

Progress on a Decision-Support System for Abdominal CT Scans

Stuart Golodetz, Irina Voiculescu* and Stephen Cameron
Oxford University Computing Laboratory
(*Work-in-Progress Paper*)

Abstract—Medical diagnosis is often made on the basis of information from visualisation instruments such as CT scanners. This information is currently used by medics only and, in most cases, browsed manually. This paper presents technical issues which are prerequisite to decision-support systems which analyse such data and provide new ways to present the data so as to make it most useful to clinicians, taking the domain of renal (kidney) cancer care as an example. In the long term such systems can also be used to improve communication between medics and patient, and may also enable communication with remote consultants.

Index Terms—Decision support systems, medical informatics, medical imaging, abdominal CT, renal cancer

I. INTRODUCTION

Although kidney cancer accounts for only 3% of all cancers [1], its survival rate continues to be fairly low. It is around 50% at one year and 75% at five years for early stage kidney cancers; worse still, the five-year survival rate is only around 10% for patients with metastases [2], [3].

It is more commonly the case (at least in the developed world) that kidney cancer is being diagnosed at smaller sizes than before because of the recent technological progress in the use of imaging techniques [4]. A vast quantity of information is being generated in this way, creating a pressing need for better data analysis methods, as well as for new ways to present the data to make it more useful.

Our medium-term goal is to provide clinicians with a tool which would help them decide whether to operate on a patient immediately or not. For kidney cancer sufferers this decision has strong quality-of-life implications as lifelong dialysis is required by patients who have lost their kidney function. In some cases it is better *not* to operate, but rather to keep such patients under active surveillance. (A longer-term goal will adapt the techniques to other forms of cancer.)

II. PROPOSED DECISION SUPPORT SYSTEM

We have been working on the requirements of a decision-support system (DSS) [5] which would help medics in the complex actions they need to take when treating kidney cancer sufferers, where the main diagnostic tool is CT (computerised tomography) images. We can summarise the main activities around the DSS in three use cases, namely *diagnosis*, *treatment*, and *follow-up*. The possible presence of kidney cancer is normally suggested by clinical problems, at which point a consultant will send the patient for a CT

scan. The radiographer will normally spot a potential tumour, and may take steps to emphasise it on the image set. The scans are subsequently viewed by a series of consultants involved in the diagnosis. Treatment may involve surgical removal of the tumour, or of the whole kidney; or the tumour may be left to see if it grows; or the tumour may be attacked via chemotherapy, immunotherapy, radiotherapy or high intensity focused ultrasound (HIFU). In the latter cases especially it is vital to keep track of the tumour, and even if a tumour is removed, follow-up scans will be scheduled to look for any sign of the cancer metastasising. For any patient who is diagnosed with cancer, therefore, there is a need to be able to catalogue sets of scans, and to be able to compare scans.

A. Segmentation

Fully-automated segmentation of clinical images is generally an unsolved problem [6], [7], [8], and is expected to remain so for the foreseeable future. Radiographers are highly-skilled individuals with a lot of responsibility, who are rightly suspicious of automated programs which may or may not produce the correct output. Whilst automated segmentation has the potential to save them a great deal of time (manual segmentation can take hours), ideally they should be able to input information interactively into the segmentation process: in other words they should be able to visualise what the program is generating and to adjust the output. This kind of interactive process is still somewhere down the line, since the segmentation information needs to be presented by the DSS in a way which is easy to understand and manipulate. The feature segmentation and identification process is necessary for a variety of applications such as 3D visualization [9] and volume estimation [10].

B. 3D visualisation

It is possible to achieve 3D visualisation by applying a marching cubes-style [9] algorithm (M3C). However, the input to this algorithm relies on individual features in each CT slice having been labelled. Hence organ identification is a central precondition to the well-functioning of a decision support system of the kind we propose.

Our DSS will allow clinicians to view a tumour and any adjacent organs or blood vessels from any chosen angle (see, for example, Figure 1), and will render slices through it (in any plane) as necessary. If the automatically generated segmentation is not satisfactory, the system may allow for

*irina.voiculescu@comlab.ox.ac.uk or Oxford University Computing Laboratory, Parks Road, Oxford OX1 3QD, UK.



