Questions and Answers

Chapter 1

Questions

(Q1) *What is the parametric model of the MR bias field proposed in this chapter?*

(Q2) *How are the parameters of the MR bias field estimated?*

(Q3) *What are alternative models for the MR bias field?*

(Q4) *How is the atlas-based geometric prior constructed?*

(Q5) *How is the atlas of priors integrated into the classification?*

(Q6) *What are the drawbacks of using such an atlas of priors and how to deal with it?*

(Q7) *What is the Markov Random Field model used in brain tissue classification?*

(Q8) *How do we limit the over-regularization of Markov Random Fields?*

Answers

(A1) We propose to estimate the bias by a parametric model. Because the main source of the bias field is a low-frequency effect, we have chosen the bias field to be modelled as a linear combination of smooth basis functions. The

choice of the appropriate parametric model is important. We use polynomial models since they describe slowly varying phenomena with a small number of parameters and are easy to compute. The maximum degree of the polynomial determines the accuracy and stability of the calculated bias field. Choosing too-large a degree results in inefficient computation times, instable coefficients, and the calculated bias field is likely to adapt itself to multiple classes. Choosing too-low a degree, will only partially remove the bias field. Our experiments showed that, for MRI with standard head-coils, polynomials up to the fourth degree sufficiently model the distortion [28, 32, 33]. We start with a zero-order bias polynomial and increase the order when convergence is detected, until the maximum order, typically 4, is reached.

(A2) The parameters of this model are computed from a weighted least-squares fit to the residue image (Fig. 8(c)), each pixel's weight (Fig. 8(d)) being inversely proportional to the variance of the class that pixel belongs to. The bias field is therefore computed primarily from pixels that belong to classes with a narrow intensity distribution, such as WM and GM. The smooth spatial model of the bias field extrapolates the bias field from regions where it can be confidently estimated from the data (WM and GM) to regions where such estimate is ill-conditioned (CSF, non-brain tissues). The algorithm of Wells et al. [4] uses the same residue and weights as our algorithm, but they compute the bias field from the residue image by spatial filtering. However, their bias estimation is a mere low-pass filtered version of the residue in large regions with the same tissue class, independent of the weights. In regions with low weights, such as the ventricular area, this can be expected to lead to errors.

(A3) We have not investigated other representations such as splines [31, 36, 50], nor did we investigate techniques to estimate the order or the stiffness of the bias field model automatically from the data itself. For 2D multi-slice sequence images, we found that correcting for a 2D polynomial bias field on each slice separately, which was also done in [31, 4], yields a more consistent classification than when a single 3D polynomial is used, even when correcting for slice-by-slice constant intensity offsets between slices.

There is also some debate as to whether the bias field should be modelled as a multiplicative or additive effect. Furthermore, should the bias field be the same for all tissue classes or tissue-class dependent? The

multiplicative model explains well the primary cause of intensity variations, namely variations in the sensitivity of the receiver coils. However, in order to make it computationally tractable, the multiplicative degradation is first logarithmically transformed into an additive one. Unfortunately, as a result, the intensity distributions are altered, making tissue separation more difficult and invalidating the Gaussian noise assumption. Therefore, the bias field is often directly modelled as an additive effect. The single multiplicative bias field model does not well approximate other bias factors such as induced currents, nonuniform excitation, biological noise (variation of electromagnetic properties of biological tissues). Marroquin et al. [50] therefore model the variable intensity for each class separately.

(A4) The initial a priori pixel class probabilities are derived from a digital brain atlas, distributed with the SPM99 package [71], which contains spatially varying prior probability maps for the location of white matter, gray matter and CSF as shown in figure 3. These probability maps, were obtained by averaging binary white matter, gray matter and CSF segmentations of MR brain images from a large number of subjects, after normalization of all images into the same space using an affine transformation [20].

(A5) To apply the a priori information of the atlas, we first normalize the atlas to the space of the study image by matching the study image to a T1 template that is already co-registered with the atlas (see figure 3) using the affine multi-modality registration technique based on maximization of mutual information of corresponding voxel intensities of Maes et al. [21].

We use six classes: three brain tissue classes for white matter, gray matter, and CSF respectively, two non-brain tissue classes and one class for the background signal. Two classes are used for the non-brain tissues because the intensity distribution of the non-brain tissue voxels can not be accurately modelled by a single Gaussian. Because the atlas provides a priori probability maps for white matter, gray matter and CSF only, we construct a priori probability maps for the 3 other classes by subtracting the sum of the atlas probability maps from a map with value 1 in all entries and dividing the residual probability equally over the three other classes. Since we are only interested in brain tissue, we confine the algorithm to the region where the atlas indicates a non-zero a priori probability for white matter, gray matter or CSF. All other pixels are of no interest and are simply discarded.

