

SCORE MATCHING FOR TRUNCATED DENSITY ESTIMATION ON A MANIFOLD

Daniel J. Williams

School of Mathematics
University of Bristol, UK
daniel.williams@bristol.ac.uk

Song Liu

School of Mathematics
University of Bristol, UK

ABSTRACT

When observations are truncated, we are limited to an incomplete picture of our dataset. Recent methods deal with truncated density estimation problems by turning to score matching, where the access to the intractable normalising constant is not required. We present a novel extension to truncated score matching for a Riemannian manifold. Applications are presented for the von Mises-Fisher and Kent distributions on a two dimensional sphere in \mathbb{R}^3 , as well as a real-world application of extreme storm observations in the USA. In simulated data experiments, our score matching estimator is able to approximate the true parameter values with a low estimation error and shows improvements over a maximum likelihood estimator.

1 INTRODUCTION

Often, data can only be observed within a certain ‘window’ of observations, such as a physical threshold (e.g., the borders of a country). This could include applications such as weather events, under-reported count data in disease modelling or undiscovered locations of wildlife habitats. Conventional statistical estimation methods are designed for datasets with full observations, even if this is not always the case. To address this complicated problem, it becomes necessary to extend existing methodology to account for this truncation of data when performing parameter estimation.

In this paper, we consider truncated densities on manifolds. Truncated density estimation on a manifold has a vast potential for applications such as in earth science applications. Estimation in a truncated setting on Earth with a Euclidean approximation may yield inaccurate results if the region is large enough due to its curvature. Additionally, artificial boundaries exist in the case of continental or national borders, which naturally span a large area of the world. When collecting observations of environmental variables, countries may only measure events that happen within their border, but not the phenomena which occur outside of their borders, for example in neighbouring countries or the ocean.

One issue with estimating a truncated density is the evaluation of the normalising constant of a statistical model, which is often intractable in such a scenario. Estimation of unnormalised models are encountered in applications such as Markov random fields for image analysis (Li, 2009), non-Gaussian undirected graphical models, and independent component analysis (Teh et al., 2003). Score matching is a statistical estimation method which bypasses the evaluation of this normalising constant, making it suitable for estimating truncated densities. Score matching has seen many developments in recent years (Hyvärinen, 2005; 2007). Recent work has proposed a generic score matching objective to estimate a truncated density function Yu et al. (2021); Liu et al. (2022). Methodology also exists for estimating a non-truncated statistical model on a manifold (Mardia et al., 2016). However, no known score matching variant currently exists which estimates a truncated density on a manifold. In this work, we propose an efficient estimator that estimates such a truncated statistical model without requiring a normalising constant.

2 BACKGROUND

2.1 SCORE MATCHING

Firstly, let us define a parametric statistical model as

$$p(\mathbf{x}; \beta) = \frac{1}{Z(\beta)} \bar{p}(\mathbf{x}; \beta), \quad Z(\beta) = \int_V \bar{p}(\mathbf{x}; \beta) d\mathbf{x},$$

for some domain V (e.g. $V = \mathbb{R}^d$), and where $p(\mathbf{x}; \beta)$ is the probability density function (pdf). We use this model to approximate an unknown data density function $q(\mathbf{x})$, where β is the parameter which would like to estimate. In many cases, the normalising constant, $Z(\beta)$, is intractable, therefore classical methods such as maximum likelihood estimation (MLE) are unsuited. Instead, one often turns to methods such as score matching (Hyvärinen, 2005; 2007), which minimises

$$J(\beta) = \int_V q(\mathbf{x}) \|\psi_q - \psi_p\|_2^2 d\mathbf{x}, \quad (1)$$

since the *score functions* $\psi_q = \nabla_{\mathbf{x}} \log q(\mathbf{x})$ and $\psi_p = \nabla_{\mathbf{x}} \log p(\mathbf{x}; \beta)$ do not depend on the normalising constant. The later inclusion of a ‘scaling function’ $g(\mathbf{x}) > 0$ yields the *generalised score matching* estimator (Lin et al., 2016; Yu et al., 2016; 2019):

$$J(\beta) = \int_V q(\mathbf{x}) g(\mathbf{x}) \|\psi_q - \psi_p\|_2^2 d\mathbf{x}. \quad (2)$$

Under some mild regularity conditions, the objective function in equation 2 can be written as

$$J(\beta) = \int_V q(\mathbf{x}) \left(g(\mathbf{x}) \left[\psi_p^2 + 2 \text{tr} \{ \nabla_{\mathbf{x}} \psi_p \} \right] + 2 \langle \nabla_{\mathbf{x}} g(\mathbf{x}), \psi_p \rangle \right) d\mathbf{x} + \text{const}, \quad (3)$$

where $\nabla_{\mathbf{x}} \psi_p$ denotes the hessian of $\log p(\mathbf{x}; \beta)$. This formulation is derived using integration by parts (see Theorem 3, Yu et al. (2019)). Now the objective no longer depends on the density of the unknown data $q(\mathbf{x})$, and thus the estimator of β is given by $\hat{\beta} := \arg \min_{\beta} \tilde{J}(\beta)$.

