Appl. Statist. (1991)
40, *No.* 3, *pp.* 461–476

The Evaluation of Fibre Transfer Evidence in Forensic Science: a Case Study in Statistical Modelling

By J. C. WAKEFIELD and A. M. SKENE†,

University of Nottingham, UK

A. F. M. SMITH

Imperial College of Science, Technology and Medicine, London, UK

and I. W. EVETT

Home Office Forensic Science Service Central Research and Support Establishment, Aldermaston, UK

[Received February 1990. Revised November 1990]

SUMMARY

Frequently, when a crime is committed, fibres are left at the scene. This paper examines the modelling aspects of evaluating the evidential content of such fibres by using a Bayesian approach. Inferences are made via the likelihood ratio, derived from bivariate colour measurements. Modelling the distribution of colour within a particular garment is discussed in detail. In addition, a large database allows an empirical prior distribution to be incorporated, utilizing kernel density estimation. Data from actual casework are analysed.

Keywords: Bayesian inference; Colour measurements; Forensic science; Kernel density estimation; Likelihood ratio; Modelling within-garment variability

1. Introduction

1.1. *The Problem*

The potential for the use of Bayesian statistical analysis in general problems of quantifying the evidential content of trace materials in criminal cases has been established by Lindley (1977), Evett (1983, 1987) and Evett *et al.* (1987), the last dealing specifically with the case of fibre transfer.

Consider the following. A crime has been committed during which several fibres have been left at the scene of the crime by the offender. A suspect has been apprehended as a result of a police investigation and a garment belonging to the suspect has been taken for examination. After the examination of the suspect's garment, the available forensic evidence consists of bivariate colour measurements, $y = (y_1, \ldots, y_n)$, from the fibres left at the scene of the crime, together with a set $x =$

†*Address for correspondence*: Department of Mathematics, University of Nottingham, University Park, Nottingham, NG7 2RD, UK.

 0035–9254/91/40461 $2.00

$(x_1, \ldots, x_m)$ of bivariate measurements from a representative sample of m fibres taken from the suspect's garment. The colour of a single fibre is expressed as a pair of *complementary chromaticity co-ordinates*. For details of how these measurements are made, see Laing *et al.* (1986).

Two mutually exclusive and exhaustive hypotheses are to be weighed against each other:

(a) C—the recovered fibres came from the suspect's garment;
(b) $\bar{C}$—the recovered fibres came from some other source.

If I denotes the rest of the information (distinct from the actual colour measurements) which has been assembled in relation to the crime, elementary use of Bayes's theorem establishes that, conditioning throughout on I,

$$\frac{p(C|x,y,I)}{p(\bar{C}|x,y,I)} = \frac{p(x,y|C,I)}{p(x,y|\bar{C},I)} \frac{p(C|I)}{p(\bar{C}|I)} \tag{1.1}$$

posterior odds on C = likelihood ratio for C × prior odds on C.

Focusing attention on the likelihood ratio as the summary of the way in which the colour measurements x, y provide additional evidence over that provided by I, we see that

$$\frac{p(x,y|C,I)}{p(x,y|\bar{C},I)} = \frac{p(x,y|C,I)}{p(x|\bar{C},I)\, p(y|\bar{C},I)} \tag{1.2}$$

if we make the obvious assumption that $\bar{C}$ implies that y is independent of x. Evaluation of this expression therefore gives the evidential import of measurements x, y (given the other information I).

1.2. *Database Information*

The modelling of the distributions involved in the likelihood ratio requires general inputs beyond those derived from any one specific case. To quote Evett *et al.* (1987): 'Forensic scientists have long recognized the need for background data collections to assist in the interpretation of evidence'. In 1982, scientists in UK forensic science laboratories started a collaborative project to amass information on fibres. At each laboratory a small representative sample of cloth was taken from every fifth fabric item submitted for examination. Many of these were garments but other articles such as carpets and bedding were included. A comprehensive questionnaire was completed for each sample, which was then sent to the Scottish College of Textiles, Galashiels, where colour measurements were obtained. In this way, by mid-1987, data including colour measurements had been collected on nearly 8000 samples. Whether such a collection can ever be truly representative of the general population of fibre sources is a moot point which we do not consider here. Certainly, however, the data collection represents the best information that is currently available to the forensic scientist. These data enable an assessment to be made of how rare or common a particular fibre colour is and this underlies the evaluation of $p(y|\bar{C},I)$ in equation (1.2). For each synthetic fibre sample in the collection five pairs of chromaticity co-ordinates were obtained by using five different fibres. For each natural fibre sample, 10 measurements were made with 10 different fibres. Colour measurements from garments made

from natural fibres exhibit greater variability. For a full description of the data collection methodology, see Laing *et al.* (1987).

The database identifies 13 different fibre types, of which six types make up 98% of all samples. The most common fibre type is cotton, of which there are 2340 samples within the database. The most common synthetic fibre is polyester, of which there are 1512 samples.

