A Proposed Decision-Support System for (Renal) Cancer Imaging

Stuart Golodetz, Irina Voiculescu and Stephen Cameron
Oxford University Computing Laboratory *

Abstract

Medical practitioners who treat cancer patients have crucial decisions to make about the best course of action for individual patients, but limited time in which to make those decisions, and limited information to help them do so. Our aim is to explore the requirements of a decision-support system which will aid them in this task. This system will be initially targeted at renal (kidney) cancer, but is intended to be extensible to other cancers in the future.

1. Introduction

Recent advances in the ability to sense the inside of the human body have provided medical practitioners ('medics') with a plethora of useful information with which to make their decisions. In particular, techniques such as those based on ultrasound imaging, magnetic resonance imaging (MRI) and computerised tomography (CT) are capable of providing high–resolution images that differentiate tissue density, and so despite the high cost of the equipment such scans are becoming routine (at least in the developed world). These, however, provide only a limited amount of information regarding the 'state' of a particular patient's situation, because only a few scans are taken for each patient, at discrete points in time, and it is difficult to infer how a tumour evolved in the time elapsed between the scans.

There are non-clinical issues involved in the routine use of medical imaging techniques:

- *Segmentation.* Deciding where a particular morphological feature — such as a tumour — begins and ends in an image is not always easy for an experienced human user ('radiologist'), let alone for a computer program. Image processing programs exist that can help radiologists in some cases, but they tend to be specialised to particular features, as the general image interpretation problem is an unsolved problem. Fully–automated segmentation is particularly useful for tasks such as the measurement of tumours, but again it is difficult to achieve reliably: it is a research topic, rather than a mature technnique.

- *Registration.* Practically all imaging systems generate two-dimensional images; even so-called three-dimensional techniques (such as MRI) take a sequence of two-dimensional slices, moving the subject in between recording slices. The process of registration is that of trying to relate points in one two-dimensional slice to neighbouring points in the adjacent slices. If the subject is completely stationary, registration is rarely a problem, but with living (breathing!) subjects it can be. Registration is always an issue with scans taken at different times, with scans taken by different generations of apparatus, as well as with scans taken under different circumstances. (For example, it is a big problem in breast imaging, when a surgeon may wish to reconcile early images taken with the patient standing with pre-surgical images taken with the patient lying down.)

- *Fusion.* Images may often be taken of the same subject using different imaging modalities (e.g. MRI and X-ray), or even variations of one or more techniques. Features will appear different in the subsequent images, and there is a very real issue of trying to fuse sets of images together to form composites, especially given the inevitable registration issues that arise under such circumstances.

- *Visualisation.* Subjects are three-dimensional, but we humans tend to have limited intuition for three-dimensional structure, and there is an obvious problem in seeing interior features in a three-dimensional image. There has been considerable progress in the use of graphics hardware to spin and slice three-dimensional data-sets, and experienced radiologists are capable of using these tools to good effect; nevertheless, there is still scope for more to be done in this area.

- *Ergonomics.* Systems that are not designed with ease-of-use in mind tend to be avoided, sometimes even if

*Authors' address: Oxford University Computing Laboratory, Wolfson Building, Parks Road, Oxford OX1 3QD, UK. stuart.golodetz@comlab.ox.ac.uk

their functionality is otherwise superior to their competitors. When designing computer-based systems for clinical use, therefore, it is important to conduct a careful requirements analysis to ensure that the end result will be as easily usable as possible. This is as yet a relatively untouched area, but even a cursory analysis suggests that it is important to consider not only the design of the software but also its presentation and hardware interface.

- *Data Management.* Finally, there are real practical issues in the storage and transmission of large three-dimensional images, together with legal and ethical issues in ensuring that the appropriate levels of security and privacy are achieved. The indexing of such datasets needs to be prioritised in a way that makes it easy to access the most relevant information first and, possibly, to discard (or at least store away) information that a radiologist is unlikely to want to see again. Such issues are also areas of research.

We are therefore proposing to develop a prototype software system that will be useful for medics specialising in one particular treatment area, in the hope that the solutions found will also be of use in other areas. We propose to apply good practice from computer science in order to help manipulate and manage imaging data, and also to use ideas from geometric modelling and finite-element modelling in order to improve the simulation of cancerous tumour growth. This paper is a discussion paper that outlines our proposal, the work for which is now getting under way.

