ESTIMATION OF TISSUE PROPORTIONS IN X–RAY CT IMAGES USING A NEW MIXED PIXEL DISTRIBUTION

CHRISTOPHER ANDREW GLASBEY[1]
AND CAROLINE DENISE ROBINSON[2]

[1]*Biomathematics and Statistics Scotland JCMB,*
King's Buildings,
Edinburgh, EH9 3JZ, Scotland
chris@bioss.sari.ac.uk

[2]*Department of Mathematics and Statistics,*
University of Edinburgh JCMB, King's Buildings,
Edinburgh, EH9 3JZ, Scotland
caroline@maths.ed.ac.uk

Abstract: A new probability distribution is obtained, termed the *mixed pixel distribution*, appropriate for the mixing proportion in digital images when the point spread function has a bivariate normal density. It is used to derive the distribution of pixel values in X–ray CT images where pixels may be a mixture of two tissue types. In a simulation study it is shown that, by fitting this distribution to histograms of pixel values, tissue proportions are estimated more accurately than using threshold–based methods.

Keywords: bivariate normal distribution, point spread function, sheep breeding, threshold estimation

1. Introduction

The primary application of X–ray computed tomography (CT) as with several other non–invasive medical imaging techniques, is diagnosis by detection of physiological abnormalities. However, X–ray CT also has the potential to be used quantitatively, to estimate tissue volumes, in human obesity studies, for example. Here, we describe a quantitative application involving estimation of sheep tissue proportions as part of a sheep breeding programme. However, the methodology is generic and could be used with other subject domains and imaging modalities.

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. For example, Figure 1a shows a lumbar image. For a discussion of the use of medical imaging devices to estimate sheep composition, see Simm (1992). In CT, X–rays are projected through a subject from different directions. and a computer reconstructs an image of the distribution of tissue types from the transmitted X–rays. As is the convention with X–ray plates, light areas in the image denote regions which transmitted less X–rays. The lightest area is the backbone.

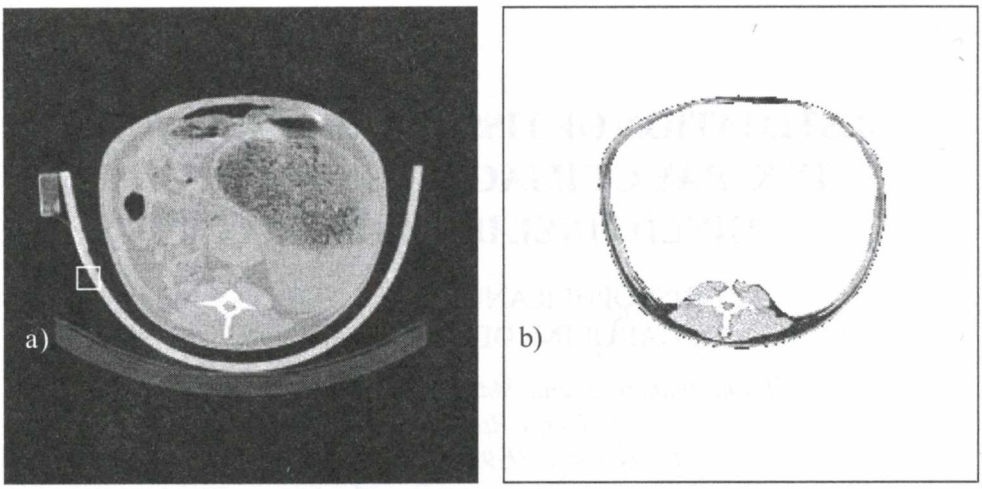

Figure 1. *X–ray CT lumbar image of a sheep: a) original (with the white square identifying the sub–image in Figure 3a), b) after processing to remove internal organs and cradle in which the sheep is lying, and with a stretched grey–scale pixel display, to emphasise differences between fat and lean tissues*

The muscles and internal organs appear slightly lighter than the fat tissue because they are slightly more opaque to X–rays. The U–shaped plastic cradle in which the sheep was lying can also be seen. X–ray attenuation is measured in Hounsfield units, which range between –1000 and about 1000. For an introduction to image analysis see, for example, Glasbey and Horgan (1995).