1.3. *Previous Statistical Analysis*

Evett *et al.* (1987) considered in detail the case where a single fibre was recovered from the scene of the crime. It is recognized that equation (1.2) can be written as

$$\frac{p(x, y|C, I)}{p(x, y|\bar{C}, I)} = \frac{p(y|C, x, I)\, p(x|C, I)}{p(y|\bar{C}, x, I)\, p(x|\bar{C}, I)} = \frac{p(y|C, x, I)}{p(y|\bar{C}, I)}, \tag{1.3}$$

by assuming that, conditional on I, x is equally likely under C and $\bar{C}$. It was assumed that the distribution of fibre colour within a particular garment was adequately described by a bivariate normal distribution. Thus the numerator of the likelihood ratio (1.3) can be written

$$p(y|C, x, I) = \int p(y|\mu, \Sigma, I)\, p(\mu, \Sigma|x, I)\, \mathrm{d}\mu\, \mathrm{d}\Sigma,$$

where $p(y|\mu, \Sigma, I)$ is the density of a bivariate normal distribution with mean μ and covariance matrix Σ, and $p(\mu, \Sigma|x, I)$ is a posterior density for these parameters, assuming $x = (x_1, \ldots, x_m)$ to have been a sample from the same bivariate normal distribution. Using conventional representations of vague prior knowledge regarding μ and Σ, the form of the numerator is shown by Evett *et al.* (1987) to be a bivariate Student t-type density, having an explicit mathematical form.

For the denominator, given I (which includes the knowledge of the type of fibre in question), the form $p(y|\bar{C}, I)$ is simply the density describing the general distribution of colour co-ordinates that we would see in whatever fibre population is defined by I, which Evett *et al.* evaluated by using a bivariate kernel density estimator.

1.4. *Motivation and Structure of this Paper*

The work described in this paper results from a collaborative project between the Home Office Forensic Science Service Central Research and Support Establishment and the University of Nottingham Statistics Group.

The project addressed numerical and modelling issues associated with the extension of the above problem to several recovered fibres. The approach described by Evett *et al.* (1987) is not tractable for more than a single fibre. Numerical procedures are necessary whatever mathematical forms are assumed and this provided an opportunity to examine more closely the conventional distributional assumptions of the original work. In addition, in the context of the possible modelling strategies available, the numerical demands of calculating the required likelihood ratio pose challenging problems within the Bayesian inference framework.

The main purpose of the present paper is to exhibit, via a case study in an important and newly emerging applications field, a class of interesting and novel modelling problems, involving both likelihood and prior specification, and some possible solutions. Our emphasis throughout will be on the statistical aspects of the problem. In particular, we restrict attention to 'single-colour' and 'single-source' problems. It

is clear that the forensic scientist has to consider multicoloured garments and that foreign fibres attached to any scene may have come from several sources. The models described here can be extended to encompass these situations but no new statistical issues arise. In the wider context of the Home Office Forensic Science Service, there are important issues of integrating the statistical methodology into intelligent knowledge-based systems for laboratory use. We shall not discuss such problems here.

In Section 2, we deal with modelling the variability of observed chromaticity co-ordinates among fibres from the same garment. In Section 3, we deal with modelling the prior (population) distribution of the parameters involved in modelling within-garment variability. Computational problems are discussed in Section 4 and extensions and refinements of the methodology are examined in Section 5. Two illustrative examples, taken from actual case work, are presented in Section 6.

2. Modelling Within-garment Variability

2.1. *General Approach*

Fig. 1 displays the 'horseshoe'-shaped region within which complementary chromaticity co-ordinates (u, v) are constrained to lie. In this figure, N denotes the neutral, or achromatic, point. The positions of typical hues, i.e. colours, are also indicated. In essence, the colour space reflects hue and saturation but does not take into account lightness. Thus all whites, greys and black map to the neutral point. If a line were drawn from the neutral point to the boundary of the colour space, movement along the line towards the boundary represents an increase in saturation (colours described by a single wavelength lie on the boundary)—the hue is constant. Most undyed cloth has a pale yellow colour which is close to the neutral point. In all dyeing processes there is some variability in uptake of the dye. When a single dye is used the colour of the fibres reflects both the undyed colour and the dye colour with saturation exhibiting rather more variability. Measurements thus appear close to a line joining the undyed point to the dye colour. Many colours are achieved via the combination of two or more dyes. In this case colour measurements for individual fibres are scattered about a line joining the two predominant dye colours. The need for multiple dyes is also more commonplace for certain colours. For example, many blues can be achieved with a single dye whereas greens often require multiple dyes.

A typical scatter diagram for the co-ordinates of fibres from a particular sample within the database is also shown in Fig. 1. The inspection of many such scatterplots indicated a preferred orientation with the long axis frequently pointing towards the yellow area, consistent with the above discussion. The direct modelling of the bivariate distribution is thus complicated by a correlation which varies stochastically with location. As an initial step therefore, for each sample of a particular fibre type, sample means were calculated and principal component axes were identified. Projections of standardized values on to these axes were then studied to try to identify possible distributional forms in the long and short directions. Fig. 2 shows the standardized plots for wool fibres.

This in turn leads to the following general modelling strategy for within-garment variability.

(a) Identify two orthogonal axes u' and v'.

