

Movies for the Visualization of Output from a Bayesian Analysis of Corbelled Domes

Nicole A. Lazar, Joseph B. Kadane and Fang Chen

Department of Statistics, Carnegie Mellon University, Pittsburgh, PA 15213 USA
nlazar@stat.cmu.edu, kadane@stat.cmu.edu, fangc@stat.cmu.edu

Abstract. The fitting of complex statistical models for archaeological data often necessitates the use of advanced computational techniques, such as Markov Chain Monte Carlo (MCMC). In this paper, we give a brief introduction to MCMC methods. The usefulness of the methodology is demonstrated on a multiple changepoint model fit to data collected on the Treasury of Atreus in Mycenae. A question that often arises out of MCMC simulations, especially with very complicated and high-dimensional models, is how to display the results in such a way that the behavior of the parameters can be assessed. We propose a statistical movie as one way of answering this question. This paper provides guidelines on building such movies and on how to watch them, as well as indicating other benefits that can be derived from the approach.

1 Archaeological Context – Corbelled Domes

Corbelling is an ancient technique, used before the development of the true dome, for spanning or roofing spaces. Examples of this construction method have been found in a number of pre-historic societies, ranging from tombs in Iberia and the British Isles, to buildings in Italy, to passages and chambers of the Great Pyramid of Kheops in Giza, Egypt. Most of these utilize corbelling to span relatively small distances, 2 to 3 meters at most, with a conservative slope.

An impressive use of the technique, which involves laying stones on top of each other, with successive layers being slightly offset (think of using rectangular Lego pieces to build an arch), comes from Mycenae in Greece, where archaeologists have found corbelled structures spanning 8 to 14 meters. As noted by Cavanagh and Laxton (1981), only the invention of the true dome enabled builders to bridge larger distances without internal support.

Due to the prevalence of corbelling in the ancient world, and the fact that many of the structures that were built with this spanning method are still standing today, thousands of years later, archaeologists have been interested in studying this technique more closely. In particular, interest has focused on finding models that will adequately describe the shape of corbelled domes, and perhaps give insight into how these structures were built in practice. One question that has drawn speculation has to do with the diffusion of the knowledge required to build a corbelled dome. For instance, was the technique developed independently in different regions of the Mediterranean where examples have been found? Or did one group learn from another?

In the first mathematical study, Cavanagh and Laxton (1981) considered the structural mechanics of the Mycenaean *tholos* tombs. By examining the structure of these tombs and the factors necessary for the vault to remain standing, they came up with the following, deterministic, model:

$$\text{radius} = \alpha \cdot \text{depth}^\beta,$$

where *depth* is the depth below the top of the tomb, *radius* is the radius of the tomb at the measured depth, and α and β are

constants to be determined for each tomb individually. Taking logarithms of each side of the equation, a simple linear equation results, and it is possible to find values for the two coefficients using least squares. Cavanagh and Laxton examined five tombs that varied in location, age and building style, and found that all five had a value of β of around $2/3$. Values of α varied, as this parameter depends on the size (specifically, the height) of each tomb. One of the tombs, the Treasury of Atreus, is almost twice the size of the other tombs in this study, and had $\alpha = 2.7$, compared to $\alpha \approx 2$ for the other four.

Their model, while providing a starting point for the analysis of the Mycenaean tombs, was lacking in two important aspects, addressed in Cavanagh, Laxton and Litton (1985). First, many of the *tholoi* were built partly below ground. The walls in this portion of the tomb may have been built according to a different plan than the part that was above the ground. For instance, they may have been straighter, with the stones lining the side of a pit constructed to hold the body. That is, instead of one value of α and one value of β sufficing to describe the entire tomb, there might be a point at which the building regime, and the parameters guiding construction, change. The changepoint, if it exists, would correspond to where the dome and the wall meet. The first simple model considered by Cavanagh and Laxton was based only on measurements from the dome, and not the entire tomb. Thus, whereas their simple form was the basic one required for stability of the dome, it did not necessarily describe adequately the Mycenaean tombs of mainland Greece or the Minoan *tholoi* on the island of Crete. Furthermore, it might not be possible to measure from the apex of the dome, that is, at depth = 0. There are two reasons for this: the top of the tomb might have collapsed, or the corbelling wasn't originally continued to the top of the arch, instead a large slab was placed down on the stones, to close off the dome. Taking under consideration these two modifications, and introducing a stochastic element to account for errors, the proposed model became