In a *truncated* setting, consider a bounded open subspace $V \subset \mathbb{R}^d$, where the boundary of the subspace is denoted by ∂V . We observe only $\mathbf{x} \in V$, thus the normalization constant $Z(\beta)$ is intractable, therefore score matching is a natural choice of method for estimation. However, the conditions which allowed the score matching objective from equation 2 to be written tractably as equation 3 are no longer valid when truncated. Instead, Liu et al. (2022) introduced an extension which enforced $g(\mathbf{x}') = 0$ for all $\mathbf{x}' \in \partial V$. By including this constraint on g , the integration by parts required for a tractable score matching objective now holds for the truncated case.

Liu et al. (2022) proposed the use of a specific family of g functions given by distance metrics. For example, the Euclidean distance given by $g(\mathbf{x}) = \min_{\tilde{\mathbf{x}} \in \partial V} \|\mathbf{x} - \tilde{\mathbf{x}}\|_2$. Distance functions naturally satisfy the criteria for $g(\mathbf{x})$.

Separately, Mardia et al. (2016) proposed a variant of score matching that is defined on a *compact oriented Riemannian manifold* M . Mardia et al. modified the classical score matching objective function given in equation 1 to include integration over the manifold:

$$J_M(\beta) := \frac{1}{2} \int_M q(\mathbf{x}) \|\psi_q - \psi_p\|_2^2 d\mathbf{x}. \quad (4)$$

They made use of the the following integration by parts as a consequence of Stokes’ Theorem (Stewart et al., 2020):

Theorem 2.1. *Let f_1 and f_2 be continuous twice differentiable functions on a compact Riemannian manifold M , then*

$$\oint_{\partial M} f_1(\mathbf{x}) \frac{\partial \nabla_{\mathbf{x}} f_2(\mathbf{x})}{\partial \mathbf{n}} \partial M = \int_M (\Delta_M f_2(\mathbf{x})) f_1(\mathbf{x}) d\mathbf{x} + \int_M \langle \nabla_{\mathbf{x}} f_1(\mathbf{x}), \nabla_{\mathbf{x}} f_2(\mathbf{x}) \rangle_M d\mathbf{x}. \quad (5)$$

where \mathbf{n} is the unit normal vector.

Additionally, let us define the Laplace-Beltrami operator, which is given by

$$\Delta_M u = \text{tr} \{ \mathbf{P} \nabla_{\mathbf{x}} \cdot (\mathbf{P} \nabla_{\mathbf{x}} \tilde{u}) \}, \quad (6)$$

and the manifold inner product, which is defined as

$$\langle u, v \rangle = \langle \mathbf{P} \tilde{u}, \mathbf{P} \tilde{v} \rangle \quad (7)$$

where u is a continuous twice differentiable real-valued function on M , \tilde{u} is the extension of u to a neighbourhood $N \subset \mathbb{R}^m$ of M , and \mathbf{P} is the $m \times m$ orthogonal projection matrix onto the m -dimensional tangent hyperplane to M .

Mardia et al. (2016) used Theorem 2.1 to rewrite equation 4 (up to a constant) as

$$J_M(\beta) = \int_M q(\mathbf{x}) [\langle \psi_p, \psi_p \rangle_M + 2(\Delta_M \log p(\mathbf{x}; \beta))] d\mathbf{x}, \quad (8)$$

which, as an expectation, no longer involves the unknown data density $q(\mathbf{x})$, see Section 3, Mardia et al. (2016) for details. This step is an analogue to using integration-by-parts in classical score matching to derive a tractable objective function.

Expectations in the objective functions in both equation 3 and equation 8 can be replaced with their Monte Carlo approximations to yield a tractable equation.

