
Normalizing Flows on Tori and Spheres

Danilo Jimenez Rezende^{*1} George Papamakarios^{*1} Sébastien Racanière^{*1} Michael S. Albergo²
Gurtej Kanwar³ Phiala E. Shanahan³ Kyle Cranmer²

Abstract

Normalizing flows are a powerful tool for building expressive distributions in high dimensions. So far, most of the literature has concentrated on learning flows on Euclidean spaces. Some problems however, such as those involving angles, are defined on spaces with more complex geometries, such as tori or spheres. In this paper, we propose and compare expressive and numerically stable flows on such spaces. Our flows are built recursively on the dimension of the space, starting from flows on circles, closed intervals or spheres.

1. Introduction

Normalizing flows are a flexible way of defining complex distributions on high-dimensional data. A *normalizing flow* maps samples from a base distribution $\pi(u)$ to samples from a target distribution $p(x)$ via a transformation f as follows:

$$x = f(u) \quad \text{where} \quad u \sim \pi(u). \quad (1)$$

The transformation f is restricted to be a *diffeomorphism*: it must be invertible and both f and its inverse f^{-1} must be differentiable. This restriction allows us to calculate the target density $p(x)$ via a change of variables:

$$p(x) = \pi(f^{-1}(x)) \left| \det \left(\frac{\partial f^{-1}}{\partial x} \right) \right|. \quad (2)$$

In practice, $\pi(u)$ is often taken to be a simple density that can be easily evaluated and sampled from, and either f or its inverse f^{-1} are implemented via neural networks such that the Jacobian determinant is efficient to compute.

^{*}Equal contribution ¹DeepMind, London, U.K. ²New York University, New York, U.S.A. ³Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, U.S.A. Correspondence to: Danilo Jimenez Rezende <danilor@google.com>, George Papamakarios <gpapamak@google.com>, Sébastien Racanière <sracaniere@google.com>.

A normalizing flow implements two operations: sampling via Equation (1), and evaluating the density via Equation (2). These operations have distinct computational requirements: generating samples and evaluating their density requires only f and its Jacobian determinant, whereas evaluating the density of arbitrary datapoints requires only f^{-1} and its Jacobian determinant. Thus, the intended usage of the flow dictates whether f , f^{-1} or both must have efficient implementations. For an overview of various implementations and associated trade-offs, see (Papamakarios et al., 2019).

In most existing implementations of normalizing flows, both u and x are defined to live in the Euclidean space \mathbb{R}^D , where D is determined by the data dimensionality. However, this Euclidean formulation is not always suitable, as some datasets are defined on spaces with non-Euclidean geometry. For example, if x represents an angle, its ‘natural habitat’ is the 1-dimensional circle; if x represents the location of a particle in a box with periodic boundary conditions, x is naturally defined on the 3-dimensional torus.

The need for probabilistic modelling of non-Euclidean data often arises in applications where the data is a set of angles, axes or directions (Mardia & Jupp, 2009). Such applications include protein-structure prediction in molecular biology (Hamelryck et al., 2006; Mardia et al., 2007; Boomsma et al., 2008; Shapovalov & Dunbrack Jr, 2011), rock-formation analysis in geology (Peel et al., 2001), and path navigation and motion estimation in robotics (Feiten et al., 2013; Senanayake & Ramos, 2018). Non-Euclidean spaces have also been explored in machine learning, and specifically generative modelling, as latent spaces of variational autoencoders (Davidson et al., 2018; Falorsi et al., 2018; Wang & Wang, 2019; Mathieu et al., 2019; Wang et al., 2019).

A more general formulation of normalizing flows that is suitable for non-Euclidean data is to take $u \in \mathcal{M}$ and $x \in \mathcal{N}$, where \mathcal{M} and \mathcal{N} are *differentiable manifolds* and $f : \mathcal{M} \rightarrow \mathcal{N}$ is a diffeomorphism between them. A difficulty with this formulation is that \mathcal{M} and \mathcal{N} are *diffeomorphic* by definition, so they must have the same topological properties (Kobayashi & Nomizu, 1963). To circumvent this restriction, Gemici et al. (2016) first project \mathcal{M} to \mathbb{R}^D , apply the usual flows there, and then project \mathbb{R}^D back to \mathcal{N} . However, a naive application of this approach can be problematic

when \mathcal{M} or \mathcal{N} are not diffeomorphic to \mathbb{R}^D ; in this case, the projection maps will necessarily contain singularities (i.e. points with non-invertible Jacobian), which may result in numerical instabilities during training. Another approach, proposed by Falorsi et al. (2019) for the case where \mathcal{N} is a Lie group, is to apply the usual Euclidean flows on the Lie algebra of \mathcal{N} (i.e. the tangent space at the identity element), and then use the exponential map to project the Lie algebra onto \mathcal{N} . Since these maps are not from \mathcal{N} to itself, they are not straightforward to compose. Also, the exponential map is not bijective and is computationally expensive in general. We discuss these issues in more detail in Section 3.

Our main contribution is to propose a new set of expressive normalizing flows for the case where \mathcal{M} and \mathcal{N} are compact and connected differentiable manifolds. Specifically, we construct flows on the 1-dimensional circle \mathbb{S}^1 , the D -dimensional torus \mathbb{T}^D , the D -dimensional sphere \mathbb{S}^D , and arbitrary products of these spaces. The proposed flows can be made arbitrarily flexible, and avoid the numerical instabilities of previous approaches. Our flows are applicable when we already know the manifold structure of the data, such as when the data is a set of angles or latent variables with a prescribed topology. We empirically demonstrate the proposed flows on synthetic inference problems designed to test the ability to model sharp and multi-modal densities.

2. Methods

We will begin by constructing expressive and numerically stable flows on the circle \mathbb{S}^1 . Then, we will use these flows as building blocks to construct flows on the torus \mathbb{T}^D and the sphere \mathbb{S}^D . The torus \mathbb{T}^D can be written as a Cartesian product of D copies of \mathbb{S}^1 , which will allow us to build flows on \mathbb{T}^D autoregressively. The sphere \mathbb{S}^D can be written as a transformation of the cylinder $\mathbb{S}^{D-1} \times [-1, 1]$, which will allow us to build flows on \mathbb{S}^D recursively, using flows on \mathbb{S}^1 and $[-1, 1]$ as building blocks. By combining these constructions, we can build a wide range of compact connected manifolds of interest to fundamental and applied sciences.

2.1. Flows on the Circle \mathbb{S}^1

Depending on what is most convenient, we will sometimes view \mathbb{S}^1 as embedded in \mathbb{R}^2 , that is $\{(x_1, x_2) \in \mathbb{R}^2; x_1^2 + x_2^2 = 1\}$, or parameterize it by a coordinate $\theta \in [0, 2\pi]$, identifying 0 and 2π as the same point. In this section, we describe how to construct a diffeomorphism f that maps the circle to itself.

