontagne^a, Joshua S Shimony^a, Tammie L. Benzinger^a, Sarah Jost Fouke^b, and Daniel S. Marcus^a

^aDepartment of Radiology, Washington University School of Medicine, St. Louis, MO, USA

^bDepartment of Neurological Surgery, Swedish Medical Center, Seattle, WA, USA

Abstract

Neuroimaging research often relies on clinically acquired magnetic resonance imaging (MRI) datasets that can originate from multiple institutions. Such datasets are characterized by high heterogeneity of modalities and variability of sequence parameters. This heterogeneity complicates the automation of image processing tasks such as spatial co-registration and physiological or functional image analysis.

Given this heterogeneity, conventional processing workflows developed for research purposes are not optimal for clinical data. In this work, we describe an approach called Heterogeneous Optimization Framework (HOF) for developing image analysis pipelines that can handle the high degree of clinical data nonuniformity. HOF provides a set of guidelines for configuration, algorithm development, deployment, interpretation of results and quality control for such pipelines. At each step, we illustrate the HOF approach using the implementation of an automated pipeline for Multimodal Glioma Analysis (MGA) as an example. The MGA pipeline computes tissue diffusion characteristics of diffusion tensor imaging (DTI) acquisitions, hemodynamic characteristics using a perfusion model of susceptibility contrast (DSC) MRI, and spatial cross-modal co-registration of available anatomical, physiological and derived patient images.

Developing MGA within HOF enabled the processing of neuro-oncology MR imaging studies to be fully automated. MGA has been successfully used to analyze over 160 clinical tumor studies to date within several research projects. Introduction of the MGA pipeline improved image processing throughput and, most importantly, effectively produced co-registered datasets that were suitable for advanced analysis despite high heterogeneity in acquisition protocols.

Corresponding author: ^aMikhail Milchenko, Ph.D., Neuroinformatics Research Group, 4525 Scott Ave, Campus Box 4225, Phone: +1-314-286-0562, mmilch@wustl.edu.

Information sharing statement

The installation package for the docker.io MGA image is available from https://bitbucket.org/mmilch01/mga_docker_install. The experimental XNAT-MGA pipeline module that uses this Docker image is available from <https://marketplace.xnat.org/plugin/hof> and from the NITRC neuroinformatics tool repository at <http://www.nitrc.org/projects/hof>.

Keywords

MRI; neuro-oncology imaging; data modeling; knowledge representation; spatial co-registration; DSC imaging; diffusion imaging

1. Introduction

Multi-modal magnetic resonance imaging (MRI) has been used with increasing frequency in clinical and research neuroimaging. Use of multiple contrast mechanisms is required in functional MRI analysis (Gholipour et al. 2007), allows for tumor boundary prediction during surgical planning (Hlaiheli et al. 2010); (Law et al. 2008; Zonari et al. 2007), enables a wide range of morphometric analyses (Buckner et al. 2004), and may improve accuracy of classification of brain tissues (Vrooman et al. 2007). Such analyses often involve complex processing schemes applied to large image data sets. These schemes typically use neuroimaging software packages such as FreeSurfer (Dale et al. 1999; Fischl et al. 1999), FSL (Fennema-Notestine et al. 2006), SPM (Friston 2007), AFNI (Cox 1996). Multiple neuroimaging pipeline environments (Rex et al. 2003); (Glauche n.d.); (Joshi et al. 2011); (Das et al. 2011); (Marcus et al. 2007) provide workflow tools and algorithm repositories for creation of meta-algorithms within a single processing framework.

The explosion of MRI data availability along with increasing complexity of the associated processing tools and pipelines makes the management of large, multi-scanner and multi-institution datasets difficult. This is especially true when the imaging to be used in research is driven by clinical considerations rather than a fixed experimental design (Hlaiheli et al. 2010). We call such datasets Clinically Acquired Research Sets (CARS) to emphasize the provenance and intended research use. In CARS, uniformity of acquisition protocols may not always be achievable, especially in multi-institutional studies. This heterogeneity can impede large-scale image analyses, including data mining and machine learning techniques employed in the current translational research such as (Aerts et al. 2014). Automatic tools that facilitate quick prototyping of complex workflows, correct selection of processing parameters and quality control of the processed data, are most critical in multispectral imaging.

Since image analysis pipelines designed for controlled research studies are hard to use on CARS directly (see Section 2.2), in this work we consider the requirements of clinical research pipelines and describe a methodology called Heterogeneous Optimization Framework (HOF) for designing and maintaining automated neuroimaging workflows for CARS data. The methodology comprises the following steps. First, a group of academic neuroradiologists, neuro-oncology researchers and MRI experts defines a set of techniques appropriate for the analysis of the imaging data. The same set of experts identifies appropriate tools for computing specific analysis steps (also referred to here as *core algorithms*), e.g., spatial co-registration, computation of physiological parameter maps, etc. HOF further represents knowledge acquired during this process by a relationship workflow model describing images, image attributes and processing tasks.

The key notion in the HOF workflow model is the image role. The image role such as “spatial alignment target in MRI study” formally describes an image characteristic in relation to a specific processing task. Using assigned roles, the workflow model controls data transfer and formatting, regularization, physiological parameter modeling, spatial co-registration and quality control. During the development of an analysis pipeline within HOF, processing results are classified according to quality and, when appropriate, reasons for failure. This classification is used to iteratively refine the formal representation of image role, increase automation, and reduce failure rates.

In this work, we illustrate HOF principles using multimodal glioma analysis (MGA) as an exemplary case. Image data were acquired by the multi-institutional Comprehensive Neuro-oncology Data Repository (CONDR) (Fouke et al. 2013). The modalities acquired in CONDR CARS include diffusion tensor (DTI) and dynamic susceptibility contrast (DSC) as well as pre- and post-Gd T1-weighted and T2-weighted structural images. MGA converts heterogeneous, multispectral MRI study into a set of uniformly structured, spatially co-registered anatomical images and physiological parameter maps suitable for multi-spectral evaluation of specific imaging biomarkers (Kumar et al. 2012). An MGA pipeline plugin to the XNAT imaging informatics platform enables MGA processing of MRI studies archived using XNAT (Marcus et al. 2007).

