

ISMRM 2002 Presentation: Partial Volume Segmentation

Should be about 10 mins : 1200 words.

Title overhead

This presentation describes a method for segmenting multi-dimensional medical images, taking into account the problem of partial volumes of multiple tissues being present in a single voxel. Ignoring this effect leads to misclassification of voxels at tissue boundaries, and so by accounting for it we can obtain significantly more accurate segmentations. The use of multiple image modalities also improves the segmentation, as it leads to better tissue separation. In brief, we build a multi-dimensional model of the image data, with contributions from both pure tissues and mixtures of tissues. We then use the expectation maximisation algorithm to iteratively adjust the parameters of the multi-dimensional model to match the data. Finally, we use Bayes theory to generate probability maps for each segmented tissue to estimate the most likely tissue volume fraction within each voxel, as opposed to previous approaches which attempt to compute how likely a certain grey level would be to be generated by a particular tissue class.

Problem definition

This work was carried out in the context of a broader project called IER-APSI, or Integrated Environment for the Rehearsal and Planning of Surgical Interventions. The aim of that project was to design a system that would allow a surgeon to plan and practise petrous bone surgical procedures, such as Mastoidectomy, Cochlear Implants, and resectioning of Acoustic Neuromas. This involved building 3D models of the area of interest, showing all of the tissues present. Therefore, a segmentation technique was required that could take multi-modal data, for instance MRI scans to identify the soft tissues, CT scans to identify the bone, and MRA scans to identify the blood vessels, and segment out these tissues. We can then build probability maps for the tissues, for later use in 3D modelling.

Data modelling 1

This shows a one-dimensional example of the multi-dimensional model we are trying to construct. In order to model the image data whilst taking account of partial voluming, we require two types of contribution. Pure tissues can be modelled using multi-dimensional Gaussian distributions, shown by the dashed lines here. If we assume linearity i.e. the grey level of a voxel containing partial volumes of tissues is due to a linear combination of the grey levels that would be generated by the pure tissues, then we can model partial volume effects using a triangular distribution convolved with a Gaussian, and this is shown by the solid lines here. The resulting distribution is shown by the dotted line.

Data modelling 2

So, the contributions from pure tissues are modelled using the standard form of a multi-dimensional Gaussian. This accounts for both inherent tissue variability and noise. The probability distribution d for grey levels g of tissue t depend on the mean tissue vector M_t , and the inverse of the covariance matrix C_t , together with some constant A_t to give unit normalisation.

A Multi-dimensional Partial Volume Distribution

Partial volume distributions are modelled as a uniform 1D distribution convolved with a Gaussian, again to account for both tissue variability and noise. This is valid for an equal prior probability of all partial volumes, and in practise was found to provide a good fit to the data. The uniform distribution can be considered as composed of two triangular distributions convolved with a Gaussian. So, we model along a line between the two centers of the pure tissue distributions, and produce a weighted average of the two pure tissue Gaussian functions, where h is the fractional distance along the line. Since the definitions of h and the inverse covariance matrix C_h are dependent, we use an iterated closest point algorithm to find the closest point on the line to the data point, which give the normal distance N_g . T_{ts} is the triangular distribution for the local density estimate of tissue t generated by partial voluming with tissue s , and B again is a constant to give unit normalisation.

Data Optimisation by Expectation Maximisation

To fit the model to the data we iteratively adjust the model parameters using the expectation maximisation algorithm, which attempts to maximise the likelihood of the data distribution given the model. The algorithm has two stages. First, we calculate the probability of each grey level being due to a certain component of the model using Bayes theory. The lower case f 's here are effectively priors, but expressed as the frequencies i.e. the number of voxels, which belong to a certain tissue type, which can be either a pure tissue or a mixture of tissues. Unknown tissues are accounted for by introducing an extra fixed term f_0 for infrequently occurring outlier data, and this enables us to separate pathological tissues from the model.