2.2 DISTRIBUTIONS ON SPHERES

We consider two primary probability distributions on the unit hypersphere (S^{d-1}) in \mathbb{R}^d : the Kent distribution (or Fisher-Bingham distribution) and the von-Mises Fisher distribution. See Chapter 9, Mardia & Jupp (2009) for an overview of spherical distributions.

2.2.1 KENT DISTRIBUTION

The pdf for the Kent distribution is

$$p(\mathbf{x}; \beta) = \frac{1}{Z(\kappa, \alpha)} \exp \{ \kappa \boldsymbol{\mu}^\top \mathbf{x} + \sum_{j=1}^{d-1} \alpha_j (\boldsymbol{\gamma}_j^\top \mathbf{x})^2 \},$$

for $\mathbf{x} \in \mathbb{R}^d$, where $Z(\kappa, \alpha)$ is the normalising constant (Kent, 1982). Here, $\boldsymbol{\mu}$ is the mean direction, whilst the $\boldsymbol{\gamma}_j$ are the axes which determine the orientation of the probability contours. The κ parameter is the concentration parameter, and the α_j are the ovalness parameters, which determines how elliptical the distribution is (Mardia & Jupp, 2009).

For S^2 in \mathbb{R}^3 , the first and second derivatives of the log pdf can be calculated as

$$\psi_p = \kappa \boldsymbol{\mu} + 2\alpha [\boldsymbol{\gamma}_1 (\boldsymbol{\gamma}_1^\top \mathbf{x}) - \boldsymbol{\gamma}_2 (\boldsymbol{\gamma}_2^\top \mathbf{x})], \quad (9)$$

$$\nabla_{\mathbf{x}} \psi_p = 2\alpha [\boldsymbol{\gamma}_1^\top \boldsymbol{\gamma}_1 - \boldsymbol{\gamma}_2^\top \boldsymbol{\gamma}_2]. \quad (10)$$

The Kent distribution is analogous to the Multivariate Normal (MVN) distribution on \mathbb{R}^d , and has similar properties, such as κ , α_j , and $\boldsymbol{\gamma}_j$ controlling the shape of the data distribution, similar to the role of the covariance matrix of the MVN.

2.2.2 VON MISES-FISHER DISTRIBUTION

The von Mises-Fisher distribution is a special case of the Kent distribution where $\alpha = 0$. It has pdf

$$p(\mathbf{x}; \beta) = \frac{1}{Z(\kappa)} \exp \{ \kappa \boldsymbol{\mu}^\top \mathbf{x} \},$$

for $\mathbf{x} \in \mathbb{R}^d$. Here, $Z(\kappa)$ is the normalising constant, $\boldsymbol{\mu}$ is the mean direction and κ is the concentration parameter (Mardia & Jupp, 2009); all defined as in the Kent distribution. The gradients of the log probability density function are calculated as

$$\psi_p = \kappa \boldsymbol{\mu}, \quad (11)$$

$$\nabla_{\mathbf{x}} \psi_p = \mathbf{0}_{d \times d}, \quad (12)$$

where $\mathbf{0}_{d \times d}$ denotes the $d \times d$ matrix of zeros. The von Mises-Fisher distribution can also be considered as analogous to the MVN with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma} = \kappa^{-1} \mathbf{I}_d$.

3 TRUNCATED MANIFOLD SCORE MATCHING

We extend the approach of both Liu et al. (2022) and Mardia et al. (2016) to combine truncated and manifold score matching. Consider a compact oriented Riemannian manifold M with boundary ∂M , and random vector $\mathbf{x} \in M$ with pdf $q(\mathbf{x})$. Similar to previous score matching implementations by Hyvärinen (2005; 2007), Mardia et al. (2016) and Liu et al. (2022), we minimise the expected squared distance between the score functions of the statistical model and the data,

$$J_S(\beta) := \frac{1}{2} \int_M q(\mathbf{x}) g(\mathbf{x}) \|\psi_q - \psi_p\|_2^2 d\mathbf{x}, \quad (13)$$

where the integration is with respect to the manifold surface M , and $g(\mathbf{x})$ is the scaling function. In the classical score matching derivation, one part of the proof involves a term which is equal to zero under the assumption that the density $q(\mathbf{x}) \rightarrow 0$ as $\mathbf{x} \rightarrow \infty$. In the *truncated* setting, this assumption is no longer valid, as the support of $q(\mathbf{x})$ ends at the boundary, ∂M , for which $q(\mathbf{x}) \neq 0$. Instead, akin to the approach of Liu et al. (2022), we instead let $g(\mathbf{x}') = 0$ for all $\mathbf{x}' \in \partial M$, the boundary of the manifold. This condition is necessary for the derivation of the following theorem.