Fig. 1. Visualisation of CT abdominal scan data set using M3C

interactive segmentation, including the ability to adjust the computer-generated results as required.

We hope to enable users to visualize the progression of a patient’s tumour (or kidney, in the case of partial nephrectomy) over time. This can be achieved by linking imaging data taken during the patient’s surveillance to geometrical models of kidney function or tumour growth. 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. Thus the DSS would ultimately provide ‘navigation’ in time, giving clinicians a way to see any changes in volume and shape.

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 [11], [12], thus helping them to improve patient care.

C. Volume calculation

Segmentation and feature identification are also key steps in the process of volume calculation. Clearly, once the measured feature has been identified on each of the slices of the scan, the volume calculation is straightforward. For example, to calculate the volume of a labelled tumour, we can use Formula (1). This sums the volumes contributed by each slice to calculate an overall volume for the tumour. The contribution from slice i is calculated as the number of pixels p in the slice (with pixel set P_i) whose label $\ell(p) = \text{FEATURE}$ (which can be chosen to be tumour, kidney, blood vessel, etc.), multiplied by the thickness th_i of the slice, calculated as the perpendicular distance between adjacent slices. In practice, we use a higher-order method (Simpson’s rule), which still only requires a linear scan of the data.

$$\sum_i th_i \left(\sum_{p \in P_i} (\ell(p) = \text{FEATURE} ? 1 : 0) \right) \quad (1)$$

The key issue here is identifying which pixels belong to the relevant feature (e.g. a tumour). Quite apart from the usual problem of identifying a tumour, the subsequent treatment may make it harder to identify; for example, if HIFU is used on it.

III. FEATURE IDENTIFICATION

We have already explained that feature identification plays an important role in the analysis of medical data. It is not an




Fig. 2. Image issues which hinder automatic segmentation

easy task to specify, let alone carry out, because sections through features can vary significantly from one slice to the next. Figures 2(a) and 2(b) show how (in a series with slices 5mm apart) consecutive sections through a vertebra have different shape. Figure 2(c) illustrates another reason why features are difficult to identify, even by experienced radiographers: the precise boundaries between features often appear unclear in the images.

A. Prerequisites

The Hounsfield Units (HU) produced in the CT scanning process are associated with each of the pixels generated in the data, and are proportional to the density of the tissue at that point. The images are displayed by any data viewers in a greyscale range, because this is what monitors can display and what the human eye can discern. This conversion causes some HU to be mapped to the same greyscale value, but nevertheless the original density is retained in the data.

So far we have been able to identify some features in the image automatically, although some continue to be labelled by hand. For the purpose of automatic feature identification we have used a series of techniques. Some of them use merely the greyscale values in the image, whereas some are based on the HU information. Unsurprisingly, these techniques appear to work well for specific parts of the body but not so well for other parts.

B. Bayesian approach

Our initial approach to feature identification has been to use a Bayesian classifier to quantify the extent to which a region resembles a feature of interest, based on some subset of the generated properties for that region. Ribs were chosen as the initial target of our algorithm because their high grey levels make them one of the easier features to identify on an image. The resulting process is fully automatic and works reasonably well. Figure 4(a) illustrates how it can pick out ribs from an abdominal scan.

The classifier was designed in terms of four properties of the region: its area, its elongatedness, the maximum grey value and the mean grey value in the region. The calculation is explained fully in our paper [13].

C. Statistical analysis

Since the HUs generated by the CT scanner are proportional to the density of the tissue, it is possible to use this




Fig. 3. Use of statistical analysis on images

information for a broad classification of the features present in the image. Typical HU tables for particular types of tissue are widely available. These tables make it easy to tell apart bone from soft tissue, and muscle from fat. What it is not possible to do easily is to distinguish the different organs from each other. For example, the respective HU ranges for liver and kidney tissue overlap significantly, as shown by the histogram in Figure 3(a).

In order to narrow down further the categorisation of these types of tissue it is possible to carry out some careful statistical analysis of the data. A pre-processing stage is used whereby one or more sample regions of a particular kind of feature (e.g. kidney) are labelled manually. The algorithm uses that information in order to label other regions in the image. In Figure 3(b) we see that it is not sufficient to apply this kind of statistical analysis on its own. Whilst the results are reasonably good, they are in need of post processing in order to label all the features correctly. The major advantage of this technique, though, is that it can identify more than one type of tissue in the same pass through the region.