Whilst the techniques we discuss should be applicable to the imaging of most types of cancers, we will eventually aim to focus on renal (kidney) cancer, which is an area that is relatively unexplored (by computer scientists) whilst, we believe, being amenable to achieving results in the relatively short term. Computational models have been emerging for some of the cancers with highest incidence (of the breast, colorectum and liver) [3]. Renal cancer will evidently have similar traits to the other types of cancer, although it is also distinctive because of the particular functioning of the kidney and its structural make-up.

Approximately half of our planned effort will consider an amalgam of the issues mentioned above and is discussed in §2. In particular, we will rely on our colleagues in the image processing community to help solve the segmentation and registration issues for this problem. (The interior of the kidney is relatively stiff and contains a number of visual cues that are useful for registration. See Figure 1.) The rest of our planned work is in simulating the growth of kidney cancers, using techniques from the computer-aided design (CAD) community. Our proposed solution of this problem is discussed in §3. Finally, we summarise our plans in §4.

Figure 1. Tumour growth on a pair of cancerous kidneys [13]: the tumour is at the top of the image. Easily-identified features of a kidney, such as the renal pelvis, prove useful for registration.

2. Proposed decision support system

In §1 we enumerated the main issues to be considered when developing a decision-support system for medical use. Figure 2 graphs the main interdependencies between them (although not every dependency is shown). Ergonomics, for example, affects mainly the features that are directly visible to the user (namely registration, segmentation, visualisation and data-management). Segmentation is done in conjunction with registration (or as a precursor to it) and is a key step in the process leading to visualisation. It is also necessary when trying to fuse images together.

We have also added an additional theme of *simulation*. As will be discussed later in this section, the idea here is to use the available images to build up a model of tumour growth over time, which can then be used to simulate the state of the tumour at times when images are unavailable (i.e. our model should make predictions about what occurred between scans and what might happen in the future).

Fully-automated segmentation of clinical images is generally an unsolved problem [4, 12]. Whilst good segmentation results can be obtained for given images (when compared to 'gold standard' segmentations which have been manually produced by a radiographer), variations in the images mean that segmentation techniques tend not to work as intended in all cases. Radiographers are highly-skilled individuals with a lot of responsibility, who are rightly suspicious of automated programs which may or may not pro-

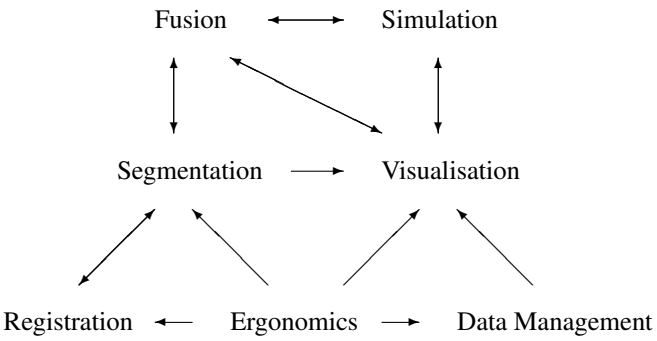

Figure 2. Interdependencies of the research topics

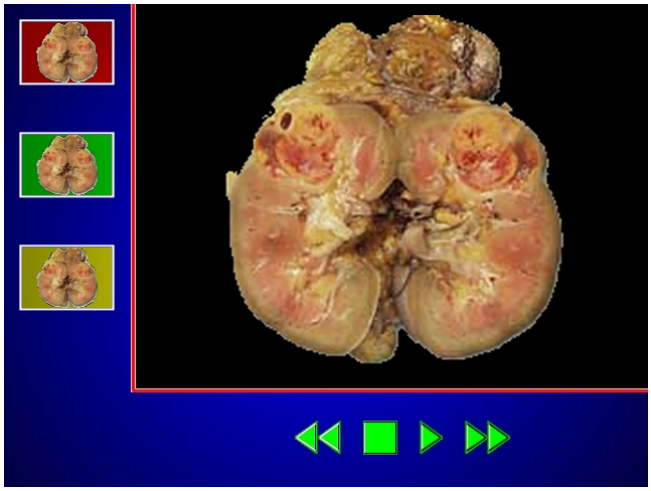

Figure 3. Rough mock-up of an interface (the tumour image is taken from [13])

duce the correct output. Whilst automated segmentation has the potential to save them a great deal of time (manual segmentation can take hours), they should be 'left in the loop' when segmentation occurs, and allowed to visualise what the program is generating and to adjust the output.

