Peter G. M. Forbes, Steffen Lauritzen and Jesper Møller

Fingerprint Analysis with Marked Point Processes

Fingerprint Analysis with Marked Point Processes

Peter G. M. Forbes
University of Oxford

Steffen Lauritzen*
University of Oxford

Jesper Møller
Aalborg University

July 22, 2014

Abstract

We present a framework for fingerprint matching based on marked point process models. An efficient Monte Carlo algorithm is developed to calculate the marginal likelihood ratio for the hypothesis that two observed prints originate from the same finger against the hypothesis that they originate from different fingers. Our model achieves good performance on an NIST-FBI fingerprint database of 258 matched fingerprint pairs.

Keywords: Bayesian alignment; complex normal distribution; forensic identification; likelihood ratio; marked point processes; von Mises distribution; weight of evidence.

1 Introduction

Fingerprint evidence has been used for identification purposes for over one hundred years. Despite this, there has been very little scientific research on the discriminatory power and error rate associated with fingerprint identification. Within the last ten years there has been a push to move fingerprint evidence towards a solid probabilistic framework, culminating in the recent paper by Neumann et al. (2012).

We discuss a novel approach for fingerprint matching using marked Poisson point processes. We develop an efficient Monte Carlo algorithm to calculate the likelihood ratio for the prosecution hypothesis that two observed prints originate from the same finger against the defence hypothesis that they originate from different fingers. Hill et al. (2012) have also considered marked Poisson point process models for fingerprints, albeit for another purpose: namely, the reconstruction of fingerprint ridges from sweat pore point patterns.

Fingerprint evidence is based on the similarity of two or more pictures, see Fig. 1. It is difficult to represent all the information from these pictures in a mathematically convenient form. Thus most fingerprint models, including the one in Neumann et al. (2012), consider only a subset of the information: namely, the points on the image where a ridge either ends or bifurcates. These points, called minutiae, generally contain sufficient information to uniquely identify an individual (Maltoni, 2009; Yager and Amin, 2004). A typical full fingerprint contains 100–200 minutiae, while a low quality crime scene fingermark may contain only one dozen (Garris and McCabe, 2000).

Lauritzen et al. (2012) note the similarity between minutia matching and the alignment problems often studied in bioinformatics. Our model exploits ideas from the model for unlabelled point set matching in Green and Mardia (2006) and applies them to the problem of fingerprint

---

*Corresponding author. Department of Statistics, University of Oxford, 1 South Parks Road, Oxford OX1 3TG, United Kingdom. email: steffen@stats.ox.ac.uk.
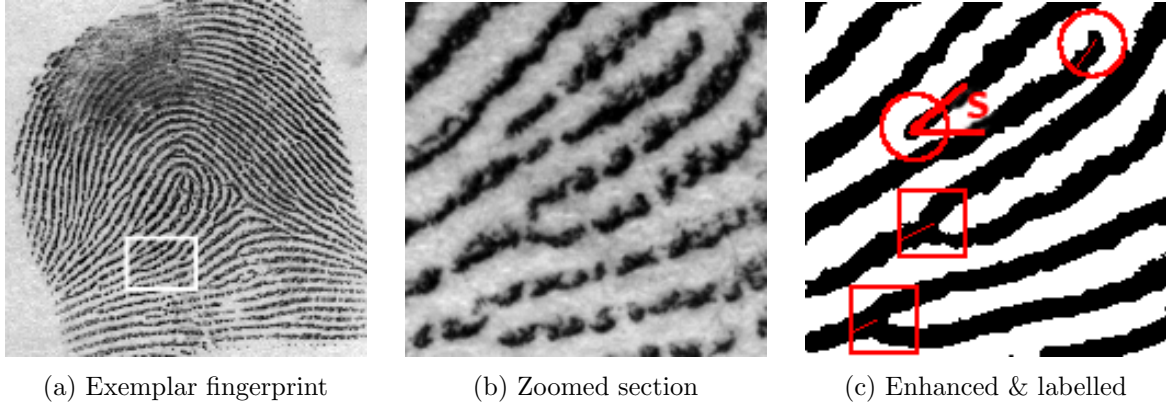
| (a) Exemplar fingerprint | (b) Zoomed section | (c) Enhanced & labelled |

Figure 1: A typical exemplar quality fingerprint from Garris and McCabe (2000). The high-lighted points in (c) are minutiae: circles are ridge endings and squares are bifurcations.

matching. Our model could be used for an automated fingerprint identification system, or it could support a courtroom presentation of fingerprint evidence.

The paper is composed as follows. After a few preliminary specifications in Section 2 we develop a generic marked Poisson point process model in Section 3 and a specific parametric version in Section 4. In Section 5 we describe our method for calculating the likelihood ratio and in Section 6 we perform an analysis using the methodology on both simulated and real data. In the appendix we give further technical details of our computational procedures.

2 Preliminaries and notation

2.1 Likelihood representation of fingerprint evidence

As in Neumann et al. (2012) we discuss the situation where we wish to compare a high-quality fingerprint A taken under controlled circumstances, with a fingermark B found on a crime scene. We consider two hypotheses

$$H_p : A \text{ and } B \text{ originate from the same finger,}$$
$$H_d : A \text{ and } B \text{ originate from different fingers,} \tag{1}$$

where H_p is referred to as the prosecution hypothesis and H_d as the defence hypothesis. Following a tradition that goes at least back to Lindley (1977), we follow standards in modern evaluation of DNA and other types of forensic evidence (Balding, 2005; Aitken and Taroni, 2004) and quantify the weight-of-evidence by calculating a likelihood ratio between H_p and H_d,

$$\Lambda = \frac{\mathrm{pr}(A, B \mid H_p)}{\mathrm{pr}(A, B \mid H_d)}. \tag{2}$$

The likelihood ratio is based on probabilistic models for the generation of the fingerprint and fingermark that shall be developed in the sequel.

2.2 Representation of fingerprints

Each minutia m consists of a location, an orientation, and a type: ridge ending, bifurcation, or unobserved; see Fig. 1c. We represent the location with a point in the complex plane $\mathbb{C}$ and the orientation with a point on the complex unit circle $\mathbb{S}^1$. The type is represented by a number in

2

$\{-1, 0, 1\}$, where -1 denotes a ridge ending, 1 a bifurcation, and 0 an unobserved type. Thus m is an element of the product space $\mathbb{M} = \mathbb{C} \times \mathbb{S}^1 \times \{-1, 0, 1\}$. We let r_m, s_m, and t_m denote the projection of m onto the location space, orientation space, and type space respectively.

A fingerprint A or a fingermark B is represented by a finite set of elements of $\mathbb{M}$. We call this representation a minutia configuration. Since A and B are observed in arbitrary and different coordinate systems, the observed minutiae are subjected to similarity transformations, which consist of translations, rotations, and scalings. These can be simply represented by algebraic operations with complex numbers,

$$(r_m, s_m, t_m) \mapsto (\psi r_m + \tau, \psi s_m / |\psi|, t_m).$$

2.3 Basic distributions

We shall use the bivariate complex normal distribution, which describes a complex random vector whose real and imaginary parts are jointly normal with a specific covariance structure (Goodman, 1963). The density with respect to the Lebesgue measure is

$$\varphi_2(r; \mu, \Sigma) = \exp\{-(\overline{r - \mu})^T \Sigma^{-1} (r - \mu)\} / (\pi^2 |\Sigma|),$$

where r and μ are two-dimensional complex numbers, Σ is a Hermitian positive definite 2×2 complex matrix with determinant $|\Sigma|$, the overline denotes the complex conjugate, and ^{T} denotes the vector transpose. The standard case of $\mu = 0$ and Σ equal to the identity matrix will be denoted $\varphi_2(r)$. When we wish to make the two arguments explicit we will write $\varphi_2(r_1, r_2; \mu, \Sigma)$ for $r_1, r_2 \in \mathbb{C}$. The univariate density will be denoted $\varphi(r; \mu, \sigma^2)$ where $r, \mu \in \mathbb{C}$ and $\sigma^2 > 0$, with the standard case denoted $\varphi(r)$.

The von Mises distribution $\mathrm{vM}(\nu_0, \kappa)$ on the complex unit circle $\mathbb{S}^1$ with position ν_0 and precision $\kappa > 0$ (Mardia and Jupp, 1999) has density

$$\upsilon(s; \nu_0, \kappa) = I_0(\kappa)^{-1} \exp\{\kappa \Re(s\overline{\nu_0})\}$$

with respect to ν, the uniform distribution on $\mathbb{S}^1$, where $\Re(z) = (z + \overline{z})/2$ is the real part of z. The normalization constant $I_0(\kappa)$ is the modified Bessel function of the first kind and order zero (Olver et al., 2010, chapter 10). The von Mises distribution can be obtained from a univariate complex Normal distribution $\varphi(s; \nu_0, 2/\kappa)$ (or equivalently $\varphi(s; \kappa\nu_0/2, 1)$) by conditioning on $|s| = 1$.

Kent (1977) shows that the von Mises distribution is infinitely divisible on $\mathbb{S}^1$ and thus it makes sense to define the root von Mises distribution $\mathrm{rvM}(\nu_0, \kappa)$ by

$$XY \sim \mathrm{vM}(\nu_0, \kappa) \text{ whenever } X, Y \text{ are independent and } X, Y \sim \mathrm{rvM}(\nu_0, \kappa).$$

The density of the root von Mises distribution is determined by a series expansion. We refrain from giving the details as we shall not need them.

3 A generic marked point process model

3.1 Model specification

We consider the observed minutia configurations $A, B \subset \mathbb{M}$ as thinned and displaced copies of a latent minutia configuration. In this paper, we use the word latent as a synonym for unobservable. This contrasts with a common usage in fingerprint forensics where a latent fingerprint refers to a fingermark which is difficult to see with the naked eye, but can still be observed via specialized techniques.

Both the observed and the latent minutia configurations are modelled as marked point processes. We assume that different fingers have independent latent minutiae configurations, whether those fingers belong to the same or different individuals. Thus we can rephrase our two model hypotheses (1) as

H_p : A and B originate from a common latent minutia configuration $M \subset \mathbb{M}$,

H_d : A and B originate from independent latent minutia configurations $M, M' \subset \mathbb{M}$.

In the notation of marked point processes, each minutia $m \in \mathbb{M} = \mathbb{C} \times \mathbb{S}^1 \times \{-1, 0, 1\}$ is a marked point. The projection of m onto the location space $\mathbb{C}$, denoted r_m, is called a point and the projection onto $\mathbb{S}^1 \times \{-1, 0, 1\}$, denoted (s_m, t_m), is called a mark. The points form a finite Poisson point process on the complex plane with intensity function $\rho : \mathbb{C} \to [0, \infty)$ such that $\rho_0 = \int_{\mathbb{C}} \rho(r) \, dr$ is positive and finite. The marks are assumed to be independently and identically distributed and independent of the points. The marks have density g with respect to the product measure $\mu = \nu \times \#$, where $\#$ is the counting measure on $\{-1, 0, 1\}$. For the latent minutiae only the types $\{-1, 1\}$ have meaning so we must insist that $g(s, 0) = 0$ for any $s \in \mathbb{S}^1$.

We write the resulting marked Poisson point process as $M \sim \text{MPPP}(\rho, g)$. The cardinality of M is Poisson distributed with mean ρ_0, and, conditionally on the cardinality $|M|$ of M, the points are independent and identically distributed with density ρ/ρ_0.

