3D Visualization of Abdominal CT Scans using Morphological Segmentation Techniques

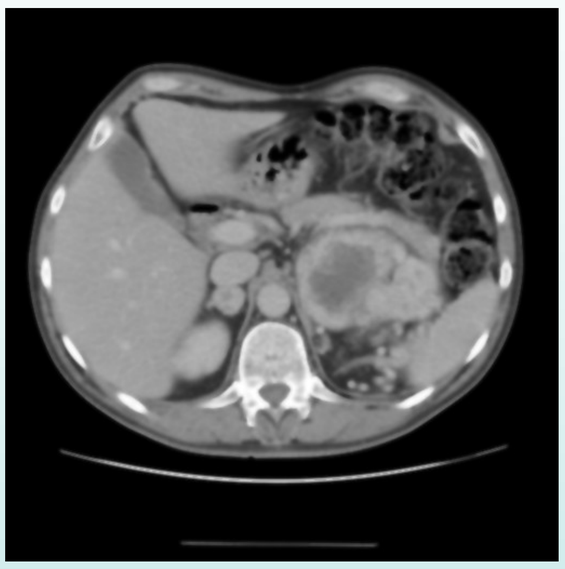

Motivation and Background

Despite major and prolonged research efforts round the world, cancer remains one of the leading causes of human mortality in the 21st century. Great progress has been made in treating certain cancers (e.g. breast cancer), partly as a result of extensive screening programmes and improvements in imaging techniques allowing the tumour to be caught early before it has metastasized to other parts of the body. However, the prognosis for other types of cancer which are harder to detect early (e.g. kidney cancer) is more bleak.

Whilst dramatic improvements to this situation are most likely to come about via the development (and, as we have seen recently, funding) of new drugs and surgical techniques within the medical community, and the use of widespread screening to catch tumours early, there are nonetheless ways in which computer scientists and mathematicians can make a valuable contribution, most notably through working with doctors to analyse medical images and produce tumour models. This helps doctors to make the best use of the images they have available, allowing them to visualize the state of a tumour more clearly before deciding how best to treat it.

To date, our own imaging work has focused on segmenting abdominal CT scans and identifying key features (e.g. the kidneys, the liver, any tumour, etc.) within them. The results of this process have applications in volume visualization (as shown here), volume estimation and landmark-based intra-series registration. Volume estimation is particularly interesting from a prognostic point of view, because it allows us to quantify the percentage of a tumour which is necrotic (i.e. the percentage of a tumour which consists of dead cells): this value has been postulated to be an indicator of the rate of tumour growth, and thus of tumour aggressiveness and patient outcome.

Volume visualization is also interesting, because it allows doctors to see the state of a tumour in 3D, aiding them in the treatment decision process.

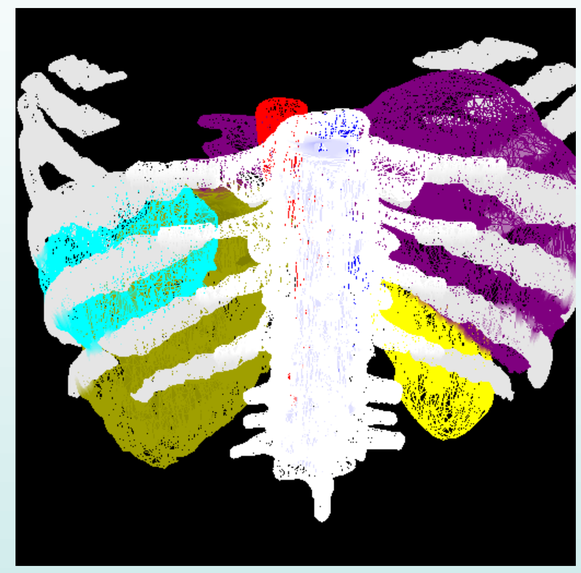

Segmentation and Feature Identification

Segmentation is the process by which we try to partition an image into regions which correspond to interesting, or salient, features therein, whereas feature identification involves assigning semantic meaning to the regions generated.

