# Volumetric Segmentation of Pigs

## - Using 3D Markov Random Fields and Graph Cuts


### Abstract

The Danish Meat Association (Slagteriernes Forsknings Institut) is currently working on using CT-scanners for virtual dissection of pigs. Since the dissection dictates earnings for the farmer (based on percentage of meat), the method used has to be robust and consistent. Furthermore, Denmark produces around 25 million pigs per year, which enforces the segmentation algorithm to be fast and accurate.
In order to achieve good classification, spatial relationship has to be included in the segmentation. In this project we have used Markov Random Fields (MRFs) and graph cuts [1]. In our implementation the spatial relationship in all three dimensions are included in the segmentation. Nevertheless, performance comparison with a per-slice-based segmentation is included as well.


### Previous Work

Currently the Danish Meat Association uses the Owen-Hjort-Mohn (OHM) algorithm [2] for classification using spatial relationship. The OHM algorithm is basically an approximation to a MRF segmentation. The reason why OHM is the current applied algorithm, is due to the computational complexity of MRF segmentation using traditionally iterative approaches. This complexity has now been removed due to the use of graph cuts, upon which this project is based.

### Problem

The figure below shows a single slice of a CT-scanned pig together with its segmented counterpart (using the MRF segmentation). Such a CT-scan can roughly be classified into the following six classes: air, skin, fat, meat, bone, and marrow. The primary class of concern is the meat. Due to this, we can narrow the classes into: fat (fat + skin), meat and bone (bone + marrow).

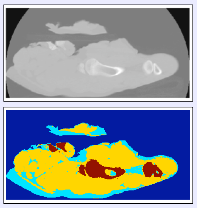

From a histogram of all voxels in the pig volume, we get a histogram as shown in the figure below . The first spike represents fat, the next contains meat, and the following flat area contains bone. By using thresholds between these classes, we are able to collect initial statistics about the classes (mean and variance) which are used for initial segmentation using our MRF implementation.

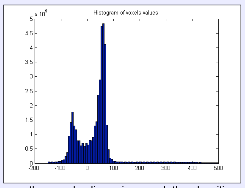


Since we know the voxels dimensions and the densities of the three classes, we are able to estimate the individual weight of all three classes, including the meat weight, and thereby answer the initial question asked by DMA.
Along with a scanned pig, the actual weight of the pig is given. Using this weight we are able to fine-tune the statistics in order to achieve an even more accurate estimate of the internal meat percentage.

### Markov Random Fields and Graph Cuts

A MRF can be seen as a multidimensional stochastic variable. Using a max-flow algorithm for solving the MRF segmentation was proposed in [1]. Here, a graph is constructed such that a single edge goes from the source to each voxel in the volume, and another edge leaves each voxel to the sink. The weight of these edges are computed using the formula below given their statistics as described in previous section.

$$U(y_{ij} \mid f_{ij} = K) = \frac{\ln(\sigma_K^2) + (y_{ij} - \mu_K)(\sigma_K^2)^{-1}(y_{ij} - \mu_K)}{2}$$

Furthermore, all internal pixels are connected in order to allow the graph flow to include their spatial information. When performing a max-flow algorithm, a graph cut is found (for a maximum flow) – thereby solving the MRF segmentation problem. The internal weights of the graph structure is set to a constant, in order to allow spatial relationship to be included. Optimal values for these internal weights will be addressed in the section "Estimation of Graph Weights". An example of such a graph can be seen in the figure below.

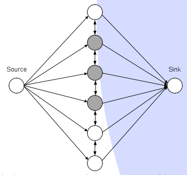


When doing multi class segmentation we repeat this process with updated source edges. This process is referred to as the alpha- expansion [1].

### Full Volumetric Segmentation

The primary focus of this project is to see if any extra information can be extracted using spatial relationship in all three dimensions. In order to do so, we have extended the graph structure – presented in previous section – into three dimensions, covering an entire pig volume. This is illustrated in the figure below.

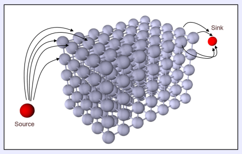


When doing this volumetric expansion we have to rethink the internal weights of the graph. The distance between each voxel in a single slice is 1 mm where the distance to the next slice is 10 mm. Knowing this, it is obvious that the amount of flow/information shared in the slice is stronger than the flow between slices. This question will be addressed in the following section and optimal values of these internal weights (or graph weights) will be found.
Another problem with this volumetric graph structure is memory consumption. In our case we have been forced to sub-sample each slice in the volume. The interesting thing about this, is that the results achieved using sub-sampling still converged extremely well towards the correct solution.

### Estimation of Graph Weights

To check the importance of the spatial relationship between a voxel and its neighbourhood, the graph weights are investigated.
A rough estimate of the optimal graph weights was achieved by segmenting three pigs at various values of the graph weights. The estimated weight of the pigs was then compared to the actual weight, and a mean squared error was calculated. A surface plot can be seen below, which illustrates the different mean squared errors.

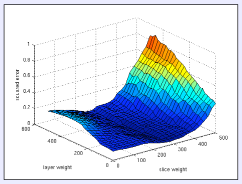


From the plot, it is found that a minimization of the error gives an in-slice-weight of 260 and a between-layer-weight of 160. This means that information, about a given voxel, may be found in the neighbouring pixel in the current slice, but information may also be extracted from the neighbour pixels in the adjacent layers.

