

USE OF DATA PROVENANCE AND THE GRID IN MEDICAL IMAGE ANALYSIS AND DRUG DISCOVERY – AN IXI EXEMPLAR

Kelvin K. Leung¹, Mark Holden¹, Rolf A. Heckemann², Nadeem Saeed³,
Keith J. Brooks³, Jacky B. Buckton⁴, Kumar Changani³, David G. Reid³,
Daniel Rueckert⁵, Joseph V. Hajnal², Derek L.G. Hill¹

¹Division of Imaging Sciences, King's College, London, UK

²Imaging Sciences Department, Imperial College (Hammersmith Hospital Campus), London, UK

³Imaging Centre, ⁴RA Disease Biology, ri-CEDD, GlaxoSmithKline, UK

⁵Department of Computing, Imperial College, London, UK

Abstract

Medical research and drug discovery rely increasingly on comparisons between magnetic resonance images from large numbers of subjects, often with multiple time points for each subject. We applied image registration and visualisation techniques for quantitative and qualitative assessment of changes over time in an experimental model of rheumatoid arthritis. We used the Grid to enable remote invocation and parallel execution of the algorithms. Data provenance was stored in a database to provide information about the validity, accuracy and timeliness of the data, which establishes a data integrity framework that is very important in the process of drug discovery and development.

1. Introduction

Medical imaging is becoming an important component of the drug discovery pipeline, and while not yet on the critical path, substantial investment [1] by the pharmaceutical industry in image acquisition technology suggests that its role is going to increase. Automated analysis of medical image data presents a challenge: it requires high data throughput and frequently employs computationally demanding algorithms. Also, because of the requirement to obtain regulatory approval for new drugs, detailed and reliable documentation of data provenance is very important, resulting in Good Laboratory Practice (GLP) and Good Clinical Practice (GCP). We show the potential of e-Science to assist in image analysis for drug discovery using an exemplar application in an experimental rheumatoid arthritis (RA) model.

The monitoring of image changes over time in RA provides information on the disease progression and on the effect of drug treatment. We apply a segmentation propagation algorithm [2] to identify and delineate two bones (the calcaneus and talus) from the *in vivo* serial magnetic resonance (MR) images of an ankle in an experimental rat model of RA. Markers of

disease progression on the bones, such as erosion, can be studied quantitatively.

Segmentation propagation is a technique that makes use of the deformation field calculated from the registration of two images to propagate the region of interest from one image to another [2;3]. We used a cubic B-spline based non-rigid registration algorithm [4] to register the images. Nonrigid registration is very computationally demanding. When applied to large cohorts, hundreds of registration tasks will take weeks to complete if they are processed sequentially on a single desktop computer. The Grid computing infrastructure [5] allows remote and parallel execution of programs, so multiple image registration tasks can be executed simultaneously.

By connecting different institutions using Grid middleware, a “virtual organisation” is formed, allowing effective and efficient collaboration between those institutions. This approach to collaborative working has great potential in image analysis for drug discovery, which often involves multiple sites within and outside a pharmaceutical company, as well as complex data logistics.

Data provenance is the description of the origins of a piece of data and the process by which it is derived [6]. It provides information about validity, accuracy and timeliness of data.