In the interest of saving radiologists from having to make similar modifications to large numbers of images, the adjustments should be applicable to many layers in a scan; the process of annotation for the purpose of segmentation should therefore be carried out in a 'segmentation module' and saved. (Segmentation modules are often adjusted by looking at things like gradient for a particular image colour; or defining an initial size for a relaxation method.) Any changes made to such a segmentation module are then propagated automatically to the slices that contain images of the affected organ or tumour 'segments'.

Segmentation has usually been a precursor to registration, as if nothing else the boundaries of the regions for registration need to be well-defined [2, 6, 5, 16], although recent work has focused on performing the two processes concurrently (a good registration of the images can help with segmentation). Three-dimensional visualisation techniques have proved useful here to show how scan layers can be built up into a three-dimensional model [9]. Finally, graphics hardware and three-dimensional algorithms have proved invaluable in allowing the resultant three-dimensional datasets to be moved, sliced, and/or the outermost layers made translucent [15, 14, 18, 19]. (One interesting aspect here is the extent to which this three-dimensional visualisation can now be carried out using off-the-shelf graphics cards, rather than the specialised equipment used previously.)

What we believe is currently weak in this mixture of techniques is (a) the ability to extrapolate from scanned data to simulated data, and (b) the organisation of the data. For point (a), we wish to look to improve mathematical and numerical modelling techniques to be able to guess how a tumour might grow; this will be considered in more detail in §3. For both (a) and (b), we wish to explore the ways in which video-editing and manipulation is now developing, as sketched out in Figure 3.

There are three major components to the interface we have sketched in the figure. The left-hand panel is intended to be an area to organise data: it will act as a 'blackboard' for ideas, images (both real and simulated) and scenarios.

The main window is where tumour visualisation will actually take place. It should allow clinicians to view the tumour from any angle and render slices through it (in any plane) as necessary. It should also allow the interactive segmentation, registration and fusion of images, including the ability to adjust the computer-generated results as required. Finally, it should allow any number of images to be either superimposed (i.e. rendered as semi-transparent layers) or viewed side-by-side simultaneously.

By linking the imaging data to mathematical/geometrical models of kidney function and tumour growth, we also hope to enable users to visualize the progression of a patient's tumour over time. Navigation would be done via the time controls beneath the main window. The idea is to use the available scans to construct a model of the situation, at discrete points in time, that can then be interpolated to give clinicians a better idea of how the tumour is growing at times in between scans being taken. It should be clear from the image context (e.g. background frame colour) whether the user is seeing a real scan or a synthetic image, but otherwise they should act the same within the system.

We believe that this sort of approach will help consolidate the data available to clinicians into a format which is easy for them to use constructively [10, 17], thus helping

them to improve patient care.

Given that a primary goal is to make the system easy for clinicians to use, it is vital that we adopt an enlightened approach to data management. In particular, we envisage a system which is in some way akin to the source control methods found in software engineering. It should keep track of changes to the imaging data (storing incremental differences each time, rather than full copies), allowing users to both backtrack to previous versions and to determine what changes have been made (including when and by whom). It should also support easy compression and decompression of the large data-sets prevalent in medical imaging, and allow users to search for three-dimensional data.

Finally, given the importance of the imaging data to patient welfare, it is important that the system should implement a robust, automated, fool-proof backup process to ensure that the data is safe. It is also necessary to ensure that the data is available when the clinicians need it.

3. Modelling Inhomogeneous Tumour Growth

Tumour growth models have traditionally been expressed using differential equations, which couple growth with the availability of nutrients and competition from other (i.e. healthy) tissue. In traditional tumour modelling these have resulted in equations over a small number of variables; for example, a spherical tumour and the distance of the cells from the centre of the tumour may be described with exactly one variable. Such models have certainly proved useful for tumours which grow in a simple fashion, but are not valid in cases in which the tumour growth is inhomogeneous. Inhomogeneous growth can be due to a variety of conditions within the tumour itself, but is more often due to the inhomogeneous nature of the tissue that the tumour is invading. Whatever the reason, predicting tumour growth in these cases proves an interesting and important challenge.