The remainder of this paper is organized as follows: in sections 2.1–2.2 we describe the CONDR imaging protocol and sources of heterogeneity in its CARS that motivated HOF; in 2.3, MGA core algorithms. Sections 2.4–2.6 develop HOF principles, using MGA as an example; section 3.1 provides technical details on MGA implementation; section 3.2 contains the evaluation of MGA atlas selection strategy; section 3.3 describes HOF software development model and its application to MGA; finally, section 3.4 describes the MGA pipeline module developed for XNAT (Marcus et al. 2007).

2. Methods and materials

2.1 CONDR imaging protocol

The CONDR study (Fouke et al. 2013) was designed to collect and integrate clinical, imaging and tissue-based data from patients with brain tumors, with an initial focus on glioblastoma multiforme (GBM). Imaging was performed at two sites: Swedish Neuroscience Institute (SNI) in Seattle, WA, and Washington University School of Medicine (WUSM) in Saint Louis, MO. The preferred CONDR imaging protocol includes pre-Gd MPRAGE and post-Gd T1 weighted high resolution image, diffusion tensor imaging (DTI), dynamic susceptibility-weighted contrast (DSC) images (also referred to as perfusion weighted imaging sequences), fluid attenuated inversion recovery (FLAIR) and susceptibility-weighted (SWI) sequences (Figure 1). Derived parametric maps include fractional anisotropy (FA) and mean diffusivity (MD) from DTI data, and cerebral blood flow (CBF), cerebral blood volume (CBV) and mean transit time (MTT) from DSC data. The actual protocol (see Figure 1 for images acquired in a representative CONDR patient) varied across scanners and institutions as detailed below.

2.2 Data heterogeneity

Because the CONDR study relied on standard of care imaging, many acquisitions did not match the preferred CONDR protocol, as described in Sections 2.2.1–2.2.4.

2.2.1 Imaging sequence availability

Incomplete acquisition: Many imaging sessions, especially during the initial stages of the CONDR study, omitted one or more sequences from the preferred protocol (Figure 2).

Although adherence improved over time, various aspects of the clinical acquisition workflow (clinician preferences, technological limitations, patient-related issues) prevent perfect matching of clinical acquisitions to the preferred research protocol. Thus, it is necessary to allow for incomplete data in the processing logic.

Incremental protocol modification: Although some sequences were missing, other sequences, such as high-resolution T2-weighted images or T2 BLADE (Figure 1), were acquired in a large fraction of cases, even though they were not described in the initial research protocol. In addition, a number of imaging sequences were upgraded and new sequences incorporated over the first four years of CONDR data collection. Because CONDR was designed to facilitate open-ended research, these sequences emerged as critical research assets; thus, the capability of incorporating new image types in the existing logic became another key requirement of HOF, described in Sections 2.3–2.5.

2.2.2. Variability of acquisition parameters—MRI acquisition parameters, such as voxel size and field of view, also varied from patient to patient. Variability was present not only across sites, but also across patients acquired at the same site. Cross-site variability was unavoidable because of differences in MRI scanners; for example, SNI used a gradient echo T1*-weighted sequence (GRE*, GE scanner platform) for identification of susceptibility artifact from blood products, whereas WUSM used a T2* susceptibility-weighted imaging (SWI) sequence (Siemens scanner platform) for the same purpose. Within-site variability at WUSM also occurred, in part because of a higher number of MRI scanners used for CONDR patients with less centralization as compared with SNI.

Acquisition details critically determine the suitability of a particular image for a specific role. For example, the initial CONDR protocol specified that a post-Gd, high-resolution, T1-weighted sequence should be used for atlas registration. However, experience showed that better results were obtained with non-contrasted enhanced T1-weighted images, even if these were of lower spatial resolution. This experience led to the development of the notion of role in HOF, as discussed in Section 2.5.

2.2.3. Variability of data representation—Although the majority of medical devices are able to store and transfer images in DICOM (dicom.nema.org) format, metadata encoding and storage of newer sequences may differ significantly across scanner manufacturers. For example, the MRI vendor at WUSM used mosaics to encode multi-frame physiological sequences, whereas the vendor at SNI used DICOM multi-frame encoding.

2.2.4. MRI artifacts—A variety of MRI artifacts, such as motion blur and wrap due to the inconsistent selection of field of view, were observed in some CONDR acquisitions. Many of these artifacts could result in data misrepresentation and, therefore, failure of processing. In our experience, many of these artifacts are hard to detect automatically once they escape the quality control procedure (if any) implemented at the clinical imaging workstation.

2.3. Core image processing algorithms of MGA

Core MGA preprocessing includes two major components designed to prepare CONDR CARS for correlation of multispectral MRI and tissue based data: (i) within-subject spatial co-registration of the available structural sequences, and (ii) generation of physiological parameter maps from DTI and DSC images.

Spatial co-registration is achieved by registering a designated study target anatomical image (typically, a high-resolution T1-weighted image) to standard atlas space, and subsequent co-registration of all other modalities to the study target, with output in target and standard spaces. Atlas space output enables standardized cross-sectional analysis and also provides a convenient space for generating quality control images. However, single-subject analyses should use native image space as registration template, since this helps to minimize partial volume error, as discussed in (Aribisala et al. 2011). The common-mode (e.g., T1W→T1W) registration algorithm used by MGA is based on a standard voxel similarity measure (maximization of source: target spatial correlation) (Hajnal et al. 1995). Cross-modal registration uses an algorithm based on alignment of signal intensity gradients (Rowland et al. 2005). Composition of appropriate transforms (e.g., DTI→T2W→T1W→atlas-template) enables resampling any image in register with any other image or in atlas space. All alignment algorithms use 12-parameter affine transform.

Diffusion processing includes motion correction and computation of the diffusion tensor from the raw diffusion measurements. Motion correction iterates through the following steps: 1) aligning each volume to the geometric mean volume of each group of images sharing the same degree of diffusion sensitization; 2) re-computing the geometric mean volume; 3) aligning each group's geometric mean to the first acquired image with $b = 0(I_0)$; 4) algebraically composing transforms between volume/group geometric mean and group/ I_0 . Three cycles of these steps yield realignments with small errors estimated by internal consistency. All transforms are computed with the same algorithm (Rowland et al. 2005). The I_0 volumes are aligned using conventional intensity correlation maximization (Hajnal et al. 1995). The motion-corrected DTI is obtained by applying the cumulative transform and averaging all data sets using cubic spline interpolation. Then, the diffusion tensor is computed from the raw motion-corrected measurements by solving diffusion equation (Basser et al. 1994).