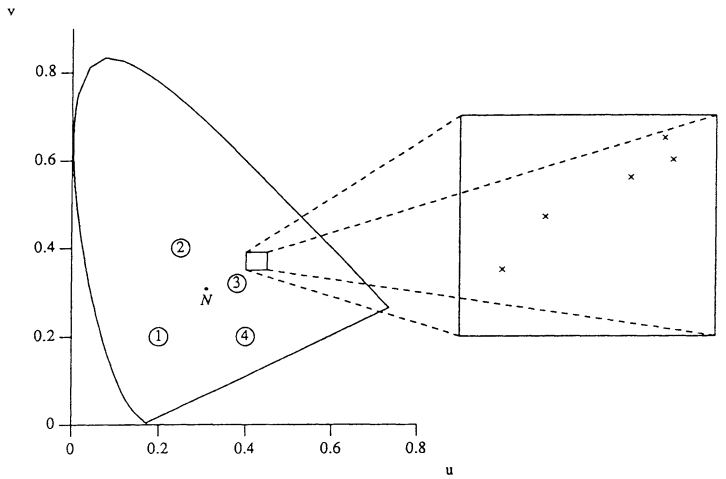

Fig. 1. Scatter of five bivariate measurements in colour space (approximate positions of the main colours are indicated: 1, yellow; 2, red; 3, blue; 4, green)

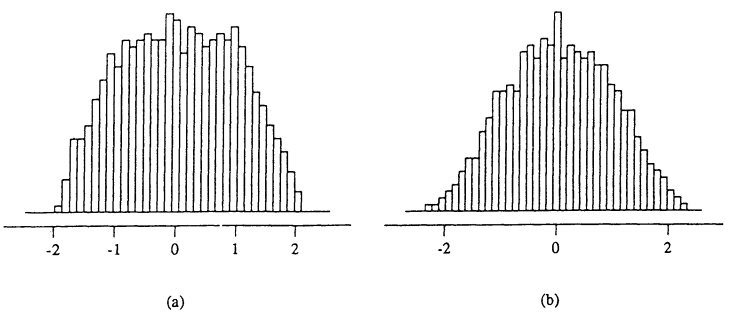

Fig. 2. Standardized values for wool fibres in (a) the long direction and (b) the short direction

(b) Identify the orientation of the u'-axis relative to the original u-axis.
(c) Assume, given orientations, locations and spreads, that the distributions of projections of the chromaticity co-ordinates on to these axes are statistically independent.

Given this strategy, the bivariate density describing fibre colour within a garment is a five-parameter mathematical family with one parameter for orientation, two for location and two for spread. The orientation is defined as the angle of the u'-axis relative to the u-axis. Additionally the spreads are constrained so that the long direction of scatter is identified as having the larger spread parameter.

The assumption made in Evett *et al.* (1987) was that the scatter was well modelled by a bivariate normal distribution. However, considering separately the projections on to the long and short axes as in Fig. 2, the normal hypothesis is clearly rejected by a conventional χ^2-test, for nearly all fibre types. This finding prompts a search for other distributional shapes providing a more satisfactory fit to the observed univariate distributions.

2.2. *Exponential Power Family*

The so-called exponential power family is a general univariate family of distributions with location τ, spread σ and shape parameter β. It is defined by

$$p(z|\tau, \sigma, \beta) \propto \sigma^{-1} \exp\left(-\frac{1}{2}\left|\frac{z-\tau}{\sigma}\right|^{2/(1+\beta)}\right).$$

This family includes both platykurtic and leptokurtic alternatives to the normal distribution. See, for example, Box and Tiao (1973).

For a particular fibre type and direction an attempt was made to find values of β which mimicked the profiles obtained via plots similar to those in Fig. 2. It was not possible, however, to identify a particular value of β which adequately modelled the tail behaviour in either direction. For example, the observed distribution in the long direction is heavy shouldered and light tailed.

2.3. *'Roof' Distribution*

A second approach considered was to describe the joint distributions via conditional distributions which were simple combinations of rectangular and triangular distributions. See Fig. 3. It was found that a good fit to the data was obtained, provided that a small proportion of fibre measurements were considered to be lying in a skirt around the base of the roof.

In addition to the usual five parameters—(μ_1, μ_2), mean locations; (σ_1, σ_2), standard deviation spreads; ϕ, orientation—further constants were required to describe the lengths of the different portions of the roof, e.g. the distance in the long direction, from the mean, to the point where the roof ridge begins to slope down and the extent of the skirt. These extra constants are easily estimated for a given fibre type by considering standardized measurements in each of the long and short directions, together with the constraints necessary for the total probability to equal 1. Although this model provided very satisfactory fits to much of the fibre data, it proved, subse-

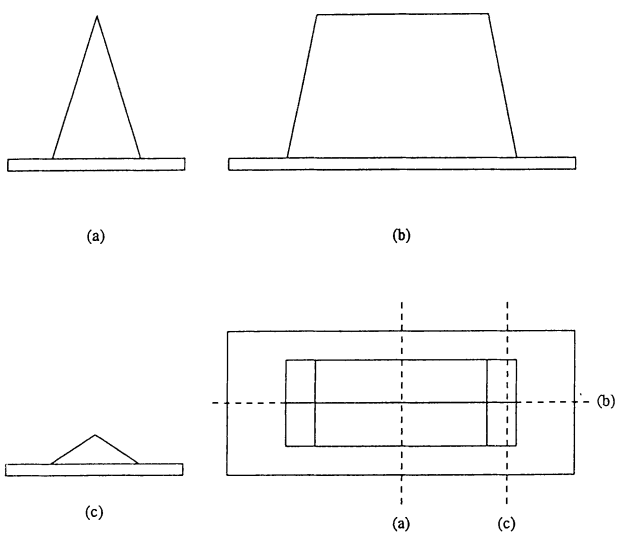

Fig. 3. Roof distribution: three sections and a plan view

quently, to be very difficult to operate with numerically and so yet another alternative was sought.