$$\log \text{radius}_i = \begin{cases} \log \alpha_1 + \beta_1 \log(\text{depth}_i + \delta) + \varepsilon_i, & \text{depth}_i \leq \gamma \\ \log \alpha_2 + \beta_2 \log(\text{depth}_i + \delta) + \varepsilon_i, & \text{depth}_i \geq \gamma. \end{cases}$$

Here, γ is the location of the changepoint, and the ε_i are assumed to be normally distributed, with mean 0 and variance σ^2 . Physical considerations require that $\beta_1 > \beta_2 > 0$. At depth γ , the two lines are required to intersect, that is, $\alpha_1(\gamma + \delta)^{\beta_1} = \alpha_2(\gamma + \delta)^{\beta_2}$, so that there are only six free parameters in the problem, instead of seven. In particular, β_2 can be expressed as $\beta_1 + [\log(\alpha_1/\alpha_2) / \log(\gamma + \delta)]$.

These first attempts at statistical modeling were all within the classical, or frequentist, framework. More recent approaches to this problem have been from the Bayesian point of view. The next section gives a brief review of some of the relevant ideas and technology.

2 An Introduction to Markov Chain Monte Carlo Methods

A simple probabilistic equality, known as *Bayes' theorem*, underlies a statistical approach that has been advocated as an alternative to classical techniques. From the Bayesian perspective, parameters of a model are random quantities, just like the observed data. Inference proceeds by specifying a likelihood for the data and a *prior distribution* for the parameters. Bayes' theorem provides statisticians with a way of combining the information in a sample with prior beliefs, to get new, updated beliefs about the parameters, in light of the observed data. The new beliefs are summarized in a *posterior distribution*. In brief, if $p(y|\theta)$ is the likelihood of some data y , which depends on parameters θ , and $p(\theta)$ is the prior distribution of θ , then Bayes' rule says that the posterior distribution of θ is given by

$$p(\theta | y) = \frac{p(\theta)p(y|\theta)}{\int p(\theta)p(y|\theta)d\theta}.$$

Historically, one of the difficulties in applying a Bayesian analysis has been the need to evaluate difficult integrals. In complicated problems, the required integrals are often over high-dimensional spaces. Note that integration over high dimensions is not uniquely a Bayesian problem, but it is perhaps more prevalent there than in other statistical contexts. With the rise of computational power (both in terms of speed and memory), it has become more feasible to use Monte Carlo simulation techniques to approximate expectations using averages. Markov Chain Monte Carlo methods in particular have become very popular in recent years. The general idea is quite simple. We want to evaluate an integral, say $E(f(\theta|y))$, the expected value of some function of the posterior distribution of θ . Monte Carlo integration approximates the desired quantity by drawing a sample x_1, x_2, \dots, x_n from $p(\theta | y)$ and setting

$$E(f(\theta | y)) \approx \frac{1}{n} \sum_{i=1}^n f(x_i).$$

Statistical limit laws guarantee that if the x_i are independent, then for large enough n , the approximation can be made arbitrarily accurate. In practice, it might not always be easy, or even possible, to draw independent samples from the posterior distribution, as this often has a complicated form. However, it is possible to draw dependent samples, and, as long as the x_i are sampled roughly according to the proportions allotted them by the poste-

rior density, the Monte Carlo estimate should be reasonable. One particular way to do this is to use Markov chains, with stationary distribution the target from which we wish to sample.