Since 0 and 2π are identified as the same point, the transformation $f : [0, 2\pi] \rightarrow [0, 2\pi]$ must satisfy appropriate boundary conditions to be a valid diffeomorphism on the circle. The following conditions are sufficient:

$$f(0) = 0, \quad (3)$$

$$f(2\pi) = 2\pi, \quad (4)$$

$$\nabla f(\theta) > 0, \quad (5)$$

$$\nabla f(\theta)|_{\theta=0} = \nabla f(\theta)|_{\theta=2\pi}. \quad (6)$$

The first two conditions ensure that 0 and 2π are mapped to the same point on the circle. The third condition ensures that the transformation is strictly monotonic, and thus invertible. Finally, the fourth condition ensures that the Jacobians agree at 0 and 2π , thus the probability density is continuous.

A restriction in the above conditions is that 0 and 2π are fixed points. Nonetheless, this restriction can be easily overcome by composing a transformation f satisfying these conditions with a phase translation $\theta \mapsto (\theta + \phi) \bmod 2\pi$, where ϕ can be a learnable parameter. Such a phase translation is volume-preserving, so it does not incur a volume correction in the calculation of the probability density.

Given a collection $\{f_i\}_{i=1, \dots, K}$ of transformations satisfying the above conditions, we can combine them into a more complex transformation f that also meets these conditions. One such mechanism for combining transformations is function composition $f = f_K \circ \dots \circ f_1$, which can easily be seen to satisfy Equations (3) to (6). Alternatively, we can combine transformations using convex combinations, as any convex combination defined by

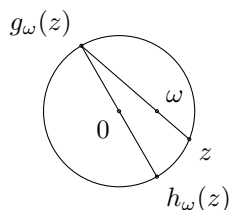
$$f(\theta) = \sum_i \rho_i f_i(\theta), \quad \text{where } \rho_i \geq 0 \text{ and } \sum_i \rho_i = 1 \quad (7)$$

also satisfies Equations (3) to (6). By alternating between function compositions and convex combinations, we can construct expressive flows on \mathbb{S}^1 from simple ones.

Next, we describe three circle diffeomorphisms that by construction satisfy the above conditions: *Möbius transformations*, *circular splines*, and *non-compact projections*.

2.1.1. MÖBIUS TRANSFORMATION

Möbius transformations have previously been used to define distributions on the circle and sphere (see e.g. Kato et al., 2008; Wang, 2013; Kato & McCullagh, 2015). We will first describe the Möbius transformation in the general case of the sphere \mathbb{S}^D , and then show how to adapt it, when $D = 1$, to create expressive flows on the circle.



Consider \mathbb{S}^D , the unit sphere in \mathbb{R}^{D+1} . Let ω be a point in \mathbb{R}^{D+1} with norm strictly less than 1. For any point z in \mathbb{S}^D , draw a straight line between z and ω (as shown on the left). This line intersects \mathbb{S}^D at exactly two distinct points. One is z . Call the other one $g_\omega(z)$.

Definition 1. We define the Möbius transformation $h_\omega(z)$ of z with centre ω to be $-g_\omega(z)$.

An explicit formula for h_ω is given by

$$h_\omega(z) = \frac{1 - \|\omega\|^2}{\|z - \omega\|^2} (z - \omega) - \omega. \quad (8)$$

When $\omega = 0$, the transformation h_ω is just the identity. In general, the variables z and ω in Equation (8) are meant to be $(D + 1)$ -dimensional real vectors. When $D = 1$, the equation also reads correctly if z and ω are taken to be complex numbers. In this case, Equation (8) and the Möbius transformations of the complex plane are related as explained in formula (6) of Kato & McCullagh (2015).

The transformation h_ω expands the part of the sphere that is close to ω , and contracts the rest. Hence, it can transform a uniform base distribution into a unimodal smooth distribution parameterized by ω . One property of the Möbius transformation is that it does not become more expressive by composing various h_ω . This is because the set of transformations Rh_ω , where R can be any matrix in $\text{SO}(D + 1)$, forms a group under function composition (see Theorem 2 of Kato & McCullagh, 2015). Since composition of two such transformations remains a member of the group, their expressivity is not increased.

In case of the circle (where $D = 1$), we would like to get expressive flows using a convex combination as in Equation (7) of Möbius transformations. The transformation h_ω defines a map from angles to angles that satisfies Equations (5) and (6), but it does not satisfy Equations (3) and (4). This can be fixed by adding a rotation after h_ω . Let R_ω be a rotation in \mathbb{R}^2 with centre $(0, 0)$ that maps $h_\omega(1, 0)$ back to $(1, 0)$. We define $f_\omega(\theta)$ to be the polar angle, in $[0, 2\pi)$, of $R_\omega \circ h_\omega(z)$, and extend f_ω by continuity to the whole range $[0, 2\pi]$ by setting $f_\omega(2\pi) = 2\pi$. This function satisfies Equations (3) to (6). We can then easily combine various f_ω via convex combinations and obtain expressive distributions on \mathbb{S}^1 —see Figure 7 in the appendix for an illustration. Unlike a single Möbius transform, a convex combination of two or more Möbius transforms is not analytically invertible, but it can be numerically inverted with precision ϵ using bisection search with $\mathcal{O}(\log \frac{1}{\epsilon})$ iterations.

2.1.2. CIRCULAR SPLINES (CS)

Spline flows is a methodology for creating arbitrarily flexible flow transformations from \mathbb{R} to itself, first proposed by Müller et al. (2019) and further developed by Durkan et al. (2019a;b). Here, we will show how to adapt the *rational-quadratic spline flows* of Durkan et al. (2019b) to satisfy the sufficient boundary conditions of circle diffeomorphisms.

Rational-quadratic spline flows define the transformation $f : \mathbb{R} \rightarrow \mathbb{R}$ piecewise as a combination of K segments, with each segment being a simple rational-quadratic function. Specifically, the transformation is parameterized by a set of $K + 1$ knot points $\{x_k, y_k\}_{k=0, \dots, K}$ and a set of K slopes

$\{s_k\}_{k=1, \dots, K}$, such that $x_{k-1} < x_k$, $y_{k-1} < y_k$ and $s_k > 0$ for all $k = 1, \dots, K$. Then, in each interval $[x_{k-1}, x_k]$, the transformation f is defined to be a rational quadratic:

$$f(\theta) = \frac{\alpha_{k2}\theta^2 + \alpha_{k1}\theta + \alpha_{k0}}{\beta_{k2}\theta^2 + \beta_{k1}\theta + \beta_{k0}}, \quad (9)$$

where the coefficients $\{\alpha_{ki}, \beta_{ki}\}_{i=0,1,2}$ are chosen so that f is strictly monotonically increasing and $f(x_{k-1}) = y_{k-1}$, $f(x_k) = y_k$, and $\nabla f(\theta)|_{\theta=x_k} = s_k$ for all $k = 1 \dots, K$ (see Durkan et al., 2019b, for more details).

We can easily restrict f to be a diffeomorphism from $[0, 2\pi]$ to itself, by setting $x_0 = y_0 = 0$ and $x_K = y_K = 2\pi$. This construction satisfies the first three sufficient conditions in Equations (3) to (5), and can be used to define probability densities on the closed interval $[0, 2\pi]$. In addition, by setting $s_1 = s_K$, we satisfy the fourth condition in Equation (6), and hence we obtain a valid circle diffeomorphism which we refer to as a *circular spline (CS)*.