Perfusion processing starts with compensation for asynchronous slice acquisition using sinc interpolation across frames and corrects for intensity errors between even and odd slices. Further, all frames are registered to a reference frame from the middle of the run. Coregistered and masked raw perfusion data is then used to estimate the perfusion model parameters. The perfusion model is based on the Bayesian tissue model (Lee et al. 2010)

where the local arterial input function is estimated iteratively from data rather than determined manually.

2.4. Identifying processing heuristics

The development of a processing stream within HOF starts from manually annotating the processing of a training set of clinical studies. Solutions to accommodate non-standard input are developed on a per-case basis. When the processing becomes stable on the training set, the accumulated solutions are generalized into a set of statements in natural language coding rules for configuring core algorithms for optimal performance. Representative HOF heuristic statements for diffusion weighted image analysis in MGA are shown in Table 1.

HOF heuristics are progressively updated to cover as much of the variability in the data as possible and to simultaneously be as accurate as possible (accuracy is favored over generality). For instance, the following statement in MGA, “For the individual study registration target, an MPRAGE (magnetization prepared rapid gradient echo) image with voxel size $1\times 1\times 1$ mm is to be used,” was replaced with a more general version accounting for variability in voxel size and availability of a (preferably pre-Gd) MPRAGE image. Thus, the rule became, “The study registration target must be a T1-weighted image with the best available spatial resolution, preferably not contrast enhanced.” Heuristics found to be incompatible with the available data are iteratively refined as development progresses. The accumulated heuristics are then formalized and incorporated into procedures for automatically configuring scripts that invoke core algorithms.

2.5. Knowledge representation

Initially, an analysis pipeline is split within HOF into multiple semi-automated pipeline tasks (e.g., see Table 2 for a list of MGA tasks). *Pipeline tasks* in HOF are automatic processing and analysis procedures that typically include input parameters whose values must be assigned based on specific acquisition’s sequence type, resolution and other quality considerations. *Roles* describe images used to perform specific tasks (see below); heuristic rules associate *image types* with roles (see Figure 3).

HOF defines *image type* as an abstract notion of a clinical image with certain characteristics. Each clinical image from CARS is assigned an image type. For instance, MGA assigns the image type, “T1w, pre-contrast, high-resolution”, to all T1-weighted images with spatial resolution higher than $2\times 2\times 2$ mm voxels acquired prior to the DTI sequence. This type is particularly useful as a study registration target, and its assignment is based on the heuristics that Gd administration cannot occur prior to DTI acquisition (for technical reasons). In this case, image type is determined from DICOM metadata, including common (e.g., study and series description), as well as modality-specific tags (e.g., in-plane spacing and slice thickness). Note that with iterative updates to heuristics (Sections 2.4 and 3.3), image types can change scope (e.g. the “high-resolution T1w” can become “T1w”).

Each image type is associated with *image instance quality*, *category* and *metadata* (see Table 3 for MGA examples). Categories are assigned to image types based on one or more properties important for core processing algorithms (e.g., T1 or T2 weighting). Instance quality is understood here as fitness of the given image type for particular pipeline tasks

within the HOF. *Metadata* is a supplementary key-value array containing all other variables required to initialize corresponding tasks.

Finally, HOF defines the notion of task specific *role* which describes an input to that task. At the final input configuration step, each task role is connected to one of available image types (see Table 2 for MGA roles). Each configuration procedure of a pipeline task fills the list of roles defined for this pipeline task, according to logic derived from the available heuristic statements.

2.6. Image-based heuristics

So far, we have described heuristic statements that use DICOM or other textual metadata for automated configuration of core algorithm steps. However, our experience is that in CARS, the key information from these sources can also be error prone or missing (for example, the DTI direction vector table is missing). In such cases, a human expert would rely on reviewing images to make a decision on the processing (in our example, determine the number of DTI directions from the image using the fact that the image with $b=0$ is much brighter than the rest, and use one of pre-saved direction tables known to be used by this scanner). It is sometimes possible to find or develop an ad hoc image analysis tool that would provide necessary information from the image voxels directly. Within HOF, such tools are called *image-based heuristics* (Figure 3). MGA employs the following image based heuristics developed during the refinement of heuristic statements (Section 2.4) using iterative development approach (Section 3.3).

1. Determining the best suitable atlas template—Due to population-specific and sequence-specific variability, a single atlas template does not always provide the best registration target even for T1-weighted brain images. MGA can use one of the four Talairach atlas templates from four different normal populations: Gd-contrast enhanced, young adult, elderly and generic. For each T1-weighted image of good instance quality, MGA computes an atlas transform for each target atlas type and a two-component score vector based on NMI (Normalized Mutual Information, (Maes et al. 2003) and atlas transform error reported by MGA's cross-modal registration algorithm (Rowland et al. 2005). A combination of T1 image and atlas with the best score vector is then selected as the study registration target and target atlas (see Section 3.2 for evaluation of this strategy).

2. Automatic detection of DTI sequence parameters (b_0 and gradient directions)—MGA extracts values for b_0 and gradient directions from vendor-specific DICOM fields using `dcm2nii`, a part of MRICron software (Rorden et al. 2007). However, in cases when the number of gradient directions cannot be determined from DICOM, it is computed using one-dimensional c-means analysis of mean DTI frame intensity, which we found to be a robust feature to classify DTI frames with $b=0$. Then, the last known gradient configuration for the detected number of directions and scanner model is assumed (certain models of Siemens and GE scanners are supported).

3. DTI slice orientation—Another common issue in automatic configuration of DTI processing is the inability to identify from DICOM metadata the correct inferior-superior

head orientation. This information is required to correctly decode Siemens mosaic images. We found that a method similar to atlas detection works well in this case. MGA determines the orientation by comparing registration to a target atlas of z-flipped and unflipped versions of the first frame of the DTI image and then selecting the co-registered image with higher NMI.

3 Implementation and Results

3.1. MGA pipeline implementation

MGA is implemented as a set of Linux shell scripts, C/C++ executables, and uses a small number of third-party utilities such as dcm2nii. Shell scripts implement the image type database that maps DICOM series descriptions to HOF image types, as well as heuristic logic and image-based techniques that govern parameter selection for core processing algorithms. The core and image-based MGA algorithms were developed mostly at Washington University School of Medicine.

