

Inference from X-ray CT Images of sheep

C.A. Glasbey

Biomathematics and Statistics Scotland

JCMB, King's Buildings, Edinburgh EH9 3JZ, Scotland

Email: chris@bioss.sari.ac.uk

C.D. Robinson

Department of Mathematics and Statistics, University of Edinburgh

JCMB, King's Buildings, Edinburgh, EH9 3JZ, Scotland

Email: caroline@maths.ed.ac.uk

Introduction

X-ray computed tomography (CT) is a non-invasive imaging technique which is ubiquitous in hospitals. It was originally developed for use in medical diagnosis, but its potential has recently been recognised in a totally different area: sheep breeding. This new area of application has generated challenging new problems in image analysis, because the objective of estimating tissue volumes is radically different from the medical purpose of detecting physiological abnormalities.




Figure: Conventional X-ray image, and CT cross-sectional images at arrowed locations, of a sheep lying on its back in a cradle. Grey levels in the images indicate the extent to which different tissues attenuate X-rays. Hand-drawn boundaries, superimposed in white, identify relevant tissue areas for estimating body composition.

The protocol established in the SAC-BioSS CT Unit in Edinburgh, is to obtain a conventional X-ray image and three anatomically-located CT cross-sectional images for each sheep, as shown in the Figure. From these data, we wish to estimate the animal's fat and muscle tissue volumes.

Our approach has three stages: first we automate the identification of relevant tissue areas in the CT images, by excluding internal and external organs; then we estimate fat and muscle proportions; and finally we infer whole-body tissue proportions by pooling information from the CT images and a conventional X-ray image. Each stage will be discussed turn.

Stage 1: Identification of tissue areas

Automatic image segmentation is a challenging task in many applications (see, for example, Glasbey and Horgan, 1995, ch.4). To develop a segmentation method for this problem, we used images from 98 sheep, together with hand-drawn boundaries as shown in white in the Figure. These boundaries separate the sheep's internal and external organs from the areas of fat, muscle and bone which we are interested in. All three boundaries enclose approximately star-shaped regions, enabling us to transform to polar coordinates (r, θ) about suitably chosen origins. To standardise the orientations of animals, the direction of zero angle was also chosen. Separately for images at each anatomical location, distributions of pixel values at a range of distances (m) from the hand-drawn boundary ($r(\theta)$) were summarised by means and standard deviations:

$$\mu_{m,\theta} = \frac{1}{N} \sum I_{r(\theta)+m,\theta}, \quad \sigma_{m,\theta} = \sqrt{\frac{1}{N-1} \sum (I_{r(\theta)+m,\theta} - \mu_{m,\theta})^2},$$

for $m = -M, \dots, M$; $\theta = -180, \dots, 180$. The summations are over N training images, M is the maximum distance considered from the boundary and I denotes arrays of polar-transformed pixel values. (Note that, in the Figure, each CT image is cropped to a 191×191 array of pixels, with each pixel 2mm square.)

We defined the automatically-chosen boundary to be the one that minimises a path across the array of outputs from a non-homogeneous filter constructed using μ and σ :

$$\hat{r} = \arg \min_r \sum_{\theta=-180}^{180} I'_{r(\theta),\theta}, \quad \text{subject to} \quad |\{r(\theta) - r(\theta-1)\} - \{\bar{r}(\theta) - \bar{r}(\theta-1)\}| \leq 1\frac{1}{2} \quad \text{for all } \theta,$$

$$\text{where} \quad I'_{r,\theta} = \log \left[\sum_{m=-M}^M \left(\frac{I_{r+m,\theta} - \mu_{m,\theta}}{\sigma_{m,\theta}} \right)^2 \right] \quad \text{for} \quad r = 0, \dots, R; \theta = -180, \dots, 180,$$

\bar{r} denotes the average boundary position in the training images and R is the maximum radius considered. If we restrict to a lattice of values of (r, θ) , then the global minimum can be obtained using dynamic programming (see, for example, Amini et al, 1990). Experimentation with a range of values of M showed a choice of 30 pixels to give the closest agreement between the hand-drawn and automatic boundaries in the training images at anatomical locations 1 and 3, and $M = 20$ was best for location 2. The algorithm was then applied to a validation set of a further 98 images, with results given in the Table. As we would expect, the fit to the validation data is not quite so good as for the training data, but it is still adequate for our purposes. The method could be extended to take account of other features of pixel distributions than simply means and variances, to allow for spatial correlations and to include a regularising term to ensure that the estimated boundary is smooth.

Table: Root-mean-square differences (mm) between hand-drawn and automatic boundaries, averaged over 98 images in each of two data sets.

anatomical location	1	2	3
training data	1.6	3.1	1.6
validation data	1.6	3.6	1.9

Stage 2: Estimation of tissue proportions

Many of the pixels in CT images are responses to mixtures of two or more tissue types. In developing an efficient estimator of tissue proportions, it is important to take this into account.