Circular splines can be made arbitrarily flexible by increasing the number of segments K . Therefore, unlike Möbius transformations, it is not necessary to combine them via convex combinations to increase their expressivity. An advantage of circular splines is that they can be inverted exactly, by first locating the corresponding segment (which can be done in $\mathcal{O}(\log K)$ iterations using binary search since the segments are sorted), and then inverting the corresponding rational quadratic (which can be done analytically by solving a quadratic equation).

2.1.3. NON-COMPACT PROJECTION (NCP)

As discussed in the introduction, Gemici et al. (2016) project the manifold to \mathbb{R}^D , apply the usual flows there, and then project \mathbb{R}^D back to the manifold. Naively applying this method can be numerically unstable. However, here we show that, with some care, the method can be specialized to \mathbb{S}^1 in a numerically stable manner. Since this method involves projecting \mathbb{S}^1 to the non-compact space \mathbb{R} , we refer to it as *non-compact projection (NCP)*.

We will use the projection $x : (0, 2\pi) \rightarrow \mathbb{R}$ defined by $x(\theta) = \tan\left(\frac{\theta}{2} - \frac{\pi}{2}\right)$. This projection maps the circle minus the point $\theta = 0$ bijectively onto \mathbb{R} . Applying the affine transformation $g(x) = \alpha x + \beta$ in the non-compact space, where $\alpha > 0$ and β are learnable parameters, defines a corresponding flow on the circle, given by

$$\begin{aligned} f(\theta) &= x^{-1} \circ g \circ x(\theta) \\ &= 2 \tan^{-1} \left(\alpha \tan \left(\frac{\theta}{2} - \frac{\pi}{2} \right) + \beta \right) + \pi, \end{aligned} \quad (10)$$

with gradient

$$\nabla f(\theta) = \left[\frac{1 + \beta^2}{\alpha} \sin^2 \left(\frac{\theta}{2} \right) + \alpha \cos^2 \left(\frac{\theta}{2} \right) - \beta \sin \theta \right]^{-1}.$$

Even though the expression for f is not defined at the endpoints 0 and 2π , the expression for the gradient is. The transformation satisfies the appropriate boundary conditions,

$$\begin{aligned} \lim_{\theta \rightarrow 0^+} f(\theta) &= 0, \\ \lim_{\theta \rightarrow 2\pi^-} f(\theta) &= 2\pi, \\ \nabla f(\theta)|_{\theta=0} &= \nabla f(\theta)|_{\theta=2\pi} = \alpha^{-1}. \end{aligned}$$

Therefore, we can extend f to $[0, 2\pi]$ by continuity such that $f(0) = 0$ and $f(2\pi) = 2\pi$, which yields a valid circle diffeomorphism. Although not immediately obvious, NCP flows and Möbius transformations are intimately related, as explained in Appendix H (see also Downs & Mardia, 2002, Section 2.1).

The above boundary conditions are satisfied when the transformation g is affine, but they are not generally satisfied when g is an arbitrary diffeomorphism. This limits the type of flow we can put on the non-compact space. Therefore, instead of making g more expressive, we choose to increase the expressivity of the NCP flow by combining multiple transformations f via convex combinations.

Similarly to Möbius transformations, NCP flows form a group. That's because affine maps with positive slope on \mathbb{R} form a group, and if g_1 and g_2 are affine maps and $f_k = x \circ g_k \circ x^{-1}$, then $f_1 \circ f_2 = x \circ g_1 \circ g_2 \circ x^{-1}$. Hence, the composition of two NCPs with parameters (α_k, β_k) , $k = 1, 2$ is another NCP with parameters $(\alpha_1 \alpha_2, \beta_1 + \alpha_1 \beta_2)$. So composing two NCP flows would produce a new NCP flow, leading to no increase in expressivity.

A potential issue with NCP is that, near the endpoints 0 and 2π , evaluating f using Equation (10) directly is numerically unstable. To circumvent this numerical difficulty, we can use equivalent linearized expressions when θ is near the endpoints. For example, for θ close to 0 we can approximate $f(\theta) \approx \theta/\alpha$, whereas for θ close to 2π we have $f(\theta) \approx 2\pi + (\theta - 2\pi)/\alpha$.

2.2. Generalization to the Torus \mathbb{T}^D

Having defined flows on the circle \mathbb{S}^1 , we can easily construct autoregressive flows on the D -dimensional torus $\mathbb{T}^D \cong (\mathbb{S}^1)^D$. Any density $p(\theta_1, \dots, \theta_D)$ on \mathbb{T}^D can be decomposed via the chain rule of probability as

$$p(\theta_1, \dots, \theta_D) = \prod_i p(\theta_i | \theta_1, \dots, \theta_{i-1}), \quad (11)$$

where each conditional $p(\theta_i | \theta_1, \dots, \theta_{i-1})$ is a density on \mathbb{S}^1 . Each conditional density can be implemented via a flow $f_{\psi_i} : \mathbb{S}^1 \rightarrow \mathbb{S}^1$, whose parameters ψ_i are a function of $(\theta_1, \dots, \theta_{i-1})$. Thus the joint transformation $f : \mathbb{T}^D \rightarrow \mathbb{T}^D$ given by $f(\theta_1, \dots, \theta_D) = (f_{\psi_1}(\theta_1), \dots, f_{\psi_D}(\theta_D))$ is an *autoregressive flow* on the torus. In the terminology of Papamakarios et al. (2019, Section 3.1), an autoregressive flow on

a torus is simply an autoregressive flow whose *transformers* are circle diffeomorphisms, i.e. obey the boundary conditions in Equations (3) to (6). As with any autoregressive flow, the Jacobian of f is triangular, therefore the Jacobian determinant in the density calculation in Equation (2) can be computed efficiently as the product of the diagonal terms.

The parameters ψ_i of the i -th transformer are a function of $(\theta_1, \dots, \theta_{i-1})$ known as the i -th *conditioner*. In order to guarantee that the conditioners are periodic functions of each θ_i , we can make ψ_i be a function of $(\cos \theta_1, \sin \theta_1, \dots, \cos \theta_{i-1}, \sin \theta_{i-1})$ instead. In our experiments, we implemented the conditioners using *coupling layers* (Dinh et al., 2017). Implementations based on *masking* (Kingma et al., 2016; Papamakarios et al., 2017) are also possible.

More generally, autoregressive flows can be applied in the same way on any manifold that can be written as a Cartesian product of circles and intervals, such as the 2-dimensional cylinder. Flows on intervals can be constructed e.g. using regular (non-circular) splines as described in Section 2.1.2. Thus, by taking f_{ψ_i} to be either a circle diffeomorphism or an interval diffeomorphism as required, we can handle arbitrary products of circles and intervals.

2.3. Generalization to the Sphere \mathbb{S}^D

The Möbius transformation (Section 2.1.1) can in principle be used to define flows on the sphere \mathbb{S}^D for $D \geq 2$, but, as noted, its expressivity does not increase by composition. Increasing the expressivity via convex combinations in $D = 1$ was only possible because we expressed diffeomorphisms on \mathbb{S}^1 as strictly increasing functions on $[0, 2\pi]$. This construction however does not readily extend to $D \geq 2$. Instead, we propose two alternative flow constructions for \mathbb{S}^D : a *recursive construction* that uses cylindrical coordinates, and a construction based on the *exponential map*. In what follows, \mathbb{S}^D will be embedded in \mathbb{R}^{D+1} as $\{(x_1, \dots, x_{D+1}) \in \mathbb{R}^{D+1}; \sum_i x_i^2 = 1\}$.