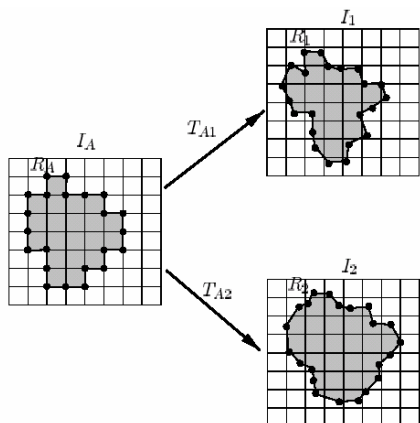


Figure 1 shows the principle of segmentation propagation. I_A represents the atlas (source) image with a segmented structure, R_A , defined by a connected set of boundary points at voxel locations, shown as dots and lines. I_1 and I_2 represent the baseline and follow-up images of a subject. Nonrigid registration of I_A to I_1 and I_2 produces the transformations T_{A1} and T_{A2} . Image I_A is transformed by T_{A1} and T_{A2} into the space of I_1 and I_2 , which results in propagated structures R_1 and R_2 . Because a transformation results generally in translations of boundary points by a non-integer number of voxels, the transformed set of boundary points does not, in general, coincide with the voxel locations of I_1 and I_2 (excerpt from [2], permission granted ©IEEE).

The U.S. Food and Drug Administration (FDA) requires maintenance of an audit trail to protect the “authenticity and integrity” of electronic records as they are created, modified, or deleted [7]. Audit trails can be derived from data provenance. Other example uses of data provenance: 1) users can determine if a certain computation has been performed previously, saving computation time if the required output is already available; 2) users can determine which data need to be recomputed in case minor errors in intermediate steps of previous calculations have been detected [8]. Chimera Virtual Data System [8] was chosen as the provenance tracking system.

In this work, we developed a prototype of the IXI (Information eXtraction from Images) workbench. The simple web interface of the workbench allows users to retrieve medical images from a data grid and to perform segmentation propagation, rigid and nonrigid image registration, and subsequent visualisation.

1.1 Virtual Data System

Virtual Data System (VDS) was designed to ‘enable documentation of data provenance, discovery of available methods and on-demand data generation (so-called “virtual data”)’ [8]. It has been applied successfully to two physics data analysis computations – analysis of data at CERN and in Sloan Digital Sky Survey (SDSS). VDS consists of a virtual data catalog and a virtual data language interpreter. The virtual data catalog is based on a relational virtual data schema, which provides a representation of computational procedures and their invocations. The virtual data language interpreter handles requests for constructing and querying the database entries. Data objects, such as input and output files, are described by logical file names (LFN), which are mapped to physical files via Globus replica catalog (RC) or Globus replica location service (RLS) [9].

Computation procedures and their invocations are described or ‘wrapped’ in virtual data language (VDL). An executable program is described by a transformation (TR) statement. An invocation of a transformation is described by a derivation (DV) statement. Compound transformations can be constructed to include multiple transformations so that workflow can be defined. Queries can be issued to search for transformations and derivations by name, version number, input and/or output filenames.

When a request for a data product (e.g. an output file) is received by VDS, an abstract directed acyclic graph (DAG) describing how to construct the data product will be generated. The abstract DAG refers only to LFN’s and is therefore not bound to any Grid site. In practice, a planner called ‘Planning for Execution in Grids (Pegasus)’ [10], converts an abstract DAG into a concrete DAG and submits it to the Globus universe of Condor® [11]. Simply speaking, it handles the mapping of the LFN to physical file name and also simplifies the DAG if a data object already exists in Globus RC or RLS. Globus Toolkit™ [9] version 2.4 was used in the implementation.

1.2 Nonrigid registration and segmentation propagation

Image registration refers to the ‘spatial alignment’ of two images so that the corresponding features in the two images are matched [12]. The result is a spatial mapping or transformation that transforms positions from one image to positions in another image. An affine registration uses 12 degrees of freedom