D. Flooding approach

We have recently begun work on an alternative flooding approach to feature identification. This is essentially a modified region-growing process. It is based on the region adjacency graph of the image which, in turn, is constructed using a waterfall algorithm [14].

Instead of doing region-growing on a per-pixel basis, we do it per-region. The results are then post-processed. We start by automatically picking a number of potential region seeds for the feature of interest based on HU, area and location, then flood outwards. The feature is ultimately identified as a collection of sub-features which entirely cover its area. The way in which location information can be used will be explained in §III-E. The algorithm grows regions according to a combination of criteria, and then post-processes the final regions to remove any artifacts. We have used this method to identify vertebrae on a significant number of very different image series with good results, and plan to make further use of this in our automatic localization work. Some preliminary results are shown in Figure 4(b). Note that the higher grey area inside the boundary of the vertebra is correctly identified as part of the vertebra (hence the double thickness border around it), whereas the lower grey area (the spinal cord) is correctly ignored.




Fig. 4. Examples of features automatically identified correctly

E. Other techniques

As explained in §III-D, it is sometimes useful to use the location of a feature in the image, as usually most of the organs are in the expected places so if the region is well away from the expected location for that type of tissue it can be discarded. (However tumours may push organs out of position.) One possible alternative is to move the origin of the coordinate system to a fixed position on the spine (or the vertebra featured in each image). One paper which proposes this scheme [15] uses it to produce interesting results.

There is a category of patients for whom such location methods are not applicable, namely those who have had an auto-transplant. In this case, their auto-transplanted kidney may be located towards the ventral side of the abdomen.

F. Use of the labelling techniques

In our work on segmentation [5], [13] we show how a segmented image can be stored as a collection of partition hierarchies. This technique summarises the information in the image in a way which allows for efficient searches. It is thus possible to use certain criteria in the search for features (as regions in the image). This data structure is central to all of our feature identification methods because the properties of each type of feature can be described in terms of certain parameters prior to any search or labelling.

IV. RESULTS

A. Feature identification

We tested the three feature identification techniques presented in §III-B, §III-C and §III-D on a number of differing image series. The first method, a Bayesian classifier approach, yielded good results for rib identification, also reported in [13]. (Other classifiers may also be used.)

Our second method used more sophisticated statistical analysis techniques to classify regions by texture, but, as expected, was incapable of dealing with the problem of regions sharing a similar texture but representing different features in the image.

Our latest method seeks to address this latter problem by incorporating localization information via the medium of a region-growing process. We have tested this approach on a wide variety of image series and achieved good automatic identification results for vertebrae.



Fig. 5. Visualisation of the kidney after partial nephrectomy

B. Volume calculation

We used our volume calculator to calculate remaining kidney volume on around 20 data sets as part of a study correlating loss of renal volume with loss of renal function in patients who had undergone a partial nephrectomy. Figure 5 shows a visualization of one of the kidneys measured: the volume of the kidney was estimated to be 253cm^3 .

Our results were verified by the surgeons engaged in the study. We are currently working on further ways of validating our results, potentially using phantoms.

V. CONCLUSIONS AND FUTURE PLANS

We have outlined the issues in producing a DSS for clinicians using abdominal scans. Rather than trying to fully-automate the process, we are looking towards a system that will integrate in with the way that health professionals work, and with their need to keep control. In discussion with our clinical partners we think that the DSS proposed will be acceptable, and we are currently focusing on some of the technological challenges, which primarily revolve around the segmentation and the activities down-stream from that.

The technical challenges are non-trivial. The above techniques, although they work well, rely on region adjacency information. This is stored in a region adjacency graph. If the image is not over-segmented, this information is easy to use. Segmentation results can be far from ideal, though, hence it being such a ‘hot’ research topic.

Ideally, the statistical analysis of the HUs should be used in conjunction with some of the other techniques described, and possibly also in conjunction with region adjacency information. We intend to perfect these combinations in the near future.