2.3.1. RECURSIVE CONSTRUCTION

In Section 2.1, we described three methods for building univariate flows on \mathbb{S}^1 . We also described how to build univariate flows on the interval $[-1, 1]$ using splines (Section 2.1.2). Using circle and interval flows as building blocks, we will construct a multivariate flow on \mathbb{S}^D for $D \geq 2$ by recursing over the dimension of the sphere.

Our construction works as follows. First, we transform the sphere \mathbb{S}^D into the cylinder $\mathbb{S}^{D-1} \times [-1, 1]$ using the map

$$T_{s \rightarrow c}(x) = \left(\frac{x_{1:D}}{\sqrt{1 - x_{D+1}^2}}, x_{D+1} \right). \quad (12)$$

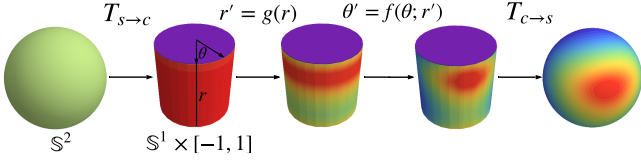


Figure 1. Illustration of the recursive flow on the sphere \mathbb{S}^2 .

Then a two-stage autoregressive flow is applied on the cylinder. The ‘height’ $r \in [-1, 1]$ is transformed first by a univariate flow $g : [-1, 1] \rightarrow [-1, 1]$, and the ‘base’ z is transformed second by a conditional flow $f : \mathbb{S}^{D-1} \rightarrow \mathbb{S}^{D-1}$ whose parameters are a function of $g(r)$, that is

$$T_{c \rightarrow c}(z, r) = \left(f(z; g(r)), g(r) \right). \quad (13)$$

Finally, the cylinder is transformed back to the sphere by

$$T_{c \rightarrow s}(z, r) = \left(z\sqrt{1-r^2}, r \right). \quad (14)$$

The flow $h : \mathbb{S}^D \rightarrow \mathbb{S}^D$ is defined to be the composition $h = T_{c \rightarrow s} \circ T_{c \rightarrow c} \circ T_{s \rightarrow c}$. When stacking multiple h (each with its own parameters), all the internal terms in $T_{s \rightarrow c} \circ T_{c \rightarrow s}$ cancel out, so it is more efficient to just not compute them. Only the first $T_{c \rightarrow s}$ and last $T_{s \rightarrow c}$ need to be computed. Figure 1 illustrates this procedure on \mathbb{S}^2 .

The above recursion continues until we reach \mathbb{S}^1 , on which we can use the flows we have already described. Unrolling the recursion up to \mathbb{S}^1 is equivalent to first transforming \mathbb{S}^D into $\mathbb{S}^1 \times [-1, 1]^{D-1}$, then applying an autoregressive flow on $\mathbb{S}^1 \times [-1, 1]^{D-1}$ as described in Section 2.2, and finally transforming $\mathbb{S}^1 \times [-1, 1]^{D-1}$ back to \mathbb{S}^D . In practice, we can use any ordering of the variables in the autoregressive flow and not just the one implied by the recursion, and we can compose multiple autoregressive flows before transforming back to \mathbb{S}^D . Figure 6 in the appendix illustrates the model with the recursion fully unrolled.

At this point, it is important to examine carefully the transformations $T_{s \rightarrow c}$ and $T_{c \rightarrow s}$. Maps from $\mathbb{S}^{D-1} \times [-1, 1]$ to \mathbb{S}^D can never be diffeomorphisms since the topology of these two spaces differ. Nevertheless, $T_{s \rightarrow c}$ and $T_{c \rightarrow s}$ are invertible almost everywhere, in the sense that we can remove sets of measure 0 from $\mathbb{S}^{D-1} \times [-1, 1]$ and \mathbb{S}^D , and the restriction of these maps to these subsets are diffeomorphisms. However, extra care must be taken to check that updates to the density caused by $T_{s \rightarrow c}$ and $T_{c \rightarrow s}$ are still numerically stable, and that the density is well-behaved everywhere on \mathbb{S}^D . In Appendix A.1, we prove that the update to a base density $\pi(z, r)$ due to $T_{c \rightarrow s}$ is

$$p(T_{c \rightarrow s}(z, r)) = \frac{\pi(z, r)}{(1-r^2)^{\frac{D}{2}-1}}. \quad (15)$$

The update due to $T_{s \rightarrow c}$ is simply the inverse of the above. Additionally, in Appendix A.1 we prove that the combined

density update due to $h = T_{c \rightarrow s} \circ T_{c \rightarrow c} \circ T_{s \rightarrow c}$ does not have any singularity whenever g is such that $g'(-1) > 0$ and $g'(1) > 0$. Since we implement g with spline flows (Section 2.1.2), this condition is satisfied by construction, so the flow density is guaranteed to be finite.

Since the density update due to $T_{s \rightarrow c}$ and $T_{c \rightarrow s}$ can be done analytically and the rest of the transformation is autoregressive, the overall density update is efficient to compute. For $D = 2$, the volume correction in Equation (15) is equal to 1, and therefore $T_{c \rightarrow s}$ and $T_{s \rightarrow c}$ preserve infinitesimal volume elements. Finally, we observe that the well-known *von Mises–Fisher* distribution on \mathbb{S}^2 (Fisher, 1953) can be obtained by transforming a uniform base density with Equations (12) to (14), where f is the identity and g has a simple functional form (for details, see Jakob, 2012).

2.3.2. EXPONENTIAL-MAP FLOWS

Ideally, we would like to specify flows directly on the manifold and avoid mapping between non-diffeomorphic sets. This motivates us to explore *exponential-map flows*, a mechanism for building flows on spheres proposed by Sei (2013).

The main idea consists of two steps. First, we define a *cost-convex* scalar field $\phi(x)$ on \mathbb{S}^D —for the definition of ‘cost-convex’, see (Sei, 2013, Formula 1). Second, we construct a flow using the exponential map of the gradient $\nabla\phi(x) \in T_x\mathbb{S}^D$, where $T_x\mathbb{S}^D$ is the tangent space of \mathbb{S}^D at x . The exponential map $\exp_x : T_x\mathcal{M} \rightarrow \mathcal{M}$ on a Riemannian manifold \mathcal{M} is defined as the terminal point $\gamma(1)$ of a geodesic $\gamma : [0, 1] \rightarrow \mathcal{M}$ that passes through $\gamma(0) = x$ with velocity $\dot{\gamma}(0) = v$ (Kobayashi & Nomizu, 1963). For a general manifold it is hard to compute exponential maps, but for the sphere \mathbb{S}^D the exponential map is straightforward:

$$\exp_x(v) = x \cos \|v\| + \frac{v}{\|v\|} \sin \|v\|, \quad (16)$$

where $\|\cdot\|$ is the Euclidean norm.