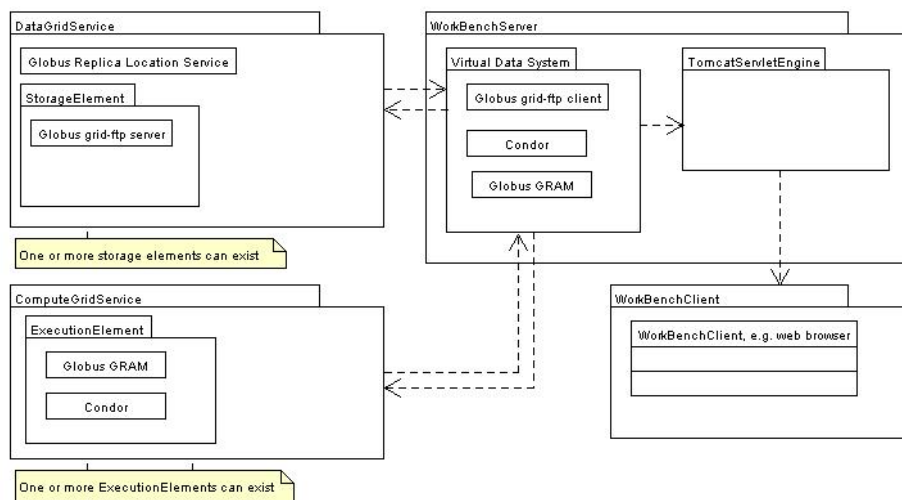


Figure 2 shows the components of our prototype.

(translation, rotation, scaling and skewing in 3-dimensional space) in determining the transformation. In contrast, a nonrigid registration can include hundreds of degrees of freedom. One example is the cubic B-spline based nonrigid registration algorithm [4]. It uses cubic B-spline based free form deformation (FFD) to model local deformations as translations of a regularly spaced grid of control points.

The principle of segmentation propagation is summarised in figure 1. The algorithm described in [2] makes use of a multi-stage registration process. The first stage involved using affine registration to compensate for the global motion and gross differences between source and target images. The second and later stages involved using cubic B-spline based nonrigid registration to compensate for local deformations.

2. Materials and methods

2.1 Prototype

Grid software used in our prototype were Globus Toolkit™ version 2.4, Condor® and Virtual Data System (VDS). A simple web interface was developed on top of VDS using the IXI workbench concept. The web interface was implemented using Java servlet technology and ran on the Apache Tomcat engine [13].

The main web page allowed users to query VDS by the name of transformation and derivation. The web page directed all the queries to VDS and displayed the returned eXtensible Markup Language (XML) results in the web page using eXtensible Style Language

Transformation (XSLT). After selecting the transformation to be invoked, a web form was generated from the XML description of the transformation using XSLT. Then, the form request was directed to VDS. A concrete DAG was generated and submitted to Condor®. When a derivation was displayed in the web page as a result of a query, data objects used in the derivation were hyperlinked to the original derivations that were used to generate the data objects or hyperlinked to the derivations that consumed the data objects as inputs.

Two web pages were developed for querying, uploading and downloading files to and from Globus RLS. A web page was developed for displaying the status of Condor® job queue.

Figure 2 shows the components of the prototype. Globus gatekeeper was run on the ComputeGridServer to receive jobs from VDS. Globus RLS and Gridftp server were run on the DataGridServer for data transfer. Both the ComputeGridServer and the DataGridServer were installed on one computer. The workbench server was installed on another computer with user access via a web browser. A Condor® pool of four computers was set up to process jobs from the ComputeGridServer.

Image registration and surface rendering services were defined using VDL. For this prototype, four services were provided: (1) rigid registration, (2) non-rigid registration, (3) segmentation propagation and (4) surface rendering.

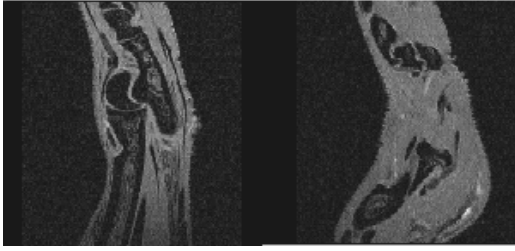


Figure 3 shows transaxial view of an ankle at day -12 (left) and at day +3 (right)

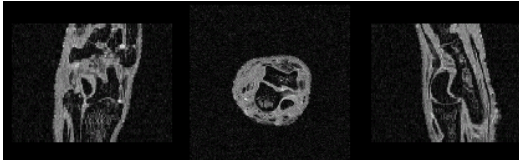


Figure 4 shows the modified atlas (subject 5 at day -12) with the part of the image removed.