We first need to know the spatial response of the X-ray instrument, which we identified by examining pixels at the edge of the cradle in which the sheep lie. An isotropic bivariate normal density, with standard deviation $\tau = 0.41$ pixels, proved to be an adequate point spread function. This can be used to derive the mixed pixel distribution. For simplicity we restrict attention to only two tissue types, though the method is extendible to more. Let us assume that a proportion p_f of pixels are pure fat, and these pixel values are normally distributed with mean ν_f , variance ω_f^2 . Similarly, assume that a proportion p_m of pixels are pure muscle, normally distributed with mean ν_m , variance ω_m^2 . From the value of τ , it is reasonable to specify that any pixel more than one pixel distant from a tissue boundary is either pure fat or pure muscle. The remainder of the pixels, proportion $(1 - p_f - p_m)$, are mixed. Let us assume that tissue boundaries are locally linear, and the perpendicular distance from a mixed pixel to a boundary is uniformly distributed between -1 and $+1$. In a mixed pixel the proportion of fat, denoted ρ , has probability density function

$$f(\rho) = \sqrt{\frac{\pi\tau^2}{2}} e^{[\Phi^{-1}(\rho)]^2/2} d\rho \quad \text{for } \Phi(1/\tau) \geq \rho \geq \Phi(-1/\tau).$$

As far as we are aware, this is a new distribution, not previously considered. Santiago and Gage (1995) also derived a mixed pixel distribution, but mistakenly assumed ρ to be uniformly distributed.

Conditional on ρ , the pixel value will be normally distributed with mean $[\rho\nu_f + (1 - \rho)\nu_m]$, variance $[\rho\omega_f^2 + (1 - \rho)\omega_m^2]$. By convolving this with the distribution for ρ , we obtain the distribution of mixed pixels:

$$f(y) = \int_{\Phi(-1/\tau)}^{\Phi(1/\tau)} \frac{e^{-[y - \rho\nu_f - (1 - \rho)\nu_m]^2/2[\rho\omega_f^2 + (1 - \rho)\omega_m^2]}}{\sqrt{2\pi[\rho\omega_f^2 + (1 - \rho)\omega_m^2]}} \sqrt{\frac{\pi\tau^2}{2}} e^{[\Phi^{-1}(\rho)]^2/2} d\rho.$$

This has no analytic solution, but can be computed by numerical integration for any specified values of the parameters. Finally, we combine this distribution with that for pure pixels, to obtain a probability density function of pixel values of

$$p_f \frac{e^{-[y - \nu_f]^2/2\omega_f^2}}{\sqrt{2\pi\omega_f^2}} + p_m \frac{e^{-[y - \nu_m]^2/2\omega_m^2}}{\sqrt{2\pi\omega_m^2}} + (1 - p_f - p_m)f(y).$$

We fitted it by maximum likelihood to histograms of pixel values, using a numerical optimisation routine, and obtained excellent agreement. An efficient estimator of the fat proportion is given by $\hat{p}_f + (1 - \hat{p}_f - \hat{p}_m)/2$.

In order to develop efficient estimators when more than two tissue types are present, it may be necessary to take into account a pixel's spatial context, i.e. values of neighbouring pixels. Also, the assumption that distances D are uniformly distributed between -1 and $+1$ may need to be generalised for images with many fine tissue structures.

Stage 3: Estimation of whole-body tissue proportions

From stages 1 and 2, we obtain detailed information on a sheep's composition at three anatomical positions. However, a conventional X-ray image also contains some information on body composition, even though it is obtained primarily as a byproduct in locating the positions for CT scanning. The information is not as detailed as in CT images; in particular, data on X-ray attenuation is integrated through the body rather than being measured at distinct locations. But, the information is available for the whole body, not just at the scan sites. It is, therefore, likely that the best predictor of body composition will involve both CT and X-ray data. For example, we could interpolate between the CT images. Another simple approach, if we denote the summed intensities by X_i for each column in the X-ray image, for columns $i = 1 \dots n$, and suppose the three anatomical sites are at columns a_j , is to include

$$\frac{1}{3} \sum_{j=1}^3 X_{a_j} \bigg/ \frac{1}{n} \sum_{i=1}^n X_i$$

as a regressor in the prediction equation for whole-body composition. We are currently exploring these ideas, both from a theoretical perspective, and empirically, using dissection data.

ACKNOWLEDGEMENTS

The first author's work was supported by funds from the Scottish Office Agriculture, Environment and Fisheries Department and the second author by an EPSRC CASE studentship.

REFERENCES

- Amini, A.A., Weymouth, T.E. and Jain, R.C. (1990). Using dynamic programming for solving variational problems in vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 12, 855-867.
- Glasbey, C.A. and Horgan, G.W. (1995). *Image Analysis for the Biological Sciences*. Wiley, Chichester.
- Santago, P. and Gage, H.D. (1995). Statistical models of partial volume effect. *IEEE Transactions on Image Processing* 4, 1531-1540.

RÉSUMÉ

La tomographie calculée (CT) à l'aide des rayons X, une technique d'imagerie a été récemment appliquée à l'élevage des moutons. On utilise une image aux rayons X et trois images CT en coupe localisées selon l'anatomie pour estimer les volumes des tissus de graisse et de muscle. Nous commençons par automatiser l'identification des zones de tissu appropriées pour les images CT en appliquant une programmation dynamique aux données produite par un filtre non homogène. Puis nous développons une nouvelle repartition des probabilités qui prend en compte le mélange des pixels, de façon à estimer les proportions de graisse et de muscle dans chaque image CT. Enfin, nous discutons quelques méthodes pour en déduire les proportions des tissus du corps entier, en mettant en commun les informations combinés des images CT et des images aux rayons X conventionnelles.