Parameterizing a general cost-convex scalar field $\phi(x)$ remains non-trivial on \mathbb{S}^D . To solve this, Sei (2013) proposes to construct $\phi(x)$ via a convex combination of simple scalar functions $\phi_i : [0, \pi] \rightarrow \mathbb{R}$ that satisfy $\phi_i'(0) = \phi_i'(\pi) = 0$ and $\phi_i'' > -1$. Specifically, $\phi(x)$ is given by

$$\phi(x) = \sum_i \alpha_i \phi_i(d(x, \mu_i)), \quad (17)$$

where $d(x, \mu_i)$ is the geodesic distance between $x \in \mathbb{S}^D$ and $\mu_i \in \mathbb{S}^D$, $\alpha_i \geq 0$ and $\sum_i \alpha_i \leq 1$. The exponential map of the gradient field of $\phi(x)$ is guaranteed to be a diffeomorphism from \mathbb{S}^D to \mathbb{S}^D (Sei, 2013, Lemma 4).

We found that it is challenging to learn concentrated multimodal densities on \mathbb{S}^D using the polynomial and high-frequency scalar fields proposed by Sei (2013). To address this limitation, we introduce a scalar field which is inspired

by radial flows (Tabak & Turner, 2013; Rezende & Mohamed, 2015) and is given by

$$\phi(x) = \sum_i \frac{\alpha_i}{\beta_i} e^{\beta_i(x^T \mu_i - 1)}, \quad (18)$$

where $\alpha_i \geq 0$, $\sum_i \alpha_i \leq 1$, $\mu_i \in \mathbb{S}^D$ and $\beta_i > 0$. It is straightforward to see that the functions

$$\phi_i(d(x, \mu_i)) = \frac{1}{\beta_i} e^{\beta_i(\cos d(x, \mu_i) - 1)} = \frac{1}{\beta_i} e^{\beta_i(x^T \mu_i - 1)}$$

satisfy the required conditions, therefore the map

$$\begin{aligned} f(x) &= \exp_x \nabla \phi(x) \\ &= x \cos \|\nabla \phi(x)\| + \frac{\nabla \phi(x)}{\|\nabla \phi(x)\|} \sin \|\nabla \phi(x)\| \end{aligned} \quad (19)$$

is a diffeomorphism from \mathbb{S}^D to \mathbb{S}^D .

As shown in Appendix A, the density update due to f is

$$p(f(x)) = \frac{\pi(x)}{\sqrt{\det(E(x)^T J(x)^T J(x) E(x))}}, \quad (20)$$

where $J(x)$ is the Jacobian of f at x when f is seen as a map from \mathbb{R}^{D+1} to \mathbb{R}^{D+1} , and the columns of $E(x)$ form an orthonormal basis of vectors tangent to the sphere at x . Unlike the recursive flow which admits an efficient density update, computing the above density update costs $\mathcal{O}(D^3)$, so the exponential-map flow is only practical for small D .

3. Related Work

The study of distributions on objects such as angles, axes and directions has a long history, and is known as *directional statistics* (Mardia & Jupp, 2009). According to the taxonomy of Navarro et al. (2017), directional statistics traditionally uses three approaches for defining distributions on tori and spheres: *wrapping*, *projecting*, and *conditioning*.

Wrapping is typically used to define distributions on \mathbb{S}^1 . The idea is to start with a distribution $p(x)$ on \mathbb{R} , and wrap it around \mathbb{S}^1 by writing $x = \theta + 2\pi k$ and marginalizing out k . A common example is the *wrapped Gaussian* (see e.g. Jona-Lasinio et al., 2012, Section 2.1). A disadvantage of wrapped distributions is that, unless $p(x)$ has bounded support, they require an infinite sum over k , which is generally not analytically tractable.

Projecting is an alternative approach where the manifold is embedded into \mathbb{R}^M with $M > D$, a distribution $p(x)$ is defined on \mathbb{R}^M , and the probability mass is projected onto the manifold. For example, if the manifold is the sphere \mathbb{S}^D embedded in \mathbb{R}^{D+1} , the probability mass in \mathbb{R}^{D+1} can be projected onto \mathbb{S}^D by writing $x = r u_x$ where $r = \|x\|$ and u_x is a unit vector in the direction of x , and then integrating out r . A common example is the projected Gaussian on

\mathbb{S}^1 (see e.g. Wang & Gelfand, 2013). A difficulty with this approach is that it requires computing an integral over r , which may not be generally tractable.

Conditioning also begins by embedding the manifold into \mathbb{R}^M and defining a distribution $p(x)$ on \mathbb{R}^M (which here can be unnormalized). Then, the distribution on the manifold \mathcal{M} is obtained by conditioning $p(x)$ on $x \in \mathcal{M}$, that is, restricting $p(x)$ on \mathcal{M} and re-normalizing. Common examples when $\mathcal{M} = \mathbb{S}^1$ are the *von Mises* (Wang, 2013, Section 1.3.3) and *generalized von Mises* distributions (Gatto & Jammalamadaka, 2007), which are obtained by taking an isotropic or arbitrary Gaussian (respectively) in \mathbb{R}^2 and conditioning on $\|x\| = 1$. For the case of $\mathcal{M} = \mathbb{S}^D$, a common example is the *von Mises–Fisher* distribution (Fisher, 1953), which is obtained by conditioning an isotropic Gaussian in \mathbb{R}^D on $\|x\| = 1$, and reduces to the von Mises for $D = 1$. The von Mises–Fisher distribution can also be extended to the *Stiefel manifold* of sets of N orthonormal D -dimensional vectors (Downs, 1972), and it can be generalized to the more flexible *Kent* distribution, also known as *Fisher–Bingham*, on \mathbb{S}^D (Kent, 1982). Finally, distributions on \mathbb{T}^D can be obtained by defining $p(x)$ in \mathbb{R}^{2D} and conditioning on $x_i^2 + x_{D+i}^2 = 1$ for all $i = 1, \dots, D$. Examples include the *multivariate von Mises* (Mardia et al., 2008) and *multivariate generalized von Mises* distributions (Navarro et al., 2017), which correspond to $p(x)$ being Gaussian with constrained or arbitrary covariance matrix respectively. A drawback of conditioning is that computing the normalizing constant of the resulting distribution requires integrating $p(x)$ on \mathcal{M} , which is not generally tractable.

In general, the above three strategies lead to distributions with tractable density evaluation and sampling algorithms only in special cases, which typically yield simple distributions with limited flexibility. One approach for increasing the flexibility of such simple distributions is via combining them into *mixtures* (see e.g. Peel et al., 2001; Mardia et al., 2007). Such mixtures could be used as base distributions for the flows on tori and spheres that we present in this work. However, unlike flows whose expressivity increases via composition, mixtures generally require a large number of components to represent sharp and complex distributions, and can be harder to fit in practice.