2.2 Automatic delineation of bones

2.2.1 Image data

The data set consisted of a group of six Lewis rats (subject 1 – 6; mean age: 60 days at the time of first scan). Animals were housed, maintained, and experiments conducted in accordance with the Home Office Animals (Scientific Procedures) Act 1986, UK. An RA inducing agent (proteoglycan polysaccharide (PGPS) from *Streptococcus pyogenes*) was injected to the right ankle of the rats at day -14. Reactivation to produce joint inflammation was carried out at day 0 by injecting PGPS into the tail vein [14]. MR scans were taken at day -12, day -4, day +3, day +10, day +14 and day +21. The T1-weighted images were acquired on a 7T 20cm bore (Bruker Biospec™) system using a 3D gradient echo sequence with the following parameters: TE = 3ms, TR = 14ms, flip angle = 30°, FOV = 15×40×15mm³, matrix = 256×256×256. This gave voxel sizes of 58.6×156×58.6 μm.

2.2.2 Atlas

As part of a separate manual analysis of the images, the calcaneus and talus bones were manually segmented from the images acquired at day -12 for all the subjects. The segmentation from subject 5 was the best under visual examination and was chosen as the atlas for the automated analysis described here. The other segmentations were used for the subsequent validation experiment. Inter- and intra-observer variability were estimated by getting a different observer to repeat the manual segmentations in subject 1 and subject 2 three more times. Figure

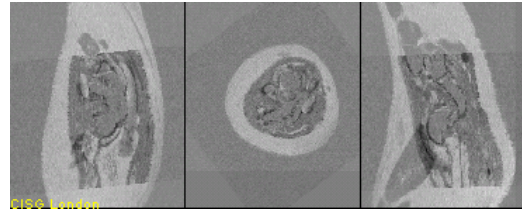


Figure 5 shows difference images of the results of rigid registration to subject 3 at day +3 using the modified atlas (subject 5 at day -12).

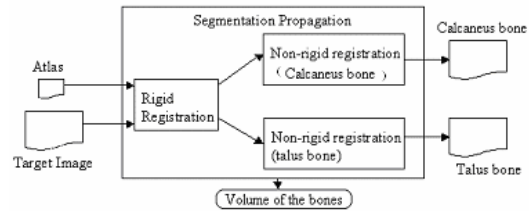


Figure 6 shows workflow of registration.

3 (right view) shows the swelling of the joint following reactivation by intravenous injection of PGPS. To enable rigid registration to align the bones of interest correctly, the original reference image was cropped to restrict the field of view to the region of the bones of interest. Figure 4 shows the modified reference images. Figure 5 shows the difference images after rigid registration to subject 3 at day +3 time point by using the modified atlas.

2.2.3 Registration and visualisation

Due to the different orientation of the atlas and the target image, they were approximately aligned manually to give a rough starting estimate for the rigid registration. Rigid registration was performed to align the atlas and the target image approximately to give a starting estimate for subsequent non-rigid registrations. For each bone of interest, non-rigid registration was used to compensate for local shape changes. Then, labels present in the atlas were transferred to the target image. This allowed the structures to be delineated and their volumes to be calculated.

Holden's algorithm provided a way to speed up the non-rigid registration [2]. It allowed the definition of a region of interest around the bone by dilating the bone outlines from the result of rigid registration. Only voxels within this region were used for the non-rigid registration. Here, the bone outlines were dilated by 5 voxels before non-rigid registration. Control point grid spacing of 0.586 mm (about the width of 10 pixels in the high resolution plane) was used in the B-spline non-rigid registration algorithm.