Essential provenance information is recorded in a detailed log file and also encoded in file names saved by MGA. For instance, raw image names follow the pattern <study label>_<DICOM series number>_<image type>, and spatially coregistered image names combine raw image or atlas template name, e.g., <raw image 1>_on_<atlas template 1>. Similar naming patterns are used for all stored output, grouped by MRI study. The information on mapping between DICOM images and MGA-generated files is also explicitly stored in variables in a shell script configuration file that can be directly sourced by post-MGA analysis routines.

The user can review MGA runtime parameters in saved configuration files that contain information on heuristic decisions and core algorithm parameters. Error reports can be extracted from execution logs using keyword search, with error and warning keywords being consistent throughout the processing flow. MGA also supports a set of custom variables that can override automatic configuration, such as explicit target T1-weighted scan and atlas template, Siemens mosaic dimensions, etc. Upon completion, MGA generates a PDF-formatted report showing slice mosaics of all co-registered sequences in isotropic 3mm atlas template space overlaid with target atlas contours (Figure 4).

3.2. Evaluation of automatic atlas selection

The ability to run and rerun MGA analysis in a standard framework on multiple CARS makes tasks such as quick core algorithm evaluation and quality assessment feasible. This allowed us, to test the performance of automatic atlas selection algorithm described in Section 2.6. A random sample of 100 MGA-processed CARS was co-registered with fixed atlas (type F processing) and automatically selected atlas (type A processing). An experienced reviewer was presented with three principal sections of atlas-registered T1-weighted brain images overlaid by atlas contours and rectangular grid. The registration quality of images co-registered using types F and A was assigned the score of 0 when outer cortical structures did not deviate from the atlas contours by more than 3 mm, 1 when the deviation did not exceed 6 mm, and 2 when it was above 6 mm. Automated metrics,

including registration algorithm target function and normalized mutual information, were also evaluated for both types. Average registration score for type A was 1.69 and 0.99 for type F, paired t-tests for the difference between type A and type F outcome for all automated metrics and reviewer scores being significant at 0.0001% level. All automated metrics correlated with reviewer scores, with NMI holding the highest correlation coefficient of 0.54 (type F) and 0.46 (type A). Thus, automatic atlas selection clearly improved the average alignment quality, although none of the considered automatic metrics was completely indicative thereof (Figure 5).

3.3. HOF development cycles and MGA data processing

As discussed in Section 2.4, the initial source of pipeline design within HOF is expert and empirical knowledge of core processing algorithms. Figure 6 details the process for refining heuristic statements at each pipeline development cycle.

When the development stage of each cycle is completed, a pilot set of clinical studies is processed with the new version of the pipeline, and the quality of output is assessed by a quality control procedure. This procedure can be more broad or specific depending on intended dataset use. In the case of MGA, where resulting datasets are intended for ROI analysis of coregistered physiological parameter maps, a reviewer estimates closeness of atlas-aligned image to overlaid atlas contours and visual quality of physiological parameter maps using mosaic views (Figure 4). The source of each error is identified, and processing errors that are due to algorithmic issues are compiled into a list, driving the development of the next cycle. A representative example of error analysis for the third HOF development cycle of MGA is shown in Table 5.

We assume that CARS coming from a single source are inherently more homogeneous, so heuristics for each imaging location or institution tend to form repeatable patterns in processing logic. It is therefore important to select the composition of CARS from all available data sources for a pilot set used at each HOF development stage. In the case of MGA, each pilot CARS included SNI and WUSTL studies whenever possible.

The first set of studies processed by MGA included data from 30 preoperative patients from CONDR study. Seven of these were acquired at SNI and 23 at WUSTL. Out of these, 19 patients had a presumptive diagnosis of glioblastoma, one and two had grade III and grade II gliomas, respectively, and the other eight had recurrent primary, metastatic or other types of intracranial tumors. The second set included the first one and added new subjects that were enrolled in the CONDR study over time (see (Fouke et al. 2013) for the CONDR study inclusion criteria). The third and current sets expanded to three other neuro-oncology research projects and include over 160 unique studies of subjects with both primary and secondary (metastatic) brain tumors, with most MRI performed before treatment and a cohort of eight studies acquired after gamma knife radiosurgery. The rate of non-unique studies processed by MGA from 2010 to 2014 is illustrated on Figure 7a. Note that the queued studies number reflects all collected studies presumed to be processed during corresponding time period. Lower number of actually processed studies in the first two years resulted mainly from insufficient analysis automation and lack of dedicated resources, whereas the much smaller difference in the last two years (when the resource dedication did

not change) is due mainly to the studies with key sequences missing or of unacceptable acquisition quality.

To quantify the dynamics of manual effort required to run MGA, we evaluated a list of 11 required or optional HOF pipeline tasks in MGA (Table 6). Each task was assigned an overall manual effort score from 1 to 10, which could change after each development cycle. Based on the development progress reports for each year, we computed the overall manual complexity as a sum of scores per each task per each year. Figure 7b shows manual complexity dynamics from 2011 to 2014 (in 2014, XNAT pipeline was used for assessment rather than the offline version). The error rate due to algorithmic failure also decreased over time. For the last set, about 2% of cases could not be processed due to algorithm or quality limitations (although some cases had to be rerun with manually set parameters for satisfactory results).

3.4. MGA virtualization and XNAT pipeline

In the effort to create a portable and maintainable installation of MGA for community use, we developed an experimental self-contained virtualized installation ('virtual appliance') of MGA in a free open source virtualization environment docker.io (Fink 2014). The Docker environment runs a virtualized image of installation that includes a core Linux operating system with application-specific libraries and packages. The runtime copy of the Docker virtualized image is called a *container*. Docker's 'light' virtualization uses only 64-bit Linux kernel functions of the host and relies on the container file system for MGA scripts, executables and libraries. The docker.io architecture initializes containers and executes a start command within milliseconds, thus making a call to any MGA command as simple as supplying the command string as arguments to a simple wrapper script on the host operating system. The Docker.io client is available for many flavors of 64-bit Linux, MacOS and Windows. The installation package for the docker.io MGA image is available from https://bitbucket.org/mmilch01/mga_docker_install.