Theorem 3.1. *Equation 13 can be expressed as*

$$\begin{aligned} J_S(\beta) = & \int_M q(\mathbf{x}) g(\mathbf{x}) \langle \psi_p, \psi_p \rangle_M d\mathbf{x} + 2 \int_M q(\mathbf{x}) g(\mathbf{x}) \Delta_M \log p(\mathbf{x}; \beta) d\mathbf{x} \\ & + 2 \int_M q(\mathbf{x}) \langle \nabla_{\mathbf{x}} g(\mathbf{x}), \psi_p \rangle_M d\mathbf{x}, \end{aligned} \quad (14)$$

where the integration is over M , $\langle \cdot, \cdot \rangle_M$ is the manifold inner product and Δ_M is the Laplace-Beltrami operator.

For proof, see Appendix A. Although we have not specified $g(\mathbf{x})$, we only need find a g such that the condition for which $g(\mathbf{x}') = 0$ for all $\mathbf{x}' \in \partial M$ holds. Any such differentiable and strictly positive function which satisfies this boundary constraint is a valid choice of g . See Liu et al. (2022) for a comprehensive study of the choice of g in a Euclidean setting. In section 3.2, we propose two distance functions that naturally fulfill this criteria, and are suited for the sphere.

The integrals can be approximated using samples from distribution $q(\mathbf{x})$ as

$$\begin{aligned} \tilde{J}_S(\beta) = & \frac{1}{n} \sum_{i=1}^n g(\mathbf{x}_i) \langle \psi_p^{(i)}, \psi_p^{(i)} \rangle_M \\ & + \frac{2}{n} \sum_{i=1}^n g(\mathbf{x}_i) \text{tr}\{\Delta_M \log p(\mathbf{x}_i; \beta)\} \\ & + \frac{2}{n} \sum_{i=1}^n \langle \nabla_{\mathbf{x}} g(\mathbf{x}_i), \psi_p^{(i)} \rangle_M, \end{aligned} \quad (15)$$

provided that n is sufficiently large, and where $\psi_p^{(i)} = \nabla_{\mathbf{x}} \log p(\mathbf{x}_i; \beta)$ denotes the score function evaluated at the i -th data point. The score matching estimator is given by $\hat{\beta} := \arg \min_{\beta} \tilde{J}_S(\beta)$, which can be optimised numerically using gradient based methods.

3.1 SCORE MATCHING ON SPHERICAL DISTRIBUTIONS

Section 2.2 describes two probability distributions defined on a unit sphere; the von Mises-Fisher distribution and the Kent distribution. The generalised manifold score matching approach detailed above is applicable to any manifold on which Stokes' theorem equation 2.1 holds, which includes the unit $(d-1)$ -hypersphere in \mathbb{R}^d where the probability distributions on a sphere are well defined.

Spherical distributions can be defined in Euclidean coordinates, denoted $\mathbf{x} \in \mathbb{R}^d$, or as spherical coordinates, denoted $\mathbf{z} \in S^{d-1}$. Whilst the methods present in this research are applicable to high dimensional manifolds (and spheres), the two dimensional sphere is primarily considered as the application is clear.

On the 2D sphere, Euclidean coordinates have components $\mathbf{x} = (x_1, x_2, x_3) \in \mathbb{R}^3$, whereas spherical coordinates have components $\mathbf{z} = (\theta, \phi) \in S^2$, where $\theta \in [0, \pi)$ is the colatitude and $\phi \in [0, 2\pi]$ is the longitude. Conversion between spherical and Euclidean coordinates is given by

$$\mathbf{x} = (r \cos \phi, r \sin \phi \cos \theta, r \sin \phi \sin \theta),$$

and conversion from Euclidean coordinates to spherical coordinates is given by

$$\mathbf{z} = (\arccos(x_1/r), \arctan(x_3/x_2)), \quad (16)$$

where $r = (x_1^2 + x_2^2 + x_3^2)^{1/2}$. The projection matrix in this case is $P = \mathbf{I}_3 - \mathbf{x}\mathbf{x}^\top$ (Mardia et al., 2016). Using this projection matrix, we can expand the Laplace-Beltrami operator equation 6 using the chain rule.