The observed fingerprint A is obtained from the latent minutia configuration M through three basic operations, thinning, displacement, and mapping, as follows:

A1: thinning. Only a subset of the latent minutiae are observed, resulting in $M_{A1} = \{m \in M : I_A(m) = 1\}$, where the indicators $I_A(m)$ are Bernoulli variables with success probabilities $\delta_A(r_m)$. Here $\delta_A : \mathbb{C} \to [0, 1]$ is a Borel function which we refer to as the selection function for A. We then have

$$M_{A1} \sim \text{MPPP}(\rho_{A1}, g_{A1}) \text{ where } \rho_{A1}(r) = \rho(r)\delta_A(r), \quad g_{A1} = g.$$

A2: displacement. The locations r_m in M_{A1} are subjected to additive errors $e_m \in \mathbb{C}$ with density f_A, the orientations s_m are subjected to multiplicative errors $v_m \in \mathbb{S}^1$ with density h_A, and the types are subjected to multiplicative classification errors $c_m \in \{-1, 0, 1\}$ with distribution d_A so that $c_m = 1$ corresponds to a correct classification, $c_m = 0$ to the type being unobserved, and $c_m = -1$ represents a misclassification. This results in $M_{A2} = \{(r_m + e_m, v_m s_m, c_m t_m) : m \in M_{A1}\}$. Consequently, $M_{A2} \sim \text{MPPP}(\rho_{A2}, g_{A2})$, where

$$\rho_{A2}(r) = f_A * \rho_{A1}(r) = \int_{\mathbb{C}} f_A(e)\rho_{A1}(r - e) \, de$$

is obtained by usual convolution in $\mathbb{C}$. The mark density is

$$g_{A2}(s, t) = \sum_{u \in \{-1,1\}} d_A(ut)h_A * g_{A1}(s, u) = \sum_{u \in \{-1,1\}} d_A(ut) \int_{\mathbb{S}^1} h_A(v)g_{A1}(s\overline{v}, u) \, d\nu(v).$$

A3: mapping. Finally, the marked points are subjected to a similarity transformation to obtain

$$A = \{(\psi_A r_m + \tau_A, \psi_A s_m/|\psi_A|, t_m) : m \in M_{A2}\}, \tag{3}$$

with $(\tau_A, \psi_A) \in \mathbb{C} \times (\mathbb{C} \setminus \{0\})$. Thus $A \sim \text{MPPP}(\rho_{A3}, g_{A3})$ where $\rho_{A3}(r) = \rho_{A2}\{(r - \tau_A)/\psi_A\}/|\psi_A|^2$ and $g_{A3}(s, t) = g_{A2}(s\overline{\psi_A}/|\psi_A|, t)$.

The model for B is specified analogously: B is the MPPP derived from a latent minutia configuration M' by three similar steps B1–B3 obtained by replacing A with B everywhere,

i.e. $B \sim \text{MPPP}(\rho_{B3}, g_{B3})$ with intensity function and the mark density defined as above, but using a new function δ_B, new indicators $I_B(m)$, new distributions f_B, h_B, d_B, new error terms e'_m, v'_m, c'_m, and new parameters τ_B, ψ_B.

Finally, we make the following independence assumptions. Under H_d we have M and M' are independent and identically distributed, while under H_p, $M = M'$. In both cases they have distribution $\text{MPPP}(\rho, g)$. Conditional on M and M', all the variables $I_A(m), e_m, v_m, c_m$ for $m \in M$, and $I_B(m), e'_m, v'_m, c'_m$ for $m \in M'$ are mutually independent with distributions which do not depend on M and M'.

3.2 Density under the defence hypothesis

The functions $\rho, g, \delta_A, \delta_B, f_A, f_B, h_A, h_B, d_A, d_B$ depend on some set of parameters denoted Θ; we describe a specific choice of these functions in Section 4. In the following we suppress the dependence on Θ for ease of presentation.

In order to obtain the densities for observed minutiae configurations we introduce the probability distribution $\zeta = \text{MPPP}(\varphi, 1/3)$ as a dominating measure. Using the fact that

$$\int_{\mathbb{C}} \rho_{A3}(r) \, dr = \int_{\mathbb{C}} \rho_{A2}(r) \, dr = \int_{\mathbb{C}} \rho_{A1}(r) \, dr = \int_{\mathbb{C}} \rho(r) \delta_A(r) \, dr,$$

the marginal density of A with respect to ζ becomes

$$\text{pr}(A \,|\, \Theta) = c(A) \exp\left\{ -\int_{\mathbb{C}} \rho(r) \delta_A(r) \, dr \right\} \prod_{a \in A} \rho_{A3}(r_a) g_{A3}(s_a, t_a), \tag{4}$$

where

$$c(A) = 3^{|A|} \exp(1) \prod_{a \in A} \varphi(r_a)^{-1}$$

depends only on the data, see e.g. Møller and Waagepetersen (2004, p. 25). Similarly, the density $\text{pr}(B \,|\, \Theta)$ of B with respect to ζ is obtained by replacing A by B everywhere in (4). Under H_d, the fingerprint A and fingermark B are independent and thus the density with respect to $\zeta \times \zeta$ is simply the product

$$\text{pr}(A, B \,|\, \Theta, H_d) = c(A) c(B) \exp\left\{ -\int_{\mathbb{C}} \rho(r) \delta_A(r) \, dr - \int_{\mathbb{C}} \rho(r) \delta_B(r) \, dr \right\}$$
$$\times \left\{ \prod_{a \in A} \rho_{A3}(r_a) g_{A3}(s_a, t_a) \right\} \left\{ \prod_{b \in B} \rho_{B3}(r_b) g_{B3}(s_b, t_b) \right\}. \tag{5}$$

3.3 Density under the prosecution hypothesis

The marginal densities of A and B are identical under both H_d and H_p, but to obtain the joint density of (A, B) under H_p we need to account for missing information, namely the matching of marked points in A and B. To handle this, we first partition M into four parts

$$M_{11} = \{ m \in M : I_A(m) = 1, I_B(m) = 1 \}, \qquad M_{10} = \{ m \in M : I_A(m) = 1, I_B(m) = 0 \},$$
$$M_{01} = \{ m \in M : I_A(m) = 0, I_A(m) = 1 \}, \qquad M_{00} = \{ m \in M : I_A(m) = 0, I_B(m) = 0 \},$$

which are independent and disjoint marked Poisson point processes, all with mark density g, and with intensity functions for the locations is

$$\rho_{11}(r) = \rho(r) \delta_A(r) \delta_B(r), \qquad\qquad \rho_{10}(r) = \rho(r) \delta_A(r) \{ 1 - \delta_B(r) \},$$
$$\rho_{01}(r) = \rho(r) \{ 1 - \delta_A(r) \} \delta_B(r), \qquad\qquad \rho_{00}(r) = \rho(r) \{ 1 - \delta_A(r) \} \{ 1 - \delta_B(r) \},$$

respectively, see Møller and Waagepetersen (2004, p.23). Note that $M_{A1} = M_{11} \cup M_{10}$ and $M_{B1} = M_{11} \cup M_{01}$, so M_{00} will play no role in the sequel. This partitioning is illustrated in Fig. 2.
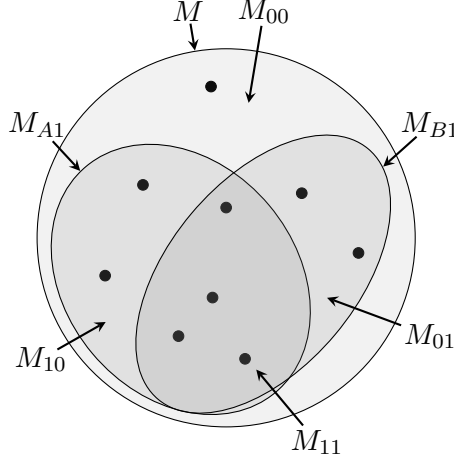

Figure 2: Partitioning the latent minutiae into those that are observed in A only (M_{10}), B only (M_{01}), both (M_{11}), and neither (M_{00}). The dots indicate minutiae locations.

Applying steps A2–A3 to M_{10} yields $M_{103} \sim \text{MPPP}(\rho_{103}, g_{A3})$, where

$$\rho_{103}(r) = f_A * \rho_{10}\{(r - \tau_A)/\psi_A\}/|\psi_A|^2. \tag{6}$$

Similarly, applying steps B2–B3 to M_{01} yields $M_{013} \sim \text{MPPP}(\rho_{013}, g_{B3})$ with

$$\rho_{013}(r) = f_B * \rho_{01}\{(r - \tau_B)/\psi_B\}/|\psi_B|^2. \tag{7}$$

Finally, for each $m \in M_{11}$ we apply steps A2–A3 to yield a marked point $a(m)$, and separately steps B2–B3 to yield a marked point $b(m)$. The set of paired marked points

$$M_{113} = \{(a(m), b(m)) : m \in M_{11}\}$$

forms an MPPP with paired points in $\mathbb{C} \times \mathbb{C}$ and corresponding marks in $(\mathbb{S}^1 \times \{-1, 0, 1\})^2$. These points have intensity function

$$\rho_{113}(r_a, r_b) = \int_{\mathbb{C}} \rho_{11}(r) f_A\{(r_a - \tau_A)/\psi_A - r\} f_B\{(r_b - \tau_B)/\psi_B - r\}/|\psi_A \psi_B|^2 \, dr. \tag{8}$$

The marks are independent and identically distributed with density

$$g_{113}(s_a, t_a, s_b, t_b) = \sum_{u \in \{-1, 1\}} d_A(ut_a) d_B(ut_b) \int_{\mathbb{S}^1} g(s, u) h_A\left(\frac{s_a \overline{s\psi_A}}{|\psi_A|}\right) h_B\left(\frac{s_b \overline{s\psi_B}}{|\psi_B|}\right) \, d\nu(s) \tag{9}$$

with respect to $\mu \times \mu$, and they are independent of the points.

The distribution of M_{113} is dominated by $\zeta_2 = \text{MPPP}(\varphi_2, 1/9)$, the MPPP whose points form a Poisson point process on $\mathbb{C} \times \mathbb{C}$ with intensity function φ_2 and whose marks are independently uniformly distributed on $(\mathbb{S}^1 \times \{-1, 0, 1\})^2$ and independent of the points. From (8) we have

$$\int_{\mathbb{C} \times \mathbb{C}} \rho_{113}(r_a, r_b) \, dr_a dr_b = \int_{\mathbb{C}} \rho_{11}(r) \, dr,$$

6

and hence the density of M_{113} with respect to ζ_2 is

$$\text{pr}(M_{113} \mid \Theta, H_p) = c_2(M_{113}) \exp\left\{-\int_{\mathbb{C}} \rho_{11}(r)\, \mathrm{d}r\right\} \prod_{(a,b) \in M_{113}} \rho_{113}(r_a, r_b) g_{113}(s_a, t_a, s_b, t_b),$$

where

$$c_2(M_{113}) = 9^{|M_{113}|} \exp(1) \prod_{(a,b) \in M_{113}} \{\varphi(r_a)\varphi(r_b)\}^{-1}.$$

Observing that $c(M_{103})c(M_{013})c_2(M_{113}) = \exp(1)c(A)c(B)$, the density for $(M_{103}, M_{013}, M_{113})$ with respect to $\zeta \times \zeta \times \zeta_2$ is

$$\begin{aligned}
\text{pr}(M_{103}, M_{013}, M_{113} \mid \Theta, H_p) = {} & c(A)c(B) \exp\left[1 - \int_{\mathbb{C}} \rho(r) \left\{\delta_A(r) + \delta_B(r) - \delta_A(r)\delta_B(r)\right\} \mathrm{d}r\right] \\
& \times \left\{\prod_{a \in M_{103}} \rho_{103}(r_a) g_{A3}(s_a, t_a)\right\} \left\{\prod_{b \in M_{013}} \rho_{013}(r_b) g_{B3}(s_b, t_b)\right\} \\
& \times \left\{\prod_{(a,b) \in M_{113}} \rho_{113}(r_a, r_b) g_{113}(s_a, t_a, s_b, t_b)\right\}.
\end{aligned} \tag{10}$$

The three marked point processes $(M_{103}, M_{013}, M_{113})$ can be identified with a labelled bipartite graph (A, B, ξ) of maximum degree one with partitioned vertex set (A, B) and edge set ξ. Specifically, we have the transformation

$$A = M_{103} \cup \Pi_A(M_{113}), \quad B = M_{013} \cup \Pi_B(M_{113}), \quad \xi = \{\langle a, b \rangle : (a, b) \in M_{113}\},$$

where we use the notation $\langle a, b \rangle$ for elements of ξ, which consist of edges between marked points, whereas the elements of $(a, b) \in M_{113}$ are the marked points themselves. Furthermore, we have the inverse transformation

$$M_{103} = A \setminus \Pi_A(\xi), \quad M_{013} = B \setminus \Pi_B(\xi), \quad M_{113} = \{(a, b) : \langle a, b \rangle \in \xi\},$$

where Π_A projects to a marked point set on $\mathbb{M}$ via

$$\Pi_A(M_{113}) = \{a : (a, b) \in M_{113} \text{ for some } b \in \mathbb{M}\}.$$

We slightly abuse notation by also writing

$$\Pi_A(\xi) = \{a : \langle a, b \rangle \in \xi \text{ for some } b \in \mathbb{M}\}.$$

The projector Π_B is defined analogously.