Markov chain theory (see, for example, Hoel, Port and Stone, 1972) implies that if a Markov chain is run long enough, it will converge to a unique *stationary distribution*, under some regularity conditions. The speed with which this happens depends on the starting condition of the chain. After the initial *burn in* period, the chain “forgets” its starting point, and all subsequent samples are assumed to be dependent draws from the stationary distribution. An average based on the observations after discarding the initial iterations, which are not from the target distribution, can then be used as an approximation to the desired integral.

The strength of this approach is that it becomes feasible to sample from the posterior distributions that result from even complex statistical models. Many analyses now routinely include an application of Markov Chain Monte Carlo (MCMC). The technique has also generated a good deal of methodological research.

We have given here only the briefest of descriptions of what is a very rich field of research in modern statistics. For more detail on both the theory and implementation issues, a good reference is Gilks, Richardson and Spiegelhalter (1996).

3 Bayesian Approaches to the Problem

Returning now to the study of corbelled domes, an important step forward was taken by Buck, Litton and Stephens (1993), who used methods developed in Stephens (1994), for the analysis of changepoint problems from a Bayesian point of view. The authors applied MCMC methods to a Minoan tholos tomb at Stylos, Crete. Their base model was as in previous work. The conditions on the model were also as described above. With the assumption of normal errors, the likelihood is easily specified, although sampling from the posterior is not simple, regardless of the prior specification, because of the parameter constraints and the piecewise nature of the model. Buck and colleagues chose non-informative (vague) priors for $\log \alpha_1$, $\log \alpha_2$, β_1 , σ^2 and γ ; the prior for δ was uniform between 0 and 0.34, to account for the thickness of the capping stone on this tomb, which was 0.34 meters.

Under this specification, the Markov Chain Monte Carlo sampling scheme is straightforward, using an algorithm called Gibbs sampling, which draws successive observations from the posterior distribution of each parameter given all of the others. The four parameters $\log \alpha_1$, $\log \alpha_2$, β_1 and σ^2 all have conditional posterior distributions from standard families, which are easy to sample from. On the other hand, the full conditional densities for δ and γ are both of non-standard form, hence are sampled using a technique called *rejection sampling* (see Buck, Litton and Stephens, 1993 and Stephens, 1994 for details).

A number of advantages of the Bayesian approach are apparent for this problem. Instead of just point estimates of the parameters of interest, entire posterior distributions are obtained. This gives archaeologists more flexibility in studying and interpreting the results of the analysis. The classical analysis did not incorporate any uncertainty about δ into the inference process, whereas

this was an outcome of the Bayesian analysis by its very nature. Furthermore, information about the main parameter of interest, the location of the changepoint, was refined. The 95% confidence interval reported for γ that arose from the classical analysis was 1.64-1.84; the 95% highest posterior density interval from the Bayesian analysis was 1.48-2.44. As pointed out by Buck *et al.*, this latter interval just includes the depth at which the lintel is below the capping stone of the tomb, 2.44 meters. This finding supported a supposition by archaeologists that the placing of the lintel would cause a break in the process of corbelling, with different structural forces coming in to play above and below that point. Using the classical approach, there was no support for this intuition, as the changepoint was placed too high, and the confidence interval was too short.

Building on this result, an obvious question to ask is whether there might be more than one changepoint. However, the appropriate number of changepoints is not known, and so it would be desirable to include this as a parameter in a new model. While seemingly a minor modification of the existing framework, in fact letting the number of building segments vary introduces non-trivial complications. The problem is that, the practical issue of how to estimate the true dimension of the model, by sampling over the space of models with differing numbers of changepoints, is not an easy one to resolve. In addition, as the number of changepoints varies across models, interpretation of the model parameters also changes.