The analysis of engineering (and other) three-dimensional rigid structures is a research area which also involves complex mathematical models. Complex structures, such as a building or a bridge, are broken down into a (large) number of pieces, each of which can be described by simple sets of differential equations. Essentially, each piece is homogeneous in terms of its properties, and so can be described by simple sets of (coupled) differential equations, together with their (time-varying) boundary values. In turn, the boundary values for the various pieces form a set of simultaneous equations, which eventually couple to the boundary conditions along the outside of the whole region of interest. Furthermore, by triangulating the interior of each homogeneous region — either to triangles in two-dimensions, or to tetrahedra in three dimensions — the solutions of the individual differential equations can be converted into *difference* equations, which can be solved as a large set of *simultaneous* equations across the system. This is, in outline, the technique known as *finite-element modelling* (FEM) [19].

A suitable finite-element triangulation is illustrated in Figure 4 for three regions (in two dimensions). The idea is that significant values pertaining to local properties are stored at the triangle vertices; simultaneous equations are then set up that relate values across adjacent vertices (i.e. vertices which are joined by triangle edges).

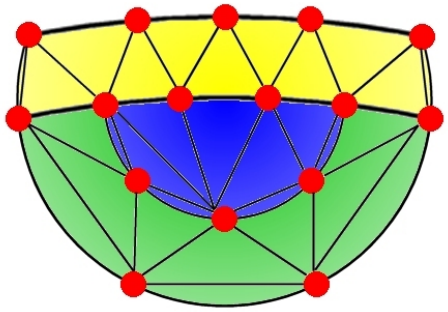

Figure 4. A sketch of a two-dimensional finite-element mesh across three homogeneous regions

This description of FEM over-simplifies a number of practical issues in applying the technique:

- how to express the differential equations in a way that admits a stable numerical solution;
- how to triangulate the individual regions in ways that admit fast and stable solution of the difference equations;
- how to break down the overall domain into (approximately) homogeneous regions anyway; and
- how to balance the speed of the solution — that may involve solving thousands of simultaneous difference equations per time step — against the achieved accuracy.

FEM in engineering has traditionally been applied using regions that form relatively simple geometric shapes, such as spheres, cylinders, polyhedra, etc. Our intention is to start with a set of simple shapes —- such as spheres, ellipsoids and cylinders (for groups of cells, blood vessels, etc.) — and to use simple templates for imprinting triangulations across these, together with additional elements to help tie the values of the varying parameters across the homogeneous region boundaries.

In order to control the execution time of the simulation, we wish to adapt and apply computer science techniques such as *data caching* and *lazy evaluation*. The former will provide a way of storing indices of images that need to be looked at frequently, whereas the latter will delay a computation until the point in time where its result is needed. In doing this, it will be essential to provide a straightforward way of trading the accuracy of the simulation against computation time.

Finally, in order to keep track of the various homogeneous regions, we will use a set of techniques from geometric modelling that have been crafted to deal with the manipulation of such regions. (These techniques exist in the geometric modelling community under the name of 'cellular modelling' [1, 8, 7, 11]. In order to avoid ambiguity we have shied away from this term here; perhaps a suitable linguistic replacement of the term would be 'partition modelling'.)

4. Summary

Medical scanners have the ability to generate vast quantities of patient data, to the extent that medics have trouble keeping up with all of the information. We plan to exploit improvements in image segmentation, registration, fusion and visualisation, and to couple these with best practice from computer science and ergonomics in order to help medics make the most of this information explosion. Furthermore, we plan to incorporate techniques from geometric modelling and analysis in order to be able to interpolate data-sets between scans, and to make predictions about the future growth of renal tumours by extrapolating from the existing images, with the aim of helping clinicians decide on how best to treat the patients under their care. This work is now underway at Oxford; we are hopeful that the results obtained can also be used in treatment for other types of cancer.

References

[1] C. Armstrong, A. Bowyer, S. Cameron, J. Corney, G. Jared, R. Martin, A. Middleditch, M. Sabin, and J. Salmon. *Djinn: A Geometric Interface for Solid Modelling (Specification and Report)*, pages 61–71, 207–227, 355–359. Information Geometers Ltd., December 2000.

[2] P. J. Besl and N. D. McKay. A Method for Registration of 3-D Shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256, 1992.

[3] M. Kelly, O. Noterdaeme, and M. Brady. Relating intra-tumor heterogeneity to morphology and its implications for assessing response to therapy. In *IEEE International Symposium on Biomedical Imaging*, 2007.