We let $\Xi(A, B)$ denote the space of all possible values for ξ, i.e. all possible edge sets for the vertex sets A and B. The cardinality of $\Xi(A, B)$ is

$$|\Xi(A, B)| = \sum_{n_\xi = 0}^{\min(n_A, n_B)} \frac{n_A!}{n_\xi!(n_A - n_\xi)!} \frac{n_B!}{n_\xi!(n_B - n_\xi)!} n_\xi!,$$

where n_A, n_B, and n_ξ denote the cardinality of A, B, and M_{113}, respectively. This reflects choosing n_ξ points each from A and B to be matched and considering all $n_\xi!$ edge sets between those points.

Let $\mathrm{pr}(A, B, \xi \,|\, \Theta, H_p)$ denote the density of (A, B, ξ) with respect to $\tilde{\zeta}$, where for fixed (A, B), $\tilde{\zeta}$ is the counting measure on $\Xi(A, B)$, i.e. it holds for $C \subseteq \Xi(A, B)$ that

$$\mathrm{d}\tilde{\zeta}(A, B, C) = |C| \, \mathrm{d}\zeta(A) \mathrm{d}\zeta(B).$$

Note that $\sum_{\xi \in \Xi(A,B)} \mathrm{d}\tilde{\zeta}(A, B, \xi) = \mathrm{d}\zeta(A) \mathrm{d}\zeta(B)$, and thus the marginal density $\mathrm{pr}(A, B \,|\, \Theta, H_p)$ of the observed points with respect to $\zeta \times \zeta$ is

$$\mathrm{pr}(A, B \,|\, \Theta, H_p) = \sum_{\xi \in \Xi(A,B)} \mathrm{pr}(A, B, \xi \,|\, \Theta, H_p). \tag{11}$$

Now let λ denote the distribution of (A, B, ξ) induced by $\zeta \times \zeta \times \zeta_2$, i.e. λ is the measure $\zeta \times \zeta \times \zeta_2$ transformed by the bijection $(M_{103}, M_{013}, M_{113}) \to (A, B, \xi)$. Using the expansion for the Poisson process measure (Møller and Waagepetersen, 2004, proposition 3.1), a long but straightforward calculation shows that $\mathrm{d}\lambda(A, B, \xi)/\mathrm{d}\tilde{\zeta} = \exp(-1)$, whence

$$\mathrm{pr}(A, B, \xi \,|\, \Theta, H_p) = \exp(-1)\mathrm{pr}(M_{103}, M_{013}, M_{113} \,|\, \Theta, H_p). \tag{12}$$

4 Parametric models

4.1 Model specification

To complete the specification of our basic point process model we need to specify parametric models for the basic elements $(\rho, g, \delta_A, \delta_B, f_A, f_B, h_A, h_B, d_A, d_B)$ introduced in Section 3 that define our marked Poisson point processes and the corresponding likelihood ratios. Clearly there are many possibilities. Below we specify a simple choice to be used in the present paper with the purpose of illustrating and investigating the methodology. We shall return to the potential for improving this choice later. Forbes (2014) provides a more detailed discussion of the issue.

We assume the intensity ρ and mark density g of M are

$$\rho(r) = \rho_0 \varphi(r; \tau_0, \sigma_0^2), \quad g(s, t) = |t| \sqrt{\chi^{|t|+t}(1-\chi)^{|t|-t}},$$

where $\rho_0 > 0$ and $\chi \in (0, 1)$ is the probability that a minutia is a bifurcation. Note that $g(s, 1) = \chi$, $g(s, 0) = 0$, and $g(s, -1) = 1 - \chi$. Without loss of generality, we assume that $\tau_0 = 0$, since this parameter can be absorbed into τ_A and τ_B, cf. (3). Similarly, we assume that $\sigma_0 = 1$, since this parameter can be absorbed into ψ_A and ψ_B. Due to the latent mark distribution $g(s, t)$ being uniform over s, we have

$$g_{A1}(s, t) = g(s, t), \quad g_{A2}(s, t) = g_{A3}(s, t) = d_A(t)\chi + d_A(-t)(1 - \chi),$$

and similarly for B.

Thinning. We assume the selection probabilities are constant with $\delta_A(r) = \delta_A \in (0, 1)$ and $\delta_B(r) = \delta_B \in (0, 1)$ so that the intensities after thinning become

$$\rho_{A1}(r) = \rho_0 \delta_A \varphi(r), \quad \rho_{B1}(r) = \rho_0 \delta_B \varphi(r).$$

Displacement. We assume the error distributions of the minutia locations and types are

$$f_A(r) = f_B(r) = \varphi(r; 0, \omega^2), \quad d_A(c) = d_B(c) = I(c = 1)\varepsilon + I(c = 0)(1 - \varepsilon)$$

for some $\varepsilon \in (0, 1)$, where I is the indicator function. Thus we assume that there are no type misclassifications, though we allow types to be unobserved. These error functions imply

$$\rho_{A2}(r) = \rho_0 \delta_A \varphi(r; 0, 1 + \omega^2), \quad \rho_{B2}(r) = \rho_0 \delta_B \varphi(r; 0, 1 + \omega^2),$$
$$g_{A2}(s, t) = g_{B2}(s, t) = (1 - |t|)\varepsilon + |t|(1 - \varepsilon)g(s, t).$$

8

The error distributions of the orientations $h_A = h_B = h$ are root von Mises distributions rvM $(1, \kappa)$ as defined in Section 2.3.

Mapping. After mapping we have

$$\rho_{A3}(r) = \rho_0 \delta_A \varphi\{r; \tau_A, (1+\omega^2)|\psi_A|^2\}, \quad \rho_{B3}(r) = \rho_0 \delta_B \varphi\{r; \tau_B, (1+\omega^2)|\psi_B|^2\},$$
$$g_{A3}(s,t) = g_{B3}(s,t) = (1-|t|)\varepsilon + |t|(1-\varepsilon)g(s,t).$$

We let $\psi = \psi_A \overline{\psi_B}/(|\psi_A||\psi_B|)$; then ψ specifies the relative rotation of A with respect to B. For simplicity we assume in the following that the minutia configurations are represented on the same scale so that $|\psi_A| = |\psi_B|$. further let $\sigma^2 = (1+\omega^2)|\psi_A|^2 = (1+\omega^2)|\psi_B|^2$.

4.2 Density under the defence hypothesis

For the defence likelihood (5) we have

$$\mathrm{pr}(A, B \,|\, \Theta, H_d) = \tilde{c}(A)\tilde{c}(B)\exp\left\{-\rho_0(\delta_A + \delta_B)\right\} \rho_0^{n_A + n_B} \delta_A^{n_A} \delta_B^{n_B}$$
$$\times \chi^{n_A^{(1)} + n_B^{(1)}}(1-\chi)^{n_A^{(-1)} + n_B^{(-1)}} \left\{\prod_{a \in A} \varphi(r_a; \tau_A, \sigma^2)\right\}\left\{\prod_{b \in B} \varphi(r_b; \tau_B, \sigma^2)\right\}, \quad (13)$$

where $n_A^{(t)} = \sum_{a \in A} I(t_a = t)$ for each $t \in \{-1, 0, 1\}$, $\tilde{c}(A) = c(A)\varepsilon^{n_A^{(0)}}(1-\varepsilon)^{n_A^{(-1)} + n_A^{(1)}}$, and similarly for $n_B^{(t)}$ and $\tilde{c}(B)$.

4.3 Density under the prosecution hypothesis

The transformed intensities (6), (7), and (8) become

$$\rho_{103}(r_a) = \rho_0 \delta_A(1-\delta_B)\varphi(r_a; \tau_A, \sigma^2), \quad \rho_{013}(r_b) = \rho_0(1-\delta_A)\delta_B\varphi(r_b; \tau_B, \sigma^2),$$
$$\rho_{113}(r_a, r_b) = \rho_0 \delta_A \delta_B \varphi_2(r_a, r_b; \tau_A, \tau_B, \Sigma_{AB}), \quad \Sigma_{AB} = \sigma^2 \begin{pmatrix} 1 & \psi/(1+\omega^2) \\ \overline{\psi}/(1+\omega^2) & 1 \end{pmatrix}. \quad (14)$$

The mark density (9) becomes

$$g_{113}(s_a, t_a, s_b, t_b) = g_{A2}(s_a, t_a)g_{B2}(s_b, t_b)T(t_a, t_b)\exp\{\kappa\Re(s_a\overline{s_b}\psi)\}/I_0(\kappa),$$

where

$$T(t_a, t_b) = (1 + t_a t_b)\left\{2^4 \chi^{|t_a| + t_a + |t_b| + t_b}(1-\chi)^{|t_a| - t_a + |t_b| - t_b}\right\}^{-t_a t_b/4}. \quad (15)$$

Note that $T(t_a, t_b) = 1$ if $t_a t_b = 0$, $T(t_a, t_b) = 0$ if $t_a t_b = -1$, $T(1, 1) = 1/\chi$, and $T(-1, -1) = 1/(1-\chi)$. Combining these basic elements with (10) and (12), we obtain

$$\mathrm{pr}(A, B, \xi \,|\, \Theta, H_p) = \tilde{c}(A)\tilde{c}(B)\exp\left\{-\rho_0(\delta_A + \delta_B - \delta_A \delta_B)\right\} \rho_0^{n_A + n_B - n_\xi}$$
$$\times \chi^{n_A^{(1)} + n_B^{(1)} - n_\xi^{(1)}}(1-\chi)^{n_A^{(-1)} + n_B^{(-1)} - n_\xi^{(-1)}} \delta_A^{n_A} \delta_B^{n_B}(1-\delta_A)^{n_B - n_\xi}(1-\delta_B)^{n_A - n_\xi}$$
$$\times \left\{\prod_{a \in A \setminus \Pi_A(\xi)} \varphi(r_a; \tau_A, \sigma^2)\right\}\left\{\prod_{b \in B \setminus \Pi_B(\xi)} \varphi(r_b; \tau_B, \sigma^2)\right\} \quad (16)$$
$$\times \left[\prod_{\langle a,b \rangle \in \xi} \varphi_2(r_a, r_b; \tau_A, \tau_B, \Sigma_{AB})\frac{\exp\{\kappa\Re(s_a\overline{s_b}\psi)\}}{I_0(\kappa)}\right],$$

where $n_\xi^{(t)} = \sum_{\langle a,b \rangle \in \xi} I(t_a = t_b = t)$.

4.4 Variability of parameters

The densities in the parametric models specified above depend on

$$\Theta = (\rho_0, \chi, \varepsilon, \delta_A, \delta_B, \tau_A, \tau_B, \sigma, \psi, \omega, \kappa),$$

where $\rho_0 > 0$, $\chi, \varepsilon, \delta_A, \delta_B \in (0,1)$, $\tau_A, \tau_B \in \mathbb{C}$, $\sigma > 0, \psi \in \mathbb{S}^1, \omega > 0$, and $\kappa > 0$ are variation independent parameters. As τ_A and τ_B are complex numbers there are thirteen real parameters in total. Of these, ρ_0 and χ relate to the latent minutiae and are common to all fingerprints and fingermarks under consideration. We shall assume the same for ε, ω, and κ. The parameters ρ_0, χ, ω, and κ will be replaced by point estimates and hence treated as being known; we suppress the dependence on these parameters in the following. Similarly ε is considered fixed; it only enters via the factors $\tilde{c}(A)$ and $\tilde{c}(B)$ which are common to both hypotheses and hence these cancel in the likelihood ratio so ε can be ignored. This would also be true if we had separate observation probabilities ε_A and ε_B for the prints and marks. The remaining parameters

$$\theta = (\delta_A, \delta_B, \tau_A, \tau_B, \sigma, \psi)$$

vary from one fingerprint or fingermark to the next, according to suitable prior distributions to be specified below. In this way, our approach takes inspiration both from empirical Bayes methods and random effect models.