$$\begin{aligned} \Delta_M \log p(\mathbf{x}; \beta) &= \text{tr}\{P \nabla_{\mathbf{x}} \cdot (P \psi_p)\} \\ &= \text{tr}\{P \nabla_{\mathbf{x}} \cdot ((\mathbf{I}_3 - \mathbf{x}\mathbf{x}^\top) \psi_p)\} \\ &= \text{tr}\{P [\nabla_{\mathbf{x}} (\mathbf{I}_3 - \mathbf{x}\mathbf{x}^\top) \psi_p + (\mathbf{I}_3 - \mathbf{x}\mathbf{x}^\top) \nabla_{\mathbf{x}} \psi_p]\} \\ &= \text{tr}\{P [-2\psi_p \mathbf{x}^\top + P \nabla_{\mathbf{x}} \psi_p]\}, \end{aligned} \quad (17)$$

In implementation, we assume an independent and identically distributed dataset given by either $\{\mathbf{z}_i\}_{i=1}^n = \{(\theta_i, \phi_i)\}_{i=1}^n$, or identically $\{\mathbf{x}_i\}_{i=1}^n = \{(x_{i,1}, x_{i,2}, x_{i,3})\}_{i=1}^n$ drawn from an unknown distribution $q(\mathbf{x})$.

3.1.1 KENT DISTRIBUTION

The derivatives of the Kent distribution are given in equations equation 9 and equation 10. Substituting these into equation 7 and equation 17 gives

$$\begin{aligned} \langle \psi_p, \psi_p \rangle_M &= 4(P\alpha [\gamma_1(\gamma_1^\top \mathbf{x}) - \gamma_2(\gamma_2^\top \mathbf{x})])^\top \cdot \\ &\quad (P\alpha [\gamma_1(\gamma_1^\top \mathbf{x}) - \gamma_2(\gamma_2^\top \mathbf{x})]), \\ \Delta_M \log p(\mathbf{x}; \beta) &= 2\text{tr}\{\alpha P (\gamma_1^\top \gamma_1 - \gamma_2^\top \gamma_2) - P \psi_p \mathbf{x}^\top\}. \end{aligned}$$

3.1.2 VON MISES-FISHER DISTRIBUTION

The derivatives of the von Mises-Fisher distribution are given in equations equation 11 and equation 12. substituting these into equation 7 and equation 17 gives

$$\begin{aligned} \langle \psi_p, \psi_p \rangle_M &= (P\kappa\mu)^\top (P\kappa\mu), \\ \Delta_M \log p(\mathbf{x}; \beta) &= -2\text{tr}\{P \psi_p \mathbf{x}^\top\}. \end{aligned}$$

3.2 SCALING FUNCTIONS FOR THE BOUNDARY

The scaling function $g(\mathbf{x})$ needs to satisfy the criteria that $g(\mathbf{x}') = 0$, $\mathbf{x}' \in \partial M$, i.e. the function takes zero at the boundary of the manifold. Choice of $g(\mathbf{x})$ is an important topic of study and is part of the focus of this research.

Similar to Liu et al. (2022), who proposed the use of a distance function for $g(\mathbf{x})$, for the 2D sphere in \mathbb{R}^3 we propose two distance functions: the spherical Haversine distance and the Euclidean distance on a projected plane. It is important to note that there is at least one central point within the truncated region where the gradient for this choice of $g(\mathbf{x})$ does not exist. Liu et al. (2022) verified the existence of this set, but also proved it to be measure zero, meaning it has no effect in implementation.

3.2.1 HAVERSINE DISTANCE

For two points on the sphere, defined by $\mathbf{z} = (\phi, \theta) \in S^2$ and $\mathbf{z}' = (\phi', \theta') \in \partial M \subset S^2$ the *Haversine*, or *great circle* distance, measures the geodesic distance between these two points along the surface of a sphere, and is given by

$$\begin{aligned} g(\mathbf{z}) &= 2r \arcsin(\sqrt{u}), \\ u &= \sin^2\left(\frac{\phi' - \phi}{2}\right) + \cos(\phi) \cos(\phi') \sin^2\left(\frac{\theta' - \theta}{2}\right). \end{aligned}$$