Several solutions to the problem of estimating the model dimension have been proposed in the Bayesian literature. A promising approach is *Reversible Jump MCMC* (Green, 1995; Richardson and Green, 1997). The reversible jump technique designs a Markov chain that jumps between models of different dimensions, while still preserving the conditions needed to eventually reach a stationary distribution. Implementation of the method can be complicated and the details tend to be unique to each application. In their study, Denison, Mallick and Smith (1998) derived a reversible jump algorithm for dealing with multiple changepoints, but they, nevertheless, made certain simplifying assumptions that were not appropriate for the problem we were working on. More recently, Fan and Brooks (2000) developed methodology specifically for the analysis of corbelled structures.

Fan and Brooks extended the basic model, to allow for two generalizations. Firstly, δ was not restricted to be the same in each linear piece. Secondly, at lower depths, the wall of the tomb might be nearly vertical, that is, governed by a single parameter. An example of a candidate two changepoint model is

$$\log \text{radius}_i = \begin{cases} \log \alpha_1 + \beta_1 \log(\text{depth}_i + \delta_1) + \varepsilon_i, & \text{depth}_i \leq \gamma_1 \\ \log \alpha_2 + \beta_2 \log(\text{depth}_i + \delta_2) + \varepsilon_i, & \gamma_1 \leq \text{depth}_i \leq \gamma_2 \\ \log \alpha_3 + \varepsilon_i & \text{depth}_i \geq \gamma_2 \end{cases}$$

The number of segments was allowed to increase to a maximum of four, by permitting as many as three changepoints. The reversible jump algorithm implemented by Fan and Brooks also explored the possibility that δ here was no constant segment at the bottom of the tomb. In addition to reanalyzing the data from the tomb at Stylos in Crete, Fan and Brooks compared the results of their reversible jump MCMC algorithm for another Minoan

tholos to those for a Mycenaean tholos and a Sardinian nuraghe. They found that for each, the method indicated different underlying models as having the highest posterior probabilities. Furthermore the estimated values of the β parameters were similar in all of the buildings taken under consideration in this comparison and as a result the authors concluded that this information alone was not enough to distinguish among structures of differing origin.

4 Visualizing MCMC Results

Reversible jump MCMC and other methods that have been suggested for exploring a space of models of differing dimensions (for example, Chib, 1995) give rise to some interesting questions about interpretation, diagnostics and visualization. As mentioned above, the parameters of models of different sizes may not be directly comparable, since the meaning of a parameter may change as the model dimension changes. Fundamental statistical issues of model selection become relevant and even pressing as statisticians and scientists start to explore more complex model spaces.

Devising visualization techniques for MCMC is an area of active research. Many of the existing methods involve displaying a trace plot of the parameters, or some function of them, along iterations of the chain. Due to the complexity of the reversible jump process, displays developed for ordinary MCMC are not applicable. Elsewhere (Lazar and Kadane, 2002), we have suggested a graphical device, in the form of a statistical movie, for visualizing the output of a Markov Chain Monte Carlo simulation and for assessing characteristics of the chain.

The status of some (but not necessarily all) MCMC runs at a given instant in the chain can be summarized in a plot. For any such summary, the stringing together of the plots to display the development of the chain over time, makes a movie. Key to the development of a good movie is to have an efficient and informative summary of the model, a “clear portrayal of complexity” (Tufte, 1983, pg. 191). Some thought therefore needs to be dedicated to finding an appropriate graphical representation, which displays as many of the parameters as possible; in particular, the movie should enable the viewer to see changes that unfold in the form of the model as the dimension switches, when this is a relevant consideration. While the proposed tool is useful for any MCMC problem where the output at a given iteration may be summarized in a single graph, we expect it to be especially helpful for examining the results of a reversible jump chain.

This approach is intrinsically a multivariate one, since it enables the examination of the posterior distributions of several parameters at once. As such, it avoids some of the problems of marginal or univariate displays. It is easy to interpret, and to extract interesting features of the MCMC simulation and the posterior distribution of the parameters from watching the movie as it runs.