Once the MGA code base matured, we implemented a dedicated XNAT pipeline with the goal of further simplifying the runtime configuration and archiving of processed data. The pipeline launch screen (Figure 7) allows the user to select the mapping of DICOM scans to MGA's HOF IDs to be used in processing. The online version of the pipeline maintains an updatable user-specific association table matching series description to HOF image type (HOF ID) and applies the user selection for previously encountered series descriptions. Since series descriptions are rarely changed manually at the point of data origin (MRI scanner, PACS, anonymization software, etc.), and because MRI studies tend to come from a few common sources per research project, the pipeline identifies most scans automatically without formally restricting series descriptions to pre-set values. In addition, manual override variables can also be initialized on the launch screen.

The pipeline uses XNAT features on progress and error tracking: the user can track execution progress on image session web pages, and receives an email link to review the processing PDF report (Figure 4) or console logs in case of failure. MGA configuration and runtime logs are stored within the associated XNAT image session for online access. The experimental XNAT-MGA pipeline module is available to install from <https://>

marketplace.xnat.org/plugin/hof and from the NITRC¹ neuroinformatics tool repository at <http://www.nitrc.org/projects/hof>.

4 Discussion

Although clinical imaging data offer research opportunities, analyzing such data is challenging: adherence to a fixed imaging protocol cannot be expected and the quality of the obtained data is inherently highly variable. In this work, we describe a methodology named Heterogeneous Optimization Framework (HOF) for building an adaptable processing flow capable of automatically configuring analysis tools to accommodate this variability. We also present a concrete implementation of this methodology in the Multispectral Glioma Analysis (MGA) pipeline.

The flexibility of HOF derives from the abstractions of image type and category. The use of these abstractions allows selecting processing patterns for large classes of images. For example, in MGA, images in the category “perfusion map” can be generated from DSC data using an in-house analysis tool. However, equivalent (but quantitatively different) “perfusion maps” may be generated by alternative commercial packages. The abstraction of category provides a formal basis for integration, co-registration and comparison of such derived maps obtained with alternative processing schemes. In addition, the ability to use images within the same type or category interchangeably allows MGA to complete analyses in certain cases of incomplete or missing data.

In the development of expert knowledge representations, it is extremely important to continually refine the heuristic statements used to configure core algorithms and use all possible data sources for pilot processing sets. We found that the validity of heuristic statements depends on the maturity of processing techniques and the availability of expertise. Maintaining heuristics generally becomes more challenging as the scope of automation becomes wider. Refining heuristics with each HOF development cycle and using different data sources allows for correction of failed assumptions concerning relevant imaging parameters for the broad range of expected input data sets.

Since HOF development is governed by incremental processing of pilot CARS and assessing the rate of successful outcomes (as defined by a specific HOF implementation), it may also be beneficial to develop a formal representation of quality control procedure. Although such criteria are to large extent core algorithm specific, a data model similar to HOF knowledge representation (Section 2.5) can provide a theoretical ground in the future for designing such criteria in a general case.

Using HOF methods in MGA substantially accelerated analyses and added flexibility. In particular, the automated MGA processing can be rerun at relatively little expense in situations where manual reprocessing would have required substantial effort. This allowed our team to quickly adapt to new scanner and image processing technology, and also enabled quick customizations of MGA for ongoing research projects. In addition, the outcomes of

¹Neuroimaging Informatics Tool and Resources Clearinghouse

automatic processing were easier to reproduce and compare, whereas manual data processing often required extra normalization steps, e.g., to resolve discrepancies in naming conventions between scanner operators and imaging groups.

The general scheme of research data flow based on HOF in application to MGA is shown on Figure 9. Here, HOF output is viewed as uniformly formatted CARS (UCARS) ready for repeatable analysis. For instance, UCARS generated by MGA were used to evaluate the performance of various perfusion software packages (Milchenko et al. 2014) and machine learning method of tumor segmentation (Prior et al. 2013), as well as in clinical research that employed Kaplan-Meier analysis evaluating the connection of blood flow and diffusion imaging metrics with survival of high grade glioma patients (LaMontagne et al. 2013).

Although the implementation of the HOF structure in MGA certainly improved processing throughput, several factors may limit the general utility of the framework. First, low quality data or incomplete image data sets can prevent successful processing, even in the setting of a strong HOF. At the other extreme, better uniformity and acquisition quality theoretically reduces the need for flexible solutions, although we do not expect that this need can be completely eliminated, given the persistent potential for human error. Additionally, HOF heuristics have the potential to introduce unintended biases because of correlations between heuristics, local imaging protocols and patient demographics. It is therefore important for researchers to review the heuristic statements used by a pipeline built within HOF prior to processing, to ensure that the experimental design is not compromised. In the latter case, manual parameter overrides can be used to overcome such biases.

Apart from general precautions due to the use of heuristics, limitations of core image analysis algorithms used by an analysis pipeline should also be carefully considered, with the most critical components in the MGA example being spatial coregistration, DSC perfusion and DTI modeling. As with many similar techniques, the MGA methods of spatial coregistration tend to perform poorly on highly distorted brains. In CONDR CARS, about 2–3% of images with large primary brain tumors were mis-registered. Regarding DSC modeling, it should be noted that DSC perfusion measures computed by different algorithms or different software packages may not be directly comparable (Orsingher et al. 2014); (Milchenko et al. 2014), and therefore only relative comparisons of measures computed by the same software may be valid.

Many sophisticated software packages for brain imaging are currently available (Gholipour et al. 2007). However, to our knowledge, the HOF methodology is one of the very few attempts to formalize the creation of automatic image analysis pipelines with emphasis on processing clinically acquired research datasets (CARS).