[4] D.-T. Lin, C.-C. Lei, and S.-W. Hung. Computer-aided kidney segmentation on abdominal CT images. *IEEE Transactions on Information Technology in Biomedicine*, 10(1):59–65, January 2006.

[5] J. Maintz and M. Viergever. A Survey of Medical Image Registration. Technical Report UU-CS-1998-20, Institute of Information and Computing Sciences, Utrecht University, 1998.

[6] J. Maintz and M. Viergever. An Overview of Medical Image Registration Methods. Technical Report UU-CS-1998-22, Institute of Information and Computing Sciences, Utrecht University, 1998.

[7] A. E. Middleditch and C. M. P. Reade. Theories of Shape and Cell Structure for Computer Aided Design. In *Proceedings: The First Korea-UK Workshop on Geometric Modeling and Computer Graphics 2000*, pages 181–187. Kyung Moon Publications, April 2000.

[8] A. E. Middleditch, C. M. P. Reade, and A. J. Gomes. Set-combinations of the mixed-dimension cellular objects of the Djinn API. *Computer-Aided Design*, 31(11):683–694(12), September 1999.

[9] J. Pitt-Francis, D. Chen, M. Slaymaker, A. Simpson, M. Brady, I. van Leeuwen, F. Reddington, P. Quirke, G. Brown, and D. Gavaghan. Multimodal imaging techniques for the extraction of detailed geometrical and physiological information for use in multi-scale models of colorectal cancer and treatment of individual patients. *Computational and Mathematical Methods in Medicine*, 7(2 & 3):177–188, June 2006.

[10] R. Procter, C. Borgman, G. Bowker, M. Jirotka, G. Olson, C. Pancake, T. Rodden, and M. C. Schraefel. Usability research challenges for cyberinfrastructure and tools. In *CHI '06: CHI '06 extended abstracts on Human factors in computing systems*, pages 1675–1678, New York, NY, USA, 2006. ACM Press.

[11] C. M. P. Reade, A. E. Middleditch, and A. J. Gomes. Least Stratifications and Cell-Structured Objects in Geometric Modelling. *International Journal of Shape Modelling*, 8(1):59–77, 2002.

[12] E. R. Ree. Segmentation of Kidneys from MR-Images. Undergraduate thesis, Norwegian University of Science and Technology, 2005.

[13] `http://resources.schoolscience.co.uk/abpi/cancer/images/kidney.jpg`.

[14] W. J. Schroeder. A Topology Modifying Progressive Decimation Algorithm. In *VIS '97: Proceedings of the 8th conference on Visualization '97*, pages 205–ff., Los Alamitos, CA, USA, 1997. IEEE Computer Society Press.

[15] W. J. Schroeder, J. A. Zarge, and W. E. Lorensen. Decimation of Triangle Meshes. In *SIGGRAPH '92: Proceedings of the 19th annual conference on Computer Graphics and Interactive Techniques*, pages 65–70, New York, NY, USA, 1992. ACM Press.

[16] T. Song, V. S. Lee, H. Rusinek, S. Wong, and A. F. Laine. An Iterative Framework for Registration and Segmentation of Dynamic Three-Dimensional MR Renography. In *Proceedings of the International Society for Magnetic Resonance in Medicine (ISMRM) 14th Scientific Meeting and Exhibition*, page 831, 2006.

[17] J. Ure, J. Geddes, C. Mackay, S. Lloyd, A. Simpson, D. Power, D. Russell, M. Jirotka, M. Katzarova, M. Rossor,

N. Fox, J. Fletcher, D. Hill, K. McLeish, Y. Chen, J. V. Hajnal, S. Lawrie, D. Job, A. McIntosh, J. Wardlaw, P. Sandercock, J. Palmer, D. Perry, R. Procter, M. Hartswood, R. Slack, A. Voss, K. Ho, P. Bath, W. Clarke, and G. Watson. Designing for e-health : Recurring scenarios in developing grid-based medical imaging systems. In *HealthGrid*, June 2006.

[18] Z. Wu and J. M. S. Jr. Multiple material marching cubes algorithm. *International Journal for Numerical Methods in Engineering*, 58(2):189–207, July 2003.

[19] H. Yin, L. Sun, G. Wang, T. Yamada, J. Wang, and M. Vannier. ImageParser: a tool for finite element generation from three-dimensional medical images. *BioMedical Engineering OnLine*, 2004.