We used the algorithm developed by Fan and Brooks (2000) to model the data collected on the Treasury of Atreus, allowing as many as four changepoints, and analyzed the output using a statistical movie. A crucial step in using the movie is to find a suitable graphical display of the underlying model connecting the

parameters. In the current situation, this is not hard to do - just plot the data points, $\log x_i$ on the x axis and $\log y_i$ on the y axis, and connect these with the fitted lines defined by the changepoints. For each iteration in the chain, we plotted the data and the fitted model. For this data set, the analysis was somewhat simplified by the fact that δ was identically zero, reducing the number of parameters to be sampled.

The plots below are summary snapshots of some of the most popular model types found by the movie. As we move from left to right along the x axis, depth increases, so the top right corner of the plot is actually the bottom of the structure. Each plot shows a representative snapshot from models of that type, for instance, two changepoint models with a constant segment at the end. Also plotted are the posterior distributions of the changepoint locations for models of that type. The changepoint placement is, in most cases, quite spread out, meaning that there is some uncertainty within a model type of where the changepoints should be.

As can be seen in the figures, models of different dimension and specification can fit the data almost equally well. Indeed, inspection of the movie showed that most fits were quite good, with only the occasional bad-fitting model. Such bad models were always quickly moved away from. We found that the chain tended to favor models of intermediate size, i.e., those with two or three changepoints. Placement of these changepoints corresponds roughly to structural features of the Treasury of Atreus - namely the lintel and the door.




Fig. 1.

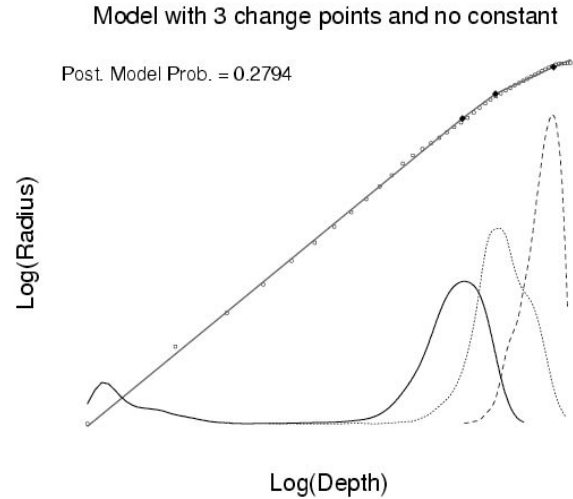


Fig. 2.

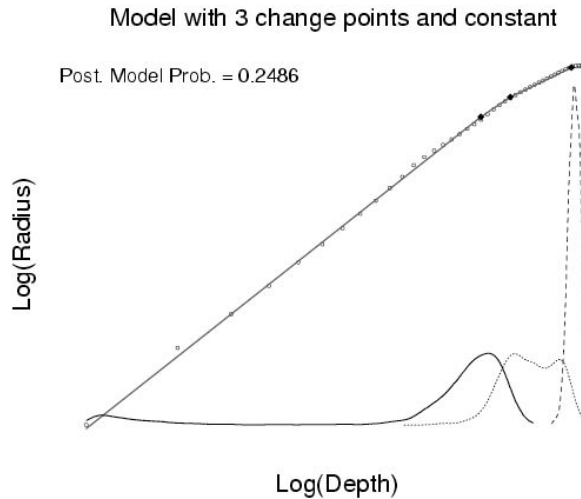


Fig. 3.

This movie made it possible to monitor the mixing and convergence of the chain - by observing changes in color, for instance, we saw that the chain favored models of particular dimensions. While the same observation could have been made by looking at a marginal trace plot for the model dimension, the movie allowed us to see in addition how the locations of the changepoints jumped from model to model, and how the slopes and intercepts of the individual lines varied to account for all of these simultaneous moves in the values of the other parameters. We noted for example that models with no changepoints resulted in poorer fits overall. Models with one or two changepoints that were generated shortly after a zero changepoint model tended to also have poor fits; on the other hand, one or two changepoint models that arose from higher dimensional models had better fits. In addition, as might be hoped, the model tended to add changepoints where they were most needed. If a changepoint was added very close to an already-existing changepoint, the fit was somewhat worse, apparently due to adding in a superfluous segment. Our movie thus offered a dynamic way of viewing the re-

sults of the MCMC run. Furthermore, we were able to present all of the parameters of interest, with the exception of the error variance, σ^2 , although we were able to get at least a qualitative impression of this parameter as well. We also created a version of the movie that included σ^2 , by plotting one standard deviation bars around the data. While this movie was informative, we felt that the result was more cluttered and harder to follow.