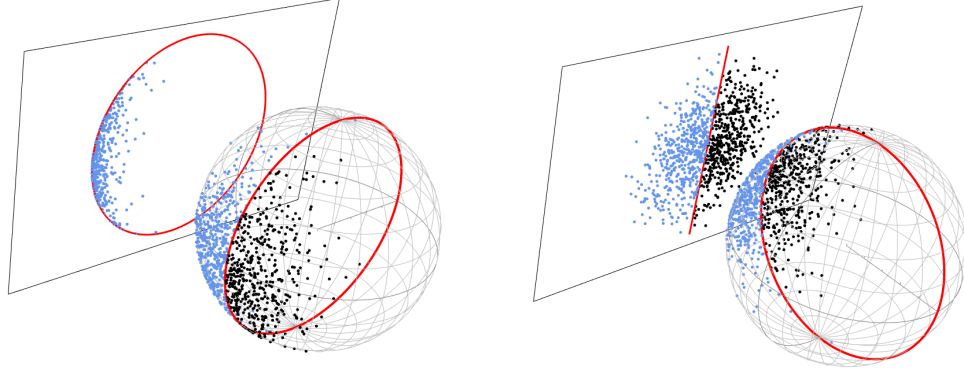


Figure 1: Visualisations of the Euclidean projection, projecting across a different axis, with simulated points from a von Mises-Fisher distribution. A boundary is visualised in red, the unobserved (truncated) points in black and the observed (non-truncated) points in blue. For the projected Euclidean method, $g(x)$ is calculated as a 2D Euclidean distance on this projected plane.




Figure 2: Example of estimation for the von Mises-Fisher distribution (left) and Kent distribution (right), where truncation happens at a constant line of colatitude (ϕ). The estimated mean directions across both cases is given for different methods, and the different distance functions as described in sections 3.2.1 and 3.2.2.

Note that the scaling function here is defined as a function of spherical coordinates z , as opposed to Euclidean coordinates x . By the chain rule we can instead use $\nabla_x g(z) = \nabla_z g(z) \nabla_x z$, where $\nabla_x z$ is the derivative of equation 16 taken with respect to x_1, x_2 and x_3 .

3.2.2 PROJECTED EUCLIDEAN DISTANCE

The points on the surface of the sphere $x \in \mathbb{R}^3$ as well as the boundary ∂M can be projected on a two dimensional plane in \mathbb{R}^2 via setting the third coordinate to zero. Suppose the truncated region is a hemisphere; then the projected plane would be a disk in \mathbb{R}^2 , with the boundary being the outer edge of the disk. See Figure 1 for a visualisation of this method. The corresponding two-dimensional Euclidean distance will be measured on this plane. For the projected point $x_e \in \mathbb{R}^2$, the distance is given by

$$g(x) = \|x_e - \tilde{x}_e\|_2,$$

where $\tilde{x}_e = \arg \min_{x'_e} \|x_e - x'_e\|_2$ such that $x'_e \in \partial M$. The gradient is easily defined as

$$\nabla_x g(x) = \frac{x_e - \tilde{x}_e}{\|x_e - \tilde{x}_e\|_2},$$

i.e. the unit direction away from the boundary.




Figure 3: Root mean squared error (with standard deviation), on a log scale, from 512 replicates between the estimates of parameters and the true parameter values, against sample size for each method. Each plot shows a different experiment, from left to right: von Mises-Fisher distribution with known κ , von Mises-Fisher distribution with unknown κ to be estimated, then Kent distribution with known κ and α .

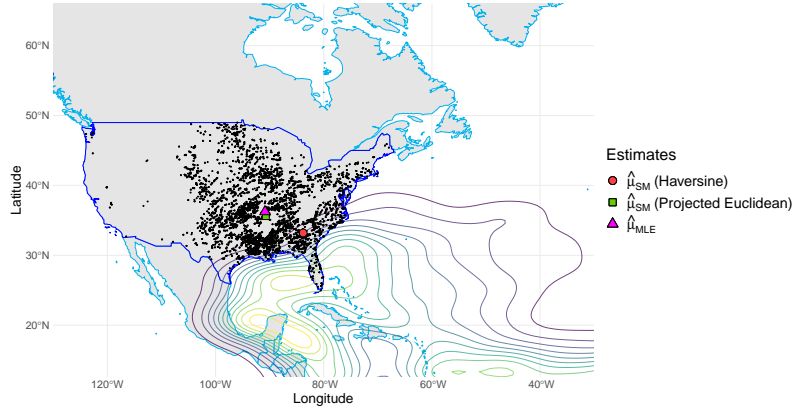


Figure 4: Observed storms as recorded by the USA from 2017 to 2020, with a kernel density estimate contour of storm origin locations over the Atlantic ocean. The mean direction is estimated by TMSM ($\hat{\mu}_{SM}$) and a maximum likelihood estimate ($\hat{\mu}_{MLE}$) for comparison.