2.4. *Beta Distribution*

It was known that the measurement procedure discarded fibres which were outlying when compared with the main cluster of a particular sample. Initially this was thought to have little influence on tail shape but subsequently this fact motivated consideration of finite range distributions. Bivariate densities defined by products of symmetric scaled beta distributions provided an adequate fit to the data and proved considerably more tractable computationally.

A univariate random variable Z has a symmetric beta distribution on $(\mu - \sigma t, \mu + \sigma t)$ if the density of $W = (z - \mu)/2\sigma t + \frac{1}{2}$ is given by

$$p(w) \propto w^{\alpha-1}(1-w)^{\alpha-1}, \qquad 0 < w < 1.$$

Thus α and t must be estimated for the two orthogonal directions for each fibre type. This was accomplished with a search procedure and a minimum chi-squared criterion, but subsequent sensitivity studies suggested that precise values were not critical.

2.5. *Comment*

Although the large amount of data available makes it possible to distinguish between alternative models for the within-garment variability by using goodness-of-fit criteria, it must be remembered that the overall objective is the computation of the likelihood ratio (1.2). It is the sensitivity of this ratio to modelling assumptions which is of primary concern and thus the eventual choice of model is mediated by sensitivity issues.

3. Modelling the Prior Distribution

Given a model for within-garment variation involving five parameters, the evaluation of the likelihood ratio requires the specification of a prior distribution $p(\mu_1, \mu_2, \sigma_1, \sigma_2, \phi)$. Given the size of the database it is technically possible to obtain estimates of the parameters from each of the samples in the database and then to construct a five-dimensional kernel estimate. However, such an estimate is computationally intensive and the numerical integration techniques to be described in Section 4 require a large number of such evaluations.

Another approach would be to evaluate the prior in advance at a regular grid of points within the five-dimensional space and then to use interpolation. However, the number of points required to characterize such a surface is large and five-dimensional interpolation is difficult.

The discussion in Section 2.1 combined with an empirical examination of the database suggested that the prior could be well approximated by a form

$$p(\mu_1, \mu_2, \sigma_1, \sigma_2, \phi) = p(\mu_1, \mu_2)\, p(\sigma_1 | r)\, p(\sigma_2 | r)\, p(\phi | \lambda), \tag{3.1}$$

where r is the radial distance from the neutral point to (μ_1, μ_2) and λ describes the radial angle of the cluster relative to the neutral point (Fig. 4), i.e. the spreads on the principal component axes are independent, given position, and depend only on the distance from the neutral point.

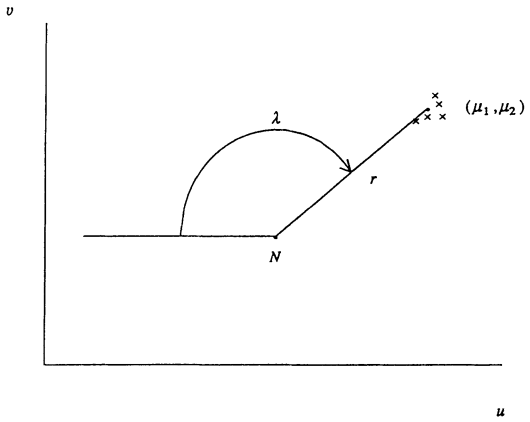

Fig. 4. r and λ for a particular data cloud

The orientation of the long axis is independent of saturation but depends on hue, i.e. on the combination of dyes involved. For example, in the blue region of the colour space, we would expect the majority of garments to have a density aligned towards the undyed point whereas with greens, which can arise from pure green dyes or from combinations of blues and yellows, we would expect the prior for ϕ to be more disperse reflecting this mixture of possibilities.

Noting that

$$p(\sigma_i|r) = p(\sigma_i, r)/p(r), \qquad i=1, 2,$$

and

$$p(\phi|\lambda) = p(\phi, \lambda)/p(\lambda),$$

we see that a prior distribution can be evaluated with four bivariate densities, $p(\mu_1, \mu_2)$, $p(\sigma_1, r)$, $p(\sigma_2, r)$ and $p(\phi, \lambda)$, each of which can be constructed by using kernel methods.

Rather than having to evaluate the four two-dimensional kernels every time that a prior probability was required, it was decided to set up four reference tables in advance. If the database were to be updated it would then be simple to recreate the required grids of points. When the prior for a particular point was required, cubic splines were used to interpolate to the desired point. The reference tables were achieved as follows: for each of the seven quantities of interest (five model parameters, the radial distance and the radial angle) sample estimates were obtained for each sample in the database of a particular fibre type. From these, kernel density estimates for each of the bivariate densities were evaluated at each point on a 20×20 grid. For σ_i and r, the kernel densities were created in terms of $\log \sigma_i$ and $\log r$. The densities were then inverted to obtain the required form in terms of the original parameters. This operation was undertaken so that density estimates close to 0 were not underestimated; see Silverman (1986). For the angles ϕ and λ, the estimates were 'wrapped around' so that, for example, λ close to 360° contributed to the kernel estimate for angles close to 0° as well as to 360°.