Using the movie technique it is very easy to make changes in the burn-in period (simply start the movie from a later or earlier frame and watch the differences), and to thin the chain (by only showing every r^{th} frame, for any value of r we might choose). Therefore multiple facets of the chain can be monitored at once. In an early implementation, for example, we noticed that there was a fair amount of “clumping” even in a multivariate sense, that is, there were portions of the chain characterized by unchanging dimension, with little movement in the existing changepoints. Plotting only every tenth frame alleviated this problem and revealed the important features of the fitted models.

The movie also displayed other characteristics that might have been difficult to spot graphing the chain parameter by parameter. We saw that the chain concentrated most of its effort on trying to fit the data well near the bottom, by placing changepoints in the lower half of the data more than in the upper half. A possible explanation for this is that the builders of the tomb had to make adjustments near the bottom to accommodate features of the ground in which the structure was placed – as previously mentioned, many Mycenaean tombs were built inside a hole dug into the earth, or inside mounds that helped provide stability, and this needed to be taken into account in the early stages of construction. Also, changes in building strategy might have been necessary around the lintel and the door, and there is some evidence for this in the fitted models. Near the top, fewer adjustments were necessary.

The flexibility of the movie makes it easy to perform sensitivity analyses under the sampling-resampling perspective of Smith and Gelfand (1992), where the importance weight for going from one prior to another could be represented by the length of the frame (equivalently, and perhaps easier to program, the number of consecutive times the frame is shown). Another use for this idea would be to sample under one prior that produces a chain that converges quickly and mixes well and then to interpret the results under another prior that more realistically summarizes the user’s beliefs about the state of nature, as was done in DiMatteo and Kadane (2001).

On a more pragmatic note, the movie can also be used for debugging. An early run of the movie revealed that the piecewise linear segments were not matching up at the changepoints, which is required by the model (and is obviously of structural importance, as the building wouldn’t stand if there were discontinuities). Simply looking at posterior distributions of the individual parameters would not have uncovered this break in the model specification, nor would any of the usual diagnostics applied to the output of the chain.

A general sketch of the method is as follows:

1. Choose or find a meaningful graphical summary of the model at a given instant in the running of the chain. Some creativity might be needed here! For particular situations, such as regression and related models, the choice is relatively straightforward. Other models might require more thought. In any case, the summary that is chosen should incorporate as many of the parameters as is feasible, and should provide the viewer (user) with a snapshot of the model at any given stage of the chain.
2. Set a burn-in period of B iterations, and a thinning factor of r . These may be determined theoretically (see for example, Raftery and Lewis, 1992), from inspection of trace plots, or from an initial viewing of the movie based on all iterations of a pilot simulation study.
3. For every r^{th} iteration after B (r might be 1), plot the graphical summary from step (1), using the parameter values for that iteration.

Programming the movie is not at all difficult - a few lines of code in any software package that can produce graphs and loop through them, are all that is required. How quickly the program cycles through the frames of the movie needs to be controlled, since if it goes too quickly, it is impossible to discern interesting features and patterns, whereas if it moves too slowly, it can be hard to form impressions about the status of the chain. Mechanisms for varying the speed of the movie will differ according to the platform. Redrawing the axes for each frame of the movie may produce a distracting flicker effect, which can make it harder to concentrate on the unfolding of the simulation. The principles of good graphics, put forth by Tukey (1977) and Tufte (1983), should guide the design not only of the individual frames of the movie, but also of the display taken as a dynamic whole.

