simultaneously given all of the data, whereas with performance measured by the percent of misclassified pixels one seeks the univariate posterior modes. Other algorithms are not driven by minimizing penalties. Professor Besag uses the maximum likelihood classifier (maximizing each pixel label given only the record there) as the starting point for *ICM*, in which one successively maximizes each label given the neighbouring labels and all the data.

ICM is a form of iterative improvement on the posterior likelihood. We have also used (local) "quenching" in experiments on segmentation and tomography, and one often arrives at a "satisfying" local maximum. In boundary-finding, our model accounts for intensities, micro-edges, and boundaries (macro-edges). The posterior is the conditional distribution on edges and boundaries given the (undergraded) intensities. We use stochastic relaxation to find a starting point and then a form of iterative improvement referred to here as "block reconstruction" to obtain a final labeling. In the case of texture segmentation we have stuck with annealing.

Obviously these issues are wide-open and ripe for further study. Professor Besag has made a most sensible start.

Professors S. Geman and D. E. McClure (Brown University, U.S.A.): Professor Besag's paper presents a timely and cogent discussion of alternative loss functions that can be used to determine algorithms for image reconstruction based on Bayesian models. The use of probabilistic models in image analysis has opened new possibilities for basing algorithms on a firm scientific foundation, and the clear connection of the methodology to the principle of minimum Bayes risk is pivotal for this foundation.

We have used the *ICM* algorithm for recent computational experiments in single photon emission computed tomography (*SPECT*; section 4.2.2) and attest to its effectiveness when the iterative reconstruction is started with a good initialization. We firmly believe that *special purpose* algorithms of this type—a judicious initialization together with simple suboptimal iterative improvement—are very important for the utility of the Bayesian methods in actual applications. By adapting the algorithm to the application, one can take advantage of special structure in the underlying model or of good alternative reconstructions to reduce the computational burden.

The *SPECT* example is an interesting departure from Professor Besag's Assumption 1 and the example in Section 5.2, where observables y have a local dependence on the true image x. In *SPECT* the dependence is manifestly global, and hence the posterior distribution $\pi(x|y)$ has nonlocal neighborhoods. Nonlocal neighbourhoods increase the computational burden, but they do not contradict the conceptual foundation of the model-based reconstructions.

While special purpose algorithms will determine the utility of the Bayesian methods, the general purpose methods—stochastic relaxation and simulation of solutions of the Langevin equation (Grenander, 1983; Geman and Geman, 1984; Gidas, 1985a; Geman and Hwang, 1986)—have proven enormously convenient and versatile. We are able to apply a single computer program to every new problem by merely changing the subroutine that computes the energy function in the Gibbs representation of the posterior distribution.

Mr J. Haslett (Trinity College, Dublin): There are, by now, a number of methods for restoring noise corrupted images. Many of these will achieve similar restorations, and not only as measured by the per pixel error rate. The emphasis by now is on the second order differences between these methods.

Professor Besag has concentrated on the differences between two 'statistical' methods, his own and that of Geman and Geman. The practical differences are computational, although there are some interesting insights into the use of Markov random fields. Both achieve very similar reconstructions to the 'non-statistical' method of relaxation, popular for many years in image processing.

The second order questions to which statisticians might be expected to contribute include the questions that are normal in any statistical classification procedure—accuracy overall (error rate in a given study) and posterior probabilities for a given case (pixel)—as well as the identification of outlying cases (pixels).

Indeed overall accuracy is the first question to be posed by any potential client thinking of using remote sensed imagery as opposed to traditional ground based surveys in, for example, an ecological inventory. At a higher level in image analysis, in terms of identifying objects—such as a cancerous growth—the same sort of questions arise: what is the posterior distribution of the size of the tumour which is suggested by image analysis?

It seems from Professor Besag's paper that *all* such questions can be answered (in principle) by sufficient simulations at a temperature $T = 1$, under Geman and Geman's method, but that *no* such questions can be derived from the *ICM* method. This latter derives from the notorious difficulty, already