We follow Dawid and Lauritzen (2000) and ensure that we use compatible prior distributions for the competing models H_d and H_p. Our compatibility condition is that the marginal distributions agree, which leads to the constraint

$$\int \text{pr}(A \,|\, \theta)\{\text{pr}(\theta \,|\, H_p) - \text{pr}(\theta \,|\, H_d)\} \, d\theta$$

for arbitrary values of A. For the parametric model described in Section 4, the constraint becomes

$$\int \{\text{pr}(\delta_A, \tau_A, \sigma \,|\, H_p) - \text{pr}(\delta_A, \tau_A, \sigma \,|\, H_d)\}$$
$$\times \exp(-\rho_0 \delta_A) \left\{ \frac{\delta_A}{\sigma^2} \exp\left(-\frac{|\tau_A|^2 - 2|\tau_A r_1| + |r_2|^2}{\sigma^2} \right) \right\}^{n_A} d(\delta_A, \tau_A, \sigma) = 0$$

for all $r_1, r_2 \in \mathbb{C}$, and all non-negative integers n_A. The fundamental lemma of the calculus of variations then implies $\text{pr}(\delta_A, \tau_A, \sigma \,|\, H_p) = \text{pr}(\delta_A, \tau_A, \sigma \,|\, H_d)$ almost everywhere. Thus $\delta_A, \delta_B, \tau_A, \tau_B$, and σ must have common priors under H_d and H_p. The remaining parameter ψ does not enter under H_d and is thus unconstrained by this consideration.

For our likelihood $\text{pr}(A, B \,|\, H_p)$ to be invariant under scale transformations, we must require that

$$\text{pr}(A, B, \xi \,|\, \theta, H_p)\text{pr}(\lambda\tau_A, \lambda\tau_B, \lambda\sigma, \psi) \, d(\lambda\tau_A, \lambda\tau_B, \lambda\sigma, \psi)$$

to be independent of the value of $\lambda > 0$. Thus, for the likelihood to be invariant under translation and rotation as well, by (16) the prior density must be of the form

$$\text{pr}(\tau_A, \tau_B, \sigma, \psi \,|\, H_p) = \sigma^{-5}.$$

A similar argument shows that $\text{pr}(\tau_A, \tau_B, \sigma, \psi \,|\, H_d) = \sigma^{-5}$. This prior density is improper, i.e. not integrable over the entire parameter domain. Normally such a prior may result in a meaningless likelihood ratio. However, in our case the improper prior is common to both models H_d and H_p under consideration and the marginal likelihood ratio is equal to the limit of likelihood ratios determined by integrals over the same large box in numerator and denominator.

Under both H_p and H_d, we also assume the following. The fingerprint selection probability δ_A has a conjugate beta distribution with parameters $(\alpha_\delta, \beta_\delta)$. Assuming that we have a database that is representative for minutiae in a fingerprint, these parameters can be estimated reliably. The fingermark selection probability δ_B has a uniform distribution on $(0, 1)$, as it will refer to a fingermark that is not taken from a well-defined population of marks. Finally we assume that δ_A, δ_B, τ_A, τ_B, σ, and ψ are mutually independent.

Thus the joint prior density of the varying parameters is the same under both H_p and H_d, and equal to

$$\mathrm{pr}(\theta) = \mathrm{pr}(\theta \mid H_d) = \mathrm{pr}(\theta \mid H_p) = \frac{\Gamma(\alpha_\delta + \beta_\delta)}{\Gamma(\alpha_\delta)\Gamma(\beta_\delta)} \delta_A^{\alpha_\delta - 1} (1 - \delta_A)^{\beta_\delta - 1} \sigma^{-5}, \qquad (17)$$

where $\Gamma(\cdot)$ is the Gamma function. We have suppressed the dependence of $\mathrm{pr}(\theta)$ on the hyper-parameters α_δ and β_δ.

Our final model contains the unknown parameters $\rho_0, \chi, \omega, \kappa, \alpha_\delta, \beta_\delta$. In the developments below we shall consider these parameters as fixed and equal to values estimated from a database of fingerprints and fingermarks as described further in Section 6.3 below.

5 Calculating the likelihood ratio

5.1 Defining the likelihood ratio

We can in principle obtain our desired likelihood ratio (2) by summing (16) over ξ, taking its expectation, and dividing by the expectation of (13), where the expectations are with respect to θ. However, under H_p the number of terms in the sum (11) is too large to compute by brute force. For example, for $n_A = n_B = 100$, $|\Xi(A, B)|$ is approximately equal to 10^{165}. We therefore proceed under H_p by approximating the expectations and the sum using a Monte Carlo sampler to be further discussed below.

Though some may prefer to call Λ a Bayes factor, integrated likelihood ratio, or marginal likelihood ratio, we use the term likelihood ratio to conform with standard terminology in forensic science.

5.2 Integrating the density under H_d

Under H_d we can analytically integrate $\mathrm{pr}(A, B \mid \theta, H_d)\mathrm{pr}(\theta)$ over θ as follows. First,

$$\int_{\mathbb{C}^2} \prod_{a \in A} \varphi(r_a; \tau_A, \sigma^2) \, \mathrm{d}\tau_A = \frac{\pi^{1 - n_A}}{n_A} \sigma^{2(1 - n_A)} \exp\left(-S_A / \sigma^2\right),$$

where $S_A = \sum_{a \in A} \|r_a - r_{A\bullet}\|^2$ is the sum of squared deviations from the average $r_{A\bullet} = n_A^{-1} \sum_{a \in A} r_a$; the integral over τ_B is analogous. Second, we can integrate over δ_A using

$$\int_0^1 e^{-\rho_0 \delta_A} \delta_A^{\alpha_\delta + n_A - 1} (1 - \delta_A)^{\beta_\delta - 1} \, \mathrm{d}\delta_A = e^{-\rho_0} \frac{\Gamma(\alpha_\delta + n_A)\Gamma(\beta_\delta)}{\Gamma(\alpha_\delta + \beta_\delta + n_A)} {}_1F_1(\beta_\delta, \alpha_\delta + \beta_\delta + n_A, \rho_0),$$

where ${}_1F_1$ is the confluent hypergeometric function (Olver et al., 2010, chapter 13). Third, for δ_B, we have

$$\int_0^1 e^{-\rho_0 \delta_B} \delta_B^{n_B} \, \mathrm{d}\delta_B = e^{-\rho_0} \frac{1}{n_B + 1} {}_1F_1(1, n_B + 2, \rho_0).$$

Fourth, the integral over σ is proportional to a gamma density for σ^{-2}:

$$\int_0^\infty \sigma^{-2n_A - 2n_B - 1} \exp\left\{-(S_A + S_B)/\sigma^2\right\} \, \mathrm{d}\sigma = \Gamma(n_A + n_B) \left(S_A + S_B\right)^{-n_A - n_B} / 2.$$

Combining these integrals with (11) and (17), the marginal likelihood under H_d is

$$\mathrm{pr}(A, B \mid H_d) = \tilde{c}(A)\tilde{c}(B)e^{-2\rho_0}\pi^2\chi^{n_A^{(1)}+n_B^{(1)}}(1-\chi)^{n_A^{(-1)}+n_B^{(-1)}}\left\{\frac{\rho_0}{\pi(S_A+S_B)}\right\}^{n_A+n_B}$$

$$\times \frac{\Gamma(\alpha_\delta+\beta_\delta)\Gamma(\alpha_\delta+n_A)\Gamma(n_A+n_B)}{2\Gamma(\alpha_\delta)\Gamma(\alpha_\delta+\beta_\delta+n_A)n_An_B(n_B+1)} {}_1F_1(\beta_\delta,\alpha_\delta+\beta_\delta+n_A,\rho_0){}_1F_1(1,n_B+2,\rho_0).$$

5.3   Approximating the likelihood under H_p

We are interested in calculating the likelihood ratio $\Lambda = \mathrm{pr}(A, B \mid H_p)/\mathrm{pr}(A, B \mid H_d)$, cf. Section 2.1, for assessing the strength of the evidence for H_p. We cannot analytically obtain $\mathrm{pr}(A, B \mid H_p)$ because the required sums and integrals are intractable. Instead we approximate the likelihood ratio using a Markov chain Monte Carlo procedure. There are a variety of possible methods but we have chosen Chib's method (Chib, 1995; Chib and Jeliazkov, 2001). Other possibilities were investigated in Forbes (2014), who found Chib's method to be superior for our specific purpose. Chib's method uses the simple relation

$$\mathrm{pr}(A, B \mid H_p) = \frac{\mathrm{pr}(A, B, \theta^*, \xi^* \mid H_p)}{\mathrm{pr}(\theta^*, \xi^* \mid A, B, H_p)},$$

which holds for any fixed values θ^* of θ and ξ^* of ξ. The numerator is simply the product of (16) and (17). Thus we can approximate $\mathrm{pr}(A, B \mid H_p)$ by approximating the denominator, which can be rewritten as

$$\mathrm{pr}(\theta^*, \xi^* \mid A, B, H_p) = \mathrm{pr}(\delta_A^* \mid A, B, H_p) \times \mathrm{pr}(\delta_B^* \mid \delta_A^*, A, B, H_p) \times \mathrm{pr}(\tau_A^*, \tau_B^* \mid \delta_A^*, \delta_B^*, A, B, H_p)$$

$$\times \mathrm{pr}(\sigma^* \mid \delta_A^*, \delta_B^*, \tau_A^*, \tau_B^*, A, B, H_p) \times \mathrm{pr}(\psi^* \mid \delta_A^*, \delta_B^*, \tau_A^*, \tau_B^*, \sigma^*, A, B, H_p)$$

$$\times \mathrm{pr}(\xi^* \mid \delta_A^*, \delta_B^*, \tau_A^*, \tau_B^*, \sigma^*, \psi^*, A, B, H_p).$$

Each of the factors on the right-hand side can be approximated with a suitable sample average of the appropriate full conditional posterior density. The accuracy of these approximations increases with the posterior probability of (θ^*, ξ^*). Our method of selecting these values and performing the approximations is detailed in the appendix.

For the final term $\mathrm{pr}(\xi^* \mid \theta^*, A, B, H_p)$, notice the following. Given a matching $\xi \in \Xi(A, B)$ and a $\beta \in B$, let the sub-matching $\xi_{<\beta} \in \Xi(A, B)$ be given by

$$\xi_{<\beta} = \{\langle a, b \rangle \in \xi : b \in B, b < \beta\},$$

where the inequality is with respect to some arbitrary total ordering on B. Given any $m \in \mathbb{M}$ and $\phi \in \mathbb{M} \setminus (A \cup B)$, we define $\Pi_{A,m} : \Xi(A, B) \to \mathbb{M}$ by

$$\Pi_{A,m}(\xi) = \begin{cases} a & \text{if } \langle a, m \rangle \in \xi, \\ \phi & \text{otherwise,} \end{cases} \tag{18}$$

which is well-defined because ξ is the edge set of a bipartite graph with maximum degree one, and hence each vertex m is incident with at most one edge $\langle a, m \rangle \in \xi$.

With this notation, we can write

$$\mathrm{pr}(\xi^* \mid \theta^*, A, B, H_p) \quad = \quad \prod_{\beta \in B} E\left\{\mathrm{pr}(\xi^* \mid \xi_{<\beta}^*, \xi_{>\beta}^*, \theta^*, A, B, H_p) \mid \xi_{\leq\beta}^*, \theta^*, A, B, H_p\right\} \tag{19}$$

12

where the expectation is over the sub-match $\xi^*_{>\beta}$. Notice that the possible values of $\xi^* \,|\, \xi^*_{<\beta}, \xi^*_{>\beta}$ differ only by which minutia is matched to β. By ignoring terms independent of the match of β, we see from (14)–(16) that

$$\text{pr}(\xi^* \,|\, \xi^*_{<\beta}, \xi^*_{>\beta}, \theta^*, A, B, H_p) \propto \exp\left[w\left\{\Pi_{A,\beta}(\xi^*), \beta \,|\, \theta^*\right\}\right] I\{\Pi_{A,\beta}(\xi^*) \notin \Pi_A(\xi^*_{<\beta} \cup \xi^*_{>\beta})\} \quad (20)$$

where w is

$$w(a, b \,|\, \theta) = I(a \in A)I(b \in B)\left[\Re\left(\kappa s_a \overline{\psi s_b} + 2\frac{\omega^2 + 1}{(\omega^2 + 1)^2 - 1}\overline{\psi}\frac{r_a - \tau_A}{\sigma}\overline{\frac{r_b - \tau_B}{\sigma}}\right)\right.$$
$$\left. -\frac{1}{(\omega^2 + 1)^2 - 1}\left(\frac{|r_a - \tau_A|^2}{\sigma^2} + \frac{|r_b - \tau_B|^2}{\sigma^2}\right) + \log\left\{\frac{T(t_a, t_b)(\omega^2 + 1)^2}{\rho_0 I_0(\kappa)(1 - \delta_A)(1 - \delta_B)\omega^2(\omega^2 + 2)}\right\}\right]. \quad (21)$$

The normalization constant of (20) can be obtained by summing over the support, which is $\xi^*_{<\beta} \cup \xi^*_{>\beta}$ and $\xi^*_{<\beta} \cup \xi^*_{>\beta} \cup \{\langle a, \beta \rangle\}$ for each $a \in A$.