4 EXPERIMENTS

We focus on the spheres S^2 in \mathbb{R}^3 , as this manifold covers many important applications in geological or earth science applications where the datasets are collected from a truncated region on a sphere (Earth).

4.1 SIMULATED EXPERIMENT

We start with a synthetic data experiment, where data are simulated from a von Mises-Fisher distribution and Kent distribution using a rejection sampling algorithm, as recommended by Wood (1994). For the von Mises-Fisher distribution experiments, samples are taken with an unknown true mean direction $\mu^* = (\pi/2, \pi)$ and an unknown concentration parameter $\kappa^* = 6$. For the Kent distribution experiments, samples are taken with an unknown true mean direction $\mu^* = (\pi/2, \pi)$ and a known concentration parameter and ovalness parameter, $\kappa^* = 10$ and $\alpha^* = 3$. The primary aim of estimation in both cases is to produce an estimate for the mean direction, $\hat{\mu}$ and in the von-Mises Fisher case, also an estimate for the concentration, $\hat{\kappa}$.

Figure 2 shows the truncated manifold score matching estimator for the von Mises-Fisher and Kent distribution, where the truncation boundary is at a constant value of colatitude ϕ . We show estimation results for using the two proposed scaling functions $g(x)$; the Haversine distance and projected Euclidean distance. These results are compared against a standard MLE approach, without

accounting for the truncation. These plots show even with half of the observations missing, the estimates of the proposed methods are closer to the true mean direction.

4.2 ESTIMATION ACCURACY

We repeat experiments to those presented in Section 4.1 across different setups. In each replicate, n samples were simulated from either the von Mises-Fisher or Kent distribution, but only those where the colatitude $\phi > \pi/2$ were observed. Figure 3 shows the root mean squared error, $\frac{1}{d} \sqrt{\sum_{j=1}^d (\hat{\mu}_j - \mu_j^*)^2}$, for different estimates $\hat{\mu}$, plotted as a function of increasing sample size n . Each method is as follows: the previous state-of-the-art, Truncated Score Matching, *TruncSM* Liu et al. (2022), with a multivariate Normal assumption (see appendix B), our method Truncated Manifold Score Matching (TMSM) using either the Haversine or projected Euclidean as g , as well as MLE for a baseline.

As expected, in all cases the error decreases as sample size increases, showing empirical consistency. Comparing the two distance functions for $g(x)$, both give near identical results across all benchmarks. Our method achieves comparable performance to TruncSM in the first experiment. When κ is also estimated, TMSM is clearly superior to TruncSM.

4.3 STORM EVENTS IN THE USA

Consider an application of estimating the probable locations of extreme storm events in the United States of America (USA) (National Climatic Data Center & of Commerce, 2022). Whilst storms are not restricted to land, observations of these events only happen on land and within the country’s borders. Furthermore, the large area spanned by the USA is significant enough that the curvature of its surface will affect any estimation, highlighting the need for considering estimating densities on a manifold.

We hypothesise that more storms are likely to occur on the East Coast of the USA, and spread across the land from there. We assume there is a true storm ‘centre’ where most storms will be spread out from beyond the country’s eastern border (Oliver, 2008).

Assuming the data can be modelled accurately by a von Mises-Fisher distribution, we estimate the mean direction μ . The estimates $\hat{\mu}$ for both MLE and generalised manifold score matching estimation are given in Figure 4. Included in the figure is a kernel density estimate of storm points of origin, which are calculated via a post-storm analysis, after a storm is observed Landsea & Franklin (2013). Not surprisingly, we see that $\hat{\mu}$ given by MLE is at the center of the dataset, whereas the truncated manifold score matching estimate is further South East, closer to the storm origins.

5 DISCUSSION

The aim of this work was to derive an estimator for truncated densities on a generic manifold. We applied Stokes’ theorem and a scaling function to obtain a tractable score matching objective, of which the estimation was possible via numerical minimisation. The method was tested on the two dimensional sphere in \mathbb{R}^3 , where we showed a superior performance comparing to a Euclidean approximation using TruncSM. A further application was presented on severe storms in the USA, in which the score matching estimator gave a mean direction estimate lying on the East Coast, where the bulk of the storms originated.

Our estimator should also be applicable to any manifold in which Stokes’ Theorem holds, but applications have only been studied on the two dimensional sphere. Other complicated manifolds would require different distribution specification and different scaling functions. It is likely there exists a general purpose scaling function on a manifold that involves projection into a Euclidean domain, which could generalise this work further, but it is beyond the scope of this work.