From images such as Figure 1a, 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 the internal organs and cradle, to produce images such as Figure 1b; 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. For a method which addresses the first stage, see Glasbey, Robinson and Young (1999). In this paper we focus on the second stage, estimation of tissue proportions from images like Figure 1b.

One way of estimating tissue proportions is by first classifying each pixel separately. However, this is unnecessary if we only want to know overall proportions. We can simply analyse the histogram of pixel values, such as that shown in Figure 2, obtained from Figure 1b. The smaller peak on the left of Figure 2 is produced by fat tissue, and the larger peak on the right is the result of muscle.

Pure tissues of either type have only a narrow range of pixel values. The many pixels falling between these peaks are caused by mixed pixels, which are part muscle and part fat (or part bone, etc.) Currently, areas are estimated simply by thresholding at the midpoint between the two peaks in the histogram. Here we propose a new mixed pixel distribution, and use it to derive a more efficient method for estimating tissue proportions. We develop the theory in section 2, and conduct a simulation study in section 3, to compare the results with those obtained by other methods.

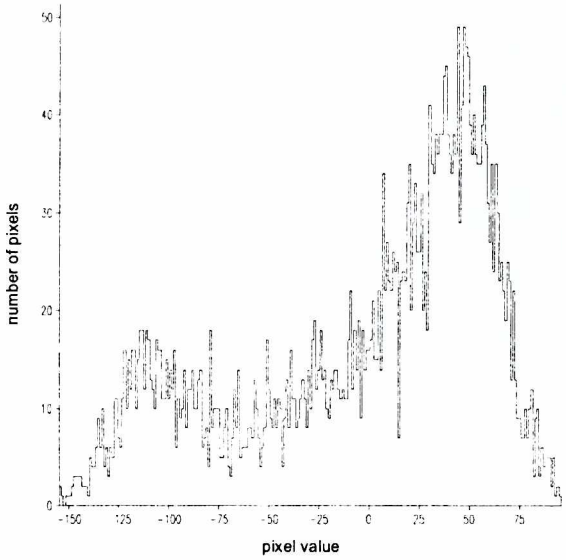

Figure 2. *Histogram of pixel values in Figure 1b, for a restricted range of Hounsfield values*

2. A new mixed pixel distribution

In order to identify the appropriate distribution to account for mixed pixels, we first need to model the spatial response of a pixel, the so–called point spread function. There are several methods for doing this (Dore, Kearney and De Guise, 1997). Here we adopt a simple approach: we chose part of the image containing the edge of the cradle in which the sheep is lying, identified by the white square in Figure 1a, which we know to be a step edge.

We modelled data near the edge of the cradle by pixel (i, j) having a predicted value of g_{ij}, given by:

$$g_{ij} = \alpha + \beta\phi\left(\frac{d_{ij}}{\tau}\right), \quad \text{where} \quad d_{ij} = \gamma - \sqrt{(\kappa_1 - i)^2 + (\kappa_2 - j)^2}. \tag{1}$$

Here, we assume that this part of the cradle can be approximated by an arc of a circle with radius γ, centred at (κ_1, κ_2), d_{ij} is the distance from pixel (i, j) to the edge, background and cradle pixels have mean values of α and $(\alpha + \beta)$ respectively, and ϕ is a 0–1 function that specifies the response of a pixel, after scaling by τ.

It is the one–dimensional integral of the two–dimensional point spread function. We estimated the 6 parameters (α, β, τ, γ, κ_1, κ_2) in (1) by minimising the sum–of–squares of the difference between the data and the model, for several choices of ϕ: a logistic function, a normal integral and a linear ramp. Table 1 summarises the results.