Thus we can evaluate and normalize (20), and therefore we can approximate (19) by approximating each expectation with a sample average. Further details are given in the appendix.

5.4 Sampling procedure

We use a Metropolis-within-Gibbs sampler to generate joint samples of (θ, ξ) from the posterior distribution $\text{pr}(A, B, \xi \,|\, \theta, H_p)\text{pr}(\theta)$, the product of (16) and (17). Our method is detailed in the appendix. Briefly, we alternate between updating δ_A, δ_B, (τ_A, τ_B), σ, ψ, and ξ. We use Gibbs updates for everything except ξ: for δ_A and δ_B this involves a rejection sampler, while the other updates are straightforward. For ξ, Green and Mardia (2006) propose using a Metropolis–Hastings sampler which creates or breaks a single, random matched pair at each iteration. However, we have developed a different sampler for ξ which considers all matches for a given minutia simultaneously and computes the probability of each match. Empirically our sampler appears to converge faster than the sampler in Green and Mardia.

6 Data analysis

6.1 Datasets

To investigate the feasibility of our model and algorithm for fingerprint analysis we now apply these to real and simulated data examples.

The real dataset originates from a small database provided by the National Institute for Standards and Technology (NIST) and the Federal Bureau of Investigation (FBI) (Garris and McCabe, 2000). This database consists of 258 fingermarks and their corresponding exemplar fingerprints. The exemplar fingerprints A are all of high quality, and the fingermarks B are of significantly lower quality. The fingerprint/fingermark pairs are partitioned into three sets based on the quality of the fingermarks: 88 pairs are of relatively good quality, 85 are bad, and 85 are ugly; see Fig. 3. All fingermarks and fingerprint images have their minutiae hand-labelled by expert fingerprint examiners. This dataset is used for estimation of unknown parameters, for model criticism, and for evaluating the performance of the calculated likelihood ratio.

For reference we also apply our method to data which are simulated from the model using the parameters estimated from the database as described below. We generated 258 fingerprint/fingermark pairs according to the model described in Section 4 and Section 4.4. To ease the comparison with the real database, we also partitioned the simulated data into a good set
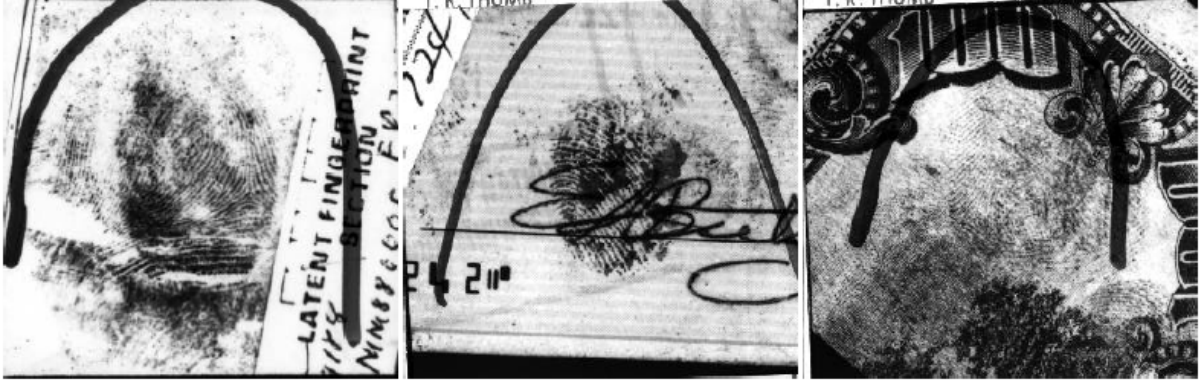
Figure 3: Example fingermarks from Garris and McCabe (2000). From left to right, the finger-mark qualities are good, bad, and ugly.

consists of those 88 pairs with the highest number of fingermark minutiae n_B, a bad set containing the next 85 pairs, and an ugly set containing those 85 pairs with the lowest n_B. By comparing our results on the NIST database to our results on the simulated data we are able to distinguish model inadequacies from algorithm errors or performance issues.

6.2 Model criticism

The question of model accuracy was investigated in Forbes (2014, chapter 7); it is apparent that some of the model features are oversimplified and the data behaviour deviates from the assumptions. For example, our model assumes the minutia are independently thinned with constant thinning frequency, have independent orientations, and have independent spatial observation errors. In fact, the thinning, orientations, and location distortions appear to be correlated amongst nearby minutiae. We abstain from giving the details here and choose to proceed with the simple model despite its apparent shortcomings.

6.3 Parameter estimation

We must find point estimates for the fixed parameters $\alpha_\delta, \beta_\delta, \rho_0, \chi, \omega$, and κ. As our real dataset contains matched fingerprint/fingermark pairs which conform with the prosecution hypothesis, we estimate all parameters under H_p.

The estimates are difficult to find without knowing the correct matching ξ. Unfortunately our dataset contains only 258 paired minutia configurations without matching the corresponding minutiae within a configuration; that is, it contains A_i and B_i but not ξ_i for $i = 1 \ldots 258$. Previous research (Mikalyan and Bigun, 2012) attempted to ameliorate this by running an automated matching algorithm on the dataset. However, we found the quality of these matchings to be extremely poor and instead we manually found and recorded what we believe to be the correct minutia matchings $\check{\xi}$ for each of the 258 fingerprint/fingermark pairs in the dataset (Garris and McCabe, 2000). With this matching $\check{\xi}$ fixed, we proceeded with the parameter estimation. We emphasize that $\check{\xi}$ is only used for estimation of the unknown parameters of the model and not otherwise for the calculation of likelihood ratios.

We estimate the fixed parameters by maximizing the likelihood function under H_p and based on matching-augmented data $(A_i, B_i, \check{\xi}_i)$, i.e.

$$\prod_{i=1}^{258} \left\{ \int \mathrm{pr}(A_i, B_i, \check{\xi}_i, \theta_i \mid H_p) \, \mathrm{d}\theta_i \right\} = \prod_{i=1}^{258} \mathrm{pr}(A_i, B_i, \check{\xi}_i \mid H_p; \alpha_\delta, \beta_\delta, \rho_0, \chi, \omega, \kappa),$$

where $\text{pr}(A_i, B_i, \check{\xi}_i, \theta \mid H_p)$ is the product of (16) and (17), and where the fixed parameters have been suppressed on the left-hand side of this equation. Each integrand on the left-hand side further factorizes into

$$\text{pr}(A_i, B_i, \check{\xi}_i, \theta \mid H_p) = f_0(A_i, B_i, \check{\xi}_i) \times f_1(A_i, B_i, \check{\xi}_i, \delta_A, \delta_B; \alpha_\delta, \beta_\delta, \rho_0)$$
$$\times f_2(A_i, B_i, \check{\xi}_i; \chi) \times f_3(A_i, B_i, \check{\xi}_i, \tau_A, \tau_B, \sigma, \psi; \omega, \kappa).$$

Here f_0 is independent of the parameters we are estimating and thus of no importance. Further

$$f_1(A, B, \check{\xi}, \delta_A, \delta_B; \alpha_\delta, \beta_\delta, \rho_0) = \exp\left\{-\rho_0\left(\delta_A + \delta_B - \delta_A\delta_B\right)\right\} \frac{\Gamma(\alpha_\delta + \beta_\delta)}{\Gamma(\alpha_\delta)\Gamma(\beta_\delta)}$$
$$\times \rho_0^{n_A + n_B - n_\xi} \delta_A^{\alpha_\delta + n_A - 1} \delta_B^{n_B} (1 - \delta_A)^{\beta_\delta + n_B - n_\xi - 1} (1 - \delta_B)^{n_A - n_\xi},$$

$$f_2(A, B, \check{\xi}; \chi) = \chi^{n_A^{(1)} + n_B^{(1)} - n_\xi^{(1)}} (1 - \chi)^{n_A^{(-1)} + n_B^{(-1)} - n_\xi^{(-1)}},$$

and

$$f_3(A, B, \check{\xi}, \tau_A, \tau_B, \sigma, \psi; \omega, \kappa) = \sigma^{-2(n_A + n_B) - 5} \left\{ \frac{(\omega^2 + 1)^2}{(\omega^2 + 1)^2 - 1} \right\}^{n_\xi} I_0(\kappa)^{-n_\xi}$$

$$\times \exp\left\{ -\left(\sum_{a \in A \backslash \Pi_A(\xi)} \frac{|r_a - \tau_A|^2}{\sigma^2} \right) - \left(\sum_{b \in B \backslash \Pi_B(\xi)} \frac{|r_b - \tau_B|^2}{\sigma^2} \right) \right\}$$

$$\times \exp\left\{ \sum_{\langle a,b \rangle \in \xi} \Re\left(\kappa s_a \overline{\psi s_b} + 2 \frac{\omega^2 + 1}{(\omega^2 + 1)^2 - 1} \overline{\psi} \frac{r_a - \tau_A}{\sigma} \overline{\frac{r_b - \tau_B}{\sigma}} \right) \right\}$$

$$\times \exp\left\{ -\frac{(\omega^2 + 1)^2}{(\omega^2 + 1)^2 - 1} \sum_{\langle a,b \rangle \in \xi} \left(\frac{|r_a - \tau_A|^2}{\sigma^2} + \frac{|r_b - \tau_B|^2}{\sigma^2} \right) \right\}.$$

Since $(\alpha_\delta, \beta_\delta, \rho_0)$ only enter into f_1, the estimates for these parameters are the maximizers of

$$\prod_{i=1}^{258} \left\{ \int f_1(A_i, B_i, \check{\xi}_i, \delta_{Ai}, \delta_{Bi}; \alpha_\delta, \beta_\delta, \rho_0) \, \mathrm{d}(\delta_{Ai}, \delta_{Bi}) \right\}.$$

The integral over δ_B can be obtained analytically as in Section 5.2. The integral over δ_A can be found numerically, and the resulting function can also be maximized numerically. We used the R package pracma for the integrals and the standard R function optim for the optimization. The resulting estimates are $\hat{\alpha}_\delta = 14\cdot67$, $\hat{\beta}_\delta = 3\cdot30$, and $\hat{\rho}_0 = 132\cdot74$.

Similarly, χ only enters into f_2 and can be found by directly maximizing $\sum_{i=1}^{258} \log f_2(A_i, B_i, \check{\xi}_i; \chi)$, yielding a linear equation for χ with the solution $\hat{\chi} = 0\cdot38$.

We estimate ω and κ by maximizing the third factor in the likelihood function

$$\prod_{i=1}^{258} \left\{ \int f_3(A_i, B_i, \check{\xi}_i, \tau_{Ai}, \tau_{Bi}, \sigma_i, \psi_i; \omega, \kappa) \, \mathrm{d}(\tau_{Ai}, \tau_{Bi}, \sigma_i, \psi_i) \right\}.$$

This function is too complicated to maximize using standard numerical techniques. We resort to a stochastic expectation-maximization algorithm (Celeux and Diebolt, 1985) based on the Monte Carlo Markov chain procedure described in the appendix. We fix $\alpha_\delta, \beta_\delta, \rho_0$, and χ to their estimated values above. Starting from some initial values for ω and κ, we generate a

posterior sample $(\tau_{Ai}, \tau_{Bi}, \sigma_i, \psi_i)$ for each fingerprint/fingermark pair $i = 1, \ldots, 258$. We then maximize

$$\prod_{i=1}^{258} f_3(A_i, B_i, \check{\xi}_i, \tau_{Ai}, \tau_{Bi}, \sigma_i, \psi_i; \omega, \kappa)$$

over ω and κ. The maximizing value for $x = (\omega^2+1)^2/\{(\omega^2+1)^2-1\}$ is a root of the polynomial equation

$$(R_2^2 - 4R_3^2)x^3 + (4R_3^2 - 2R_1R_2 - R_2^2)x^2 + (R_1^2 + 2R_1R_2 - R_3^2)x - R_1^2 = 0,$$

where $R_1 = \sum_{i=1}^{258} |\check{\xi}_i|$, $R_2 = \sum_{i=1}^{258} \left\{ \sigma_i^{-2} \sum_{\langle a,b \rangle \in \check{\xi}_i} |r_a - \tau_{Ai}|^2 + |r_b - \tau_{Bi}|^2) \right\}$, and

$$R_3 = \sum_{i=1}^{258} \left[\sigma_i^{-2} \sum_{\langle a,b \rangle \in \check{\xi}_i} \Re \left\{ (r_a - \tau_{Ai})\overline{\psi_i(r_b - \tau_{Bi})} \right\} \right].$$

This can be solved using the cubic formula. The maximizing value for κ solves

$$R_1 \frac{I_1(\kappa)}{I_0(\kappa)} = \sum_{i=1}^{258} \left\{ \sum_{\langle a,b \rangle \in \check{\xi}_i} \Re \left(s_a \overline{\psi_i s_b} \right) \right\},$$

where I_1 is the modified Bessel function of the first kind and first order. The ratio $I_1(\kappa)/I_0(\kappa)$ is always between zero and one, so this equation is simple to solve numerically.