Property Tree Generation

Our approach uses the morphological watershed and waterfall algorithms [Meijster98, Marcotegui05] to generate a hierarchy of partitions of the image, which we represent using a tree structure (or, more precisely, a forest). A set of useful properties are then calculated for each region in the tree (in practice, we calculate the properties for the lowest layer of the tree and propagate the results upwards), resulting in a *property tree* which can be fed into the feature identification process.

Feature Identification

We currently identify features in the property trees interactively using a mouse, but the long-term goal is to identify features as automatically as possible. We have had some preliminary success in using Bayesian classifiers to identify ribs in the images, and plan to continue working on this in the near future.

The actual process of identifying a region as a particular feature requires structural changes to the property tree and we have developed an algorithm for this, as well as algorithms to dynamically rearrange the partitions within the tree.

Volume Estimation and Necrosis Quantification

The output of the feature identification process is a set of labelled slices, which together form a labelling of the CT volume. (Experiments by our MSc student, Qingnan Zhang, have shown that the slices are already very well registered, so no additional work is necessary to combine the slices into a volume.) This can be used for volume visualization (see right), but is also immediately useful for volume estimation. As mentioned above, there is value in being able to estimate both the total volume of a tumour and the volume of necrotic cells within a tumour. In particular, this allows us to calculate the percentage of a tumour which is necrotic as:

% necrosis = volume of necrotic cells / volume of tumour (incl. necrotic cells)

By determining the percentage of a tumour which is necrotic, doctors can gain an insight into the rate of the growth of a tumour, since one potential cause of tumour cell death is the rapid proliferation of a tumour cutting off the blood supply to the cells within it. A relatively high percentage of necrosis is thus hypothesised to correlate with a relatively high rate of tumour growth. This has implications for treatment, since slowly-growing tumours can sometimes be monitored within the body, whereas rapidly-growing tumours will generally need to be resected (surgically removed) sooner rather than later.

Given a labelled CT volume, the actual calculation of the relevant volumes is fairly straightforward. Some care does need to be taken, however, since adjacent CT slices can overlap (e.g. each slice can represent a volume 5mm thick, but the centres of adjacent slices may be only 2.5mm apart). Our EPSRC Vacation Bursary scholar, Clarice Poon, has recently tackled this task.