SearchVDC

Type:

Namespace:

Name:

Version:

Search Results:

type=transformation, namespace=null, name=null, version=null

Transformation	Version	Arguments
kkl:vtkParam	1.0	bins(none) iterations(none) lambda(none) length(none) param(output) res(none) similarity(none) steps(none)
kkl:areg	1.0	do(output) param(input) source(input) target(input)
kkl:nreg	1.0	do(output) dof_in(input) ds(none) param(input) source(input) target(input)
kkl:vtkSegPropParam	1.0	dilate_01(none) dilate_02(none) dofout(none) ds_01(none) ds_02(none) padding(none) param(output) param_in_01(input) param_in_02(input) propagation(none)
kkl:vtkSegProp	1.0	dofin(input) dofout(output) param(input) seg(output) seg_source(input) source(input) target(input)
kkl:lds-areg	1.0	bins(none) do(output) ds(none) iterations(none) lambda(none) length(none) param(input) res(none) similarity(none) source(input) steps(none) target(input) temp_dof(input)
kkl:areg	2.0	do(output) dof_in(input) padding(none) param(input) source(input) target(input)
kkl:vtkSurRender	1.0	seg1(input) seg2(input) wrf(output)
kkl:vtkFindBones	1.0	aregdof(input) bone_areg_param(input) bone_nreg_param(input) cal_dof(output) cal_dof_name(none) cal_param(input) cal_seg(input) calcaneus(output) calcaneus_name(none) dilate(none) dof_in(input) ds(none) mask(none) padding(none) reference_image(input) stage(none) tal_dof(output) tal_dof_name(none) talus(output) talus_name(none) talus_param(input) talus_seg(input) target(input)

Figure 7 shows some of the available transformations in VDS.

Transformation:kkl:vtkFindBones:1.0

Derivation name:

Source Pool:

Destination Pool:

Run Pool:

aregdof (input)

bone_areg_param (input)

bone_nreg_param (input)

cal_dof (output)

cal_dof_name (none)

cal_param (input)

cal_seg (input)

calcaneus (output)

calcaneus_name (none)

dilate (none)

dofin (input)

ds (none)

mask (none)

padding (none)

reference_image (input)

stage (none)

tal_dof (output)

tal_dof_name (none)

talus (output)

talus_name (none)

talus_param (input)

talus_seg (input)

target (input)

Figure 8 shows the form used to invoke the transformation vtkFindBones version 1.0.

Visualization Toolkit (VTK) [15] was used to render the surfaces of the delineated bones in 3D to studying the shape of the bones. Figure 6 summarises the workflow of the whole process.

2.2.4 Intra-subject registration

The registration algorithm was applied to different time points of one subject (subject 5). Change over time was quantified from the

volumes of the calcaneus and talus bones determined by segmentation propagation.

2.2.5 Inter-subject registration

The registration algorithm was also applied to the first time point of subject 1 and subject 2 to propagate the labels of the talus bone. Since four manual segmentations of the first time points of the two subjects were available, the