### Volume Statistics Adjustments

All CT-scanned pigs will also have their weight measured. Using the physical weight, we are able to tweak the segmentation of all voxels to equal the actual physical weight. In terms of adjusting the segmentation we have two possibilities: move/adjust the threshold (between classes) or adjust class statistics (i.e. mean value, variance, or both). We have found that the easiest and most stable solution is simply to adjust the class mean value.
An average pig body approximately consists of 65% meat, 25% fat, and 10% bone. The density for meat and bone is above 1 kg/dm³ and the density of fat is below. If the total estimated weight of the pig is above the actual weight, we have to enhance the chance of a voxel to be fat, in order to lower the total weight. If the weight is under estimated the chances for a voxel being meat or bone should be enhanced.
In order to control the percentage the mean is shifted, the step size is weighted by the a priori class percentage (100% = fat [the only class for lowering weight] and 87% = meat and 13% = bone [the internal distribution of the two classes]). How the actual mean values are shifted is illustrated in the figure below.

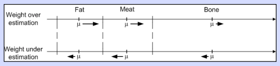


This approach may seem to destroy the statistics for the pig. Nevertheless, the initial result of the volumetric segmentation is so good (as shown in the "Results" section) that the mean values will only be slightly modified and still result in very reasonable values.
The figure below shows how the total estimated weight converges towards the actual physical weight during 10 iterations of the algorithm.

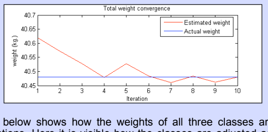


The figure below shows how the weights of all three classes are adjusted during iterations. Here it is visible how the classes are adjusted according to their percentage (i.e. meat is modified more then bone).

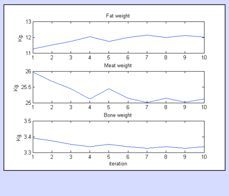


### Results

Now we have to see how the algorithm performs. In this section we will show how well the total weight is estimated when performing MRF segmentation on a per slice basis, and on the entire pig volume. Lastly we will compare the results with the OHM approach by comparing the $R^2$ measure.
The figure below shows the fit, of the estimated weights compared to the actual weights, using a slice based segmentation. The $R^2$, using this approach, is found to be 0.9991 – clearly a very good result.

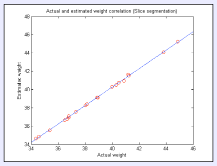


The figure below shows the result using the full volumetric segmentation with the graph weights found in the "Estimation of Graph Weights" section. Here the $R^2$ equals 0.9988, which is close to the results achieved using the slice based segmentation, however, slightly better results were achieved using small values of the graph weights. The current inclusion of depth information did not seem to have significant impact of the final result.

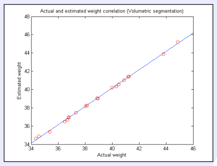


Finally, we have to test the volumetric segmentation when performed in conjunction with the iterative update of the statistics. Using this we get the result shown in the figure below. This approach results in a perfect match between the estimated total weight and the physical weight. The $R^2$ is in this case equal to 1.00 showing perfect a match – at least concerning the total weight.

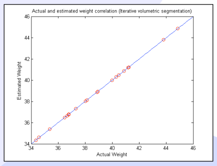


The results compared with the OHM approach seems to be nearly the same. The weights found, using the OHM algorithm, gives raise to $R^2$ values in the range of 0.999 as well. Only our iterative approach seemed to triumph the OHM approach.

### Conclusion

During this project we have successfully investigated and implemented a segmentation method, using MRFs and graph cuts, as an alternative method for the pig segmentation problem addressed by DMA. We have investigated the impact of including spatial relationship in all three dimensions during the segmentation process. Based on our results, this spatial relationship did not seem to reveal any new information, since no improvement was seen. Nevertheless, the estimation of the graph weights was not comprehensive enough to make any final conclusions, and further investigation in this area is heavily suggested.
Using the concept of including a prior physical weight of the pigs, an iterative approach was used to adjust the volume statistics, and thereby improving the segmentation results in positive way. The only thing left, is to find out if the estimated meat percentages are correct as well.
Overall the results using the MRFs and graph cuts did perform just as well as the existing solutions, and the more widespread usage of MRFs in general, may talk in favour for this type of segmentation.

### Future work and optimizations

Apart from the clear problem about the rough estimation of the graph weights, other improvements have been considered.
The entire implementation is based on C and MatLab, so a speed up would be achieved if the implementation was based entirely on C instead.
At the moment, the volumetric segmentation of a pig can be carried out in around 30 seconds (without using the iterative statistics adjustment). This result, however, is based on a volume where each of the slices is a sub-sampled version of the original slices.
If sub-sampling was not used, the size of the current data structure for the MRF-volume of each pig, would be in the scope of 1.5 to 2.0 Gb, a size which is hard for our implementation to handle. Changing the data types (i.e. floating point values with lower precision) in the data structure could probably help this, but it was not investigated further.
Also in each of the CT-slices there is a large amount of air (that does not contribute to the weight estimation), which is put into the MRF-volume. By defining a region of interest, the amount of air put into the graphs could be reduced, and thereby reducing the entire MRF-volume significantly.
Finally, since we are sub-sampling the original data, it should be investigated if a blurring of the image, prior to sub-sampling, can spread out image information in order to improve the segmentation result.

### References


[1] Vladimir Kolmogorov, "Graph Based Algorithms for Scene Reconstruction from Two or More Views", Sep. 2003

[2] Rasmus Larsen, "3-D Contextual Bayesian Classification", March 2000.



**Technical University of Denmark**
Course 02503 - Advanced Image Analysis - Spring 2008

Anders Rong s032263@student.dtu.dk
Jacob Lercke Skytte s031956@student.dtu.dk

Supervisor: Martin Vester-Christensen



DTU Informatik
Institut for Informatik og Matematisk Modellering