Bivariate normal kernels used after initial experimentation indicated that the

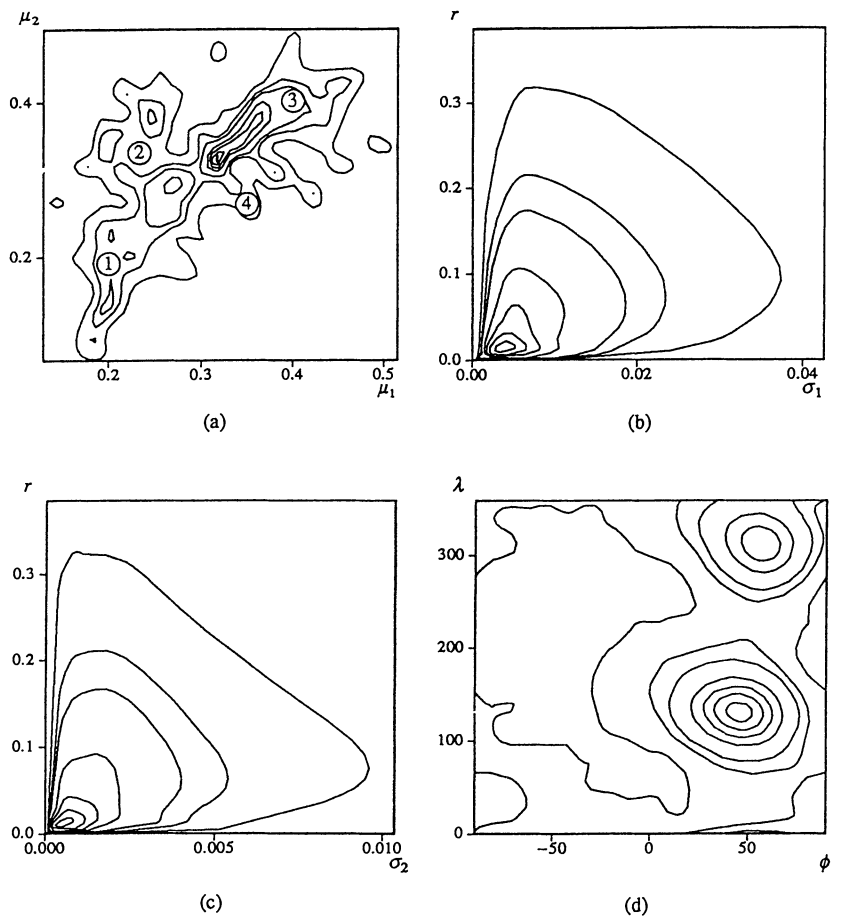

Fig. 5. Kernel density estimates for wool fibres for (a) $p(\mu_1, \mu_2)$, (b) $p(\sigma_1, r)$, (c) $p(\sigma_2, r)$ and (d) $p(\phi, \lambda)$: 1, yellow; 2, red; 3, blue; 4, green

density estimates were insensitive to the choice of kernel shape. Initially, window widths were chosen proportional to $n^{-1/5}$ by using guidelines suggested by Silverman (1986). However, for fibre types where the amount of data was large this resulted, for $p(\mu_1, \mu_2)$ especially, in a very spiky density estimate. Consequently the final density estimates used slightly larger window widths giving a smoothness which was consistent with the considerable prior knowledge communicated to us by the Home Office colour chemists. Contour plots of the four bivariate density estimates for wool are shown in Fig. 5. Displayed contours are at 1%, 5%, 10%, 25%, 50%, 75% and 95% of the mode.

Fig. 5(a) shows that a large amount of the data lies in the blue and the red regions. The triangular shape of the contour plot is typical and is similar for other fibre types. Each of the multiple modes represents popular colours. The large central mode occurs at the neutral point. Figs 5(b) and 5(c) confirm that the standard deviations increase with radius—the clusters show more variability as we move away from the neutral point. Fig. 6(a) confirms this and displays the prior for $p(\sigma_1 | r)$ for two particular

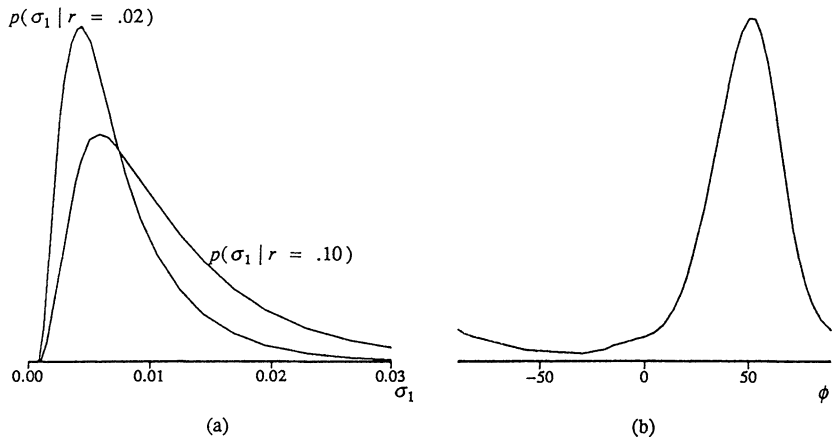

Fig. 6. (a) $p(\sigma_1 \mid r)$ and (b) marginal distribution $p(\phi)$

choices of r, one small and one large. Referring now to Fig. 5(d), the mode occurring at approximately $\lambda = 300°$ corresponds to the yellows, which all tend to point towards the neutral point. In this region of the colour space, this orientation corresponds to $\phi = 50°$. The other mode at $\lambda = 130°$ represents the blues, which again are predominantly oriented towards the neutral point. ϕ is defined between $-90°$ and $90°$ and so even though the blues and yellows lie on opposite sides of the neutral point they have similar values of ϕ. Greens, lying at $\lambda = 240°$, have a flatter distribution, reflecting the fact that greens often arise as mixtures of two different colours; see the discussion in Section 2.1. If the greens tended to point towards the neutral point we would expect the mode of $p(\phi \mid \lambda = 240°)$ to lie at $\phi = -60°$. The mode of this conditional distribution actually lies at $\phi = 50°$ indicating that the greens are often combinations of blues and yellows. Fig. 6(b) shows the marginal distribution $p(\phi)$. This distribution was obtained from $p(\phi, \lambda)$ via discrete approximations to the required integrals. The mode at approximately 50° reflects the previous discussion of the blues, the yellows and the greens.