				<pre> talus_dofin=rarfe3_10_ib1Torarfe3_8_ib1_tal1_dilated.dof talus_name=vtkFindBone06_rarfe3_10_ib1_tal.gipl talus_seg=rarfe3_8_ib1_4_1_talus1_refined1.gipl target=rarfe3_10_ib1_4_1.gipl </pre>
vtkFindBone06_rarfe3_2_ia2	3.0	vtkFindBones		<pre> affine_dilate=2 affine_param=bone_areg_param.txt bone_nreg_param_01=bone_nreg_param_01.txt bone_nreg_param_02=bone_nreg_param_02.txt cal_dof=vtkFindBone06_rarfe3_2_ia2_cal.dof cal_dof_name=vtkFindBone06_rarfe3_2_ia2_cal.dof cal_dofin=rarfe3_2_ialTorarfe3_8_ib1_cal1_dilated.dof cal_seg=rarfe3_8_ib1_4_1_cal1.gipl calcaneus=vtkFindBone06_rarfe3_2_ia2_cal.gipl calcaneus_name=vtkFindBone06_rarfe3_2_ia2_cal.gipl dilate_01=3 dilate_02=3 ds_01=1.172 ds_02=0.586 padding=0 reference_image=rarfe3_8_ib1_4_1.gipl tal_dof=vtkFindBone06_rarfe3_2_ia2_tal.dof tal_dof_name=vtkFindBone06_rarfe3_2_ia2_tal.dof talus=vtkFindBone06_rarfe3_2_ia2_tal.gipl talus_dofin=rarfe3_2_ialTorarfe3_8_ib1_tal1_dilated.dof talus_name=vtkFindBone06_rarfe3_2_ia2_tal.gipl talus_seg=rarfe3_8_ib1_4_1_talus1_refined1.gipl target=rarfe3_2_ia2_4_1.gipl </pre>
vtkFindBone06_rarfe3_3_ial	3.0	vtkFindBones		<pre> affine_dilate=2 affine_param=bone_areg_param.txt bone_nreg_param_01=bone_nreg_param_01.txt bone_nreg_param_02=bone_nreg_param_02.txt cal_dof=vtkFindBone06_rarfe3_3_ial_cal.dof cal_dof_name=vtkFindBone06_rarfe3_3_ial_cal.dof cal_dofin=rarfe3_3_ialTorarfe3_8_ib1_cal1_dilated.dof cal_seg=rarfe3_8_ib1_4_1_cal1.gipl calcaneus=vtkFindBone06_rarfe3_3_ial_cal.gipl calcaneus_name=vtkFindBone06_rarfe3_3_ial_cal.gipl dilate_01=3 dilate_02=3 ds_01=1.172 ds_02=0.586 padding=0 reference_image=rarfe3_8_ib1_4_1.gipl tal_dof=vtkFindBone06_rarfe3_3_ial_tal.dof tal_dof_name=vtkFindBone06_rarfe3_3_ial_tal.dof talus=vtkFindBone06_rarfe3_3_ial_tal.gipl talus_dofin=rarfe3_3_ialTorarfe3_8_ib1_tal1_dilated.dof talus_name=vtkFindBone06_rarfe3_3_ial_tal.gipl talus_seg=rarfe3_8_ib1_4_1_talus1_refined1.gipl target=rarfe3_3_ial_4_1.gipl affine_dilate=2 affine_param=bone_areg_param.txt </pre>

Figure 9 shows some derivations in VDS.

Search Results:
type=derivation, namespace=null, name=null, version=null

Transformation	Version	Arguments
Derivation	Version	Transformation Arguments
boneNregParameters_02	1.0	vtkparam bins=64 iterations=100 lambda=0 length=0.293 param=bone_nreg_param_02.txt res=1 similarity=NMI steps=4

Figure 10 shows the derivation, which was used to generate the file bone_nreg_param_02.txt. The file was used in the two derivations shown in figure 9.

result was compared to them to give an indication of the accuracy of the registration by using similarity index [16], which is the ratio of the intersection of the two segmentations and the union of the two segmentations. It is a measure of segmentation error as well as registration error. Manual segmentations were also compared to each other to give estimates of inter- and intra-observer variability.

3. Results

3.1 Prototype

Figure 7 shows some of the available transformations in VDS. Figure 8 shows the form to invoke the transformation vtkFindBones version 1.0. Figure 9 shows some derivations in VDS. Figure 10 shows the result of clicking on an input file of a derivation. It shows the derivation, which was used to generate the input file “bone_nreg_param_02.txt”.

3.2 Automatic delineation of bones

3.2.1 Intra-subject registration

An example result is shown in Figure 11. The volume of the calcaneus bone and the talus bone are summarised in Table 1. The running time for the rigid registration was about 0.5 hour using step sizes of 1 mm, 0.5 mm, 0.25 mm and 0.125 mm. The running time for the non-rigid registration was about 11 – 12 hours for each bone at a time point. Figure 12 shows an example rendering of the delineated bones.