Data Optimisation by Expectation Maximisation 2

The next step of the EM algorithm is to update the parameters in order to maximise the likelihood function. The basic idea is to make some original guess of the model parameters, calculate the multi-dimensional probability density in the expectation step, and then calculate new estimates of the model parameters from this probability density in a maximum likelihood fashion. These new model

parameters can be plugged back into the expectation step, and the algorithm can be iterated. I won't go through the derivation in detail, as it is described in the standard text books, but the model parameters are updated using the equations given here, for the new frequencies of the tissues, the tissue means, and the inverse covariance matrices. After some number of iterations, when the algorithm has converged, the probabilities can be used to produce probability maps for each tissue class.

The EM algorithm has been shown to be guaranteed to converge, but it suffers from several problems common to all such optimisation techniques. In particular, it can converge only to a local minimum and can be badly affected by outliers. It is therefore important to ensure that the initial estimates of the model parameters are close to the true distribution of the data, and to deal with outliers in some way, for instance by including an outlier class in the model.

MRI Image Sequences

To illustrate the algorithm, we have applied it to the problem of segmenting multi-dimensional MR images of the head. Here are the sequences used: inversion-recovery turbo-spin echo, variable echo T2, variable echo proton density, and fluid attenuated inversion recovery. These sequences were chosen because they provide good separation of the tissues of interest, which include air and bone as a single class, because they both feature as signal voids in the MR images, fat, soft tissue (including skin and muscle), CSF, grey matter and white matter. Voxel dimensions for these images were 0.89 by 0.89 by 3.5 mm. Using greater slice thicknesses gives better signal to noise ratios, but increases the problem of partial voluming as there is a higher chance of a tissue boundary passing through a voxel.

Scatter Plots

This shows a scatter plot of the data for two of the modalities used, inversion recovery turbo spin echo and variable echo T2. You can see the data clusters due to the pure tissues, with the partial volume distributions stretching between them. The second scatter plot shows the initial model before application of the EM algorithm, where again you can see the same features, and the final plot shows the model after ten iterations of the EM algorithm, and you can see that the model has converged to the data successfully.

Histogram Plots

To see this in more detail, we can plot histograms of the images used, where the red line is the original image data, the blue line is the total model, the green lines show the pure tissue contributions, and the pink lines show the partial volume contributions. This is the histogram for one of the images before

application of EM...

Animation

...and this is the histogram after 10 iterations of EM, and you can see that the model has successfully converged to the data within ten iterations.

Histogram Plots 2

Remember that we are working with 4D data here: so we can plot the same thing for all four of the image sequences used. So the top row shows the histograms before application of EM, and the bottom row shows histograms after 10 iterations of EM, and you can again see that the model has converged on the data: the red and blue lines are in the same place.

Probability Maps

We can take the Bayesian probabilities output by the algorithm and build probability maps for each tissue type included in the model, picking out the tissue boundaries at the point where the probability drops to 0.5 between each pair of tissues. Here you can see maps for bone and air, fat, soft tissue, CSF, grey matter and white matter. You can see that the method has successfully segmented out the six tissue types. There are a few misclassifications, which you can see most easily in the white matter map. These are due to genuine lack of differentiation in the data, which could be addressed in a number of ways, for example by applying an additional step of probability relaxation labelling. Remember that we are estimating the partial volume contributions to each voxel, not the most likely tissue class, and so we can get accurate tissue volume measurements that include partial volume voxels

Conclusions

To conclude, about half of all the voxels in the data I have shown here are accounted for by partial voluming, mainly due to the slice thickness. By building a data model that accounts for this effect, we can obtain more accurate segmentations. This in turn means that we obtain a better visual appearance for the segmented tissues, which is important for the target application of visualisation and simulation. The software also supports tissue volume and volume change measurement, and we have published a paper on its application to measuring tumor volume changes. Finally, it could be extended to any number of image modalities for which the linearity assumption holds, and we are particularly interested in including CT to enable segmentation of the bone.

Acknowledgements

Finally, I would like to acknowledge the support of the European Union for funding this work as part of the IERAPSI project, and mention that all of our machine vision software and publications can be downloaded as open source from our website.