Bibliography

- Aerts HJWL, et al. Decoding tumour phenotype by noninvasive imaging using a quantitative radiomics approach. *Nature communications*. 2014; 5:4006. Available at: <http://www.nature.com/ncomms/2014/140603/ncomms5006/full/ncomms5006.html>.
- Aribisala BS, He J, Blamire AM. Comparative study of standard space and real space analysis of quantitative MR brain data. *Journal of magnetic resonance imaging : JMRI*. 2011; 33(6):1503–9. Available at: <http://www.ncbi.nlm.nih.gov/pubmed/21591021>. [PubMed: 21591021]

- Basser PJ, Mattiello J, LeBihan D. MR diffusion tensor spectroscopy and imaging. *Biophysical journal*. 1994; 66(1):259–67. Available at: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=1275686&tool=pmcentrez&rendertype=abstract>. [PubMed: 8130344]
- Buckner RL, et al. A unified approach for morphometric and functional data analysis in young, old, and demented adults using automated atlas-based head size normalization: reliability and validation against manual measurement of total intracranial volume. *NeuroImage*. 2004; 23(2):724–38. Available at: <http://www.ncbi.nlm.nih.gov/pubmed/15488422>. [PubMed: 15488422]
- Cox RW. AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and biomedical research, an international journal*. 1996; 29(3):162–73. Available at: <http://www.ncbi.nlm.nih.gov/pubmed/8812068>.
- Dale AM, Fischl B, Sereno MI. Cortical surface-based analysis. I. Segmentation and surface reconstruction. *NeuroImage*. 1999; 9(2):179–94. Available at: <http://www.ncbi.nlm.nih.gov/pubmed/9931268>. [PubMed: 9931268]
- Das S, et al. LORIS: a web-based data management system for multi-center studies. *Frontiers in neuroinformatics*. 2011; 5:37. Available at: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3262165&tool=pmcentrez&rendertype=abstract>. [PubMed: 22319489]
- Fennema-Notestine C, et al. Quantitative evaluation of automated skull-stripping methods applied to contemporary and legacy images: effects of diagnosis, bias correction, and slice location. *Human brain mapping*. 2006; 27(2):99–113. Available at: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2408865&tool=pmcentrez&rendertype=abstract>. [PubMed: 15986433]
- Fink, J. Docker : a Software as a Service, Operating System-Level Virtualization Framework; *Code4lib Journal*. 2014. p. 3-5. Available at: <http://journal.code4lib.org/articles/9669>
- Fischl B, Sereno MI, Dale AM. Cortical surface-based analysis. II: Inflation, flattening, and a surface-based coordinate system. *NeuroImage*. 1999; 9(2):195–207. Available at: <http://www.ncbi.nlm.nih.gov/pubmed/9931269>. [PubMed: 9931269]
- Fouke, SJ., et al. [Accessed October 14, 2013] The Comprehensive Neuro-oncology Data Repository (CONDR): A Research Infrastructure to Develop and Validate Imaging Biomarkers. *Neurosurgery*. 2013. Available at: <http://www.ncbi.nlm.nih.gov/pubmed/24089052>
- Friston, K. *Statistical parametric mapping : the analysis of functional brain images*. Amsterdam;;Boston: Elsevier/Academic Press; 2007.
- Gholipour A, et al. Brain functional localization: a survey of image registration techniques. *IEEE transactions on medical imaging*. 2007; 26(4):427–51. Available at: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=4141192. [PubMed: 17427731]
- Glauche, V. [Accessed August 17, 2014] MATLAB Batch System. Available at: <http://sourceforge.net/p/matlabbatch/wiki/Home/>
- Hajnal JV, et al. A registration and interpolation procedure for subvoxel matching of serially acquired MR images. *Journal of computer assisted tomography*. 1995; 19(2):289–96. Available at: <http://www.ncbi.nlm.nih.gov/pubmed/7890857>. [PubMed: 7890857]
- Hlaihel C, et al. Predictive value of multimodality MRI using conventional, perfusion, and spectroscopy MR in anaplastic transformation of low-grade oligodendrogliomas. *Journal of neuro-oncology*. 2010; 97(1):73–80. Available at: <http://www.ncbi.nlm.nih.gov/pubmed/19727562>. [PubMed: 19727562]
- Joshi A, et al. Unified framework for development, deployment and robust testing of neuroimaging algorithms. *Neuroinformatics*. 2011; 9(1):69–84. Available at: <http://www.springerlink.com/content/m763304211482j8r/>. [PubMed: 21249532]
- Kumar V, et al. Radiomics: the process and the challenges. *Magnetic resonance imaging*. 2012; 30(9):1234–48. Available at: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3563280&tool=pmcentrez&rendertype=abstract>. [PubMed: 22898692]
- LaMontagne, P., et al. *Neuro-Oncology*. Oxford University Press; 2013. Reliability of Quantitative Biomarkers of Tumor Progression Based on Multispectral MRI in Glioblastoma Patients; p. iii191-iii205. Available at: <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC3823904/> [Accessed September 17, 2015]
- Law M, et al. Gliomas: predicting time to progression or survival with cerebral blood volume measurements at dynamic susceptibility-weighted contrast-enhanced perfusion MR imaging.

- Radiology. 2008; 247(2):490–8. Available at: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3774106&tool=pmcentrez&rendertype=abstract>. [PubMed: 18349315]
- Lee JJ, et al. Dynamic susceptibility contrast MRI with localized arterial input functions. *Magnetic resonance in medicine : official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine*. 2010; 63(5):1305–14. Available at: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3014609&tool=pmcentrez&rendertype=abstract>.
- Maes F, Vandermeulen D, Suetens P. Medical image registration using mutual information. *Proceedings of the IEEE*. 2003; 91(10):1699–1722. Available at: <http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=1232201>.
- Marcus DS, et al. The Extensible Neuroimaging Archive Toolkit: an informatics platform for managing, exploring, and sharing neuroimaging data. *Neuroinformatics*. 2007; 5(1):11–34. Available at: <http://www.ncbi.nlm.nih.gov/pubmed/17426351>. [PubMed: 17426351]
- Milchenko, MV., et al. [Accessed August 1, 2014] Comparison of Perfusion- and Diffusion-Weighted Imaging Parameters in Brain Tumor Studies Processed Using Different Software Platforms. *Academic Radiology*. 2014. Available at: <http://www.sciencedirect.com/science/article/pii/S1076633214002219>
- Orsingher L, Piccinini S, Crisi G. Differences in dynamic susceptibility contrast MR perfusion maps generated by different methods implemented in commercial software. *Journal of computer assisted tomography*. 2014; 38(5):647–54. Available at: <http://www.ncbi.nlm.nih.gov/pubmed/24879459>. [PubMed: 24879459]
- Prior, FW., et al. Predicting a multi-parametric probability map of active tumor extent using random forests. *Conference proceedings : ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Conference; 2013; 2013*. p. 6478-6481. Available at: <http://www.ncbi.nlm.nih.gov/pubmed/24111225>
- Rex DE, Ma JQ, Toga AW. The LONI Pipeline Processing Environment. *NeuroImage*. 2003; 19(3):1033–1048. Available at: [http://dx.doi.org/10.1016/S1053-8119\(03\)00185-X](http://dx.doi.org/10.1016/S1053-8119(03)00185-X). [PubMed: 12880830]
- Rorden C, Karnath H-O, Bonilha L. Improving lesion-symptom mapping. *Journal of cognitive neuroscience*. 2007; 19(7):1081–8. Available at: <http://www.ncbi.nlm.nih.gov/pubmed/17583985>. [PubMed: 17583985]
- Rowland DJ, et al. Registration of [18F]FDG micropET and small-animal MRI. *Nuclear medicine and biology*. 2005; 32(6):567–72. Available at: <http://www.ncbi.nlm.nih.gov/pubmed/16026703>. [PubMed: 16026703]
- Vrooman HA, et al. Multi-spectral brain tissue segmentation using automatically trained k-Nearest-Neighbor classification. *NeuroImage*. 2007; 37(1):71–81. Available at: <http://www.ncbi.nlm.nih.gov/pubmed/17572111>. [PubMed: 17572111]
- Zonari P, Baraldi P, Crisi G. Multimodal MRI in the characterization of glial neoplasms: the combined role of single-voxel MR spectroscopy, diffusion imaging and echo-planar perfusion imaging. *Neuroradiology*. 2007; 49(10):795–803. Available at: <http://www.ncbi.nlm.nih.gov/pubmed/17619871>. [PubMed: 17619871]