(A6) A drawback of this approach is that it assumes such an atlas to be available, which currently limits the practical use of this method to brain images only. Also, while the affine registration approach is able to compensate for global differences in brain morphology, the classification will fail in the presence of large, pathology induced abnormalities in brain shape, for instance in subjects with dramatically increased ventricles. The fuzziness of the atlas currently used helps to temper the influence of local mis-registration not corrected for by the affine registration. However, while the fuzziness of the atlas has no significant influence on the bias correction itself, it may occasionally introduce voxel classification errors. These problems may be overcome by using non-rigid rather than affine registration for atlas matching.

The brain atlas template used here is a group-averaged template [20], created by averaging MR images (and the corresponding tissue segmentation maps) from 305 young, normal subjects, following affine spatial normalization. For global spatial normalization, involving mainly translation, rotation and scaling, such an average brain template avoids any spatial biases arising when a single brain template were used. It is not adequate, though, when high resolution spatial normalization is required to remove anatomical differences between individual 3D brain images. Holmes et al. [80] created a high-resolution, high signal-to-noise-ratio brain MRI template. This template provides excellent resolution and contrast and has been manually segmented and anatomically labelled. However, matching of an image to this template will inevitably lead to bias in the matching quality because of anatomical features unique to this particular brain. Kochunov et al. [81], therefore, developed a method to define, for a particular multi-subject study, a "minimal deformation target" (MDT). The MDT is defined as that brain image that minimizes the overall deformation between it and all brain images in the study.

(A7) Within the context of brain tissue classification, Markov Random Field (MRF) models incorporate spatial and anatomical constraints, such as "neighboring voxels tend to belong to the same class" or "a voxel surrounded by non-brain tissue cannot belong to gray matter". A MRF model assumes that the probability that a particular voxel belongs to a certain tissue type depends on the tissue type of its neighbors. A simple MRF Potts

model (the extension of the binary Ising model [41] to more than 2 classes) is used. It is defined on a first order neighborhood system, i.e. only the six nearest neighbors on the 3-D-image lattice are used. The MRF parameters are 4×4 matrices G and H, the (k, j)-th element of which describes the cost associated with transitions from class k to class j among neighboring voxels (Eq. 0.3.1). We use four classes: white matter (class 1); gray matter (class 2); CSF (class 3); and other (class 4).

Due to the fact that the in-plane resolution is generally different from the out-plane resolution in MR images, we have modelled in-plane class transition costs with different parameters G than between-slice class transition costs, parameterized by H. Since the resolution of MR images can largely vary, we re-estimate G and H for every image separately. It should be noted, however, that the MRF parameter estimation step is responsible for almost half of the total computational burden. To speed up the algorithm, one could neglect the out-plane interactions H and precalculate the in-plane interactions G just once on a normal brain data set. Since the in-plane resolution in MR is fairly constant, this precalculated G could then be used for the segmentation of all following data sets. However, this approach does not make use of the full 3D nature of the MR images and can therefore be expected to yield less powerful discrimination between brain and non-brain tissues. On the contrary, if a large set of similar images with equal voxel sizes has to be segmented, both G and H could be precalculated on one image and applied to all other images of the set. Although we have not validated this approach, we expect this to speed up the process without loss of accuracy.

(A8) Application of the genuine Potts model, leads to the loss of small details in the resulting segmentations, as shown in Figure 12(e). This overregularization is a well-known effect which other models try to overcome. In particular, the so-called Chien-model, first proposed by Descombes et al. [42], seems better adapted to medical images since it better preserves fine structures and linear shapes. Unfortunately, generalizing this 2D model to 3D induces neighborhoods of 124 voxels and leads in practice to intractable computations. A more efficient extension to 3D has been proposed in [82], but this still involves 60 neighbors. Instead, we propose a modification that penalizes impossible combinations, such as a gray matter voxel

surrounded by voxels belonging to the class "other", while at the same time preserving edges between tissues that are known to border each other. We impose that a voxel surrounded by white-matter and gray-matter voxels must have the same probability to be white matter as to be gray matter. With the class numbers as defined above, this can be achieved by imposing the constraints that $G_{11} = G_{21}$ and $G_{22} = G_{12}$, the same for H. As a result, voxels surrounded by brain tissues have a low probability for CSF and "other" and a high, but equal, probability for white and gray matter. The actual decision between white and gray matter is therefore only based on the intensity, so that the interface between white and gray matter is unaffected by the MRF. The same rationale applies for the interface between gray matter and CSF ($G_{22} = G_{32}$) and ($G_{23} = G_{33}$), the same for H, and to the interface between CSF and other ($G_{33} = G_{43}$) and ($G_{34} = G_{44}$), the same for H. This reduces the number of MRF parameters to be estimated from $2 \times 4^2 = 32$ to 20.