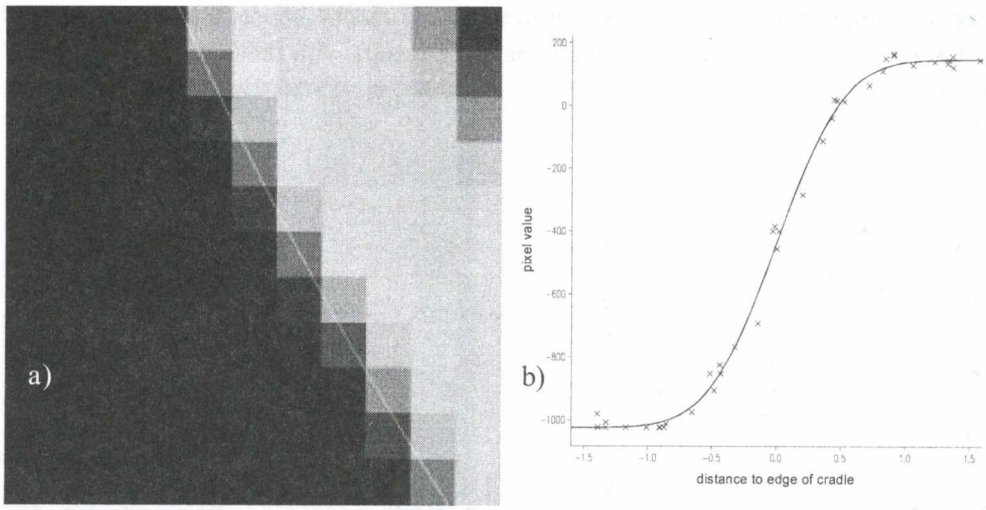

Figure 3. *Data used to estimate point spread function: a) subimage from Figure 1a together with estimated edge of the cradle, b) data near the edge of the cradle plotted against that distance, together with fit of model (1) using a normal integral function*

Table 1. *Least squares fit of model (1) to data in Figure 3b*

ϕ	*ramp*	*logistic*	*normal integral*
sum of squares	19.4	18.8	18.1

The normal integral gave the best fit, and is also the only one of the three that corresponds to an analytically simple two–dimensional point spread function, in this case a bivariate normal density. Therefore, we chose:

$$\phi(x) = \Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{x} e^{-y^2/2} dy,$$

where Φ is the standard notation for a normal integral. Figure 3a shows the estimated edge of the cradle superimposed on the sub–image, and Figure 3b shows data and fitted values plotted against D, from which we see that the agreement between the model and the data is very good. The most important parameter is τ, which we estimate to be 0.41.

The standard use of the point spread function in image analysis is in deconvolution, to restore the preblurred image. However, this is not our purpose, because we are not interested in the individual classification of each pixel in the

image. Instead, we use the point spread function to derive the appropriate distribution to fit to histograms such as Figure 2. 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 μ_f, variance σ_f^2. Similarly, assume that a proportion p_m of pixels are pure muscle, normally distributed with mean μ_m, variance σ_m^2. From the value of τ, it is reasonable to specify that any pixel more than one unit 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. We will assume that tissue boundaries are sufficiently smooth to be approximated locally by straight lines, and that a negligible number of pixels are affected by more than one boundary. We claim that the distance, D, from a random line to a randomly chosen pixel within one unit distance of it, is uniformly distributed between –1 and +1. The justification is as follows:

The perpendicular distance from a pixel at (i, j) to the straight line, $y = \lambda\,(x - c)$, is:

$$D_{i,j} = \frac{\lambda(i-c)-j}{\sqrt{1+\lambda^2}}.$$

Without loss of generality we can assume that the slope $|\,\lambda\,| \geq 1$, because if not we simply need to exchange x and y. The line intersects the x–axis at c, and we can choose the x origin to be such that $c \in [-\frac{1}{2},+\frac{1}{2}]$. Further, we assume that the positioning of the sheep in the X–ray machine has sufficient variability that we can assume the boundary is positioned randomly with respect to the pixel lattice, and therefore that c is uniformly distributed between $-\frac{1}{2}$ and $+\frac{1}{2}$, i.e.:

$$c \sim U\left(-\frac{1}{2},+\frac{1}{2}\right).$$

It follows that:

$$D_{i,j} \sim U\left(\frac{\lambda(i-\frac{1}{2})-j}{\sqrt{1+\lambda^2}},\frac{\lambda(i+\frac{1}{2})-j}{\sqrt{1+\lambda^2}}\right).$$

If we restrict to points on the x–axis (i.e. $j = 0$), and choose i at random from the set $(-n, -(n-1), \ldots, n)$, for $n \geq 1$, subject to the restriction that $D_{i,j} \in [-1,+1]$ then it is not difficult to see that:

$$D_{\cdot,0} \sim U(-1,+1).$$

The same result applies for other values of $j \in (-(n-1), -(n-2), \ldots, (n-1))$ provided that $|\,\lambda\,| \geq 1$. Therefore, if we also choose j at random from the set $(-(n-1), -(n-2), \ldots, (n-1))$ then:

$$D_{\cdot,\cdot} \sim U(-1,+1),$$

which is the result we require.