4. Implementation

4.1. *Evaluation of Likelihood Ratio*

Each term in equation (1.2) represents the joint distribution of a set of bivariate colour measurements from a single source. Reference to C and $\bar{C}$ can therefore be dropped and, if z denotes a collection $(z_1, \ldots, z_k)$ of k bivariate measurements,

$$p(z \mid I) = \int \prod_{i=1}^{k} p(z_i \mid \theta)\, p(\theta \mid I)\, \mathrm{d}\theta, \tag{4.1}$$

with $p(z_i \mid \theta)$ modelled by one of the forms discussed in Section 2 and assumed to be independent of I, and $p(\theta \mid I)$ specifying a prior of the form discussed in Section 3. Note that the integral still depends on I, the information additional to the actual colour measurements, including fibre type.

4.2. *Computation*

For any choice of $p(z_i|\theta)$ and $p(\theta|I)$, the evaluation of equation (4.1) requires the evaluation of a five-dimensional numerical integral.

Naylor and Smith (1982, 1983) and Smith *et al.* (1985, 1987) describe several novel numerical integration strategies for the implementation of Bayesian methods. Among these are adaptive Cartesian product Gauss–Hermite methods which we outline very briefly here.

The Gauss–Hermite strategy is motivated by first noting that, in one dimension, if an integrand can be well approximated by a polynomial × normal density then the integral will be well estimated by a Gauss–Hermite rule (with an exact answer from an n-point rule if the integrand is actually of polynomial × normal form, with the polynomial component of degree up to $2n-1$). We then note that, possibly after suitable reparameterization, many of the marginal posterior densities arising in Bayesian inference can plausibly be regarded as of polynomial × normal form. A joint posterior density for many parameters, however, is likely to exhibit several dependencies, which would make the application of Gauss–Hermite product rules highly inefficient. Moreover, the implementation of the Gauss–Hermite rule in each parameter direction requires the specification of a location and scale for that direction, and correct values for these are not known.

This suggests the following iterative strategy for finding a product grid at which the evaluations of the likelihood × prior function provide the basis for the accurate evaluation of the normalizing constant and hence any aspects of interest of the joint posterior density.

(a) Reparameterize individual parameters so that the resulting working parameters can be well represented by a polynomial × normal form. Often this parameterization will allow the working parameters to take any value on the real line.
(b) Using initial estimates of the joint posterior mean vector and covariance matrix for the working parameters, transform further to a centred, scaled, more orthogonal set of parameters.
(c) Using the derived initial location and scale estimates for these orthogonal parameters, perform Cartesian product integration of functions of interest using suitably dimensioned grids.
(d) Iterate, successively updating the mean and covariance estimates, until stable results are obtained both within and between grids of specified dimension.

The parameterization adopted for the problem described here,

$$\psi_1 = \mu_1,$$
$$\psi_2 = \mu_2,$$
$$\psi_3 = \log \sigma_1^2,$$
$$\psi_4 = \log\left(\frac{\sigma_2^2}{\sigma_1^2 - \sigma_2^2}\right),$$
$$\psi_5 = \log\left(\frac{90+\phi}{90-\phi}\right),$$

proved adquate for normal, exponential power and roof densities. However, further constraints in the range of μ_1, μ_2, σ_1 and σ_2 are implied by $z_1, \ldots, z_k$ when the bivariate beta distribution is considered. This point is discussed in more detail in Wakefield *et al.* (1989). A further point is worth mentioning. The mean parameters μ_1 and μ_2 must lie within the horseshoe-shaped region of Fig. 1 and so could be constrained accordingly. In practice, however, they never fall close to the boundary. Fig. 5(a) bears this out. The variances σ_1^2 and σ_2^2 are similarly constrained but again, in practice, this is not a problem.

4.3. *Sensitivity Studies*

Elements of the likelihood ratio (1.2) were calculated with a variety of combinations of the within-garment variability model (normal, double exponential, roof and beta) and prior specifications (reflecting our uncertainty in the choice of window width). Detailed findings of the sensitivity are not given here. The results showed, however, that although the values of individual components were clearly influenced by the choice of model and prior, the resulting likelihood ratios were within a factor of 2 in most cases.

Although all such combinations are computationally feasible, the computational cost is influenced by the choice of model to some extent with bivariate normality being the most efficient given the use of the iterative Gauss–Hermite strategy. Furthermore, likelihood ratios accurate to within a factor of 10 are currently viewed as sufficient to represent evidential content. Thus the modelling and sensitivity exercise shows that the adoption of bivariate normality as the default model is adequate at present. The examples which follow in Section 6 have been analysed under this assumption. Implementation of the integration procedures is possible on the current range of personal computer systems although more powerful single-user workstations are more appropriate. Further developments in computer technology will minimize the argument for choosing bivariate normality over a more realistic distributional form. Bivariate beta distributions are our preferred choice.