Figure 1. Sequences acquired in a CONDR participant: (a) MPRAGE; (b) high-resolution T2; (c) SWI; (d) perfusion EPI; (e) diffusion weighted EPI; (f) T2 FLAIR; (g) T2 BLADE; (h) high-resolution post-gadolinium T1w. Note that T2w BLADE (Siemens) and high-resolution T2w images were not part of the standard CONDR protocol.

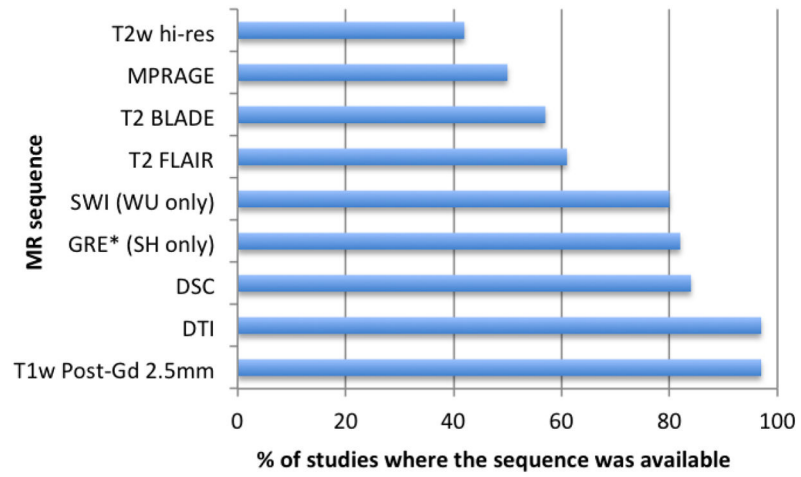


Figure 2.

Availability of selected sequences in the training set of 21 CONDR imaging studies. T2 BLADE (WU only) and GRE* (T1w, GE scanner platform) are examples of sequences not included in initial CONDR protocol but have potential research value.




Figure 3. HOF data model. Circled H indicates the use of heuristics based on textual data; circled I indicates the use of image-based heuristics. The output of HOF is termed Uniform CARS (UCARS).




Figure 4.

A page showing T1 weighted image (grey) overlaid by atlas contours (red) from MGA QC report. QC mosaics allow to check orientation and coarse spatial match to target across the entire imaged volume.




Figure 5. NMI coefficient (Y axis) distribution for different groups of registration quality (X axis) for type A processing. NMI distribution for type F processing was similar.

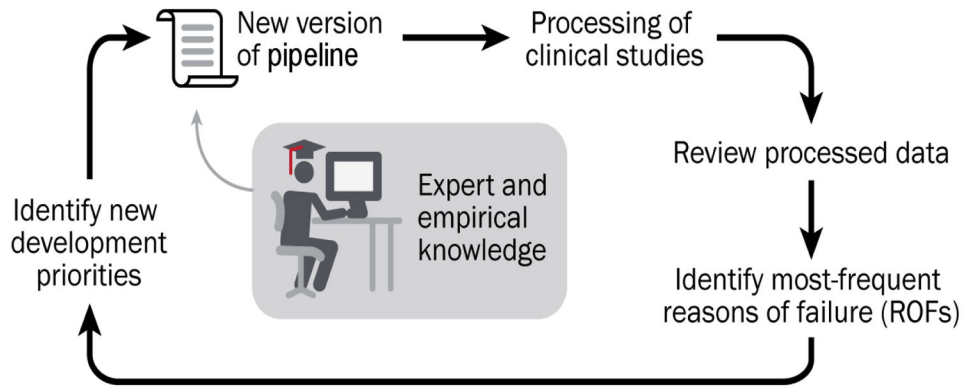


Figure 6. Pipeline development cycle within HOF

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

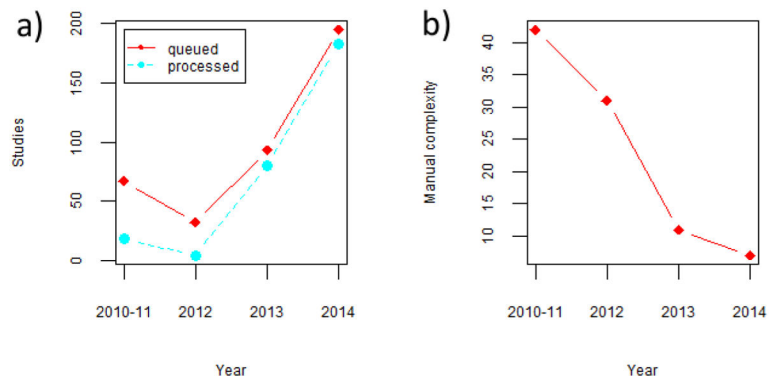


Figure 7.

a) MRI studies of patients with primary and metastatic brain tumors queued for processing by MGA and actually processed (non-unique) each year; b) overall complexity of manual MGA configuration by development year.