In a mixed pixel, the proportion of fat, denoted ρ, is related to D by:

$$\rho = \Phi\left(\frac{D}{\tau}\right). \tag{2}$$

Differentiating this, we obtain:

$$\frac{d\rho}{dD} = \frac{e^{-D^2/2\tau^2}}{\sqrt{2\pi\tau^2}}.$$

As D has probability density function:

$$f(D) = \frac{1}{2}dD \quad \text{for} \quad 1 \geq D \geq -1,$$

it follows that the probability density function for ρ is:

$$\begin{aligned} f(\rho) &= \frac{\sqrt{2\pi\tau^2}}{2e^{-D^2/2\tau^2}}d\rho \\ &= \sqrt{\frac{\pi\tau^2}{2}}\, e^{[\Phi^{-1}(\rho)]^2/2}d\rho \quad \text{for} \quad \Phi(1/\tau) \geq \rho \geq \Phi(-1/\tau). \end{aligned} \tag{3}$$

Figure 4 shows this distribution. As far as we are aware, it has never previously been proposed. It is not included in the many distributions given by Johnson, Kotz and Balakrishnan (1994), derived from the normal distribution. We term it the *mixed pixel distribution*. Santago and Gage (1995) also derived a mixed pixel distribution but, in our view, mistakenly assumed ρ to be uniformly distributed.

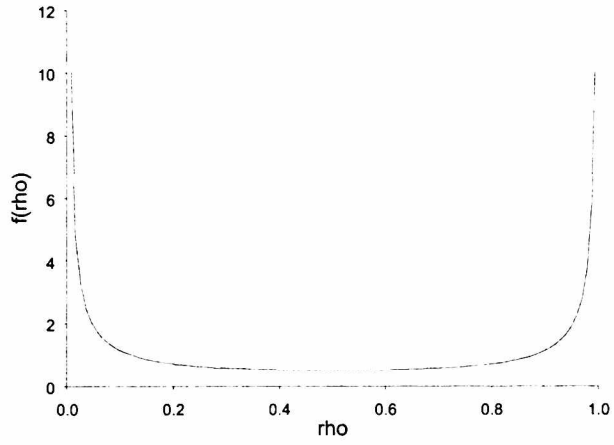

Figure 4. *The new mixed pixel distribution for ρ, the proportion of fat in a mixed pixel, given by (3)*

Conditional on ρ the pixel value is normally distributed, as follows:

$$y \mid \rho \sim \mathrm{N}(\{\rho\mu_f + (1-\rho)\mu_m\}, \{\rho\sigma_f^2 + (1-\rho)\sigma_m^2\}). \tag{4}$$

By convolving this with (3), we obtain the distribution of mixed pixels:

$$f(y) = \int_{\Phi(-1/\tau)}^{\Phi(1/\tau)} \frac{e^{-[y-\rho\mu_f-(1-\rho)\mu_m]^2/2[\rho\sigma_f^2+(1-\rho)\sigma_m^2]}}{\sqrt{2\pi[\rho\sigma_f^2+(1-\rho)\sigma_m^2]}} \sqrt{\frac{\pi\tau^2}{2}} e^{[\Phi^{-1}(\rho)]^2/2}\, d\rho. \tag{5}$$

This has no analytic solution, but can be computed by numerical integration for any specified values of the parameters.