5. Extensions and Refinements

We now consider in a little more detail the conditioning event I appearing in equation (1.1). I represents all additional information available about the fibres apart from the actual complementary chromaticity co-ordinate measurements.

As a simple illustration of a typical event I, consider the situation where foreign fibres have been found on a victim. An initial examination of the fibres reveals that they

(a) are of a particular type (e.g. wool or nylon),
(b) may have characteristics such as delustrant etc. and
(c) to the naked eye are of a particular colour (e.g. red or blue) and possibly can be described as light or dark etc.

Then I is the event that a suspect has been found with a garment containing fibres which *match* this description.

The precise measurement of the colour of both recovered and control fibres is then undertaken and y and x are obtained.

As we see from equation (1.1), the overall contribution to the posterior odds on the recovered fibres having come from the suspect's garment is a combination of the likelihood ratio (1.2), based on y and x and I, and an odds ratio conditioned just on I. Detailed quantification of the latter term, $p(C|I)/p(\bar{C}|I)$, clearly depends on the specific circumstances surrounding the identification of a suspect and a garment. For example, was the suspect primarily identified because of the garment, or was the garment found subsequently to the identification of the suspect on other grounds? We shall not pursue such issues here, but instead concentrate on the effect of the information in (a)–(c) on $p(\theta|I)$.

Taking account of an I based on (a) and/or (b) is straightforward in theory, though it raises some issues in practice. Such information simply identifies the subset of the database on which the kernel methodology of Section 3 is to be applied.

For natural fibres such as wool and cotton, there is usually no additional information of type (b) available and thus the present database is sufficient to obtain good kernel density estimates. For rare fibre types or situations where a combination of synthetic type, cross-sectional shape and radius of fibres and delustrant leads to a precise initial match, the present database is not yet sufficiently large to yield precise descriptions of colour distribution. In such cases the likelihood ratio can still be calculated using the near uniform priors obtained from the few data points available. The evidential value of the colour measurements will thus be small but the initial matching process will usually have high evidential content. Information of the kind described by (c) and by the fact that an apparent colour match has been made is not straightforward to incorporate. In effect, it imposes a further conditioning, which implies that the kernel, for the appropriate part of the database implied by (a) and (b), is defined only over a subset of the colour space, rather than over the entire horseshoe-shaped region described earlier (see Fig. 1). What is difficult to quantify is precisely which subregion of the colour space is defined by the description in (c) and the fact that a *prima facie* visual match has been obtained.

Our solution to this problem is to define the appropriate subset of the colour space to be a circle centred at the mean location of the recovered fibres. The size of this circle was chosen after consulting forensic science case officers. The region of a match does, to some extent, depend on the mean location but our analyses proved to be reasonably insensitive to the exact value chosen. The implied form of $p(\theta|I)$ is then modified from that discussed in Section 3 by replacing the kernel density estimate of the $p(\mu_1, \mu_2)$ component by one constrained to the circular region. The use of a circle rather than some other shaped neighbourhood has negligible impact on the values of the likelihood ratio obtained and this approach has been used in the following case studies.

6. Illustrative Analyses

6.1. *Example 1*

The victim was dragged into a garden and was robbed of her handbag. The victim was wearing a green wool and synthetic fibre jumper. The assailant was wearing a grey–green checked cotton shirt. Three green wool fibres were found on the suspect's shirt. Seven grey–green cotton fibres were found on the victim's jumper. 10 control fibres were taken from both the suspect's shirt and the victim's jumper.

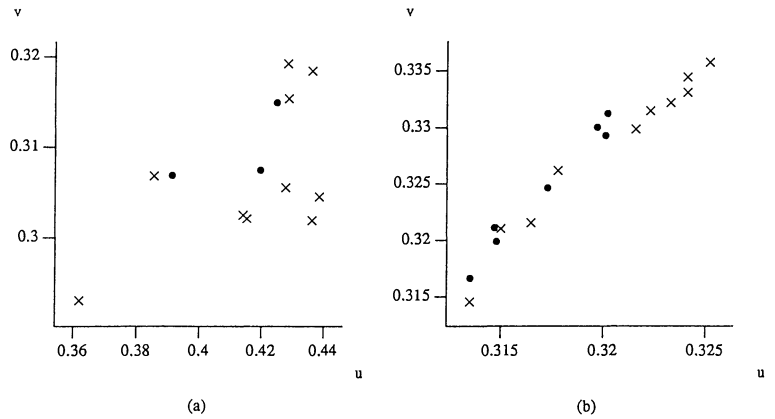

Fig. 7. Control (×) and recovered (•) fibre measurements for (a) the green wool and (b) the grey–green cotton

Consider first the green wool. Fig. 7(a) gives a scatterplot showing measurements from both recovered and control fibres. The value of the likelihood ratio based on these data alone is 700.

Considering only the grey–green cotton, the value of the likelihood ratio is 30. The control and recovered fibre measurements are shown in Fig. 7(b).

In most cases, it can be argued that, conditional on I, the likelihood ratio for two-way transfer is the product of the separate ratios. Overall, therefore, the likelihood ratio in this case is 21000.