We repeat the process of generating new values of $(\tau_{Ai}, \tau_{Bi}, \sigma_i, \psi_i)$ and updating ω and κ until the latter stabilize. After they stabilize we run 500 more iterations while saving the maximizing values of ω and κ. Our point estimates for ω and κ are the average of these maximizing values, yielding $\hat{\omega} = 0 \cdot 047$ and $\hat{\kappa} = 35$.

6.4 Results

The Monte Carlo Markov chain algorithm was programmed in C# version $4 \cdot 51$. We chose this language due to its multi-thread support for multiple parallel fingerprint comparisons and advanced data visualization capabilities. Our algorithm generates approximately 5000 joint samples of θ and ξ per thread per second on a 3GHz Intel Xeon processor.

For both simulated and real data we set the initial value of ξ to the empty match. Within 2000 iterations the variable traces appeared to be stationary. We used 5000 samples for burn-in and generated another 5000 samples to estimate the likelihood ratio. In our experience this sample size is sufficient to reduce the Monte Carlo error in the log likelihood ratio estimate to less than $0 \cdot 2$.

We computed the log likelihood ratio for all possible 258×258 fingerprint/fingermark pairs in our simulated dataset. The log likelihood ratios for the 258 pairs that originate from the same finger are shown in the blue histogram with solid lines at the top of Fig. 4. The remaining 258×257 log likelihood ratios (the false matches) are shown in the red histogram with dashed lines. The inset receiver-operating characteristic curve describes our discrimination of true matches from false matches based on any chosen cutoff point for the log likelihood ratio.

The other three histograms subdivide the pairs into the 88×258 pairs where the fingermark is good, the 85×258 pairs where the fingermark is bad, and the 85×258 pairs where the fingermark is ugly. We achieve perfect separation for the good and bad fingermarks, and worse separation for ugly fingermarks, reflecting that these have fewer minutiae and thus are less informative.

The same type of histograms for our real dataset are displayed in Fig. 5. The discrimination here is not as good as for the simulated data; this could be another indication that our model does not completely describe the variability in real fingermarks.

Similarly, the likelihood ratios appear to be slightly more extreme than they should be for the false pairings; for example, the maximal value of $\log_{10} \Lambda$ is equal to 12·2, which appears too high to occur by chance, and higher than the similar value for simulated data, which is 9·6.

7   Discussion

We have described a marked Poisson point process model for paired minutia configurations in fingerprints and fingermarks, and the corresponding matching between these minutia configurations. We can efficiently sample from the distribution of the unknown matching and parameters in this model using a Markov chain Monte Carlo method. The resulting sample can be used to compute likelihood ratios for comparing the hypothesis that the two configurations originate from the same finger against the hypothesis that they originate from different fingers.

The method provides excellent discrimination on simulated data. Using the method on a specific NIST-FBI database indicate that the model yields good discrimination between these two hypotheses as long as the fingermark is of reasonable quality. However, some inaccuracies are apparent for the simplistic model discussed in the present paper, in particular concerning the model for selection of observed fingermarks which appears to be non-constant; also, the minutiae do tend to occur along fingerprint ridges which means that orientations of nearby minutiae are not independent, as assumed.

This may result in the likelihood ratios calculated to be more extreme than what can be justified. The ratios can still be used as a sensible model based method for discrimination between true and false matches, but they would have to be calibrated against a large real dataset along the lines described in Forbes (2014, chapter 9) before they can be interpreted as an accurate measure of the strength of evidence. In any case, we believe the framework can be used to establish a sound and model-based foundation for the analysis of fingerprint evidence.

8   Acknowledgements

Appendix

Overview of the sampling procedures

Chib's method as discussed in Section 5.3, including our choice of θ^* and ξ^*, is detailed in Algorithm 1. The Metropolis-within-Gibbs sampler discussed in Section 5.4 is described in Algorithm 2. To generate our samples, we use Marsaglia and Tsang (2000b) for Gaussian variables, Marsaglia and Tsang (2000a) for Gamma variables, Dagpunar (1978) for truncated Gamma variables, Cheng (1978) for Beta variables, and Best and Fisher (1979) for von Mises variables. For those variables whose full conditionals are not one of the above type, we give a detailed sampling algorithm below. All samplers use Marsaglia (2003) as source of pseudo-random integers.
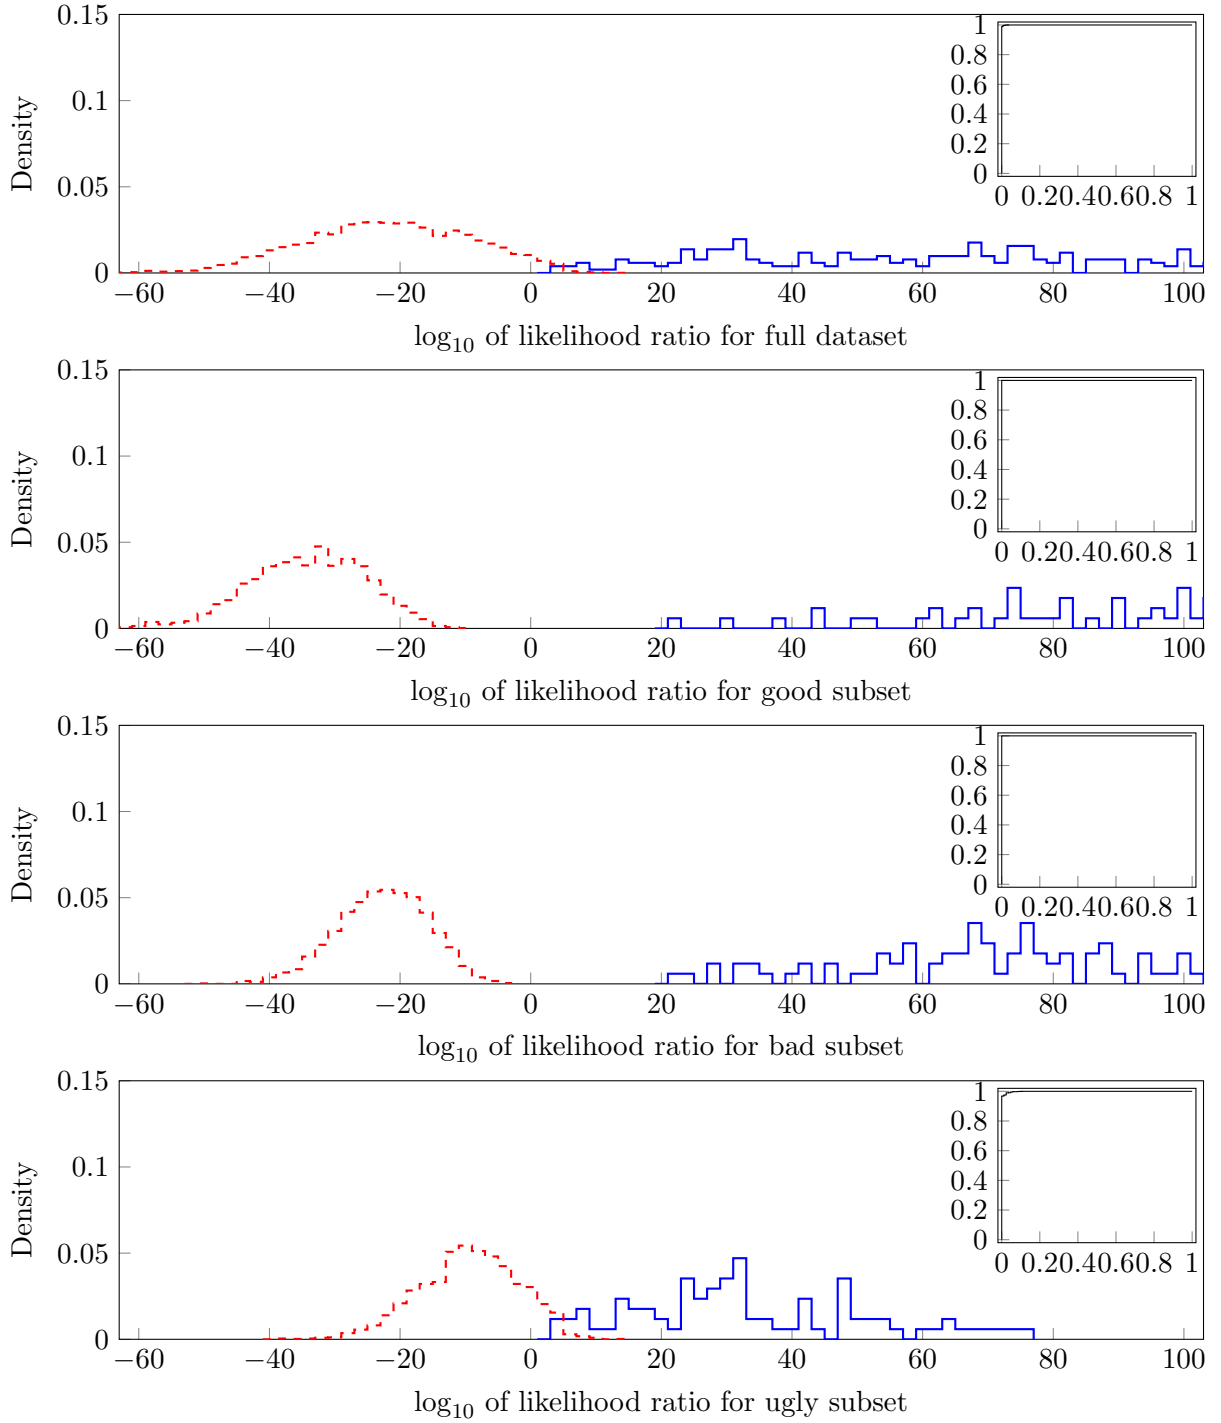
Figure 4: Histogram of the log-likelihood ratios for simulated data. Log-likelihood ratios corresponding to false matches are dashed and red, and true matches are solid and blue. Inset is a receiver-operating characteristic curve with the rate of false positives (i.e., the type 1 error rate) on the x-axis and the rate of true positives (i.e., one minus the type 2 error rate) on the y-axis.
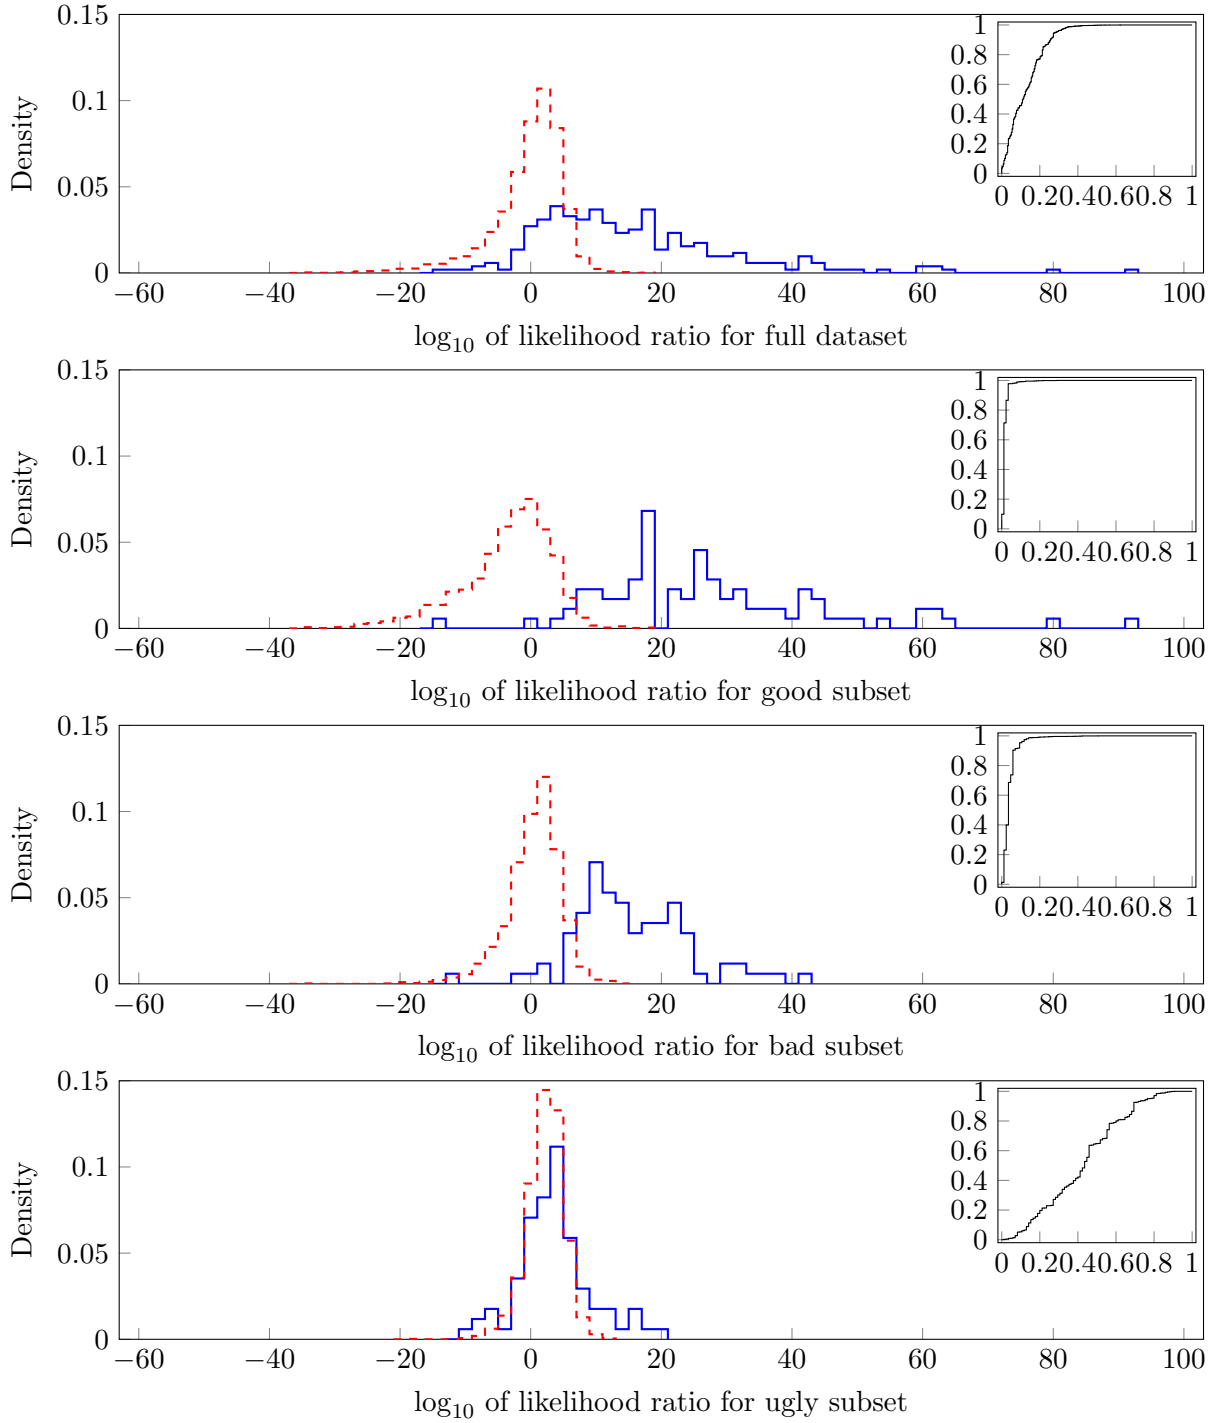
Figure 5: Histogram of the log-likelihood ratios for NIST-FBI data. Log-likelihood ratios corresponding to false matches are dashed and red, and true matches are solid and blue. Inset is a receiver-operating characteristic curve with the rate of false positives (i.e., the type 1 error rate) on the x-axis and the rate of true positives (i.e., one minus the type 2 error rate) on the y-axis.

Algorithm 1: Chib's method for approximating $\text{pr}(A, B \mid H_p)$.

```
/* All samples Sᵢ are generated by holding the starred variables constant
   while sampling the non-starred variables as described in Algorithm 2.
   */
```
$r \leftarrow 1$;

Generate sample S_1 of size N from $\delta_A, \delta_B, \tau_A, \tau_B, \sigma, \psi, \xi \mid A, B, H_p$;

$\delta_A^* \leftarrow \hat{E}_{S_1}(\delta_A)$;

Generate sample S_2 of size N from $\delta_B, \tau_A, \tau_B, \sigma, \psi, \xi \mid \delta_A^*, A, B, H_p$;

$\delta_B^* \leftarrow \hat{E}_{S_2}(\delta_B)$;

$r \leftarrow r \times \hat{E}_{S_2}\{\text{pr}(\delta_A^* \mid \delta_B, \tau_A, \tau_B, \sigma, \psi, \xi, A, B, H_p)\}$;

Generate sample S_3 of size N from $\tau_A, \tau_B, \sigma, \psi, \xi \mid \delta_A^*, \delta_B^*, A, B, H_p$;

$(\tau_A^*, \tau_B^*) \leftarrow \hat{E}_{S_3}(\tau_A, \tau_B)$;

$r \leftarrow r \times \hat{E}_{S_3}\{\text{pr}(\delta_B^* \mid \delta_A^*, \tau_A, \tau_B, \sigma, \psi, \xi, A, B, H_p)\}$;

Generate sample S_4 of size N from $\sigma, \psi, \xi \mid \delta_A^*, \delta_B^*, \tau_A^*, \tau_B^*, A, B, H_p$;

$\sigma^* \leftarrow \hat{E}_{S_4}(\sigma)$;

$r \leftarrow r \times \hat{E}_{S_4}\{\text{pr}(\tau_A^*, \tau_B^* \mid \delta_A^*, \delta_B^*, \sigma, \psi, \xi, A, B, H_p)\}$;

Generate sample S_5 of size N from $\psi, \xi \mid \delta_A^*, \delta_B^*, \tau_A^*, \tau_B^*, \sigma^*, A, B, H_p$;

$\psi^* \leftarrow \hat{E}_{S_5}(\psi)$;

$r \leftarrow r \times \hat{E}_{S_5}\{\text{pr}(\sigma^* \mid \delta_A^*, \delta_B^*, \tau_A^*, \tau_B^*, \psi, \xi, A, B, H_p)\}$;

Generate sample S_6 of size N from $\xi \mid \delta_A^*, \delta_B^*, \tau_A^*, \tau_B^*, \sigma^*, \psi^*, A, B, H_p$;

$r \leftarrow r \times \hat{E}_{S_6}\{\text{pr}(\psi^* \mid \delta_A^*, \delta_B^*, \tau_A^*, \tau_B^*, \sigma^*, \xi, A, B, H_p)\}$;

$\xi^* \leftarrow \text{argmax}_\xi\{\text{pr}(\theta^*, \xi, A, B, H_p)\}$;   `/* An efficient method for finding the`
`maximizer over ξ is given in Forbes (2014, chapter 3) */`

for $\beta \in B$ **do**
  Generate sample S_β of size N_ξ from $\xi \mid \xi_{\leq\beta} = \xi_{\leq\beta}^*, \theta^*, A, B, H_p$;
  $r \leftarrow r \times \hat{E}_{S_\beta}\{\text{pr}(\xi \mid \xi_{<\beta}, \xi_{>\beta}, \theta^*, A, B, H_p)\}$;   `/* See (20) */`

return r, an estimate of $\text{pr}(A, B \mid H_p)$;

Algorithm 2: Metropolis-within-Gibbs sampler for the posterior of θ and ξ.

Input: θ^0, ξ^0 set to some initial value

for $n = 1, \dots, N$ **do**
  $\delta_A^n \leftarrow \text{Sample}\left(\delta_A \mid A, B, \delta_B^{n-1}, \tau_A^{n-1}, \tau_B^{n-1}, \sigma^{n-1}, \psi^{n-1}, \xi^{n-1}\right)$;
  $\delta_B^n \leftarrow \text{Sample}\left(\delta_B \mid A, B, \delta_A^n, \tau_A^{n-1}, \tau_B^{n-1}, \sigma^{n-1}, \psi^{n-1}, \xi^{n-1}\right)$;
  $(\tau_A^n, \tau_B^n) \leftarrow \text{Sample}\left(\tau_A, \tau_B \mid A, B, \delta_A^n, \delta_B^n, \sigma^{n-1}, \psi^{n-1}, \xi^{n-1}\right)$;
  $\sigma^n \leftarrow \text{Sample}\left(\sigma \mid A, B, \delta_A^n, \delta_B^n, \tau_A^n, \tau_B^n, \psi^{n-1}, \xi^{n-1}\right)$;
  $\psi^n \leftarrow \text{Sample}\left(\psi \mid A, B, \delta_A^n, \delta_B^n, \tau_A^n, \tau_B^n, \sigma^n, \xi^{n-1}\right)$;
  $\xi^n \leftarrow \xi^{n-1}$;
  for $j = 1, \dots, n_A$ **do**
    `/* Sample repeatedly to reduce autocorrelation` `*/`
    $\xi^n \leftarrow \text{Sample}\left(\xi \mid A, B, \delta_A^n, \delta_B^n, \tau_A^n, \tau_B^n, \sigma^n, \psi^n, \xi^n\right)$;

Sampling δ_A, δ_B

Define the distribution $D(\alpha, \beta, \lambda)$ to have density

$$f_D(\delta) \propto \delta^{\alpha-1}(1-\delta)^{\beta-1}e^{-\lambda\delta}$$

for $\delta \in (0, 1)$, where $\alpha > 0, \beta > 0$, and $\lambda \in \mathbb{R}$. If $\lambda = 0$ this is a Beta distribution, and if $\beta = 1$ it is a Gamma distribution right-truncated at one. The full conditionals for δ_A and δ_B are

$$\delta_A \sim D(\alpha_\delta + n_A, \beta_\delta + n_B - n_\xi, \rho_0 - \rho_0\delta_B), \quad \delta_B \sim D(n_B, n_A - n_\xi, \rho_0 - \rho_0\delta_A).$$

We describe an algorithm to sample from D in Algorithm 3.

Algorithm 3: Rejection sampler for $D(\alpha, \beta, \lambda)$.

---

if $\lambda = 0$ **then**
   | $\delta \leftarrow$ Sample $(\text{Gamma}(\alpha, \beta) \,|\, \delta \leq 1)$;
else
   | $\delta_0 \leftarrow \text{argmax}\{x^{\alpha-1}(1-x)^{\beta-1}e^{-\lambda x} : x \in [0, 1]\}$;
   | **if** $\delta_0 > 0.5$ **and** $\lambda < (\alpha - 1)/(1 - \delta_0)$ **then**
   | | $\delta \leftarrow 1 -$ Sample $(D(\beta, \alpha, -\lambda))$;
   | **else**
   | | **repeat**
   | | | $\delta \leftarrow$ Sample $(\text{Gamma}\{\alpha, \lambda + (\beta - 1)/(1 - \delta_0)\} \,|\, \delta \leq 1)$;
   | | | $U \leftarrow$ Sample $(\text{Uniform}(0, 1))$;
   | | **until** $U^{1/(\beta-1)} < (1 - \delta)\delta_1 \exp\{(\delta - \delta_0)/(1 - \delta_0)\}$;
return δ;

Briefly, let δ_0 be the mode of D, which can be easily computed by applying the quadratic formula to $\mathrm{d}\log f_D(\delta)/\mathrm{d}\delta = 0$. To sample from D, first notice that $\log(1-\delta) \approx 1 - \delta_0 - \delta/(1 - \delta_0)$ for small δ_0 ($\delta_0 \leq 0.5$ in Algorithm 3), and hence $(1 - \delta)^{\beta-1} \approx C\exp\{-(\beta - 1)\delta/(1 - \delta_0)\}$ where C is a constant independent of δ. Plugging this approximation into $f_D(\delta)$ yields a Gamma$\{\alpha, (\beta - 1)/(1 - \delta_0) + \lambda\}$ density right-truncated at one. Thus when $\delta_0 \leq 0.5$, we can use rejection sampling with proposals drawn from this distribution. Similarly, when the mode $\delta_0 \geq 0.5$, let $\tilde{\delta} = 1 - \delta$ so that $\tilde{\delta} \sim D(\beta - 1, \alpha - 1, -\lambda)$ with mode $\tilde{\delta}_0 = 1 - \delta_0$. Using the same approximation as before, we can use rejection sampling on $\tilde{\delta}$ with a Gamma$(\beta, (\alpha - 1)/(1 - \tilde{\delta}_0) - \lambda)$ proposal, right-truncated at one, provided $(\alpha - 1)/(1 - \tilde{\delta}_0) - \lambda > 0$. In practice we achieve acceptance rates greater than 0.9.

Sampling τ_A, τ_B

The full conditional is bivariate complex normal with mean $(K_d + n_\xi\Sigma_{AB}^{-1})^{-1}(K_dr_d + n_\xi\Sigma_{AB}^{-1}r_p)$ and inverse variance $K_d + n_\xi\Sigma_{AB}^{-1}$, where

$$r_d = \begin{pmatrix} n_A^{-1}\sum_{a\in A}r_a \\ n_B^{-1}\sum_{b\in B}r_b \end{pmatrix}, \quad r_p = \frac{1}{n_\xi}\sum_{\langle a,b\rangle\in\xi}\begin{pmatrix} r_a \\ r_b \end{pmatrix}, \quad K_d = \sigma^{-2}\begin{pmatrix} n_A & 0 \\ 0 & n_B \end{pmatrix}.$$

Sampling σ

We make the change of variables $u = \sigma^{-2}$. The improper prior $\mathrm{pr}(\sigma) = \sigma^{-5}$ becomes $\mathrm{pr}(u) \propto u$. The full conditional of u is a Gamma distribution with shape parameter $n_A + n_B + 2$ and inverse

scale parameter

$$\sum_{a \in A \setminus \Pi_A(\xi)} |r_a - \tau_A|^2 + \sum_{b \in B \setminus \Pi_B(\xi)} |r_b - \tau_B|^2$$

$$+ \frac{(\omega^2 + 1)^2}{(\omega^2 + 1)^2 - 1} \sum_{\langle a,b \rangle \in \xi} \left[|r_a - \tau_A|^2 + |r_b - \tau_B|^2 - \frac{2}{\omega^2 + 1} \Re \left\{ (r_a - \tau_A) \overline{(r_b - \tau_B)\psi} \right\} \right].$$

Sampling ψ

The full conditional of ψ is a von Mises distribution with location parameter $\nu_0/|\nu_0|$ and concentration parameter $|\nu_0|$, where

$$\nu_0 = \sum_{\langle a,b \rangle \in \xi} \kappa s_a \overline{s_b} + 2 \frac{\omega^2 + 1}{(\omega^2 + 1)^2 - 1} \frac{r_a - \tau_A}{\sigma} \frac{\overline{r_b - \tau_B}}{\sigma}.$$

Sampling ξ

Finally we sample the matching ξ. A possible Metropolis–Hastings sampler for ξ is described in Green and Mardia (2006). They propose creating or breaking a single, random matched pair at each iteration. In contrast, our Algorithm 4 considers all matches for a given minutia simultaneously and computes the probability of each match.

We need one more piece of notation. In analogy with $\Pi_{A,m}$ in (18), for each $m \in \mathbb{M}$, define $\Pi_{B,m} : \Xi(A, B) \to \mathbb{M}$ by $\Pi_{B,m}(\xi) = b$ if $\langle m, b \rangle \in \xi$ for some $b \in B$, and ϕ otherwise.

We will sample ξ with the help of an auxiliary random variable β that takes values uniformly on B. Consider the following transition kernel for moving in the augmented state space from (ξ, β) to (ξ', β'):

$$q(\xi', \beta' \,|\, \xi, \beta) \propto \mathrm{pr}(\xi' \,|\, \theta, A, B, H_p) I \left[\xi' \setminus \{ \langle \Pi_{A,\beta}(\xi'), \beta \rangle \} = \xi \setminus \{ \langle \Pi_{A,\beta}(\xi), \beta \rangle \} \right],$$

where $\mathrm{pr}(\xi' \,|\, \theta, A, B, H_p)$ is proportional to (16). This transition kernel allows transitions to any $\xi' \in \Xi(A, B)$ which differs from ξ only in its match for β. The states ξ' which are accessible from the state ξ are illustrated in Fig. 6. We can move from any state ξ to any other state ξ'

(a) Add   (b) Swap on A (c) Remove   (d) Swap on B (e) Add/remove   (f) No change
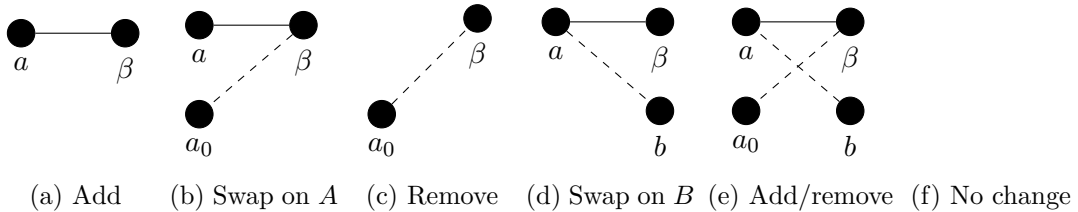
Figure 6: Illustration of which states for ξ' are accessible from a given state ξ. Dashed edges are removed matches, solid edges are added matches. Edges that are common to both ξ' and ξ are not shown. We write a for $\Pi_{A,\beta}(\xi')$, a_0 for $\Pi_{A,\beta}(\xi)$ and b for $\Pi_{B,a}(\xi)$, assuming that none of these are equal to ϕ. Hence $\langle a_0, \beta \rangle \in \xi$, $\langle a, b \rangle \in \xi$, and $\langle a, \beta \rangle \in \xi'$.

in at most n_B steps, so the Markov chain with this transition kernel is irreducible. Clearly it is also aperiodic and therefore ergodic. Its stationary distribution is $\mathrm{pr}(\xi \,|\, \theta, A, B, H_p)$ as desired.

The densities of the allowed states have many terms in common. By ignoring these common terms, we obtain from (20)

$$q(\xi', \beta' \,|\, \xi, \beta) \propto \exp[w(a, \beta \,|\, \theta) - w\{a, \Pi_{B,a}(\xi) \,|\, \theta\}],$$

where $a = \Pi_{A,\beta}(\xi')$ and w is given by (21). Thus the proposal function can be computed very quickly, and it can be normalized over ξ by simply summing over the permitted moves. There are $n_A + 1$ such moves, one for each possible value of $a \in A \cup \phi$. The full algorithm is described in Algorithm 4.

Algorithm 4: Sampler for ξ using the auxiliary variable β.

Input: Previous value ξ
$\beta \leftarrow \text{Sample}(\text{Uniform over } B)$;
$\xi' \leftarrow \xi \setminus \{\langle \Pi_{A,\beta}(\xi), \beta \rangle\}$ # remove the old match of β;
$\alpha \leftarrow \text{Sample}(\text{pr}(a) \propto \exp[w(a, \beta \,|\, \theta) - w\{a, \Pi_{B,a}(\xi) \,|\, \theta\}])$ for $a \in A \cup \phi$;
if $\alpha = \phi$ **then**
   | **return** ξ';
else
   | **return** $\xi' \cup \{\langle \alpha, \beta \rangle\}$;

References

Aitken, C. C. G. and F. Taroni (2004). *Statistics and the Evaluation of Evidence for Forensic Scientists* (2nd ed.). Statistics in Practice. Chichester, UK: Wiley.

Balding, D. J. (2005). *Weight-of-evidence for Forensic DNA Profiles*. Statistics in Practice. Chichester, UK: Wiley.

Best, D. J. and N. I. Fisher (1979). Efficient simulation of the von Mises distribution. *Journal of the Royal Statistical Society Series C 28*(2), 152–157.

Celeux, G. and J. Diebolt (1985). The SEM algorithm: a probabilistic teacher algorithm derived from the EM algorithm for the mixture problem. *Computational Statistics Quarterly 2*, 73–82.

Cheng, R. C. H. (1978). Generating beta variates with nonintegral shape parameters. *Communications of the ACM 21*(4), 317–322.

Chib, S. (1995). Marginal likelihood from the Gibbs output. *Journal of the American Statistical Association 90*(432), 1313–1321.

Chib, S. and I. Jeliazkov (2001). Marginal likelihood from the Metropolis-Hastings output. *Journal of the American Statistical Association 96*(453), 270–281.

Dagpunar, J. (1978). Sampling of variates from a truncated gamma distribution. *Journal of Statistical Computation and Simulation 8*, 59–64.

Dawid, A. P. and S. L. Lauritzen (2000). Compatible prior distributions. In *Bayesian Methods with Applications to Science, Policy and Official Statistics*, pp. 109–118. International Society for Bayesian Analysis.

Forbes, P. G. M. (2014). *Quantifying the Strength of Evidence in Forensic Fingerprints*. Ph. D. thesis, University of Oxford.

Garris, M. and R. McCabe (2000). NIST special database 27: Fingerprint minutiae from latent and matching tenprint images. Technical report, NIST, Gaithersburg, MD, USA.

Goodman, N. R. (1963). Statistical analysis based on a certain multivariate complex Gaussian distribution (an introduction). *Annals of Mathematical Statistics 34*(1), pp. 152–177.

Green, P. J. and K. V. Mardia (2006). Bayesian alignment using hierarchical models, with applications in protein bioinformatics. *Biometrika 93*(2), 235–254.

Hill, B. J., W. S. Kendall, and E. Thönnes (2012). Fibre-generated point processes and fields of orientations. *The Annals of Applied Statistics 6*(3), 994–1020.

Kent, J. (1977). The infinite divisibility of the von Mises–Fisher distribution for all values of the parameter in all dimensions. *Proceedings of the London Mathematical Society 35*(3), 359–384.

Lauritzen, S., R. G. Cowell, and T. Graversen (2012). Discussion on the paper by Neumann et al. (2012). *Journal of the Royal Statistical Society Series A 175*(2), 405–406.

Lindley, D. V. (1977). A problem in forensic science. *Biometrika 64*(2), 207–213.

Maltoni, D. (2009). *Handbook of Fingerprint Recognition* (2nd ed.). New York: Springer-Verlag.

Mardia, K. V. and P. E. Jupp (1999). *Directional Statistics* (2nd ed.). Chichester, UK: Wiley.

Marsaglia, G. (2003). Xorshift RNGs. *Journal of Statistical Software 8*(14), 1–6.

Marsaglia, G. and W. W. Tsang (2000a). A simple method for generating gamma variables. *ACM Transactions on Mathematical Software 26*(3), 363–372.

Marsaglia, G. and W. W. Tsang (2000b). The ziggurat method for generating random variables. *Journal of Statistical Software 5*(8), 1–7.

Mikalyan, A. and J. Bigun (2012). Ground truth and evaluation for latent fingerprint matching. In *CVPR Workshop on Biometrics*.

Møller, J. and R. P. Waagepetersen (2004). *Statistical Inference and Simulation for Spatial Point Processes*. Boca Raton: Chapman and Hall/CRC.

Neumann, C., I. W. Evett, and J. E. Skerrett (2012). Quantifying the weight of evidence from a forensic fingerprint comparison: a new paradigm (with discussion). *Journal of the Royal Statistical Society Series A 175*(2), 371–415.

Olver, F. W. J., D. W. Lozier, R. F. Boisvert, and C. W. Clark (Eds.) (2010). *NIST Handbook of Mathematical Functions*. New York, NY: Cambridge University Press.

Yager, N. and A. Amin (2004). Fingerprint verification based on minutiae features: a review. *Pattern Analysis and Applications 7*(1), 94–113.