Finally, we combine (5) with that for pure pixels, to obtain a probability density function of pixel values of:

$$f^*(y) = p_f \frac{e^{-[y-\mu_f]^2/2\sigma_f^2}}{\sqrt{2\pi\sigma_f^2}} + p_m \frac{e^{-[y-\mu_m]^2/2\sigma_m^2}}{\sqrt{2\pi\sigma_m^2}} + (1-p_f-p_m) f(y). \tag{6}$$

We can estimate the 6 parameters (μ_f, μ_m, σ_f^2, σ_m^2, p_f, p_m) in (6) numerically by maximising the log–likelihood:

$$L = \sum_y n_y \log f^*(y),$$

where n_y is the number of pixels taking value y in histograms such as Figure 2. Or, we could assume μ_f, μ_m, σ_f^2 and σ_m^2 are known and only estimate p_f and p_m. Figure 5 shows the fit to the histogram in Figure 2, with all six parameters estimated. For comparison, the histogram of pure pixels alone is also shown for the same values of μ_f, μ_m, σ_f^2 and σ_m^2, from which the effect of mixed pixels is evident. The agreement between the data and the distribution given by (6) can be seen to be excellent.

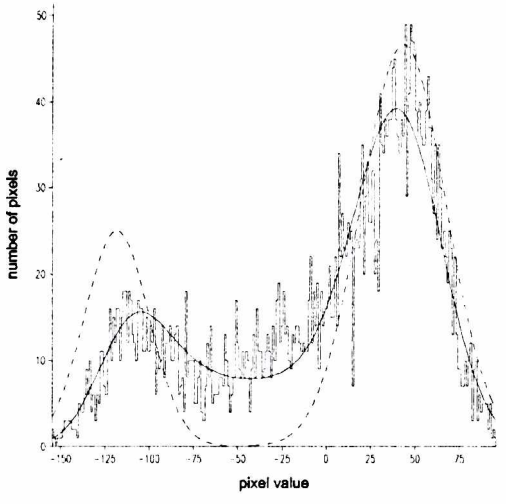

Figure 5. *Histogram of pixel values in Figure 1b, together with: (—) maximum likelihood fit of distribution (6) and, (– – –) the same distribution but with the mixed pixel component omitted*

3. Simulation

We compared three methods for estimating tissue areas from a histogram of pixel values:

1. Choose a threshold midway between the means of the two normal densities:

$$t_1 = \frac{\mu_f + \mu_m}{2},$$

and estimate the fat area by the number of pixels less than t_1. This is the method in current usage in the SAC–BioSS CT unit.

2. Choose as a threshold the point of minimum misclassification, where the two normal probability densities cross. This is the solution of:

$$\frac{1}{\sigma_f}\exp\left\{\frac{-(y-\mu_f)^2}{2\sigma_f^2}\right\} = \frac{1}{\sigma_m}\exp\left\{\frac{-(y-\mu_m)^2}{2\sigma_m^2}\right\},$$

which leads to:

$$t_2 = \frac{(\sigma_f^2\mu_m - \sigma_m^2\mu_f) \pm \sigma_f\sigma_m\sqrt{(\mu_f-\mu_m)^2 + 2(\sigma_f^2-\sigma_m^2)\log(\sigma_f/\sigma_m)}}{(\sigma_f^2-\sigma_m^2)},$$

where the sign of ± is chosen to ensure $\mu_f < t_2 < \mu_m$. Again we estimate the fat area by the number of pixels less than t_2.

3. Fit (6) numerically by maximum likelihood to the histogram, and estimate the fat area to be proportional to $\{\hat{p}_f + \frac{1}{2}(1-\hat{p}_f-\hat{p}_m)\}$.

Further, we either assumed that values of μ_f, μ_m, σ_f^2 and σ_m^2 were correctly known, or had to be estimated from the histogram by maximum likelihood. For the two threshold methods, we fitted a mixture of two normal distributions, excluding the mixture component.

Images were simulated of size 60 × 60 to approximate the number of pixels in histograms such as Figure 2. A set of seven randomly–positioned parallel lines were simulated which crossed each image, to represent alternate layers of fat and muscle tissue. By choosing seven boundaries, the number of mixed pixels matched that in typical sheep CT images. True tissue areas were computed using standard formulae from coordinate geometry. Perpendicular distances of pixels were computed from their nearest boundary. Pixels more than a unit distance from a boundary and lying in a fat region, were generated from normal distributions with mean $\mu_f = -125$, variance $\sigma_f^2 = 300$, chosen to represent typical values from sheep CT images. Similarly, muscle pixels were generated from normal distributions with mean $\mu_m = 45$, variance $\sigma_m^2 = 500$. The remainder of the pixels were represented as mixed pixels. Based on the distance, D to the nearest boundary, ρ was computed using (2), and then a pixel value was drawn from the distribution given in (4). Finally, a histogram of pixel values was computed from each simulated image.