We also intend to devise a location method based on a coordinate system centred on the spine. This strategy would make it possible to use symmetry information more easily.

Once the feature identification can work reliably, it can be applied for a variety of uses. Figure 6, based on manual labelling, shows how accurate 3D reconstruction can be.

VI. ACKNOWLEDGEMENTS

We would like to thank medics and technicians from the Churchill Hospital, Oxford, for their invaluable help and support: Zoë Traill Consultant Radiologist, David Cranston Consultant Urological Surgeon, Mark Sullivan Consultant Urological Surgeon, Andrew Protheroe Consultant Medical




Fig. 6. Use of manual feature identification for visualisation

Oncologist, Anthony McIntyre Superintendent Radiographer, Nilay Patel Specialist Registrar, Rob Ritchie Academic Clinical Fellow in Urology. We are also grateful to Clarice Poon who carried out some of the statistical analysis.

REFERENCES

- [1] “Office for national statistics, registrations of cancer diagnosed in 2005,” England 2008.
- [2] M. P. Coleman, B. Rachet, L. M. Woods, E. Mitry, M. Riga, N. Cooper, M. J. Quinn, H. Brenner, and J. Estève, “Trends and socioeconomic inequalities in cancer survival in England and Wales up to 2001,” *British Journal of Cancer*, vol. 90, no. 7, pp. 1367–73, 2004.
- [3] SEER., Cancer Statistics Review 1975-2004, 2007.
- [4] B. Ljungberg, D. Hanbury, M. Kuczyk, A. Merseburger, P. Mulders, J.-J. Patard, and I. Sinescu, “Guidelines on renal cell carcinoma,” European Association of Urology, 2007.
- [5] S. Golodetz, I. Voiculescu, and S. Cameron, “A proposed decision-support system for (renal) cancer imaging,” in *Proceedings of FBIT 2007*, October 2007. [Online]. Available: <http://dphil.gxstudios.net/golodetz-RenalCancerDecisionSupport.pdf>
- [6] 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*, vol. 10, no. 1, pp. 59–65, January 2006.
- [7] E. R. Ree, “Segmentation of Kidneys from MR-Images,” Undergraduate thesis, Norwegian University of Science and Technology, 2005.
- [8] D. L. Pham, C. Xu, and J. L. Prince, “Current methods in medical image segmentation,” *Annual Review of Biomedical Engineering*, vol. 2, no. 1, pp. 315–337, 2000. [Online]. Available: <http://dx.doi.org/10.1146/annurev.bioeng.2.1.315>
- [9] Z. Wu and J. M. S. Jr., “Multiple material marching cubes algorithm,” *International Journal for Numerical Methods in Engineering*, vol. 58, no. 2, pp. 189–207, July 2003.
- [10] R. Lu, P. Marziliano, and C. H. Thng, “Liver tumor volume estimation by semi-automatic segmentation method,” in *IEEE-EMBS Engineering in Medicine and Biology Society, 27th Annual International Conference, 2005*, pp. 3296 – 3299.
- [11] R. Procter, C. Borgman, G. Bowker, M. Jirotko, 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*. New York, NY, USA: ACM Press, 2006, pp. 1675–1678.
- [12] J. Ure, J. Geddes, C. Mackay, S. Lloyd, A. Simpson, D. Power, D. Russell, M. Jirotko, 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.
- [13] S. Golodetz, I. Voiculescu, and S. Cameron, “Region Analysis of Abdominal CT Scans using Image Partition Forests,” in *Proceedings of CSTST 2008*, October 2008, pp. 432–7. [Online]. Available: <http://dphil.gxstudios.net/golodetz-CIMI08.pdf>
- [14] B. Marcotegui and S. Beucher, “Fast Implementation of Waterfall Based on Graphs,” in *Mathematical Morphology: 40 Years On*. Springer Netherlands, 2005.
- [15] H. Fujimoto, L. Gu, and T. Kaneko, “Recognition of abdominal organs using 3d mathematical morphology,” *Systems and Computers in Japan*, vol. 33, pp. 75–83, 2002. [Online]. Available: <http://www3.interscience.wiley.com/search/allsearch?mode=quicksearch&WISindexid1=WISall&WISsearch1=93520774>