6.2. *Example 2*

Following a complex series of armed robberies, the forensic scientist was asked to investigate links between three suspects and a number of getaway cars. In particular, it was required to investigate the link between a knife found in the possession of one of the suspects and several purple nylon-66 cash bags which had been split open. We shall consider

(a) purple fibres from the cash bags and the knife and
(b) blue wool fibres found in one of the cars and matched with a suit owned by one of the suspects.

6.2.1. *Purple cash bags*

The manufacturers were contacted and only two batches of this type of fibres had ever been made. 73 fibres were found on the knife and chromaticity co-ordinates were calculated from a randomly selected 15 of these. Fibres from all the bags were visually indistinguishable but they could be split into two sets by visible spectroscopy and thin layer chromatography. 13 fibres matched the first batch and the remaining two fibres matched the second batch. There were five control measurements for each of the batches. The scatterplot of the control and recovered fibres for the second batch is shown in Fig. 8(a). Although there are only two recovered fibres, the likelihood ratio is 217. Posterior moments of model parameters indicated that, in this example, both

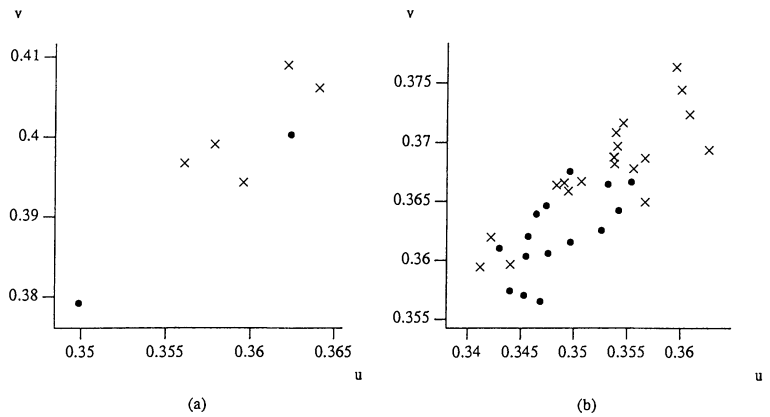

Fig. 8. Control (×) and recovered (•) fibre measurements for (a) the purple cash bags and (b) the blue wool suit

the orientation and the dispersion of the recovered and control data were very similar. The high value is *not* a consequence of the rarity of the purple colour as the ratio is conditional on I. The fact that a knife was recovered with fibres that matched the rare colour of the cash bags would normally increase the strength of evidence. However, this aspect of the evidence must be modelled via $p(C|I)/p(\overline{C}|I)$.

6.2.2. *Blue wool suit*

15 blue wool fibres were found and these matched fibres taken from a blue suit owned by one of the suspects. In this case 20 control measurements were taken and the subsequent measurements, along with the recovered co-ordinates, are displayed in Fig. 8(b).

In this case a likelihood ratio of 0.8 was obtained, the data being sufficiently numerous to suggest differences in dispersion between the recovered and control samples.

Acknowledgements

The work reported here was made possible by a Home Office grant supporting J. C. Wakefield. We also acknowledge invaluable assistance from P. E. Cage, A. W. Hartshorne and D. K. Laing of the Home Office Central Research and Support Establishment and R. Cook of the Metropolitan Police Forensic Science Laboratory.

References

Box, G. E. P. and Tiao, G. C. (1973) *Bayesian Inference in Statistical Analysis*, p. 157. New York: Addison-Wesley.

Evett, I. W. (1984) A quantitative theory for interpreting transfer evidence in criminal cases. *Appl. Statist.*, **33**, 25–32.

——(1987) Bayesian inference and forensic science: problems and perspectives. *Statistician*, **36**, 99–106.

Evett, I. W., Cage, P. E. and Aitken, C. G. G. (1987) Evaluation of the likelihood ratio for fibre transfer evidence in criminal cases. *Appl. Statist.*, **36**, 174–180.

Laing, D. K., Hartshorne, A. W., Cook, R. and Robinson, G. (1987) A fibre data collection for forensic scientists—collection and examination methods. *J. Forens. Sci.*, **32**, 364–369.

Laing, D. K., Hartshorne, A. W. and Harwood, R. J. (1986) Colour measurements on single textile fibres. *Forens. Sci. Int.*, **30**, 65–77.

Lindley, D. V. (1977) A problem in forensic science. *Biometrika*, **64**, 207–213.

Naylor, J. C. and Smith, A. F. M. (1982) Applications of a method for the efficient computation of posterior distributions. *Appl. Statist.*, **31**, 214–225.

——(1983) A contamination model in clinical chemistry: an illustration of a method for the efficient computation of posterior distributions. *Statistician*, **32**, 82–87.

Silverman, B. W. (1986) *Density Estimation for Statistics and Data Analysis*. London: Chapman and Hall.

Smith, A. F. M., Skene, A. M., Shaw, J. E. H. and Naylor, J. C. (1987) Progress with numerical and graphical methods for practical Bayesian statistics. *Statistician*, **36**, 75–82.

Smith, A. F. M., Skene, A. M., Shaw, J. E. H., Naylor, J. C. and Dransfield, M. (1985) The implementation of the Bayesian paradigm. *Communs Statist. Theory Meth.*, **14**, 1079–1102.

Wakefield, J. C., Skene, A. M., Smith, A. F. M. and Evett, I. W. (1989) A case study in forensic science: the evaluation of fibre transfer evidence. *Research Report*. Nottingham Statistics Group, University of Nottingham.