The six methods of estimation were compared on the basis of the root mean square error (RMSE) of estimated fat area, averaged over 100 independent simulations. If, in the ith simulation, the true fat area is F_i and the estimated area using a particular method is $\hat{F}_i$, then

$$\text{RMSE} = \sqrt{\frac{1}{100}\sum_{i=1}^{100}(F_i - \hat{F}_i)^2}.$$

Results are given in the middle row of Table 2. We see that the new method outperforms both of the threshold–based ones, irrespective of whether parameters are known or have to be estimated. To check the sensitivity of the results to the particular choices of parameters, simulations were repeated for other values of σ_f^2 and σ_m^2, it not being necessary to also vary μ_f and μ_m. These are also given in Table 2. The threshold methods sometimes give biased results and this accounts for the RMSEs sometimes being large. The new method, both when parameters are known or estimated, consistently estimates fat area more accurately than the other methods, except when $\sigma_f^2 = 600$, $\sigma_m^2 = 1000$, often by a considerable margin.

Table 2. *Root mean square errors of estimated fat area, by applying six methods to data from 100 simulations for each of a range of values of σ_f^2 and σ_m^2*

		parameters known			parameters estimated		
σ_f^2	σ_m^2	t_1	t_2	*new*	t_1	t_2	*new*
150	250	6.0	25.2	4.7	9.5	51.4	5.0
	500	13.5	57.0	5.4	8.5	74.6	8.1
	1000	51.4	79.7	7.3	27.0	95.0	18.6
300	250	5.8	11.2	5.2	8.4	13.4	5.5
	500	10.9	23.5	5.8	9.5	38.3	8.5
	1000	48.6	47.6	7.7	30.3	64.5	22.3
600	250	21.1	39.0	6.6	19.6	33.9	8.7
	500	13.3	9.4	7.1	18.8	18.5	9.9
	1000	30.9	26.4	8.9	22.1	46.1	24.9

4. Discussion

We have developed a new mixed pixel distribution, and demonstrated its use in estimating tissue proportions in images more accurately than using threshold–based methods. Although the application used X–ray CT images and involved estimation of sheep tissue proportions, the methodology is generic and could be used with other imaging modalities and subject domains, such as human obesity studies.

Further work is needed 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. Methods have already been developed for cases in remote sensing where the image is multivariate (Foody and Cox, 1994).

Also, the assumption that distances D are uniformly distributed between -1 and $+1$ may need to be modified for images with many fine tissue structures, by using nonparametric density estimation.

Acknowledgements

The authors wish to thank Chris Theobald and Mark Young for advice with this work, which was funded by the Ministry of Agriculture, Fisheries and Food, the Meat and Livestock Commission and the Scottish Executive Rural Affairs Department as part of the LINK Sustainable Livestock Production Programme (Project LK 0607), and by EPSRC funding the second author.

References

[1] Dore S., Kearney R.E. and De Guise J.A. (1997) Experimental correlation–based identification of X–ray CT point spread function. Part 1: method and experimental results, *Medical & Biological Engineering & Computing*, **35**, 2–8

[2] Foody G.M. and Cox D.P. (1994) Sub–pixel land cover composition estimation using a linear mixture model and fuzzy membership functions, *International Journal of Remote Sensing*, **15**, 619–631

[3] Glasbey C.A. and Horgan G.W. (1995) *Image Analysis for the Biological Sciences*, Wiley, Chichester

[4] Glasbey C.A., Robinson C.D. and Young M.J. (1999) Segmentation of X–ray CT images using stochastic templates, *ICIAP'99 10th International Conference on Image Analysis and Processing*, Venice, 746–751

[5] Johnson N.L., Kotz S. and Balakrishnan N. (1994) *Continuous Univariate Distributions*, Volume 1 (2nd edition), Wiley, New York

[6] Santago P. and Gage H.D. (1995) Statistical models of partial volume effect, *IEEE Transactions on Image Processing*, **4**, 1531–1540

[7] Simm G. (1992) Selection for lean meat production in sheep. In *'Recent Advances in Sheep and Goat Research'*, (ed A.W. Speedy), CAB International, 193–215