Please set the parameters required for the HOF diffusion, perfusion and spatial registration processing of your study.

Session: W001_MR_01122010_preop Project: CONDR

Siemens scanner

Override automatic parameters (advanced):

Scan	Description	Frames	HOF Id
<input type="checkbox"/> 1-MR1	localizer	3	Select
<input type="checkbox"/> 1-OT1	MRI Brain wwoBH	0	Select
<input checked="" type="checkbox"/> 2	TRAN 3D T1 GRE 2MM	88	MPRAGE
<input checked="" type="checkbox"/> 3	TRA SWI	56	SWI
<input checked="" type="checkbox"/> 4	DIFFUSION TENSOR std	13	DWI
<input checked="" type="checkbox"/> 5	DIFFUSION TENSOR std_FA	20	FA
<input type="checkbox"/> 6	DIFFUSION TENSOR std_ColFA	0	Select
<input checked="" type="checkbox"/> 8	TRAN FLAIR T2	26	T2FLAIR
<input type="checkbox"/> 9	OEF_qBOLD_fieldmap	96	Select
<input checked="" type="checkbox"/> 10	TRA PERFUSION	800	DSC
<input checked="" type="checkbox"/> 11	TRA PERFUSION_TTP	16	TTP
<input checked="" type="checkbox"/> 12	TRA PERFUSION_PBP	16	PBP
<input type="checkbox"/> 13	TRA PERFUSION_GBP	0	Select
<input type="checkbox"/> 14	localizer	3	Select

Figure 8.
XNAT configuration screen of MGA pipeline.

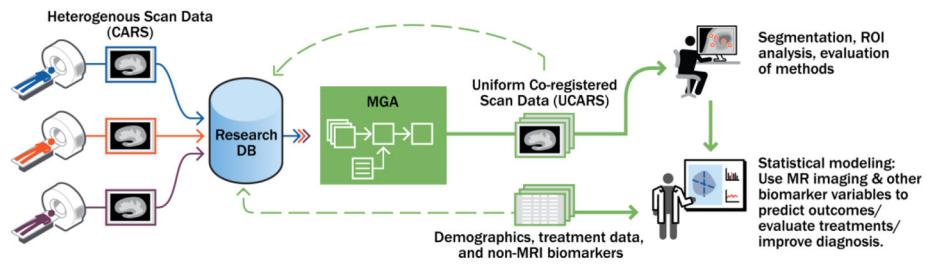


Figure 9.
Data flow in research based on MGA.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 1

HOF heuristic statements for MGA's DTI processing.

The processing requires T1, T2 and DTI sequences to run

Scans of higher spatial resolution typically yield better registration results

The processing requires the encoding of DTI sequences from WUSTL to be specified manually

Direction vectors for DTI sequences from WUSTL should be based on pre-calculated tables rather than values found in DICOM

DTI sequences from WUSTL can have both z orientations

DTI sequences from WUSTL can have more than one "run" (number of images per direction). This number has to be explicitly specified.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 2

Common HOF processing tasks in MGA, with associated roles and metadata.

Task	Task image roles	Metadata
DTI processing	atlas image, T1 image, T2 image, DTI sequence	direction vectors, z orientation, B0
DSC processing	atlas image, T1 image, T2 image, DSC sequence	repetition time, z orientation
Co-registration of original sequences	atlas image, study target, images of all available types	

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 3

Common mappings between HOF image type, category, quality and series description in CONDR CARS.

Image type	Category	Instance quality ²	Sample of mapped DICOM series descriptions
MPRAGE	T1, high-res	excellent	MPRAGE, SAGMPRAGE, SAGT1MPR
TRA_T1_2.5mm	T1, high-res	good	TRAN3DT1GRE2MM, TRA3DGRE2.5MM, TRAGRE2.0
SWI	susceptibility-weighted	good	SWI, TRASWI
DTI	diffusion	poor	TRANDIFFUSION, TRADTI++NOANGLE+++, AXEPIDTIDTPT
T2FLAIR	T2, low-res	adequate	FLAIR, TRAFLAIR, TRANFLAIRT2
T2	T2, high-res	good	TRANTSET2, TRANTSET23MM, TRANST2TSE1X1X2ORSEQ
DSC	perfusion	poor	PERFUSION, PWI, TRAPERFUSION
FA	diffusion derived	poor	N/A
MD	diffusion derived	poor	N/A
CBF	perfusion derived	poor	N/A
CBV	perfusion derived	poor	N/A
MTT	perfusion derived	poor	N/A

²Suitability of image instance for accurate registration.

Table 4

HOF role assignment procedures in MGA

Role	Procedure
Registration target	Pick an image with the “T1, high-resolution” image type of highest instance quality, preceding DTI scan in the series (if available).
Perfusion/diffusion target T1 image	Pick T1 image of highest instance quality.
Perfusion/diffusion target T2 image	Pick T2 image of highest instance quality.
Diffusion image	Pick DTI sequence closest to research protocol specifications (B0=1000, 12 directions), if several are available.
Perfusion image	Pick longest sequence, if several are available.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 5

Cause of Failure analysis on CONDR CARS used in MGA development cycle 3

Identified primary cause of failure	Category	% of all errors
Failed automatic detection of DTI sequence orientation	Algorithmic	40
Poor registration	Registration	13
Perfusion processing did not complete	Algorithmic	10
Perfusion/diffusion sequence parameters are incorrect	Administrative	8.6
High EPI distortion	MR artifact	7.1
Failed atlas registration	Registration	7.1
Missing sequence	Administrative	5.7
Registration of T2 image to atlas failed	Registration	4.3
High distortion of brain anatomy	Registration	2.9
Unable to identify DTI directions	Administrative	1.4

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 6

MGA manual tasks

MGA manual task	Complexity in 2011, 1..10	Year fully automated
Copy/download DICOM study	1	2014
Assign HOF scan types	3	N/A
Diffusion scan z orientation	1	2013
DSC scan z orientation	1	2012
Select target atlas template	3	2013
Select target T1w scan	3	2013
Configure diffusion algorithm	5	2014
Configure DSC algorithm	1	2012
Integrate images from other studies	5	N/A
Quality control/error tracking	10	N/A
Archive/upload processing results	1	2014
Integrate results into other studies	5	N/A

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript