

A BAYESIAN HIERARCHICAL MODEL FOR EVALUATING FORENSIC FOOTWEAR EVIDENCE

BY NEIL A. SPENCER¹ AND JARED S. MURRAY²

¹*Department of Statistics and Data Science, and Machine Learning Department, Carnegie Mellon University, nspencer@andrew.cmu.edu*

²*Department of Information, Risk, and Operations Management and Department of Statistics and Data Science, University of Texas at Austin, jared.murray@mcombs.utexas.edu*

When a latent shoeprint is discovered at a crime scene, forensic analysts inspect it for distinctive patterns of wear such as scratches and holes (known as accidentals) on the source shoe's sole. If its accidentals correspond to those of a suspect's shoe, the print can be used as forensic evidence to place the suspect at the crime scene. The strength of this evidence depends on the random match probability—the chance that a shoe chosen at random would match the crime scene print's accidentals. Evaluating random match probabilities requires an accurate model for the spatial distribution of accidentals on shoe soles. A recent report by the President's Council of Advisors in Science and Technology criticized existing models in the literature, calling for new empirically validated techniques. We respond to this request with a new spatial point process model (code and synthetic data is available as Supplementary Material) for accidental locations, developed within a hierarchical Bayesian framework. We treat the tread pattern of each shoe as a covariate, allowing us to pool information across large heterogeneous databases of shoes. Existing models ignore this information; our results show that including it leads to significantly better model fit. We demonstrate this by fitting our model to one such database.

1. Introduction. Forensic footwear analysis encompasses a suite of techniques used to analyze latent shoeprints as part of forensic investigations. A principal goal of these investigations is to link a suspect's shoe to a crime scene print, providing evidence to place the suspect at the scene of the crime. Figure 1(a) provides an example of a latent crime scene shoeprint.

As described by Bodziak (2017), the procedure for determining the source of a latent print typically consists of two stages. First, the examiner inspects the tread of the latent print to identify class characteristics (brand, model and size) of the source shoe. This identification can be carried out manually, or automated using tread matching algorithms (e.g., Kong, Ramanan and Fowlkes (2017), Richetelli et al. (2017a), Srihari and Tang (2014)).

Manufacturers routinely produce thousands of shoes of the same make and model, meaning that class characteristics alone are often insufficient for determining a print's source. For this reason examiners regularly turn to a second stage of analysis, the inspection of *accidentals*. Accidentals, also known as randomly acquired characteristics, are the postmanufacturing cuts, scrapes, holes and debris that accumulate on a shoe sole. Examiners are trained to identify accidentals on a shoe by inspecting both the shoe's sole and test impressions—high quality prints created using the shoe in a controlled laboratory setting. Figure 1(b), Figure 1(c) and Figure 1(d) depict a shoe sole, test impression, and accidentals locations, respectively. These images all correspond to the same shoe obtained from the JESA database (Yekutieli et al. (2012)) (we describe the JESA database in Section 2.2).

Received June 2019; revised November 2019.

Key words and phrases. Random match probability, forensic statistics, forensic footwear analysis, random measures, hierarchical Bayes, point processes, spatial statistics.




FIG. 1. (a)–(d) represent objects pertaining to the same shoe from the JESA database: (a) is a photograph of a latent crime scene print, (b) a photo of the shoe’s sole, (c) is a raw image of a test impression and (d) is the contact surface obtained from standardizing the test impression. The superimposed blue points in (d) correspond to accidental locations.

In theory, if both the class characteristics and the accidentals of a suspect’s shoe coincide exactly with those detected from the crime scene print, then the suspect’s shoe is almost certainly the source of the print. In practice, the comparison is less clear-cut. Latent crime scene prints are typically of low quality, making it difficult to pick out all of the individual accidentals. Furthermore, accidental locations are known to vary slightly from test print to test print due to variability in the impression-taking process (Shor et al. (2018)), so there is some uncertainty on their exact locations on the source shoe. As a result, accidental comparisons typically involve comparing a subset of approximate accidental locations on the test impression to those detected on the crime scene print. This uncertainty leaves the possibility of a false positive due to chance, especially for partial prints and tread patterns on which accidentals are very likely to occur in certain regions.

To account for the possibility of a false positive, shoeprint analysts are encouraged to provide a measure of the uncertainty of the match when testifying in court (Edwards and Gotsonis (2009)). One popular summary for communicating this uncertainty is the random match probability (RMP) (Thompson and Newman (2015)). The RMP is the probability that a randomly sampled shoe would produce a print matching the observed features at the crime scene. For instance, if 15 out of 10,000 relevant shoes were consistent with the crime scene print, the RMP would be 0.0015.

The standard approach for evaluating RMPs decomposes into three terms: the evidence given by the class characteristics, the evidence based on general wear and accidental-based evidence (Evet, Lambert and Buckleton (1998), Skerrett, Neumann and Mateos-Garcia (2011)). In this work we focus on accidental-based evidence, inspired by the recent report on forensic science by The President’s Council for Advisors on Science (PCAST (2016)) that criticized existing work in the area.

We address the concerns of PCAST (2016) by developing and estimating the parameters of a model for the distribution of accidental configurations on a shoe. Specifically, we model the spatial distribution of accidentals on a shoe sole as a point process, treating the sole’s tread pattern as a covariate. We fit and evaluate our model using the JESA database (Yekutieli et al. (2012)), a ground truth dataset of 386 accidental-annotated shoeprints compiled by the Israeli Police Department’s Division of Forensic Science. The JESA database is one of the largest existing databases of its kind (Speir et al. (2016)), consisting of shoes with a variety of tread patterns.

We define our model within a hierarchical Bayesian framework, pooling information across JESA to infer general trends spanning a broad variety of shoes. Our model is a finite resolution version of the normalized compound random measure framework of [Griffin and Leisen \(2017\)](#), modified to incorporate spatial covariates and dependency of the intensity across space. We develop the computational tools to fit our model, evaluate it and demonstrate that it outperforms existing approaches by a wide margin.

The remainder of this paper is organized as follows. In Section 2 we review the literature related to random match probabilities, formalize the link between evaluating random match probabilities and modeling spatial distributions of accidentals, describe the JESA database of annotated shoeprints collected by [Yekutieli et al. \(2012\)](#) and review the relevant literature pertaining to vectors of dependent probability measures. In Section 3 we provide the details of our hierarchical Bayesian model for spatial configurations of accidentals. In Section 4 we propose a Markov chain Monte Carlo algorithm for inferring the parameters of the model and an importance sampling algorithm for evaluating marginal likelihoods. In Section 5 we showcase the results of fitting our model to the JESA dataset and compare its performance to other candidate models. Section 6 contains some concluding remarks.

2. Preliminaries.

2.1. *Random match probabilities.* A theory to evaluate RMPs for footwear evidence was laid out in [Evelt, Lambert and Buckleton \(1998\)](#) in the context of evaluating likelihood ratios. The framework is equally applicable to evaluating raw RMPs. Let y denote a crime scene print and \mathbb{A} denote the relevant population of plausible sources of the crime scene print. For instance, \mathbb{A} could be all shoes belonging to residents of a particular city or town. As per [Evelt, Lambert and Buckleton \(1998\)](#), the random match probability for footwear evidence is given by

$$(2.1) \quad \text{RMP} = p(y \equiv s \mid s \sim \mathbb{A}),$$

where $y \equiv s$ indicates that shoe s exhibits features consistent with those of the print y , and $s \sim \mathbb{A}$ is shorthand for s being chosen uniformly at random from all shoes in the set \mathbb{A} . Further discussion of random match probabilities with examples from forensic science is available in [Srihari and Su \(2011\)](#).

Following the classical two-step process of forensic footwear analysis, [Evelt, Lambert and Buckleton \(1998\)](#) suggested that the RMP be calculated using the factorization $\text{RMP} = \text{rmp}_M \text{rmp}_U$. Here, rmp_M denotes the probability that a randomly chosen shoe in \mathbb{A} has class characteristics matching the latent crime scene print, and rmp_U denotes the probability that the shoe is also consistent with the wear patterns and accidentals, given that it matches on the class characteristics. [Skerrett, Neumann and Mateos-Garcia \(2011\)](#) refined this representation by further decomposing rmp_U into rmp_W and rmp_V , corresponding to separate conditional probabilities of matching on general wear and accidentals, respectively.

Let y_M, y_W and y_V denote the class characteristics, general wear and accidentals observed on the latent print y with s_M, s_W, s_V denoting the same features as observed on a shoe $s \in \mathbb{A}$. The factorization proposed by [Skerrett, Neumann and Mateos-Garcia \(2011\)](#) can be formally expressed as:

$$(2.2) \quad \text{RMP} = \text{rmp}_M \cdot \text{rmp}_W \cdot \text{rmp}_V,$$

$$(2.3) \quad \text{rmp}_M = p(y_M \equiv s_M \mid s \sim \mathbb{A}),$$

$$(2.4) \quad \text{rmp}_W = p(y_W \equiv s_W \mid s \sim \{s' \in \mathbb{A} : y_M \equiv s'_M\}),$$

$$(2.5) \quad \text{rmp}_V = p(y_V \equiv s_V \mid s \sim \{s' \in \mathbb{A} : y_M \equiv s'_M, y_W \equiv s'_W\}),$$

where $y_M \equiv s_M$ denotes the class characteristics of s being consistent with those of y , and $y_W \equiv s_W$ and $y_V \equiv s_V$ defined similarly. Implicit in this decomposition is the assumption that $y \equiv s$ is characterized by $y_M \equiv s_M$, $y_W \equiv s_W$ and $y_V \equiv s_V$, a reasonable choice given that these features form the basis of forensic footwear analysis (Bodziak (2017)). Strategies for evaluating rmp_M and rmp_W based on relevant databases (e.g., Evett, Lambert and Buckleton (1998), Champod, Evett and Jackson (2004) for rmp_M and Fruchtenicht, Herzig and Blackledge (2002), Facey, Hannah and Rosen (1992), Bodziak et al. (2012) for rmp_W) were discussed in Skerrett, Neumann and Mateos-Garcia (2011). However, evaluating the accidental-based component rmp_V was left as a subject for future work. In this work we focus on the remaining accidental-based component. We begin by making two simplifying assumptions.

First, we follow Petraco et al. (2010) in assuming that the evidence present in a configuration of accidentals on a crime scene print y_V is characterized by the locations (e.g., the blue points shown in Figure 1(d)). We omit secondary characteristics, such as shape or size of the accidental, as they are difficult to reliably glean from latent prints. We use x^s to denote the accidental locations on shoe s and x^y to denote the locations detected on print y . Employing a standardized coordinate system (details provided in Section 2.2), we have $x^s \in ([0, 100] \times [0, 200])^{N_s}$, $x^y \in ([0, 100] \times [0, 200])^{N_y}$, where N_s denotes the number of accidentals on shoe s and N_y denotes the number detectable on print y . We use $x_n^s = (x_{n,1}^s, x_{n,2}^s)$ to denote the n th row of x^s .

Because examiners are adept at recovering y_M and y_W from a shoeprint y , our second assumption is that a shoe's class characteristics and wear are characterized by its *contact surface*. A shoe's contact surface refers to the portion of its sole that typically touches the ground when worn—the part responsible for leaving prints. An example contact surface is provided in Figure 1(d). We provide a more detailed definition of contact surface in Section 2.2.1. Letting \mathcal{C}^s denote the contact surface of shoe s , this assumption can be formalized as $s, s' \in \mathbb{A}$, $\mathcal{C}^s = \mathcal{C}^{s'}$ if and only if $s_M \equiv s'_M$ and $s_W \equiv s'_W$.

After characterizing y_V using accidental locations and y_W, w_M using the contact surface, we can now reexpress the accidental-based random match probability in (2.5) in a form that is more tractable for statistical inference. The relation $y_V \equiv s_V$ reduces to a comparison of the point clouds x^y and x^s (denoted $x^y \equiv x^s$). The set $\{s' \in \mathbb{A} : y_M \equiv s'_M, y_W \equiv s'_W\}$ reduces to the set of relevant shoes with the given contact surface (i.e., $\mathbb{A}_{\mathcal{C}^y} = \{s' \in \mathbb{A} : \mathcal{C}^{s'} = \mathcal{C}^y\}$, where \mathcal{C}^y denotes the contact surface as determined from y). Thus, the accidental-based random match probability given in (2.5) reduces to

$$(2.6) \quad \text{rmp}_V = p(x^y \equiv x^s \mid s \sim \mathbb{A}_{\mathcal{C}^y}).$$

In theory, computing rmp_V using (2.6) is straightforward. One would simply inspect all shoes in \mathbb{A} with contact surface \mathcal{C}^y to determine the ratio that also have accidentals consistent with x^y . Even if \mathbb{A} were not completely accessible, a large random sample would suffice to provide a sufficiently accurate approximation. Figure 2 illustrates this strategy for a small example.

In practice, the computation of rmp_V is complicated by two issues:

1. In many cases, no shoes in $\mathbb{A}_{\mathcal{C}^y}$ (other than the suspect's shoe) are accessible by the examiner. Examiners are left to rely on previous experience and limited data (e.g., a small convenience sample from \mathbb{A} or a related database) to make inferences regarding the conditional distribution of $x^s \mid s \sim \mathbb{A}_{\mathcal{C}^y}$. Historically, these inferences have been based on heuristics that lack empirical support (PCAST (2016)).

2. Determining if $x^y \equiv x^s$ is complicated by three phenomena: (i) a shoe's detected accidental locations are known to vary slightly each time it is printed (Shor et al. (2018)), meaning that the locations in x^y may only approximate those in x^s , (ii) some accidentals do not reliably show up on crime scene prints (Richetelli et al. (2017b)), meaning that the accidentals in




FIG. 2. (a) depicts the accidental locations (blue) and contact surface (orange) for eight synthetic draws from the population \mathbb{A}_{C^y} corresponding to the crime scene print y shown in Figure 1(a). (b) depicts the contact surface C^y (orange) and accidental locations x^y (blue). (c) illustrates the close correspondence between x^y (blue) and x^s (red) given by the accidental locations from the rectangle enclosed shoe in (a).

x^y could be a thinned version of x^s and (iii) test impressions may not be obtained until long after the crime was committed, leaving the opportunity for new accidentals to arise (Wyatt, Duncan and Trimpe (2005)) or existing accidentals to change (Sheets et al. (2013)) in the meantime.

We concentrate on issue 1 in this paper, developing a more principled approach to inferring the distribution $x^s | s \sim \mathbb{A}_{C^y}$ using the JESA database. Issue 2 is beyond the scope of this paper, as determining an appropriate definition of $x^y \equiv x^s$ would require much richer data than is currently available in the literature. However, given a definition of $x^y \equiv x^s$, our model can compute the RMP via Monte Carlo. Figure 2 demonstrates this process with Figure 2(a) depicting the samples drawn from the distribution $x^s | s \sim \mathbb{A}_{C^y}$.

2.2. JESA. The Jerusalem Shoeprint Accidentals Database (JESA) is one of a series of datasets created by the Israel Police Department's Division of Forensic Science. It pertains to 386 men's shoes collected as evidence through casework. A full description of the database is available in Yekutieli et al. (2012). For each shoe there are two data structures relevant to our work—the *standardized shoeprint image* (contact surface) and the *accidentals*.

2.2.1. *Standardized shoeprint image.* Test impressions for each shoe were obtained by applying orange powder to their soles, pressing them onto clear films, then digitally photographing the residual orange impressions on the films. An example impression image is shown as Figure 1(c).

For consistency across shoes, each image was standardized onto a 200 by 100 grid. Standardization involved translating, aligning and scaling the images so the prints were centered, pointed upward and of the same length. The axes for the alignment were designated through point-and-click software by trained examiners. All left shoes were mirrored to appear as right shoes. Alignment of the images facilitates the pooling of information across shoes, even if they differ in size or chirality (i.e., left shoe or right shoe).




FIG. 3. (a), (b), (c) are standardized test impressions from the JESA database; (d) is the mean test impression across the entire JESA database.

After standardization, the images were smoothed and denoised to isolate the *contact surface*—the areas of the shoe sole that typically touch the ground. The smoothing was performed to preserve the shoe’s tread pattern and general wear while filtering out small breaks due to accidentals or imperfections in the impression. These contact surfaces take the form of 200 by 100 binary arrays, with each bit defining contact or noncontact of a region of the shoe. Figure 1(d) illustrates the positive values of contact surface for the shoe in Figure 1(b). The superimposed points are the locations of accidentals.

Additional example contact surfaces are shown in Figures 3(a), 3(b) and 3(c), demonstrating the variety of tread patterns in the JESA database. No two contact surfaces in the JESA database are exactly alike, although those that correspond to the same brand of shoe are similar (differences in wear patterns as well as variation in test impressions, account for the differences).

Figure 3(d) depicts the average contact surface across the entire database. It shows that it is far more common for regions of the shoe corresponding to the heel and toes to be part of the contact surface than regions corresponding to the shoe arch. This discrepancy drives home the importance of conditioning on contact surface when evaluating accidental-based RMPs; shoes with arches that do make contact with the ground (the minority) would likely have different accidental distributions than those that do not. We use $\mathbb{C} = \{0, 1\}^{100 \times 200}$ to denote the space of values that a contact surface can take and $\mathcal{C}^s \in \mathbb{C}$ to denote the contact surface of shoe s .

2.2.2. Accidentals. For each shoe, examiners identified accidentals by inspecting the shoeprint image and the shoe sole itself. The locations of the centroids of the accidentals were recorded using a computerized system. These locations were stored as real numbers in $[0, 100] \times [0, 200]$ corresponding to the standardized space of the contact surface. The region $[0, 1] \times [0, 1]$ corresponds to the bottom left-hand corner of the standardized grid, and $[99, 100] \times [199, 200]$ corresponds to the top right. Figure 1(d) gives an example of the locations accidentals as points on the shoeprint image.

The number of accidentals and their locations varies from shoe to shoe. Figure 4(a) provides a histogram of the number of accidentals on each shoe. The distribution is heavily skewed to the right—the median number of accidentals is 20, whereas the mean is 33 and the maximum is 268.

Figure 4(b) aggregates the coordinates of all accidentals recorded in the JESA database. Its similarity to that of Figure 3(d) is consistent with the intuition that accidentals should appear




FIG. 4. (a) is a histogram summarizing the number of accidentals on each shoe in JESA. (b) illustrates the locations of these accidentals with points.

more frequently in areas of the shoe which are part of the contact surface. However, not all accidental locations fall on the sole with contact surface. Of the accidentals in JESA, about 12 percent of them occur in grid points without contact. Therefore, a robust model should be able to assign probability to situations in which accidentals do not occur directly on the contact surface. Examples of shoes in JESA for which accidentals occur away from the contact surface are available as Figure 4 of the Supplementary Material (Spencer and Murray (2020)). Following Damary et al. (2018), we exclude rift-type accidentals from our analysis because they occur only on specific type of shoe tread, making their spatial distribution markedly different than the more frequently occurring types of accidentals (e.g., hole or scratch).

2.3. Existing models for the distribution of accidentals. Going forward, we use the shorthand $x^s | \mathcal{C}^s$ to refer the distribution of accidental locations x^s on a shoe s with contact surface \mathcal{C}^s , with $\mathcal{C}^s = \mathcal{C}^y$ referring to the distribution required to compute the RMP (2.6). For easy comparison of existing models in the literature with the approach we develop, we use a unified notation.

We begin by treating each $x^s | \mathcal{C}^s$ as a draw from a two-dimensional spatial point process (Daley and Vere-Jones (2008)) over the standardized space $[0, 100] \times [0, 200]$. For our model we make three additional assumptions regarding the structure of these point processes: (1) the individual accidentals $(x_n^s)_{n=1, \dots, N_s}$ are exchangeable, (2) the marginal distribution of each x_n^s is independent of the total number of accidentals N_s and (3) the distribution of x^s depends on s only through the contact surface \mathcal{C}^s . The first two assumptions are common in the literature, whereas the third is unique to our model because we are first to incorporate the contact surface.

Following assumptions (1) and (2), $(x_n^s)_{n=1, \dots, N_s}$ can be treated as independent draws from a random probability measure Λ_s on $[0, 100] \times [0, 200]$. The literature has mostly focused on universal models for Λ_s , assuming a fixed Λ that is common to all shoes $s \in \mathbb{A}$. Stone (2006) proposed a uniform model for Λ , that is, $\Lambda \propto 1$. This assumption has been criticized for its lack of empirical support, as noted by PCAST (2016). Yekutieli et al. (2012) instead inferred Λ using a kernel density estimator on the accidentals in JESA (Section 2.2). Speir et al. (2016) applied a similar histogram estimator to a different database, yielding comparable results.

Because estimating a single Λ does not allow for conditioning on class characteristics or wear, these approaches implicitly assume that a shoe's accidental locations are independent of its contact surface. Evidence against this assumption was provided by Damary et al. (2018); their analysis of multiple replicates of three different tread patterns appearing in the

JESA database revealed that different tread patterns tend to yield different accidental distributions. Therefore, having distinct Λ_s that depend on C^s seems more appropriate, serving as the motivation for our assumption (3) above.

We encode assumption (3) in our model by explicitly treating each Λ_s as a draw from a distribution G_{C^s} . As the notation suggests, $G_{C^s} = G_{C^{s'}}$ if $C^s = C^{s'}$, but the distributions of Λ_s and $\Lambda_{s'}$ can differ otherwise. Other works have followed a similar line of thought by restricting analysis to a single type of shoe at a time (Adair et al. (2007), Petraco et al. (2010), Wilson (2012)). In each of those studies, several replicates of the exact same pair of shoes were worn independently for a period of time, after which their accidental locations were annotated, analyzed and compared. This allowed for the identification of common trends for one specific type of shoe. Though such data is ideal for modeling G_{C^s} , the approach cannot be practically scaled to all types of shoes. Collecting multiple annotated observations for all given tread patterns is prohibitively expensive. In addition, the project would have to continue in perpetuity, continually updating the database to account for the ever-growing list of footwear styles and brands.

For this reason we propose a more general and scalable approach in our modeling of Λ_s . Instead of developing independent models G_{C^s} for each unique contact surface, we propose a Bayesian hierarchical model to pool information across many contact surfaces at once. Let \mathbb{C} denote the space of possible contact surfaces. Our goal is to infer the entire family of distributions $G = (G_C)_{C \in \mathbb{C}}$ as a single model, treating a shoe's contact surface C as a covariate. This joint modeling approach helps to leverage the information available in heterogeneous databases—in our case the JESA database—to identify the relationship between the contact surface and accidental locations and to capture commonalities that span across many shoe types.

Let \mathbb{J} denote a set of available shoes (e.g., JESA) used to infer G . Then, $(\Lambda_s)_{s \in \mathbb{J}}$ is a vector of dependent random probability measures, with the dependence between them induced by a hierarchical model on G . We now review existing approaches for modeling vectors of dependent probability measures, limiting our discussion to that which is most relevant to our model. We defer discussion of additional related work to Section 2 of the Supplementary Material (Spencer and Murray (2020)).

2.4. Random vectors of dependent probability measures. Over the years, there has been a broad interest in modeling dependent probability measures, especially via nonparametric Bayes (Hjort et al. (2010), Foti and Williamson (2015)). The approach we use to model $(\Lambda_s)_{s \in \mathbb{J}}$ in this paper is not fully nonparametric, but it is a finite-resolution approximation of one. Thus, it is natural to frame our review within the nonparametric Bayesian literature.

The canonical Bayesian nonparametric approach to modeling a measure μ on a space Ω is to treat it as a random draw from some subclass of measures on Ω . Completely random measures (Kingman (1967)) are an especially tractable subclass of random measures that are composed of a (possibly countably infinite) collection of weighted atoms in Ω . We use $(\theta_i)_{i=1, \dots, \infty} \in \Omega^\infty$ to denote the locations of the atoms of the completely random measure μ and $(w_i)_{i=1, \dots, \infty} \in \mathbb{R}_+^\infty$ to denote the corresponding (nonnegative) atom weights. The defining feature of a completely random measure is that, for any disjoint subsets $\Omega_1, \Omega_2 \subset \Omega$, $\mu(\Omega_1)$ is independent of $\mu(\Omega_2)$ (complete randomness). An accessible review of completely random measures as they pertain to statistical modeling is available in Chen et al. (2010).

For our purposes we are interested in atomic measures that do not necessarily satisfy the complete randomness assumption. In particular, we are interested in atomic random *probability* measures—random measures μ consisting of atoms such that $\mu(\Omega) = 1$. Any finite atomic random measure can be converted to a probability measure via normalization. For

instance, a normalized completely random measure takes the form

$$(2.7) \quad \bar{\mu}(\cdot) = \frac{\sum_{i=1}^{\infty} w_i \delta_{\theta_i}(\cdot)}{\sum_{i=1}^{\infty} w_i},$$

where w_i, θ_i are defined analogously to above. The strength of atomic probability measures is that they can be convolved with probability kernels to define mixture models for densities (e.g., Escobar and West (1995), Rasmussen (2000)). Each atom acts as its own mixture component, providing a framework that is flexible and computationally tractable.

Rather than a single normalized random measure, we are concerned with a vector of dependent random probability measures $(\Lambda_s)_{s \in \mathbb{J}}$ that can capture commonalities across all shoes in JESA. Particularly relevant to our work is the recently proposed normalized compound random measure framework (NCoRM) of Griffin and Leisen (2017), which formulates the vector of random probability measures μ_1, \dots, μ_K on Ω as

$$(2.8) \quad \mu_k(\cdot) = \frac{\sum_{i=1}^{\infty} m_i^k w_i \delta_{\theta_i}(\cdot)}{\sum_{i=1}^{\infty} m_i^k w_i},$$

where $(\theta_i, w_i)_{i=1, \dots, \infty}$ are drawn as in a single completely random measure and $(m_i^k)_{i=1, \dots, \infty}$ are i.i.d. random “score” variables for $k = 1, \dots, K$, following a distribution ρ , that up-weight or down-weight the shared set of atoms defined by the (θ_i, w_i) for each of the μ_k ’s. The distribution of the scores controls the strength of the dependence, with much of the exposition in Griffin and Leisen (2017) devoted to gamma distributions due to their computational tractability. We use the idea of scoring in normalized atomic random measures to develop our model. However, modifications must be made.

The NCoRM approach, as described in Griffin and Leisen (2017), was developed for exchangeable vectors of random probability measures. However, exchangeability does not hold when each measure has an associated covariate (as we have in the contact surfaces \mathcal{C}^s). For this reason we generalize the idea of “scoring” from NCoRMs to the nonexchangeable setting, allowing us to incorporate covariate information. It is worth noting that Griffin and Leisen (2018) also generalizes the NCoRM framework to a nonexchangeable regression framework but differently than we do here.

3. Model. Recall that for a given shoe $s \in \mathbb{J}$, we have assumed each accidental location x_n^s is drawn independently from a probability measure Λ_s on $[0, 100] \times [0, 200]$ where Λ_s itself is randomly drawn from a distribution $G_{\mathcal{C}^s}$ that depends on the contact surface $\mathcal{C}^s \in \mathbb{C}$. Because it is impractical to independently model $G_{\mathcal{C}}$ for all possible $\mathcal{C} \in \mathbb{C}$, we develop a hierarchical model to jointly infer all entries of G , treating each $\mathcal{C} \in \mathbb{C}$ as a high-dimensional spatial covariate.

Before specifying how we model the family of distributions G , it is useful to first address the limited precision of the data. As per Section 2.2, the contact surface variables $\mathcal{C} \in \mathbb{C}$ are defined on a discrete 200 by 100 equally-spaced grid over $[0, 100] \times [0, 200]$. We use \mathcal{A} to denote the set of entries in this grid,

$$(3.1) \quad \mathcal{A} = \{(a_1, a_2) : a_1 \in \{1, \dots, 100\}, a_2 \in \{1, \dots, 200\}\},$$

with gridpoint $(a_1, a_2) \in \mathcal{A}$ corresponding to the area $(a_1 - 1, a_1] \times (a_2 - 1, a_2]$ in $[0, 100] \times [0, 200]$. We restrict our model for Λ_s to have the same resolution as \mathcal{A} by discretizing Λ_s to be a piecewise constant over each gridpoint in \mathcal{A} . This reduced resolution provides computational advantages, simplifies interpretation and guards against overfitting. Further discussion of the discretization is available in Section 1 of the Supplementary Material (Spencer and Murray (2020)).

After discretization, each Λ_s can be characterized by the values it takes at the grid points in \mathcal{A} , and each $G_C \in G$ can be characterized by the multivariate distribution it assigns to those grid points. This provides a natural representation for parametrizing our model. We view G as a family of distributions over the 20,000-dimensional simplex indexed by \mathbb{C} , with each G_{C^s} characterized by the joint distribution it defines over the vector of values in the probability measure $\Lambda_s \mid \mathcal{C}^s$. It is most straightforward to describe G in terms of the generative process it assigns to a generic $\Lambda_s \mid \mathcal{C}^s$, as we do below.

3.1. *Parameterization of Λ_s .* We model each measure $\Lambda_s \sim G_{C^s}$ as the convolution of a normalized random atomic measure μ_s with a two dimensional piecewise constant probability kernel k . We define μ_s to consist of 20,000 atoms at fixed locations, one for each gridpoint in \mathcal{A} . To model the weights of each of these atoms, we generalize the NCoRM scoring technique of Griffin and Leisen (2017) to incorporate the covariate information in \mathcal{C}^s and to allow for spatial dependence between atom weights.

For each $a \in \mathcal{A}$, we define the distribution of $\mu_s \mid \mathcal{C}^s$ as

$$(3.2) \quad \mu_s(a) = \frac{w_a m_a^s}{\sum_{b \in \mathcal{A}} w_b m_b^s} = \frac{w_a \epsilon_a^s \phi_a^s}{\sum_{b \in \mathcal{A}} w_b \epsilon_b^s \phi_b^s}.$$

Here, $(w_a)_{a \in \mathcal{A}}$ are parameters common to all G , and $(m_a^s)_{a \in \mathcal{A}}$ are random shoe-specific location-specific scores applied to the weights of the atoms. The scores further decompose into two components: $m_a^s = \epsilon_a^s \phi_a^s$, with ϵ_a representing “traditional” scores as in NCoRM (assumed to be independent for all shoes and all locations), and ϕ_a^s representing contact-dependent score—variables that depend on the nearby configuration of \mathcal{C}^s . We model the traditional scores as independent draws from $\rho_q = \text{Gamma}(q, 1)$. The contact-dependent scores ϕ_b^s are treated as parameters, defined as follows.

Let $\phi \in [0, 1]^{32}$. For all $a \in \mathcal{A}$, $s \in \mathbb{J}$ define

$$(3.3) \quad \phi_a^s = \phi_{r_a^s} \quad \text{where}$$

$$(3.4) \quad r_a^s = 1 + \sum_{i=-1}^1 \sum_{j=-1}^1 2^{3+i+2j} \mathcal{C}_{a+(i,j)}^s I(\|(i, j)\|^2 \leq 1).$$

By this formulation, ϕ_a^s takes one of $2^5 = 32$ values depending on the value of the contact surface at the gridpoints surrounding a . For instance, if a is completely surrounded by contact surface, that is,

$$(3.5) \quad \mathcal{C}_{a+(-1,0)} = \mathcal{C}_{a+(0,-1)} = \mathcal{C}_a = \mathcal{C}_{a+(1,0)} = \mathcal{C}_{a+(0,1)} = 1,$$

then $\phi_a^s = \phi_{32}$. Similarly, if a is in an area devoid of contact surface, that is,

$$(3.6) \quad \mathcal{C}_{a+(-1,0)} = \mathcal{C}_{a+(0,-1)} = \mathcal{C}_a = \mathcal{C}_{a+(1,0)} = \mathcal{C}_{a+(0,1)} = 0,$$

then $\phi_a^s = \phi_1$. A demonstration of the possible configurations is provided in Figure 5(a) along with an depiction of r_a^s for two $a \in \mathcal{A}$ in Figure 5(b).

Before specifying the functional form for the kernel k (which smooths the atom weights), let us first interpret of the various components that define the atoms weights for μ_s in the context of the shoe sole and accidentals.

The weights are the normalized product of three components:

1. ϕ , which specifies the impact of a gridpoint’s surrounding contact surface on the relative likelihood of accidental occurrence,
2. w , which specifies the impact of the position of a gridpoint’s spatial coordinates on the relative likelihood of accidental occurrence, and

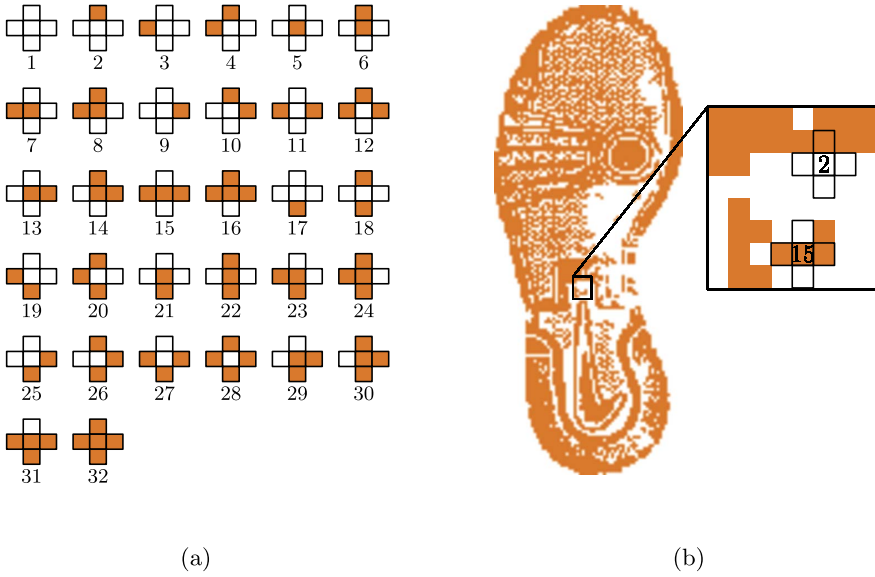


FIG. 5. (a) provides a list of the possible shapes the contact surface can take around an atom, accompanied by the index in $\phi \in [0, 1]^{32}$ to which it corresponds. (b) zooms in on an example shoe’s contact surface (zoomed region outlined in black) to demonstrate the ϕ_a^s value of two example locations.

3. ρ_q (parameterized by q), which specifies the variability across shoes of each gridpoint’s relative probability of accidental occurrence, controlling for position and contact surface.

Essentially, the parameters ϕ and w control the mean of μ_s , whereas its variance depends on the ϵ_a^s scores, distributed according to ρ_q . These choices are in-line with a common belief in forensic footwear analysis—that the locations of accidentals tend to follow a spatially inhomogeneous distribution across the shoe sole (captured by w) and that some areas are more likely to be affected than others, depending on their contact with the ground (captured by ϕ). We model each of ϕ , w and q as global parameters, assuming they take the same value for all shoes JESA.

The random shoe-specific errors ϵ^s capture deviations from this common trend. The coefficient of variation of ρ_q , given by $q^{-1/2}$, indicates the strength of the deviations. The smaller the value of q , the larger the variation of μ_s around its mean.

Finally, we convolve all atoms in all μ_s with a kernel k to obtain Λ_s . The kernel is parameterized to smooth the weights across nearby atoms. Recognizing that the smoothing should be local, we define the kernel k to have finite support, symmetrically redistributing the mass over a window extending three grid points from a in all four axis-aligned directions. Figure 6(a) illustrates the shape of the probability kernel. We refer to this parameterization as the *tiered cake* representation, due to the resultant kernel resembling a tiered cake, with p^α controlling the size of each tier.

We parameterize k as a function $k : \{-3, \dots, 3\}^2 \rightarrow [0, 1]$ such that

$$(3.7) \quad k(i, j) = \kappa_{1+|i|}^h \kappa_{1+|j|}^v.$$

Here, $\kappa^h, \kappa^v \in [0, 1]^4$ define independent symmetric kernels in the horizontal and vertical directions, and k is their composition. To ensure that κ^h and κ^v are unimodal probability kernels, we further reparameterize them as

$$(3.8) \quad \kappa_i^\alpha = \frac{\sum_{j=i}^4 \exp(p_j^\alpha) / (2j - 1)}{\sum_{j=1}^4 \exp(p_j^\alpha)}$$

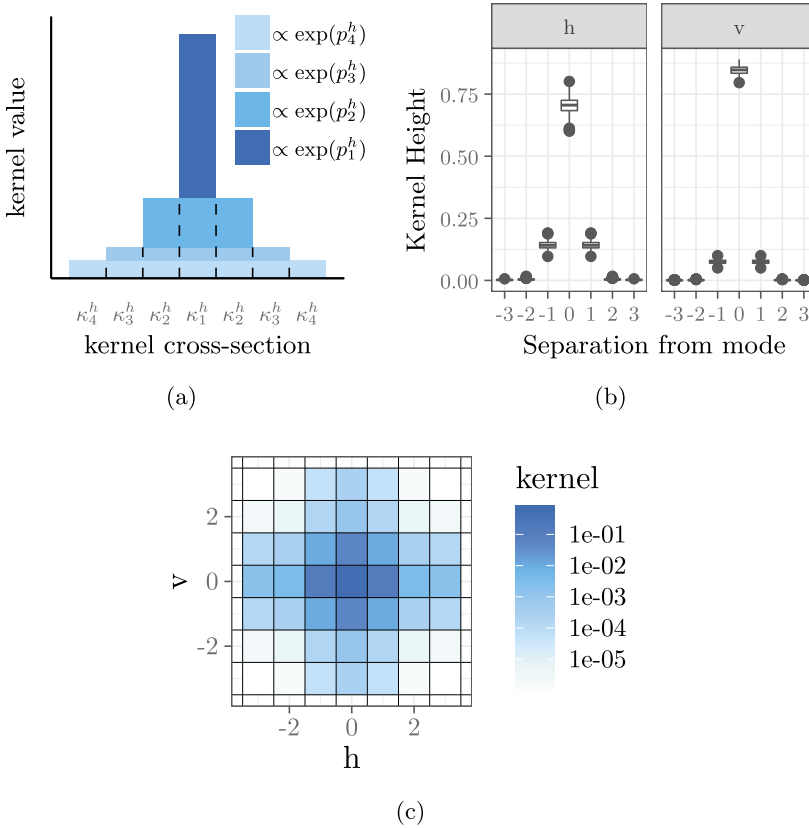


FIG. 6. (a) illustrates the tiered cake parametrization of κ^h . Each uniquely colored tier is proportional to the corresponding $\exp p_i^h$, with the dotted lines depicting how the cake is sliced that form each κ_i^h . (b) demonstrates the posterior fit of κ^h and κ^v using symmetrically arranged boxplots. (c) depicts the posterior mean of $k(i, j) = \kappa_i^h \kappa_j^v$ centered $(0, 0)$. The decay in the h direction controlled by κ^h and the decay in the vertical direction (v) controlled by κ^v , with the hue changing according to a logarithmic scale.

for $i = 1, \dots, 4, \alpha = v, h$ and each $p^\alpha \in \mathbb{R}^4$. Note that our fitted results (Figure 6(b)) indicate that extending the window for three grid points appears to be excessive, but parameterizing three allowed for such a discovery. Going forward, we will often suppress the p^α parameterization to make the presentation more concise, instead relying on the κ^α representation.

3.2. *Model summary and prior.* Having parametrized G , we now formulate the full hierarchical Bayesian model. Let Θ denote the concatenation of the global parameters ϕ, w, q, p^h and p^v . Our prior distribution on Θ is the composition of independent priors on its entries. Letting $MVN(0, 4I_4)$ denote four-dimensional isotropic Gaussian distribution with variance 4 and $MVLN(0, \Sigma)$, denote a multivariate log normal distribution with mean parameter 0 and precision matrix Σ ; the following provides a bird’s eye view of the model via the generative process of the JESA data given $(C_s, N_s)_{s \in \mathbb{J}}$:

Step 1: Generate global parameters:

$$\begin{aligned}
 q &\sim \text{Gamma}(2, 2), & w_E &\sim \text{MVLN}(0, \Sigma), \\
 \phi &\sim \text{unif}([0, 1]^{32}), & p^h, p^v &\sim \text{MVN}(0, 4I_4).
 \end{aligned}$$

Step 2: Generate the densities $\Lambda_s \sim G_{C^s}$ for $s \in \mathbb{J}$:

For $a \in \mathcal{A}$:

$$\epsilon_a^j \sim \text{Gamma}(q, 1),$$

$$\Lambda_s(a) = \sum_{-3 \leq i, j \leq 3} \kappa_{1+|i|}^h \kappa_{1+|j|}^v \frac{w_{a+(i,j)} \epsilon_{a+(i,j)}^s \phi_{a+(i,j)}^s}{\sum_{a' \in \mathcal{A}'} w_{a'} \epsilon_{a'}^s \phi_{a'}^s}.$$

Step 3: Generate the accidental Locations x^s for $s \in \mathbb{J}$.

$$\text{For } n = 1, \dots, N_s: \text{ generate } x_n^s \sim \Lambda_s.$$

Our prior on w in Step 1 uses a coarsened representation, the details of which are provided in Section 3 of the Supplementary Material (Spencer and Murray (2020)). The vector w_E denotes the subvector of unique values after the coarsening. The scales for all of the priors were chosen based on the range of possible behaviors we expected in the model. For example, the variances of 4 for p^h and p^v were chosen to strike a balance: too small of a variance concentrates the kernel around its geometric decaying mean, and too large of a variance places most of the prior density kernels that are essentially step functions. For $\phi \in [0, 1]^{32}$, the upper bound on the uniforms is arbitrary; the likelihood in (4.2) is invariant to scalings of ϕ due to the normalization of μ_s . The rate of ρ_q is fixed at 1 for the same reason.

4. Computation. There are two key computational challenges associated with our model:

1. How do we efficiently compute the posterior of Θ ?
2. How do we efficiently compute the density of an observed set of accidentals x^s given C^s ?

Task 1 (addressed in Section 4.1) arises when fitting our model to the JESA data, and task 2 (addressed in Section 4.2) arises when evaluating models. Before describing our strategies for addressing these tasks, we develop a trick to compute the likelihood of $x^s \in ([0, 100] \times [0, 200])^{N_s}$ given C^s for a given Θ .

The raw likelihood takes the form

$$(4.1) \quad p(x^s | C^s; \Theta) = \int \prod_{n=1}^{N_s} \Lambda(x_n^s) G_{C^s}(d\Lambda)$$

$$(4.2) \quad = \int \prod_{n=1}^{N_s} \sum_{-3 \leq i, j \leq 3} \kappa_{1+|i|}^h \kappa_{1+|j|}^v \frac{w_{x_n^s+(i,j)} \epsilon_{x_n^s+(i,j)}^s \phi_{x_n^s+(i,j)}^s}{\sum_{a \in \mathcal{A}} w_a \epsilon_a^s \phi_a^s} d\rho(\epsilon^s).$$

In a slight abuse of notation, we have overloaded x_n^s to also denote the atom $a \in \mathcal{A}$ to which the real-valued $x_n^s \in [0, 100] \times [0, 200]$ is associated. At first glance the $|\mathcal{A}|$ -dimensional integral over the ϵ^s variable in (4.2) appears to be computationally intractable. It has no closed form and is too high dimensional to efficiently compute using quadrature or generic Monte Carlo algorithms. To overcome this problem, we introduce auxiliary variables.

For each accidental location x_n^s on shoe $s \in \mathbb{J}$, we define Z_n^s by

$$(4.3) \quad \mathbb{P}(Z_n^s = x_n^s + (i, j) | \kappa^h, \kappa^v) = k(i, j) = \kappa_{1+|i|}^h \kappa_{1+|j|}^v,$$

with κ^v, κ^h being the kernel parameters as defined in (3.7) and each $x_n^s \in \mathcal{A}$. We use the shorthand Z^s to refer to the collection $(Z_n^s)_{1 \leq n \leq N_s}$ and use C_a^s to denote the number of times each $a \in \mathcal{A}$ occurs in Z^s . We also introduce the auxiliary variables

$$(4.4) \quad u^s \sim \text{Gamma}\left(N_s, \sum_{a \in \mathcal{A}} w_a \epsilon_a^s \phi_a^s\right),$$

with $\text{Gamma}(\alpha, \beta)$ denoting a gamma distribution with shape α and rate β . We can now analytically marginalize the ϵ^s variables to obtain

$$(4.5) \quad p(x^s \mid \mathcal{C}^s; \Theta) = \int_0^\infty \frac{u^{N_s-1}}{\Gamma(N_s)} \mathbb{E} \left(\frac{1}{\Gamma(q)^{|\mathcal{A}|}} \prod_{a \in \mathcal{A}} \frac{\Gamma(q + C_a^s)(w_a \phi_a^s)^{C_a}}{(u^s w_a \phi_a^s + 1)^{q+C_a^s}} \right) du^s,$$

where $\Gamma(\cdot)$ denotes the gamma function and \mathbb{E} denotes an expectation taken with respect to the distribution of Z^s as given in (4.3). By swapping (4.2) for (4.5), we have exchanged an $|\mathcal{A}|$ -dimensional integral over ϵ^s for a more tractable one-dimensional integral. The full derivation of moving from (4.2) to (4.5) is provided in Section 3 of the Supplementary Material (Spencer and Murray (2020)).

This new expression for the marginal likelihood (4.5) enables us to address challenges (1) and (2) using Monte Carlo algorithms, relying on Markov chain Monte Carlo (MCMC) and importance sampling, respectively. For background information regarding MCMC and importance sampling, we refer the reader to Brooks et al. (2011) and Tokdar and Kass (2010).

4.1. *Computing the posterior for Θ .* We consider an augmented version of the posterior that instantiates the auxiliary variables $Z = (Z^s)_{s \in \mathbb{J}}$ and $U = (u^s)_{s \in \mathbb{J}}$. We use $\mathcal{L}(\Theta, Z, U)$ to denote the augmented likelihood

$$(4.6) \quad \mathcal{L}(\Theta, Z, U) = \frac{1}{\Gamma(q)^{|\mathbb{J}||\mathcal{A}|}} \prod_{s \in \mathbb{J}} \frac{u^{N_s-1}}{\Gamma(N_s)} \prod_{a \in \mathcal{A}} \frac{\Gamma(q + C_a^s)(w_a \phi_a^s)^{C_a^s}}{(u^s w_a \phi_a^s + 1)^{q+C_a^s}} \prod_{n=1}^{N_s} k(\Delta_n^s),$$

where Δ_n^s and C_a^s are defined as in (4.5). Our target is the posterior distribution Θ, U, Z , with density $p(\Theta, U, Z \mid (x^s, \mathcal{C}^s)_{s \in \mathbb{J}})$ satisfying

$$(4.7) \quad p(\Theta, U, Z \mid (x^s, \mathcal{C}^s)_{s \in \mathbb{J}}) \propto \mathcal{L}(\Theta, Z, U)p(\Theta).$$

Our MCMC algorithm consists of sequential updates of the parameters, akin to Metropolis within Gibbs, with most of the components being updated according to slice sampling (Neal (2003), Murray, Adams and MacKay (2010)). The updates are repeatedly performed in the following sequence:

- Each auxiliary variable $(u^s)_{s \in \mathbb{J}}$ is updated one-by-one using slice sampling. These updates can be performed in parallel.
- The entire vector w is updated jointly using elliptical slice sampling.
- Each entry in $(\psi_i)_{i=1, \dots, 32}$ is updated one-by-one using slice sampling.
- The parameter q is updated using a slice sampler.
- Each entry in p^h then p^v is updated one-by-one using slice sampling.
- Each auxiliary variable (z_n^s) is updated one-by-one by Gibbs sampling.

The details and conditional distributions for these updates are available in the Section 4 of the Supplementary Material (Spencer and Murray (2020)). This algorithm provides a sequence of draws of Θ from its posterior that can be used to approximate posterior expectations. Notably, we can use these to approximate the posterior marginal probability of a configuration of accidentals (Task 2), as we now detail in Section 4.2.

4.2. *Computing marginal densities via importance sampling.* A natural metric for assessing the performance of our model is to split \mathbb{J} into a training set T and test set T' , then evaluate the held-out density of the accidental locations on each shoe in T' (given T). Doing this requires computing

$$(4.8) \quad p(x^\tau \mid \mathcal{C}^\tau, T) = \mathbb{E}_\Theta(p(x^\tau \mid \mathcal{C}^\tau, \Theta) \mid (x^s, \mathcal{C}^s)_{s \in T})$$

for each $\tau \in T'$, where $p(\cdot | \mathcal{C}^\tau, T)$ denotes the posterior density. Here, $\mathbb{E}_\Theta(\cdot | (x^s, \mathcal{C}^s)_{s \in T})$ denotes the expected value under the posterior of Θ given the contact surfaces and accidentals in T . Note that the nested integrals in the expression in (4.8) can be separated into an outer integral and an inner integral. The outer integral is the posterior expectation over the global parameters Θ and can be approximated using MCMC draws, as described above. The inner integral, computed for each posterior draw, is over the local auxiliary variables u^τ and Z^τ , as shown in (4.5).

We approximate this integral using importance sampling. Specifically, given a draw of Θ , we define an importance distribution given by

$$(4.9) \quad u | \Theta, N_s, \mathcal{C}^\tau \sim \text{Gamma}\left(N_s, q \sum_{a \in \mathcal{A}} w_a \phi_a^\tau\right),$$

$$(4.10) \quad \mathbb{P}(Z_n = x_n^\tau + a | \Theta, x_n^\tau) = \frac{w_{a+x_n^\tau} \phi_{a+x_n^\tau}^\tau k(a)}{\sum_{b \in B} w_{b+x_n^\tau} \phi_{b+x_n^\tau}^\tau k(b)},$$

where $B = \{-3, \dots, 3\}^2$ and $a \in B$ for all $n \in \{1, \dots, N_s\}$. After drawing $M > 0$ importance samples $u_1, \dots, u_M \in \mathbb{R}_+$ by (4.9) and $Z^1, \dots, Z^M \in \mathcal{A}^{N_\tau}$ by (4.10), the inner integral can be approximated as

$$p(x^\tau | \mathcal{C}^\tau, \Theta) = \frac{\prod_{n=1}^{N_\tau} (\sum_{b \in B} w_{b+x_n^\tau} \phi_{b+x_n^\tau}^\tau k(b))}{\Gamma(q)^{|\mathcal{A}|} (q \sum_{a \in \mathcal{A}} w_a \phi_a^\tau)^{N_\tau}} \sum_{m=1}^M \frac{\exp(u^m q \sum_{a \in \mathcal{A}} w_a \phi_a^\tau)}{\prod_{a \in \mathcal{A}} \frac{(u^m w_a \phi_a^\tau + 1)^{q+C_a^m}}{\Gamma(C_a^m + q)}},$$

where C_a^M denotes the number of times $a \in \mathcal{A}$ occurs as an entry in Z^M .

Thus, using one importance sample ($M = 1$) for each MCMC draw $\Theta^\ell = (\phi^\ell, w^\ell, q^\ell, (p^h)^\ell, (p^v)^\ell)$ yields the approximation

$$(4.11) \quad p(x^\tau | \mathcal{C}^\tau, T) \approx \sum_{\ell=1}^L \frac{\prod_{n=1}^{N_\tau} (\sum_{b \in B} w_{b+x_n^\tau} (\phi^\ell)_{b+x_n^\tau}^\tau k^\ell(b))}{L \Gamma(q^\ell)^{|\mathcal{A}|} (q^\ell \sum_{a \in \mathcal{A}} w_a (\phi^\ell)_a^\tau)^{N_\tau}} \frac{\exp(u^\ell q^\ell \sum_{a \in \mathcal{A}} w_a (\phi^\ell)_a^\tau)}{\prod_{a \in \mathcal{A}} \frac{(u^\ell w_a (\phi^\ell)_a^\tau + 1)^{q^\ell + C_a^\ell}}{\Gamma(C_a^\ell + q^\ell)}},$$

where L is the total number of MCMC draws and the $(u^\ell, Z^\ell)_{1 \leq \ell \leq L}$ are each drawn according to the respective importance distribution for Θ^ℓ . Detailed derivations and discussion of this strategy are available the Section 5 of the Supplementary Material (Spencer and Murray (2020)).

5. Comparisons to competitors and summary of fit.

5.1. *Comparison to competitors.* To demonstrate that efficacy of our model, we compare its performance to three competitor models. The first two models we consider—the uniform model of Stone (2006) and the kernel density estimator of Yekutieli et al. (2012)—rely on fitting a single fixed density Λ for all shoes. Recall from Section 2.3 that the kernel density estimator does not make use of contact surface information when estimating Λ and that the uniform model does not rely on any data at all.

For this reason we introduce a third competitor called the *contact model*. In the contact model each $G_{\mathcal{C}^s}$ is defined as a point mass at $\Lambda_{\mathcal{C}^s}$ with

$$(5.1) \quad \Lambda_{\mathcal{C}^s}(a) \propto \exp(\alpha r_a^s).$$

Here, $\alpha \in \mathbb{R}^{32}$ are shared amongst all of G , similar to ϕ with r_a^s following the same set-up as defined as in (3.4). The parameters α are straightforward to infer using maximum likelihood (fixing $\alpha_1 = 1$ to obtain identifiability).



FIG. 7. Comparison of the performance of four models: the contact model (red) the kernel density estimate (green), our model (blue) and the uniform model (purple) on 50 held out shoes across four data splits. The solid lines depict the metric given in (5.2) for each of 50 shoes (sorted by our model's performance). The dotted lines depict the mean for each model.

We fit our model and the three competitor models to four test/train splits of the JESA data, with each training set consisting of 336 randomly selected shoes. The remaining 50 serve as the test set. For our model the posterior was computed by running the MCMC algorithm outlined in Section 4.1 for 30,000 full sweeps and discarding the first 10,000 iterations as warm-up.

Let T denote a training set and T' denote the test set. As a metric of performance, we used our importance sampling technique to evaluate the held-out density of the accidental locations x^τ on each shoe $\tau \in T'$ given T . Figure 7 depicts the held-out likelihood per accidental on each held-out shoe for each of the four models fit to each of the four splits. Specifically,

$$(5.2) \quad 20,000 \times p(x^\tau | \mathcal{C}^\tau, T)^{1/N_\tau}$$

is reported for each $\tau \in T'$. The scaling by 20,000 is performed for readability of the y -axis (it is equivalent to transforming \mathcal{A} to the unit square), and the N_τ th root is taken to facilitate comparison of average performance on shoes with different numbers of accidentals. This metric is equivalent to comparing the per-accidental average log loss of each shoe. The held-out shoes were sorted according to our model's performance for each of the four splits. Note that, for the uniform model, only those atoms in \mathcal{A} were given positive density, hence the constant density of 1.743 rather than 1.

It is evident from Figure 7 that the two models that account for contact surface (our model and the simple contact surface model (5.1)) vastly outperform the two that do not. Notably, the kernel density estimator assigns 0 density to a shoe in splits 3 and 4, showing an alarming lack of robustness. The performance of our model and the contact model tend to track together across shoes, suggesting that the incorporation of the contact surface is the major driver of both models' success.

We also checked whether the other components of the model (w , κ and ϵ) contribute positively to the model's performance. We fit an additional five variants of our model to the training data and summarized their results in Table 1, along with the performance of the four original competitors. The variant models are defined as follows. "Without scores" refers to

TABLE 1
The mean predictive performance (measured by (5.3)) of our model, five variants on our model and three competitor models. The best performing result is bolded for each split

	Method	Split 1	Split 2	Split 3	Split 4
Existing models	Uniform (Stone (2006))	1.743	1.743	1.743	1.743
	KDE (Yekutieli et al. (2012))	2.266	2.182	0.000	0.000
Other models	Contact	3.954	3.823	4.106	3.995
Our model and variants	Full	4.060	3.832	4.272	4.144
	Without scores	4.052	3.831	4.260	4.131
	Without kernel	4.041	3.794	4.244	4.081
	Without scores and kernel	4.039	3.791	4.238	4.072
	Without w	3.981	3.860	4.131	4.070
	Without ϕ	2.217	2.124	2.187	2.144

our model with all ϵ_a^s variables are fixed at one, “without kernel” refers to our model but without k smoothing, “without scores and kernel” excludes both ϵ_a^s and k , “without w ” fixes $w_E = 1$ and “without ϕ ” fixes all ϕ at 1. Posterior computation for all variant models were performed using appropriate analogs of the MCMC algorithm given in Section 4.1.

For each model and test set T' , Table 1 reports the geometric mean of (5.2) across all held-out shoes, that is,

$$(5.3) \quad 20,000 \times \left(\prod_{\tau \in T'} p(x^\tau | \mathcal{C}^\tau, T)^{1/N_\tau} \right)^{1/|T'|}.$$

This metric is equivalent to the mean per-accidental log loss across shoes.

Table 1 demonstrates that our full model outperforms all competitors, with variants on Splits 1, 3 and 4, being edged out only by “without w ” on Split 2. Nearly all variants perform close to comparably to the full model; the notable exception is “without ϕ .” It performs far worse, highlighting the importance of the contact surface. The persisting decrease in performance of the other variants across splits indicates that each component provides a small gain and is worth keeping in the model.

Note that the superior performance of “without w ” in Split 2 is explained by the presence of an atypical shoe in the test set. It possesses only two accidentals, both of which are located at the left side of the heel. As illustrated in Figure 9(w), w is small toward the heel, especially on the left-hand side. Consequently, including w leads to far lower predictive posterior probability for this particular shoe. Excluding this shoe from test set 2 results in the full model regaining its spot as the top performer.

5.2. *Summary of inferred model parameters.* To investigate our fitted model, we consider the posterior of Θ from Split 1 in Section 5.1. Components of the posterior distribution are summarized in Figures 6, 8 and 9.

Figure 6 summarizes the posterior fit for the kernel k . Figure 6(a) uses boxplots to demonstrate the posterior distribution of both κ^h and κ^v , arranged symmetrically to facilitate visualization of the kernel. For both h and v , the kernel’s mass is mostly concentrated on its mode and immediate neighbours. The smoothing is also more diffuse in the horizontal direction than the vertical direction, suggesting that the accidental distributions are smoother in the horizontal direction than vertical direction. Figure 6(c) demonstrates the composition of the vertical and horizontal kernel into the bivariate kernel.

Figure 8 displays the marginal posterior distributions of each ϕ_1, \dots, ϕ_{32} using boxplots. Here, the larger the associated posterior value, the more likely an accidental is to occur nearby




FIG. 8. Posterior distribution boxplots of the parameters of the 32 possible shapes (listed in Figure 5). Boxplot color indicate the amount of contact surface present in each, with vertical lines partitioning the levels.

contact surface taking on the shape. There is a stark difference in accidental proclivity between gridpoints surrounded mostly by contact surface (shapes 32, 31, 30, 28, 24, 16 as depicted in Figure 5(a)) and those with little contact surface present (shapes 1, 2, 3, 5, 9, 17). This difference supports the intuition among shoeprint examiners that regions which rarely make contact with the ground are typically less likely to accumulate accidentals. Also notable is the discrepancy between different shapes containing the same amount of contact surface. For example, accidentals appear to be nearly to twice as likely to be associated with gridpoints exhibiting shape 31 than those exhibiting shape 24, even though both shapes consist of four of five possible contact components. This inference suggests the shape of the contact surface—and not just the amount of contact surface—also plays a role in a region’s likelihood of being marked with an accidental. However, we caution against over-interpreting such differences due to ϕ being just one component of the larger model.

Figure 9 illustrates the posterior predictive distribution of an accidental for four separate contact surfaces. The first panel is synthetic, consisting entirely of contact surface to demonstrate w . The inward facing side of the toe tends to exhibit more accidentals than the in-outward facing portion, and the front of the heel tends to exhibit more accidentals than the rear of the heel. A depiction of the fit and uncertainty of the raw w parameter is available as Figure 1(b) in the Supplementary Material (Spencer and Murray (2020)).

The second through fourth panels of Figure 9 (Shoe A, Shoe B, Shoe C) demonstrate the posterior mean of Λ^s for three example contact surfaces in JESA. The difference in the magnitude of the density between Shoe B and Shoe C demonstrates that the density associated with a particular location is heavily contingent on the total amount of contact surface present for the shoe; because Shoe C demonstrates relatively little contact surface, the density is much higher in locations where contact surface is present.

6. Discussion. In this work we made progress on a problem put forth by the President’s Council of Advisors in Science and Technology (PCAST (2016)). Namely, we formalized the problem of modeling accidental distributions for random match probabilities, developed a modeling framework for the spatial distribution of accidentals on shoe soles, fit our hierarchical Bayesian model to real data within a Bayesian nonparametric setting to pool information across a variety of shoes and demonstrated that our model vastly outperforms existing models in the literature on a held-out data task.

A key takeaway from this endeavor was the importance of explicitly incorporating the contact surface when modeling accidental distributions. We were the first to do so, and it

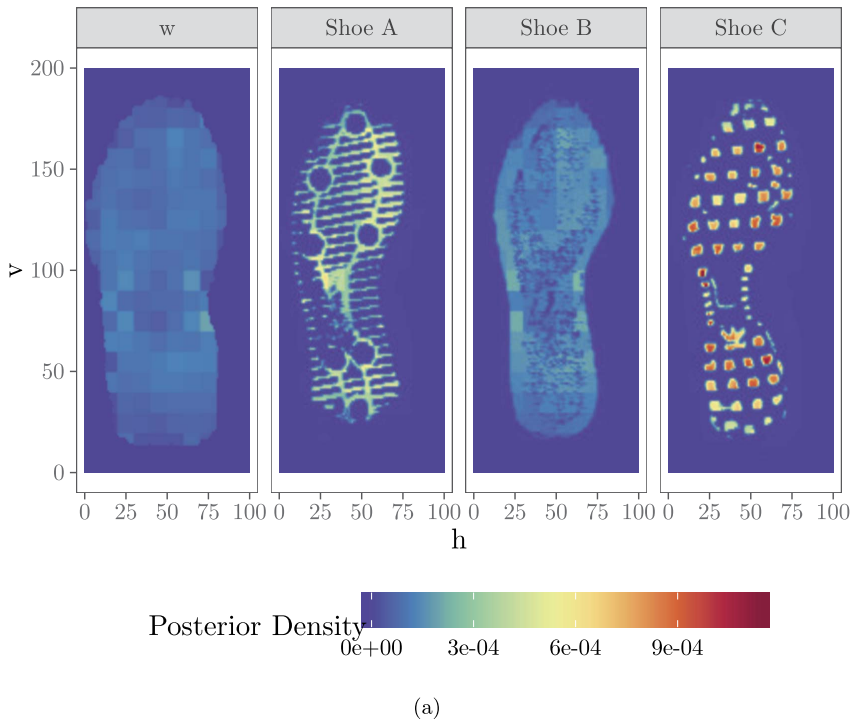


FIG. 9. Panels *w*, *Shoe A*, *Shoe B* and *Shoe C* demonstrate the posterior predictive distribution of accidental locations for four contact surfaces. Panel *w* is synthetic (entirely contact surface). *Shoe A* corresponds to the shoe shown in Figure 1. Shoes *B* and *C* are other contact surfaces from JESA.

resulted in a major improvement over the traditional models. We took care to develop our model hierarchically, allowing for the pooling of information across shoes of different types to capture commonalities in how the contact surface influences accidental distributions. As data sources grow and new data collection efforts are undertaken (CSAFE (2019)), we anticipate the opportunity for more sophisticated models to better capture the relationship between contact surface and accidentals.

Along these lines, a natural extension of our model would be to allow the w , ϕ , and q parameters to differ across shoes according to a nonparametric mixture model. Another possibility would be to extend the model to a spatiotemporal setting, using the temporal data being collected by CSAFE (2019) to model how accidentals accumulate over time.

A possible limitation of our model stems from treating the contact surface parameter ϕ and spatial location parameter w separately. It is plausible that a shoe's intensity would involve dependence between the contact surface and the spatial location. For instance, accidentals could be more likely to occur in high-contact areas when on the toe, but more likely to occur in low-contact areas when on the heel. In such instances a model including an interaction effect would outperform our current model.

Another issue we briefly touched on without addressing was the open problem of formally defining when two impressions “match” ($x^s \equiv x^y$). Given a similarity metric defining when $x^s \equiv x^y$, our model is tailored to computing the RMP. Draws from the posterior distribution of $x^s \mid C^s$ can serve as a surrogate for sampling from \mathbb{A}_{C^y} in (2.6), providing a straightforward Monte Carlo strategy for evaluating the RMP. It is worth noting that, although our exposition focused on RMPs, our approach is equally applicable to calculating other related summaries of uncertainty, such as likelihood ratios or Bayes factors (Evelt, Lambert and Buckleton (1998)).

Finally, we would like to highlight uses of our model outside of direct evaluation of random match probabilities. Recently, the National Institute for Standards in Technology has started development of a multipurpose software tool for forensic footwear examiners (Herman (2016)). One of the tools in development is ShoeGuli, a program for developing synthetic footwear impressions complete with accidentals. As our framework results in an accurate generative model, it is a natural choice for simulating accidental patterns.

Acknowledgments. Thank you very much to Yaron Shor and Sarena Wiesner of the Israeli Police Department and Yoram Yekutieli of Hadassah Academic College for sharing and explaining the details of the JESA database. Thank you to our colleagues in Center for Statistics and Applications in Forensic Evidence and Martin Herman, Hari Iyer, Steven Lund, and others at NIST for advice and feedback. Special thanks to Steve Fienberg for his encouragement and involvement in the early stages of this project.

N. A. Spencer is supported in part by the Center for Statistics and Applications in Forensic Evidence (CSAFE) through Cooperative Agreement #70NANB15H176 between NIST and Iowa State University, which includes activities carried out at Carnegie Mellon University, University of California Irvine and the University of Virginia as well as a fellowship from the Natural Science and Research Council of Canada.

SUPPLEMENTARY MATERIAL

Supplement A (DOI: [10.1214/20-AOAS1334SUPPA](https://doi.org/10.1214/20-AOAS1334SUPPA); .pdf). Additional discussion, related work, derivations, and details of the Monte Carlo algorithms.

Supplement B (DOI: [10.1214/20-AOAS1334SUPPB](https://doi.org/10.1214/20-AOAS1334SUPPB); .zip). Code and synthetic data.

REFERENCES

- ADAIR, T. W., LEMAY, J., McDONALD, A., SHAW, R. and TEWES, R. (2007). The Mount Bierstadt study: An experiment in unique damage formation in footwear. *J. Forensic Identif.* **57** 199–205.
- BODZIAK, W. J. (2017). *Forensic Footwear Evidence*. CRC Press, Boca Raton, FL.
- BODZIAK, W. J., HAMMER, L., JOHNSON, G. M. and SCHENCK, R. (2012). Determining the significance of outsole wear characteristics during the forensic examination of footwear impression evidence. *J. Forensic Identif.* **62** 254–278.
- BROOKS, S., GELMAN, A., JONES, G. L. and MENG, X.-L., eds. (2011). *Handbook of Markov Chain Monte Carlo*. Chapman & Hall/CRC Handbooks of Modern Statistical Methods. CRC Press, Boca Raton, FL. MR2742422 <https://doi.org/10.1201/b10905>
- CHAMPOD, C., EVETT, I. W. and JACKSON, G. (2004). Establishing the most appropriate databases for addressing source level propositions. *Sci. Justice* **44** 153–164.
- CHEN, M.-H., DEY, D. K., MÜLLER, P., SUN, D. and YE, K., eds. (2010). *Frontiers of Statistical Decision Making and Bayesian Analysis: In Honor of James O. Berger*. Springer, New York. MR2766461 <https://doi.org/10.1007/978-1-4419-6944-6>
- CSAFE (2019). Longitudinal shoe outsole impression study. Center for Statistics and Applications in Forensic Evidence. Available at <http://forensicstats.org/shoeoutsoleimpressionstudy/>.
- DALEY, D. J. and VERE-JONES, D. (2008). *An Introduction to the Theory of Point Processes, Vol. II: General Theory and Structure*, 2nd ed. *Probability and Its Applications (New York)*. Springer, New York. MR2371524 <https://doi.org/10.1007/978-0-387-49835-5>
- DAMARY, N. K., MANDEL, M., WIESNER, S., YEKUTIELI, Y., SHOR, Y. and SPIEGELMAN, C. (2018). Dependence among randomly acquired characteristics on shoeprints and their features. *Forensic Sci. Int.* **283** 173–179.
- EDWARDS, H. T. and GOTSONIS, C. (2009). Strengthening forensic science in the United States: A path forward. Statement before the United States Senate Committee on the Judiciary.
- ESCOBAR, M. D. and WEST, M. (1995). Bayesian density estimation and inference using mixtures. *J. Amer. Statist. Assoc.* **90** 577–588. MR1340510
- EVETT, I. W., LAMBERT, J. A. and BUCKLETON, J. S. (1998). A Bayesian approach to interpreting footwear marks in forensic casework. *Sci. Justice* **38** 241–247.

- FACEY, O. E., HANNAH, I. D. and ROSEN, D. (1992). Shoe wear patterns and pressure distribution under feet and shoes, determined by image analysis. *J. Forensic Sci. Soc.* **32** 15–25. [https://doi.org/10.1016/s0015-7368\(92\)73042-4](https://doi.org/10.1016/s0015-7368(92)73042-4)
- FOTI, N. J. and WILLIAMSON, S. A. (2015). A survey of non-exchangeable priors for Bayesian nonparametric models. *IEEE Trans. Pattern Anal. Mach. Intell.* **37** 359–371.
- FRUCHTENICHT, T. L., HERZIG, W. P. and BLACKLEDGE, R. D. (2002). The discrimination of two-dimensional military boot impressions based on wear patterns. *Sci. Justice* **42** 97–104.
- GRIFFIN, J. E. and LEISEN, F. (2017). Compound random measures and their use in Bayesian non-parametrics. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **79** 525–545. MR3611758 <https://doi.org/10.1111/rssb.12176>
- GRIFFIN, J. and LEISEN, F. (2018). Modelling and computation using NCoRM mixtures for density regression. *Bayesian Anal.* **13** 897–916. MR3807871 <https://doi.org/10.1214/17-BA1072>
- HERMAN, M. (2016). Measurements and scoring procedures for footwear impression comparisons. National Institute for Standards in Technology.
- HJORT, N. L., HOLMES, C., MÜLLER, P. and WALKER, S. G. (2010). *Bayesian Nonparametrics. Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge.
- KINGMAN, J. F. C. (1967). Completely random measures. *Pacific J. Math.* **21** 59–78. MR0210185
- KONG, B., RAMANAN, D. and FOWLKES, C. (2017). Cross-domain forensic shoeprint matching. In *British Machine Vision Conference (BMVC)*.
- MURRAY, I., ADAMS, R. and MACKAY, D. (2010). Elliptical slice sampling. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics* 541–548.
- NEAL, R. M. (2003). Slice sampling. *Ann. Statist.* **31** 705–767. MR1994729 <https://doi.org/10.1214/aos/1056562461>
- PCAST (2016). Report to the President, Forensic science in criminal courts: Ensuring scientific validity of feature-comparison methods. President’s Council of Advisors on Science and Technology.
- PETRACO, N. D. K., GAMBINO, C., KUBIC, T. A., OLIVIO, D. and PETRACO, N. (2010). Statistical discrimination of footwear: A method for the comparison of accidentals on shoe outsoles inspired by facial recognition techniques. *J. Forensic Sci.* **55** 34–41. <https://doi.org/10.1111/j.1556-4029.2009.01209.x>
- RASMUSSEN, C. E. (2000). The infinite Gaussian mixture model. In *Advances in Neural Information Processing Systems* 554–560.
- RICHETELLI, N., LEE, M. C., LASKY, C. A., GUMP, M. E. and SPEIR, J. A. (2017a). Classification of footwear outsole patterns using Fourier transform and local interest points. *Forensic Sci. Int.* **275** 102–109.
- RICHETELLI, N., NOBEL, M., BODZIAK, W. J. and SPEIR, J. A. (2017b). Quantitative assessment of similarity between randomly acquired characteristics on high quality exemplars and crime scene impressions via analysis of feature size and shape. *Forensic Sci. Int.* **270** 211–222. <https://doi.org/10.1016/j.forsciint.2016.10.008>
- SHEETS, H. D., GROSS, S., LANGENBURG, G., BUSH, P. J. and BUSH, M. A. (2013). Shape measurement tools in footwear analysis: A statistical investigation of accidental characteristics over time. *Forensic Sci. Int.* **232** 84–91. <https://doi.org/10.1016/j.forsciint.2013.07.010>
- SHOR, Y., WIESNER, S., TSACH, T., GUREL, R. and YEKUTIELI, Y. (2018). Inherent variation in multiple shoe-sole test impressions. *Forensic Sci. Int.* **285** 189–203. <https://doi.org/10.1016/j.forsciint.2017.10.030>
- SKERRETT, J., NEUMANN, C. and MATEOS-GARCIA, I. (2011). A Bayesian approach for interpreting shoemark evidence in forensic casework: Accounting for wear features. *Forensic Sci. Int.* **210** 26–30. <https://doi.org/10.1016/j.forsciint.2011.01.030>
- SPEIR, J. A., RICHETELLI, N., FAGERT, M., HITE, M. and BODZIAK, W. J. (2016). Quantifying randomly acquired characteristics on outsoles in terms of shape and position. *Forensic Sci. Int.* **266** 399–411. <https://doi.org/10.1016/j.forsciint.2016.06.012>
- SPENCER, N. A. and MURRAY, J. S. (2020). Supplement to “A Bayesian hierarchical model for evaluating forensic footwear evidence.” <https://doi.org/10.1214/20-AOAS1334SUPPA>, <https://doi.org/10.1214/20-AOAS1334SUPPB>
- SRIHARI, S. N. and SU, C. (2011). Generative models and probability evaluation for forensic evidence. In *Pattern Recognition, Machine Intelligence and Biometrics* 533–559. Springer, Berlin.
- SRIHARI, S. N. and TANG, Y. (2014). Computational methods for the analysis of footwear impression evidence. In *Computational Intelligence in Digital Forensics: Forensic Investigation and Applications* 333–383. Springer, Berlin.
- STONE, R. S. (2006). Footwear examinations: Mathematical probabilities of theoretical individual characteristics. *J. Forensic Identif.* **56** 577–599.
- THOMPSON, W. C. and NEWMAN, E. J. (2015). Lay understanding of forensic statistics: Evaluation of random match probabilities, likelihood ratios, and verbal equivalents. *Law Hum. Behav.* **39** 332–349. <https://doi.org/10.1037/lhb0000134>
- TOKDAR, S. T. and KASS, R. E. (2010). Importance sampling: A review. *Wiley Interdiscip. Rev.: Comput. Stat.* **2** 54–60.

- WILSON, H. D. (2012). Comparison of the individual characteristics in the outsoles of thirty-nine pairs of Adidas Supernova Classic shoes. *J. Forensic Identif.* **62** 194–203.
- WYATT, J. M., DUNCAN, K. and TRIMPE, M. A. (2005). Aging of shoes and its effect on shoeprint impressions. *J. Forensic Identif.* **55** 181–188.
- YEKUTIELI, Y., SHOR, Y., WIESNER, S. and TSACH, T. (2012). Expert assisting computerized system for evaluating the degree of certainty in 2d shoeprints. Technical Report TP-3211, National Institute of Justice.