3.2.2 Inter-subject registration

An example result is shown in Figure 13. The similarity indices including estimates of inter- and intra-observer variability are summarised in Table 2. The similarity index for the automatic segmentation is notably less than the intra-observer variability, but similar to the inter-observer variability.

4. Discussion and conclusions

We integrated Grid middleware and data provenance tracking tools with medical image

processing software to build a prototype system. The system was used to perform automatic delineation of multiple bones in an experimental model of RA. If all the image processing tasks had been performed sequentially on a single desktop computer, the total processing time to segment 12 bones would have been about 132 hours. This would be about 4 times longer than our processing time (about 33 hours). These results suggest that the Grid has the potential to provide easily accessible computing power for processing large amounts of image data. Data provenance of our results was stored and could be queried easily from the web interface. This is particularly important and useful in the drug discovery process because an “audit trail” of electronic records is required.

The volumes of bones shown in Table 1 provide quantitative information about disease progression. Visualisation provides additional qualitative information about the disease progression. Initial results demonstrate the potential of our approach. In ongoing work, we are implementing a better atlas, and refining the parameters of the registration algorithm. Because of the differences between the tasks of inter- and intra-subject registration, similarity indices in Table 2 are not necessarily an indication of the accuracy of the intra-subject registration. Further validation, and application of the technique to a larger number of subjects is in progress to determine the sensitivity of this technique to the disease process.

For this application, Globus Toolkit version 2.4 was stable, with only a few technical problems encountered. Intermittent errors were reported by the Globus gatekeeper when ‘too many’ jobs were submitted to it. Globus RLS sometimes hung and required restarting. The RLI server of Globus RLS sometimes lost a few entries and required an updating command to be issued manually. These problems will be investigated in the near future.

In future, more feedback is required from users to evaluate the ‘value’ of the system in helping them to speed up medical image processing and to keep track of data provenance. This feedback will also help to improve any inadequacy of the system and provide requirements for future development.

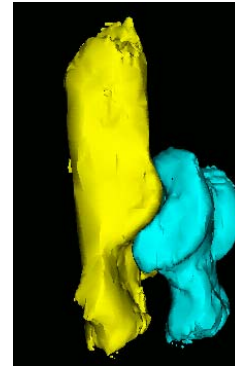


Figure 12 shows the surface rendering of the delineated bones of subject 5 at day +21

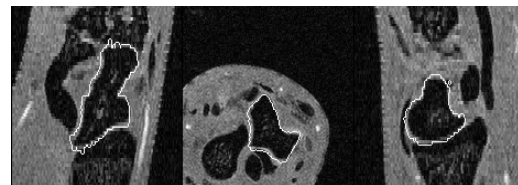


Figure 13. Inter-subject registration: automatic delineation of the talus bone of subject 3 at day -12; the boundaries of the bone from the automatic delineation are shown outlined in white.

	Index A	Index B	Index C
Subject 1	0.9259	0.9421	0.9797
Subject 2	0.9337	0.9452	0.9775

Table 2. Similarity indices (SI) calculated by comparing the manual segmentations and the calculated results for the first time points of subject 1 and subject 2. Index A is the average SI by comparing the calculated result to the manual segmentations. Index B is the average SI by comparing the manual segmentations done by two different observers. Index C is the average SI by comparing the manual segmentations done by the same observer.

ACKNOWLEDGEMENTS

This work is part of the UK e-Science project “Information eXtraction from Images” (IXI), with additional support from GlaxoSmithKline.

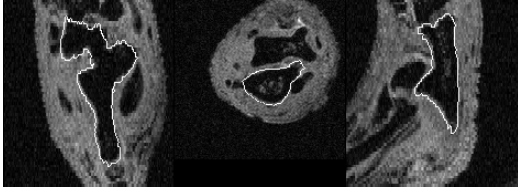


Figure 11. Intra-subject registration: automatic delineation of the calcaneus bone of subject 5 at day +21; the boundaries of the bone from the automatic delineation are outlined in white.