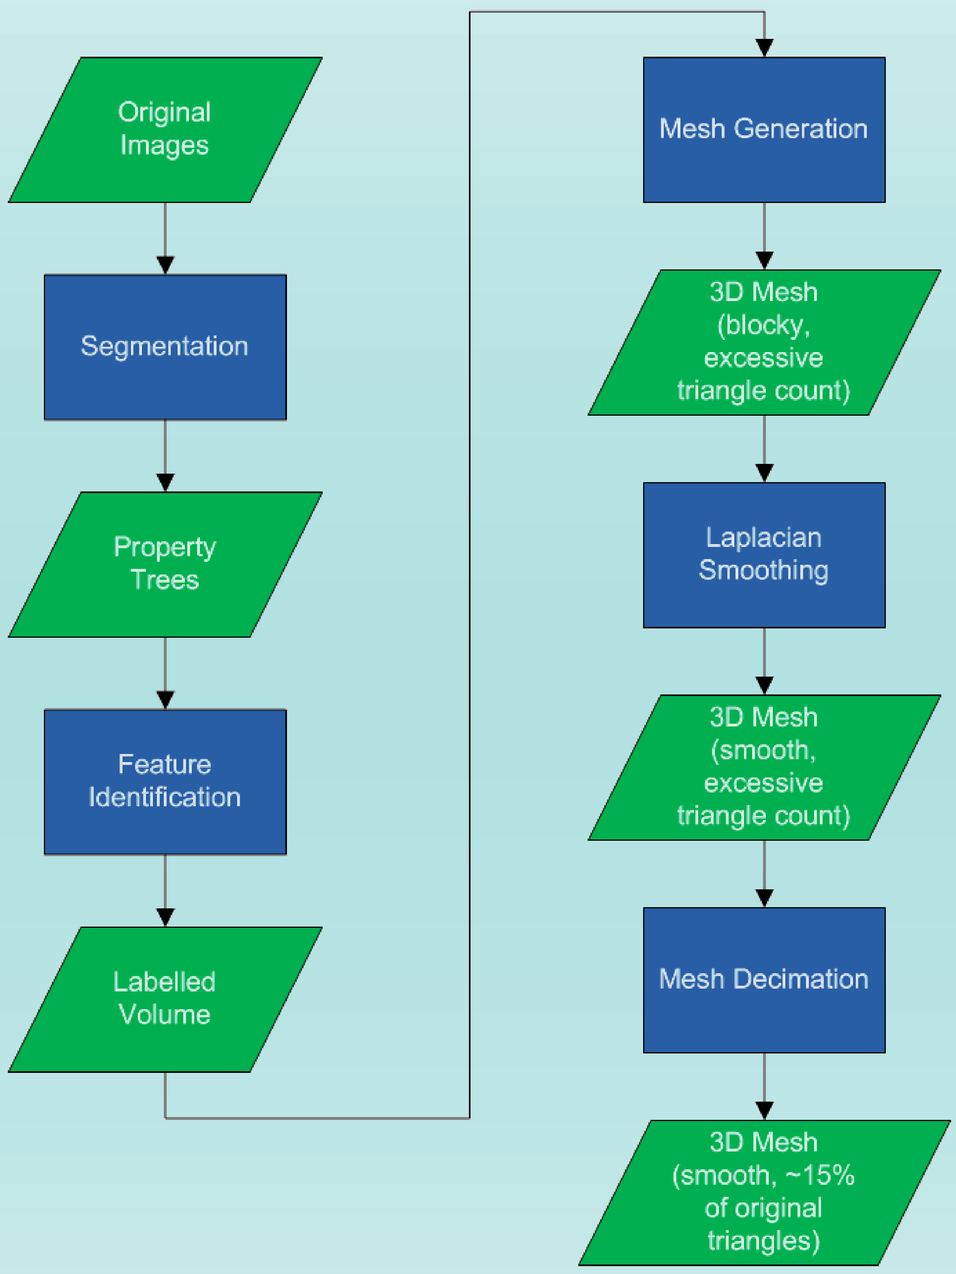

Volume Visualization

In general terms, volume visualization is the process of rendering a volume of data in a manner that aids human understanding. An example can be seen in the image above, which shows an abdominal section through a patient: it's much clearer to see what's going on in the 3D visualization than it would have been in a sequence of segmented slices like those shown in the image at the bottom-left of the poster.

Our work uses the Multiple Material Marching Cubes (M3C) algorithm [Wu03] to visualize the results of identifying features in the slices. This is essentially a 3-stage process: first, a mesh is generated for the 3D labelled volume on a cube-by-cube basis, in such a way that the mesh pieces in adjacent cubes match up on the cube faces (this ensures a valid resultant mesh).

Laplacian Smoothing

Since the mesh produced by this process tends to be a bit blocky and suffer from 'stair-stepping' (due to the inter-slice spacing in the z direction being much greater than the intra-slice spacing in the x and y directions), the mesh is next smoothed using a process known as *Laplacian smoothing*. This involves iteratively moving vertices a small amount towards each of their neighbours: it is effectively a parameterised averaging process.

Mesh Decimation

Finally, the mesh undergoes a process known as *decimation* in order to reduce its triangle count. This is essential, since generated meshes can consist of millions of triangles, making them impossible to render as is at a reasonable frame-rate. The goal of the decimation process is to reduce the triangle count by a specified amount whilst preserving the *topology* of the mesh to the greatest extent possible. It is possible to reduce triangle counts by upwards of 80% without the topology being visibly affected.

Roughly speaking, the decimation process works as follows: firstly, an 'average plane' is calculated for each decimable vertex from its surrounding triangles. Its distance from that plane is used as a metric, so that flat areas of the mesh get decimated first. All the decimable vertices are placed in a priority queue, in increasing order by metric. Vertices are then iteratively popped off the queue, and the mesh locally decimated around them. The process terminates when the desired triangle reduction has been achieved.

Examples

As can be seen in the images (above and below), the output of the M3C algorithm is a valid, good-looking mesh that makes it easy to see what's going on. The image above shows an abdominal section through a patient with a large tumour on his left-hand side (shown in beige). The contralateral kidney (in yellow) is not cancerous. The image also shows part of the spine (in white), the liver (in purple), the spleen (in cyan) and the aorta (in red).