SOFTWARE AND DATA

A fully packaged implementation in R, with examples and a tutorial, can be found at the GitHub repository here: <https://github.com/dannyjameswilliams/truncsm>.

Additionally, the arxiv version of this paper, with full appendices, can be found here: <https://arxiv.org/abs/2206.14668>.

ACKNOWLEDGEMENTS

We would like to thank Dr. Farhad Babaei for the generous help with the geometric algebra portion of this work. Additionally, we would like to thank the two anonymous reviews for their feedback. Daniel J. Williams was supported by a PhD studentship from the EPSRC Centre for Doctoral Training in Computational Statistics and Data Science (COMPASS).

REFERENCES

- Aapo Hyvärinen. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(Apr):695–709, 2005.
- Aapo Hyvärinen. Some extensions of score matching. *Computational statistics & data analysis*, 51(5):2499–2512, 2007.
- John T Kent. The fisher-bingham distribution on the sphere. *Journal of the Royal Statistical Society: Series B (Methodological)*, 44(1):71–80, 1982.
- C. W. Landsea and J. L. Franklin. Atlantic hurricane database uncertainty and presentation of a new database format, 2013. *Mon. Wea. Rev.*, 141, 3576–3592. Date accessed: 10th May 2022.
- John M Lee. Smooth manifolds. In *Introduction to Smooth Manifolds*, pp. 1–31. Springer, 2013.
- Stan Z Li. *Markov random field modeling in image analysis*. Springer Science & Business Media, 2009.
- Lina Lin, Mathias Drton, and Ali Shojaie. Estimation of high-dimensional graphical models using regularized score matching. *Electronic journal of statistics*, 10(1):806, 2016.
- Song Liu, Takafumi Kanamori, and Daniel J. Williams. Estimating density models with truncation boundaries using score matching. *Journal of Machine Learning Research*, 23(186):1–38, 2022. URL <http://jmlr.org/papers/v23/21-0218.html>.
- Kanti V Mardia and Peter E Jupp. *Directional statistics*, 2009.
- Kanti V Mardia, John T Kent, and Arnab K Laha. Score matching estimators for directional distributions. *arXiv preprint arXiv:1604.08470*, 2016.
- NOAA National Climatic Data Center, NESDIS and US Department of Commerce. Ncdc storm events database, 2022. Date accessed: 10th May 2022.
- John E Oliver. *Encyclopedia of world climatology*. Springer Science & Business Media, 2008.
- Manos Papadakis, Michail Tsagris, Marios Dimitriadis, Stefanos Fafalios, Ioannis Tsamardinos, Matteo Fasiolo, Giorgos Borboudakis, John Burkardt, Changliang Zou, Kleanthi Lakiotaki, and Christina Chatzipantsiou. *Rfast: A Collection of Efficient and Extremely Fast R Functions*, 2021. URL <https://CRAN.R-project.org/package=Rfast>. R package version 2.0.3.
- R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2021. URL <https://www.R-project.org/>.
- James Stewart, Daniel K Clegg, and Saleem Watson. *Calculus: early transcendentals*. Cengage Learning, 2020.

- Yee Whye Teh, Max Welling, Simon Osindero, and Geoffrey E Hinton. Energy-based models for sparse overcomplete representations. *Journal of Machine Learning Research*, 4(Dec):1235–1260, 2003.
- Michail Tsagris, Giorgos Athineou, Anamul Sajib, Eli Amson, and Micah J. Waldstein. *Directional: A Collection of R Functions for Directional Data Analysis*, 2021. URL <https://CRAN.R-project.org/package=Directional>. R package version 5.0.
- Hadley Wickham. *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York, 2016. ISBN 978-3-319-24277-4. URL <https://ggplot2.tidyverse.org>.
- Andrew TA Wood. Simulation of the von mises fisher distribution. *Communications in statistics-simulation and computation*, 23(1):157–164, 1994.
- Ming Yu, Mladen Kolar, and Varun Gupta. Statistical inference for pairwise graphical models using score matching. *Advances in Neural Information Processing Systems*, pp. 2829–2837, 2016.
- Shiqing Yu, Mathias Drton, and Ali Shojaie. Generalized score matching for non-negative data. *Journal of Machine Learning Research*, 20(76):1–70, 2019. URL <http://jmlr.org/papers/v20/18-278.html>.
- Shiqing Yu, Mathias Drton, and Ali Shojaie. Generalized score matching for general domains. *Information and Inference: A Journal of the IMA*, 01 2021.