A general method for defining distributions on *Lie groups*, which are groups with manifold structure, was proposed by Falorsi et al. (2019). The tangent space of a D -dimensional Lie group at the identity element is isomorphic to \mathbb{R}^D and is known as the *Lie algebra* of the group. The method is to define a distribution $p(x)$ on the Lie algebra (e.g. using standard Euclidean flows), and then map the Lie algebra onto the group using the exponential map. If the Lie group is a compact connected manifold, the exponential map is surjective but not injective, and the method can be thought

of as a generalization of *wrapping*. When the Lie group is $U(1) \cong \mathbb{S}^1$, we recover wrapped distributions on \mathbb{S}^1 , since in this case the exponential map is $x \mapsto (\cos x, \sin x)$. When the Lie group is \mathbb{T}^D , the exponential map wraps \mathbb{R}^D around the torus along each dimension. This involves an infinite sum, but even if the sum is restricted to be finite (e.g. by bounding the support of $p(x)$), the number of terms to be summed scales exponentially with D . In addition to the above, Falorsi et al. (2019) provide concrete examples for $SO(3)$ and $SE(3)$. The main drawback of this method is that the exponential map cannot be easily composed with other maps (since it is not generally bijective), and can be expensive to compute in high-dimensions.

Finally, a general method for defining normalizing flows on a Riemannian manifold \mathcal{M} was proposed by Gemici et al. (2016). This method relies on the existence of two maps: an embedding $g : \mathcal{M} \rightarrow \mathbb{R}^M$ and a coordinate chart $\phi : \mathcal{M} \rightarrow \mathbb{R}^D$, where $D \leq M$. The idea is to apply the usual Euclidean flows on \mathbb{R}^D , and then transform the density onto the (embedded) manifold via the map $g \circ \phi^{-1}$. The density update associated with this transformation is $\sqrt{\det J^T J}$, where J is the Jacobian of $g \circ \phi^{-1}$ (as shown in Appendix A). This method leads to composable transformations, but is better suited for manifolds that are homeomorphic to \mathbb{R}^D (so that ϕ is well-defined). Compact manifolds such as tori and spheres are not homeomorphic to \mathbb{R}^D , hence ϕ cannot be defined everywhere on \mathcal{M} . In such cases, the transformed density may not be well-behaved, which can create numerical instabilities during training.

4. Experiments

For Möbius transforms, we reparameterized the centre ω , where $\|\omega\| < 1$, in terms of an unconstrained parameter $\omega' \in \mathbb{R}^2$ via $\omega = \frac{0.99}{1 + \|\omega'\|} \omega'$. The constant 0.99 is used to prevent $\|\omega\|$ from being too close to 1. For circular splines, we prevent slopes from becoming smaller than 10^{-3} by adding 10^{-3} to the output of a softplus of unconstrained parameters. For autoregressive flows, the parameters of each conditional transformation are computed by a ReLU-MLP with layer sizes $[N_i, 64, 64, N_o]$, where N_i is the number of inputs and N_o the number of parameters for each transform. All flow models use uniform base distributions.

We evaluate and compare our proposed flows based on their ability to learn sharp, correlated and multi-modal target densities. We used targets $p(x) = e^{-\beta u(x)}/Z$ with inverse temperature β and normalizer Z . We varied β to create targets with different degrees of concentration/sharpness. The models q_η were trained by minimizing the KL divergence

$$\text{KL}(q_\eta \| p) = \mathbb{E}_{q_\eta(x)}[\ln q_\eta(x) + \beta u(x)] + \ln Z. \quad (21)$$

For an additional evaluation of how well the flows match the target, we used $S = 20,000$ samples x_s from the *trained*

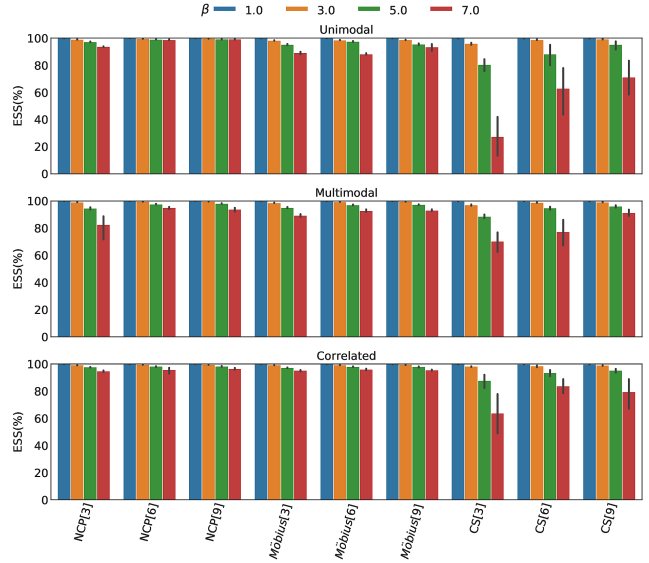


Figure 2. Comparing flows on \mathbb{T}^2 using ESS, for the targets in Figure 3 and various inverse temperatures β . A larger value of β makes the target densities more concentrated and therefore harder to learn. Numbers in square brackets indicate the number of components K used for each transform. All values are averages across 10 runs with different neural-network initialization. See also Figure 8 (appendix) for KL values and further comparisons. **Top row:** Unimodal target. **Middle row:** Multi-modal target with 3 mixture components. **Bottom row:** Correlated target.

flows to estimate the normalizer via importance sampling:

$$Z = \mathbb{E}_{q_\eta(x)} \left[\frac{e^{-\beta u(x)}}{q_\eta(x)} \right] \approx \frac{1}{S} \sum_{s=1}^S w_s, \quad (22)$$

where $w_s = e^{-\beta u(x_s)}/q_\eta(x_s)$. The effective sample size, (Arnaud et al., 2001; Liu, 2008), can be estimated by

$$\text{ESS} = \frac{\text{Var}_{\text{unif}} [e^{-\beta u(x)}]}{\text{Var}_q \left[\frac{e^{-\beta u(x)}}{q_\eta(x)} \right]} \approx \frac{\left(\sum_{s=1}^S w_s \right)^2}{\sum_{s=1}^S w_s^2}. \quad (23)$$

Higher ESS indicates that the flow matches the target better (when reliably estimated). We report ESS as a percentage of the actual sample size.

All models were optimized using Adam (Kingma & Ba, 2015) with learning rate 2×10^{-4} , 20,000 iterations, and batch size 256. The reported error bars are the standard deviation of the average, computed from independent replicas of each experiment with identical hyper-parameters, but with different initialization of the neural network weights. All shown model densities are kernel density estimates using 20,000 samples and Gaussian kernel with bandwidth 0.2.

Figures 2 and 3 show results for \mathbb{T}^2 . The target densities are: a unimodal von Mises, a multi-modal mixture of von

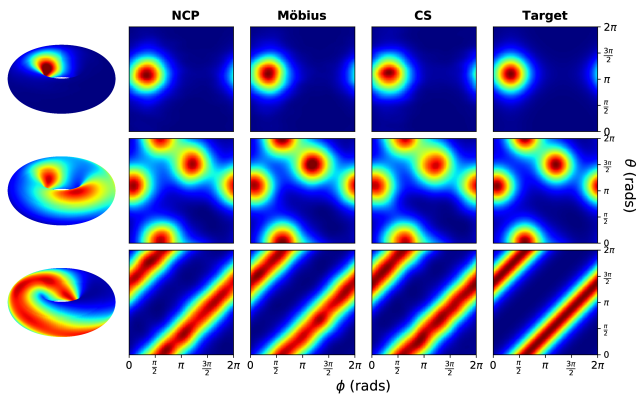


Figure 3. Learned densities on \mathbb{T}^2 using NCP, Möbius and CS flows. Densities shown on the torus are from NCP.