The image below shows a close-up view of part of the spine to illustrate the results of a 75% decimation of the original mesh. Note that none of the important details have been lost, in spite of such a serious reduction in the triangle count.

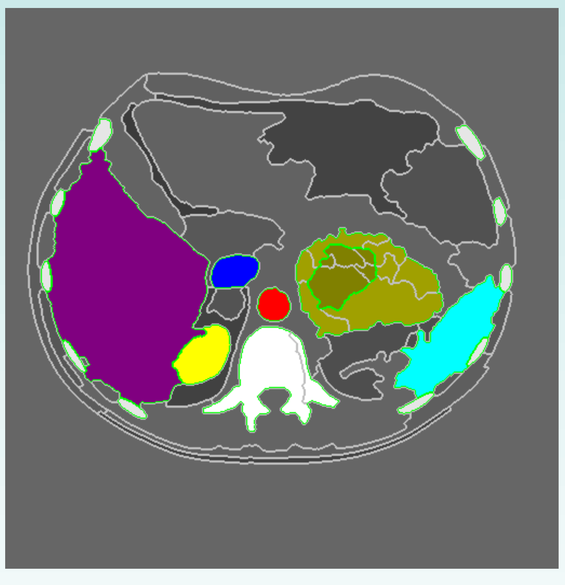

Summary

- Feature identification in CT scans is a vital preliminary process that allows us to do more directly useful things like volume visualization (rendering the inside of a patient's body to make it easier to see what's going on) and volume estimation (calculating the size of a tumour, etc.).
- Volume estimation (in particular) is important: if we can estimate the extent of necrosis within a tumour, doctors may be able to make deductions about its rate of growth and thus how best to treat the patient.
- We are currently trying to identify features automatically using a partition tree-based approach and have developed novel tree-modification algorithms which will help with this by allowing us to dynamically rearrange the partitions within the tree to obtain a better segmentation.

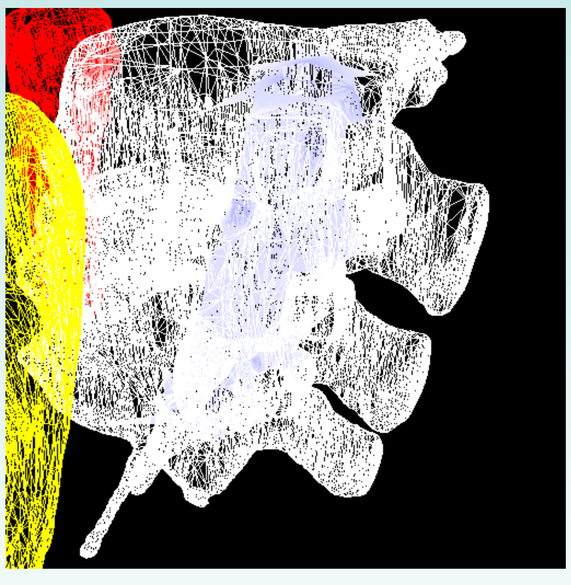

Acknowledgements

We gratefully acknowledge the invaluable help and support of our medical colleagues at the Churchill Hospital Oxford:

Dr Andrew Protheroe
Dr Zoe Traill
Mr Mark Sullivan (FRCS)
Mr David Cranston (FRCS)

References

[Marcotegui05] *Fast Implementation of Waterfall Based on Graphs*, Beatriz Marcotegui and Serge Beucher, Mathematical Morphology: 40 Years On (Springer Netherlands), 2005.
[Meijster98] *A Disjoint Set Algorithm for the Watershed Transform*, Arnold Meijster and Jos B T M Roerdink, Proceedings of the 9th European Signal Processing Conference (EUSIPCO '98), 1998.
[Wu03] *Multiple material marching cubes algorithm*, Ziji Wu and John M Sullivan Jr, International Journal for Numerical Methods in Engineering, volume 58, pages 189-207, 2003.

Stuart Golodetz, Irina Voiculescu and Stephen Cameron
stug@comlab.ox.ac.uk, irina@comlab.ox.ac.uk, cameron@comlab.ox.ac.uk
Oxford University Computing Laboratory, OX1 3QD, UK