	Calcaneus bone	Talus bone
Day -12	36.59*	18.14*
Day -4	36.31	17.87
Day +3	36.02	19.09
Day +10	36.33	18.16
Day +14	34.29	16.99
Day +21	37.61	18.44

Table 1. Bone volumes of subject 5 in mm^3 .

*Notice that the time point day-12 is used as the atlas. The two numbers are the volume of the manual segmentations of the bones in the atlas.

References

- [1] £76 million research centre to make the UK a global centre of excellence for clinical imaging, <http://www.ic.ac.uk/P4980.htm>; accessed 16-3-2004.
- [2] M. Holden, J. A. Schnabel, and D. L. Hill, "Quantification of small cerebral ventricular volume changes in treated growth hormone patients using nonrigid registration," *IEEE Trans. Med Imaging*, vol. 21, no. 10, pp. 1292-1301, Oct.2002.
- [3] G. Calmon and N. Roberts, "Automatic measurement of changes in brain volume on consecutive 3D MR images by segmentation propagation," *Magn Reson. Imaging*, vol. 18, no. 4, pp. 439-453, May2000.
- [4] D. Rueckert, L. I. Sonoda, C. Hayes, D. L. G. Hill, M. O. Leach, and D. J. Hawkes, "Nonrigid registration using free-form deformations: application to breast MR images," *Medical Imaging, IEEE Transactions on*, vol. 18, no. 8, pp. 712-721, 1999.
- [5] I. Foster, C. Kesselman, and S. Tuecke, "The Anatomy of the Grid: Enabling Scalable Virtual Organizations.," *International J. Supercomputer Applications*, vol. 15, no. 3 2001.
- [6] P. Buneman, S. Khanna, and W. C. Tan, "Why and Where: A Characterization of Data Provenance," *Lecture notes in computer science.*, no. 1973, p. 316, 2001.
- [7] Guidance for Industry on Part 11, Electronic Records; Electronic Signatures--Scope and Application, <http://www.fda.gov/>; accessed 10-11-2003.
- [8] I. Foster, J. Vockler, M. Wilde, and Y. Zhao, "Chimera: A Virtual Data System for Representing, Querying, and Automating Data Derivation," in *Proceedings of 14th International Conference on Scientific and Statistical Database Management 2002*.
- [9] The Globus Alliance, <http://www.globus.org/>; accessed 10-11-2003.
- [10] Planning for Execution in Grids, <http://pegasus.isi.edu/>; accessed 10-11-2003.
- [11] The Condor® Project, <http://www.cs.wisc.edu/condor/>; accessed 10-11-2003.
- [12] D. L. Hill, P. G. Batchelor, M. Holden, and D. J. Hawkes, "Medical image registration. [Review] [126 refs]," *Physics in Medicine & Biology.*, vol. 46, no. 3, pp. R1-45, Mar.2001.
- [13] The Jakarta Site - Apache Tomcat, <http://jakarta.apache.org/tomcat/>; accessed 10-11-2003.
- [14] R. E. Esser, S. A. Stimpson, W. J. Cromartie, and J. H. Schwab, "Reactivation of streptococcal cell wall-induced arthritis by homologous and heterologous cell wall polymers," *Arthritis Rheum.*, vol. 28, no. 12, pp. 1402-1411, Dec.1985.
- [15] The Visualization ToolKit, <http://public.kitware.com/VTK/>; accessed 10-11-2003.
- [16] B. M. Dawant, S. L. Hartmann, J. P. Thirion, F. Maes, D. Vandermeulen, and P. Demaerel, "Automatic 3-D segmentation of internal structures of the head in MR images using a combination of similarity and free-form transformations: Part I, Methodology and validation on normal subjects," *IEEE Trans. Med Imaging*, vol. 18, no. 10, pp. 909-916, Oct.1999.