Model	KL [nats]	ESS
MS ($N_T = 1, K_m = 12, K_s = 32$)	0.05 (0.01)	90%
EMP ($N_T = 1$)	0.50 (0.09)	43%
EMSRE ($N_T = 1, K = 12$)	0.82 (0.30)	42%
EMSRE ($N_T = 6, K = 5$)	0.19 (0.05)	75%
EMSRE ($N_T = 24, K = 1$)	0.10 (0.10)	85%

Table 1. Comparing baseline and proposed flows on \mathbb{S}^2 using KL and ESS. The target density is the mixture of 4 modes shown in Figure 4. We compare recursive Möbius-spline flow (MS), exponential-map polynomial flow (EMP) and exponential-map sum-of-radial flow (EMSRE). Brackets show error bars on the KL from 3 replicas of each experiment. N_T is the number of stacked transformations for each flow; K_m is the number of centres used in Möbius; K_s is the number of segments in the spline flow; K is the number of radial components in the radial exponential-map flow. The polynomial scalar field is shown in Appendix E.

Mises, and a correlated von Mises (their precise definitions are in Table 2 of the appendix). The results demonstrate that we can reliably learn multi-modal and correlated densities with different degrees of concentration. Among the circle flows, NCP and Möbius performed the best, whereas CS performed less well for high β . An additional experiment is shown in Appendix I, where flows on \mathbb{T}^6 are used to learn the posterior density over joint angles of a robot arm.

Table 1 shows results for \mathbb{S}^2 , where we compare the recursive flow to the exponential-map flow with polynomial and radial scalar fields. The recursive flow uses Möbius for the circle and splines for the interval. For the exponential-map flow, the polynomial scalar field was proposed by Sei (2013) and is shown in Appendix E, whereas the radial scalar field is the new one we propose in Equation (18). The target density is the mixture of 4 modes shown in Figure 4. We demonstrate substantial improvement relative to the polynomial scalar field, but we find that the exponential-map

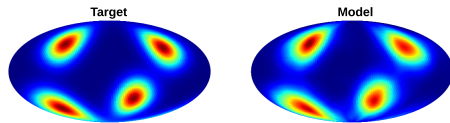


Figure 4. Learned multi-modal density on \mathbb{S}^2 using exponential-map flows, using the Mollweide projection for visualization. The model is a composition of 24 exponential-map transforms, using the radial scalar field with 1 component.

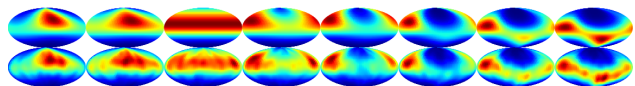


Figure 5. Learned multi-modal density on $SU(2) \cong \mathbb{S}^3$ using the recursive flow. Each column shows an \mathbb{S}^2 slice of the \mathbb{S}^3 density along a fixed axis using the Mollweide projection. **Top row:** target density. **Bottom row:** learned density. We used a Möbius transform with $K_m = 32$ for the circle, and spline transforms with $K_s = 64$ for the two intervals (ESS = 84%, KL = 0.14).

flow with either scalar field is not yet competitive with the recursive flow. Figure 5 shows an example of learning a density on $SU(2) \cong \mathbb{S}^3$ using the recursive flow. Finally, Appendix J shows an example of training a recursive flow (using splines for both the circle and the interval) on data sampled from a ‘map of the world’ density on \mathbb{S}^2 .

5. Discussion

This work shows how to construct flexible normalizing flows on tori and spheres of any dimension in a numerically stable manner. Unlike many of the distributions traditionally used in directional statistics, the proposed flows can be made arbitrarily flexible, but have tractable and exact density evaluation and sampling algorithms. We conclude with a comparison of the proposed models, a discussion of their limitations, and some preliminary thoughts on how to extend flows to other manifolds of interest to fundamental physics.

5.1. Comparison, Scope and Limitations

Among the flows on the circle, Möbius and NCP performed the best, with CS performing less well for highly concentrated target densities. However, increasing the expressivity of Möbius and NCP required convex combinations, whereas CS can be made more expressive by adding more spline segments. As a result, CS is the cheapest to invert (it can be done analytically), whereas Möbius and NCP (with more than one component) require a root-finding algorithm such as bisection search. Therefore, in practice it may be preferable to use CS if both density evaluation and sampling are required, and use Möbius or NCP otherwise.

On \mathbb{S}^D , the recursive flow performed better than the

exponential-map flow. In addition, the recursive flow scales better to high dimensions since its density can be computed efficiently, whereas the density of the exponential-map flow has a computational cost of $\mathcal{O}(D^3)$. The theoretical advantage of the exponential-map flow is that it is intrinsic to the sphere, but this advantage did not result in a practical benefit in our experiments.

5.2. Towards Normalizing Flows on $SU(D)$ and $U(D)$

The unitary Lie groups $U(1)$, $SU(2)$ and $SU(3)$ are of particular interest to fundamental physics, because the symmetry groups of particle interactions are constructed from them (Woit, 2017). We have shown that we can construct expressive flows on $U(1) \cong S^1$ and $SU(2) \cong S^3$.

Our recursive construction for S^D provides a starting point for flows on $SU(D)$ for $D \geq 3$ and $U(D)$ for $D \geq 2$, via recursively building $SU(D)$ from $SU(D-1)$ and $U(1)^{2D-1}$, and $U(D)$ from $SU(D)$ and extra angles. These decompositions provide coordinate systems on $SU(D)$ and $U(D)$ in terms of Euler angles as shown by Tilma & Sudarshan (2002; 2004). Working with explicit coordinate systems on $SU(D)$ and $U(D)$ is more involved than on spheres; figuring out the necessary details such as appropriate boundary conditions for each coordinate/angle is a direction we plan to explore in the future in order to build flows on a wider variety of compact connected manifolds.

Acknowledgements

We thank Heiko Strathmann and Alex Botev for discussions and feedback. PES is partially supported by the National Science Foundation under CAREER Award 1841699 and PES and GK are partially supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics under grant Contract Number DE-SC0011090. KC is supported by the National Science Foundation under the awards ACI-1450310, OAC-1836650, and OAC-1841471 and by the Moore-Sloan data science environment at NYU.