5 Conclusion

As archaeologists begin to fit more complex models to their data, using simulation techniques such as Markov Chain Monte Carlo, visualization of results will take a more central role. We have presented here a dynamic tool for visualizing the output of an MCMC simulation, which allows the user to monitor multiple model parameters simultaneously, and control aspects of the simulation such as length of burn-in and amount of thinning. This provides the user the opportunity to examine interplay and interconnections among parameters (how does a change in the value of one parameter affect the value of the others) and the ability to study mixing of the Markov chain, in a multivariate sense. The statistical movie is flexible, is easy to program, and can be used with a wide variety of models. Principles of good graphical display should be considered at all stages of implementation.

In the particular example presented here, of fitting piecewise linear models to the Treasury of Atreus, the movie has given us some insight into how the tomb might have been built. Models of intermediate size – two or three changepoints - were favored, although there is a fair amount of uncertainty regarding the location of the changepoint. Watching the movie run made it clear that many possible types of models could fit the data well, a fact which could be taken advantage of in future archaeological analyses of Mycenaean tholoi.

References

- BUCK, C.E., LITTON, C.D. and STEPHENS, D.A. (1993) Detecting a change in the shape of a prehistoric corbelled tomb. *The Statistician*, **42**, 483-490.
- CAVANAGH, W.G. and LAXTON, R.R. (1981) The structural mechanics of the Mycenaean tholos tomb. *Annual of the British School at Athens*, **76**, 109-140.
- CAVANAGH, W.G., LAXTON, R.R. and LITTON, C.D. (1985) An application of change-point analysis to the shape of prehistoric corbelled domes; I. The maximum likelihood method. *PACT*, **11** -- **III.4**, 191-203.
- CHIB, S. (1995) Marginal likelihood from the Gibbs output. *Journal of the American Statistical Association*, **90**, 1313-1321.
- DENISON, D.G.T., MALLICK, B.K. and SMITH, A.F.M. (1998) Automatic Bayesian curve fitting. *Journal of the Royal Statistical Society, Series B*, **60**, 333-350.
- DIMATTEO, I. and KADANE, J.B. (2001) Vote tampering in a district justice election in Beaver County, PA. *Journal of the American Statistical Association*, **96**, 510-518.
- FAN, Y. and BROOKS, S.P. (2000) Bayesian modelling of prehistoric belled tombs. *The Statistician*, **49**, 339-354.
- GILKS, W.R., RICHARDSON, S. and SPIEGELHALTER, D.J. (1996) *Markov Chain Monte Carlo in Practice*, Chapman & Hall, Boca Raton.
- GREEN, P.J. (1995) Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika*, **82**, 711-732.
- HOEL, P.G., PORT, S.C. and STONE, C.J. (1972) *Introduction to Stochastic Processes*, Waveland Press, Illinois.
- LAZAR, N.A. and KADANE, J.B. (2002) Movies for the visualization of MCMC output. *Journal of Computational and Graphical Statistics, In Press*.
- RAFTERY, A. and LEWIS, S. (1992) How many iterations in the Gibbs sampler? In *Bayesian Statistics 4*, Bernardo, J.M., Berger, J., Dawid, A.P. and Smith, A.F.M. (eds), Oxford, U.K., Oxford University Press, 765-776.
- RICHARDSON, S. and GREEN, P.J. (1997) On Bayesian analysis of mixtures with an unknown number of components. *Journal of the Royal Statistical Society, Series B*, **59**, 731-792.
- SMITH, A.F.M. and GELFAND, A.E. (1992) Bayesian statistics without tears: A sampling-resampling perspective. *The American Statistician*, **46**, 84-88.
- STEPHENS, D.A. (1994) Bayesian retrospective multiple-change-point identification. *Applied Statistics*, **43**, 159-178.
- TUFTE, E.R. (1983) *The Visual Display of Quantitative Information*, Cheshire, Graphics Press.
- TUKEY, J.W. (1977) *Exploratory Data Analysis*. Reading, Addison-Wesley.