References

- Arnaud, D., de Freitas, N., and Gordon, N. *Sequential Monte Carlo methods in practice*. Information Science and Statistics. Springer New York, 2001.
- Boomsma, W., Mardia, K. V., Taylor, C. C., Ferkinghoff-Borg, J., Krogh, A., and Hamelryck, T. A generative, probabilistic model of local protein structure. *Proceedings of the National Academy of Sciences*, 105(26):8932–8937, 2008.
- Davidson, T. R., Falorsi, L., De Cao, N., Kipf, T., and Tomczak, J. M. Hyperspherical variational auto-encoders. *Conference on Uncertainty in Artificial Intelligence*, 2018.
- Dinh, L., Sohl-Dickstein, J., and Bengio, S. Density estimation using Real NVP. *International Conference on Learning Representations*, 2017.
- Downs, T. D. Orientation statistics. *Biometrika*, 59(3): 665–676, 1972.
- Downs, T. D. and Mardia, K. V. Circular regression. *Biometrika*, 89(3):683–698, 2002.
- Durkan, C., Bekasov, A., Murray, I., and Papamakarios, G. Cubic-spline flows. *ICML workshop on Invertible Neural Networks and Normalizing Flows*, 2019a.
- Durkan, C., Bekasov, A., Murray, I., and Papamakarios, G. Neural spline flows. *Advances in Neural Information Processing Systems*, 2019b.
- Falorsi, L., de Haan, P., Davidson, T. R., De Cao, N., Weiler, M., Forré, P., and Cohen, T. S. Explorations in homeomorphic variational auto-encoding. *ICML workshop on Theoretical Foundations and Applications of Deep Generative Models*, 2018.
- Falorsi, L., de Haan, P., Davidson, T. R., and Forré, P. Reparameterizing distributions on Lie groups. *International Conference on Artificial Intelligence and Statistics*, 2019.
- Feiten, W., Lang, M., and Hirche, S. Rigid motion estimation using mixtures of projected Gaussians. *International Conference on Information Fusion*, 2013.
- Fisher, R. Dispersion on a sphere. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 217(1130):295–305, 1953.
- Gatto, R. and Jammalamadaka, S. R. The generalized von Mises distribution. *Statistical Methodology*, 4(3):341–353, 2007.
- Gemici, M. C., Rezende, D. J., and Mohamed, S. Normalizing flows on Riemannian manifolds. *arXiv preprint arXiv:1611.02304*, 2016.
- Hamelryck, T., Kent, J. T., and Krogh, A. Sampling realistic protein conformations using local structural bias. *PLoS Computational Biology*, 2(9), 2006.
- Jakob, W. Numerically stable sampling of the von Mises Fisher distribution on S^2 (and other tricks). Technical report, Interactive Geometry Lab, ETH Zürich, 2012.
- Jona-Lasinio, G., Gelfand, A., and Jona-Lasinio, M. Spatial analysis of wave direction data using wrapped Gaussian processes. *The Annals of Applied Statistics*, 6(4):1478–1498, 2012.
- Kato, S. and McCullagh, P. Möbius transformation

- and a Cauchy family on the sphere. *arXiv preprint arXiv:1510.07679*, 2015.
- Kato, S., Shimizu, K., and Shieh, G. S. A circular–circular regression model. *Statistica Sinica*, 18:633–645, 2008.
- Kent, J. T. The Fisher–Bingham distribution on the sphere. *Journal of the Royal Statistical Society. Series B (Methodological)*, 44(1):71–80, 1982.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *International Conference for Learning Representations*, 2015.
- Kingma, D. P., Salimans, T., Jozefowicz, R., Chen, X., Sutskever, I., and Welling, M. Improved variational inference with inverse autoregressive flow. *Advances in Neural Information Processing Systems*, 2016.
- Kobayashi, S. and Nomizu, K. *Foundations of differential geometry*, volume 1. Interscience Publishers, 1963.
- Liu, J. S. *Monte Carlo strategies in scientific computing*. Springer Science & Business Media, 2008.
- Mardia, K. V. and Jupp, P. E. *Directional statistics*. Wiley Series in Probability and Statistics. John Wiley & Sons, 2009.
- Mardia, K. V., Taylor, C. C., and Subramaniam, G. K. Protein bioinformatics and mixtures of bivariate von Mises distributions for angular data. *Biometrics*, 63(2):505–512, 2007.
- Mardia, K. V., Hughes, G., Taylor, C. C., and Singh, H. A multivariate von Mises distribution with applications to bioinformatics. *The Canadian Journal of Statistics*, 36(1):99–109, 2008.
- Mathieu, E., Le Lan, C., Maddison, C. J., Tomioka, R., and Teh, Y. W. Continuous hierarchical representations with Poincaré variational auto-encoders. *Advances in Neural Information Processing Systems*, 2019.
- Müller, T., McWilliams, B., Rousselle, F., Gross, M., and Novák, J. Neural importance sampling. *ACM Transactions on Graphics*, 38(5):1–19, 2019.
- Navarro, A. K. W., Frellsen, J., and Turner, R. E. The multivariate generalised von Mises distribution: Inference and applications. *AAAI Conference on Artificial Intelligence*, 2017.
- Papamakarios, G., Pavlakou, T., and Murray, I. Masked autoregressive flow for density estimation. *Advances in Neural Information Processing Systems*, 2017.
- Papamakarios, G., Nalisnick, E., Rezende, D. J., Mohamed, S., and Lakshminarayanan, B. Normalizing flows for probabilistic modeling and inference. *arXiv preprint arXiv:1912.02762*, 2019.
- Peel, D., Whiten, W. J., and McLachlan, G. J. Fitting mixtures of Kent distributions to aid in joint set identification. *Journal of the American Statistical Association*, 96(453):56–63, 2001.
- Rezende, D. J. and Mohamed, S. Variational inference with normalizing flows. *International Conference on Machine Learning*, 2015.
- Sei, T. A Jacobian inequality for gradient maps on the sphere and its application to directional statistics. *Communications in Statistics—Theory and Methods*, 42(14):2525–2542, 2013.
- Senanayake, R. and Ramos, F. Directional grid maps: modeling multimodal angular uncertainty in dynamic environments. *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2018.
- Shapovalov, M. V. and Dunbrack Jr, R. L. A smoothed backbone-dependent rotamer library for proteins derived from adaptive kernel density estimates and regressions. *Structure*, 19(6):844–858, 2011.
- Tabak, E. G. and Turner, C. V. A family of nonparametric density estimation algorithms. *Communications on Pure and Applied Mathematics*, 66(2):145–164, 2013.
- Tilma, T. and Sudarshan, E. Generalized Euler angle parametrization for $SU(N)$. *Journal of Physics A: Mathematical and General*, 35(48):10467, 2002.
- Tilma, T. and Sudarshan, E. Generalized Euler angle parametrization for $U(N)$ with applications to $SU(N)$ coset volume measures. *Journal of Geometry and Physics*, 52(3):263–283, 2004.
- Wang, F. and Gelfand, A. E. Directional data analysis under the general projected normal distribution. *Statistical Methodology*, 10(1):113–127, 2013.
- Wang, M. *Extensions of probability distributions on torus, cylinder and disc*. PhD thesis, Graduate School of Science and Technology, Keio University, 2013.
- Wang, P. Z. and Wang, W. Y. Riemannian normalizing flow on variational Wasserstein autoencoder for text modeling. *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2019.
- Wang, Z., Cheng, S., Li, Y., Zhu, J., and Zhang, B. A Wasserstein minimum velocity approach to learning unnormalized models. *Symposium on Advances in Approximate Bayesian Inference*, 2019.
- Woit, P. *Quantum theory, groups and representations: An introduction*. Springer, 2017.