Numerically stable sampling of the von Mises Fisher distribution on $\mathcal{S}^2$ (and other tricks)

Wenzel Jakob

July 21, 2012 (revised Jun 28, 2015)

1 Introduction

The *von Mises-Fisher* (vMF) distribution is a popular all-purpose distribution [1, 4] for statistical inference involving directional data. On the 2-sphere, it is defined as

$$f_{\mathrm{vMF}}(\omega) = \frac{\kappa}{4\pi \sinh \kappa} \exp(\kappa \mu^T \omega)$$

where $\mu \in S^2$ is the mean direction and κ denotes the *concentration* parameter ($\kappa \to 0$ approaching the uniform distribution). A recent application [2] of this distribution in computer graphics entailed fitting mixture models composed of vMF functions to arbitrary spherical data using the expectation maximization procedure.

Unfortunately, many basic operations involving this distribution are prone to severe numerical issues when implemented in finite precision computer arithmetic. There is a surprising lack on information on how these can be circumvented, and hence the purpose of this document is to serve as a collection of numerically-well behaved recipes for common operations.

2 Evaluation

Evaluation of the vMF distribution easily overflows single precision arithmetic even for moderate concentration values (for instance, $\sinh 100 = 1.34406 \cdot 10^{43}$), and double precision fails shortly thereafter. The following equivalent expression follows from exponential function identities and works reliably over a much larger range of concentrations:

$$f_{\mathrm{vMF}}(\omega) = \begin{cases} \dfrac{1}{4\pi}, & \kappa = 0 \\ \dfrac{\kappa}{2\pi(1 - \exp(-2\kappa))}\, e^{\kappa(\mu^T\omega - 1)}, & \kappa > 0 \end{cases}$$

3 Sample generation

Several prior works have investigated how independent samples can be drawn so that they are distributed according to the vMF distribution [5, 6, 3]. The following is a brief summary of [3], which leads to a simple but numerically ill-behaved method:

Observe that the following random vector with mean direction $\mu = (0, 0, 1)$ is distributed according to f_{vMF} [5]:

$$\omega_\kappa = (\sqrt{1 - W^2}\, V, W)^T$$

where V and W are independent random variables, $V \in \mathbb{R}^2$ is a uniformly distributed vector on the unit circle, and $W \in [-1, 1]$ follows the density

$$f_W(w) = \frac{\kappa}{2 \sinh \kappa} \exp(\kappa w).$$

All that is needed for a computer implementation is a way to generate realizations of W. Applying the inversion method results in

$$F_W^{-1}(\xi) = \kappa^{-1} \log \left(\exp^{-\kappa} + 2\, \xi \sinh \kappa \right)$$

To handle other values of μ, one can simply apply a rotation to directions obtained in this manner.

3.1 Numerically stable variant

Again, we can apply exponential function identities to arrive at an equivalent and numerically well-behaved expression, which avoids overflow for large values of κ:

$$F_W^{-1}(\xi) = 1 + \kappa^{-1} \log \left(\xi + (1 - \xi) e^{-2\kappa} \right)$$

4 Finding κ such that $f_{\text{vMF}}(\mu) = c$

One very useful tool is the ability to create distributions that have a specified solid angle density into a certain direction. In the case of the van Mises-Fisher distribution, we can see that f_{vMF} takes on its maximum into direction μ, where

$$g(\kappa) := \frac{\kappa}{4\pi}(1 + \coth \kappa).$$

gives the maximum as a function of the concentration. Unfortunately, it is not convenient to analytically invert this expression for κ. However, note that

$$\coth \kappa = \frac{e^{2\kappa} + 1}{e^{2\kappa} - 1}$$

rapidly approaches 1. For instance, $\coth 5$ is already approximately equal to 1.0009. Assuming that there are no particularly stringent accuracy requirements on the inversion, we can use the following approximate scheme:

$$g^{-1}(x) \approx \begin{cases} 2\pi x, & x > g(5) \approx 0.795 \\ g_{\text{rat}}^{-1}(x), & \text{otherwise} \end{cases}$$

where we have approximated $\coth \kappa \approx 1$ for $\kappa > 5$ and make use the following rational interpolant elsewhere:

$$g_{\text{rat}}^{-1}(x) := \max \left\{ 10^{-5}, \frac{168.479x^2 + 16.4585x - 2.39942}{-1.12718x^2 + 29.1433x + 1} \right\}.$$

On the interval $[1/4\pi, g(5)]$, the function g_{rat}^{-1} has an absolute error of < 0.007 The relative error is infinite, as $g^{-1}(x) \to 0$ $(x \to 1/4\pi)$. Figure 1 shows an illustration of the fit.

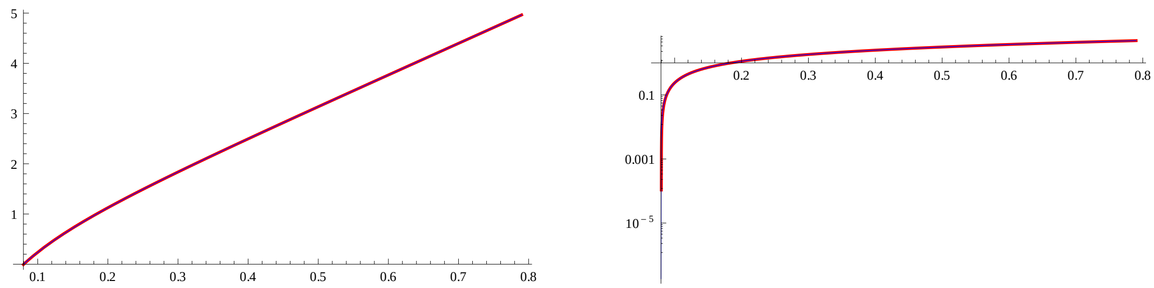

(a) Linear plot

(b) Logarithmic plot

Figure 1: Fit of the rational function g_{rat}^{-1} (blue) to g^{-1} (red).

5 Convolution

The convolution of two vMF distributions does not generally produce another vMF distribution. However, the result of this operation can be well-approximated by a vMF distribution with a suitably chosen value of κ. Mardia and Jupp [4] describe one approach to obtain this parameter, which entails approximating the distributions to be convolved by wrapped normal distributions, convolving them instead, and transforming the result back into a vMF distribution. A C implementation of this is given below:

```
float A3(float kappa) {
    return 1 / std::tanh(kappa) - 1 / kappa;
}

float dA3(float kappa) {
    float csch = 2.0f / (std::exp(kappa) - std::exp(-kappa));
    return 1 / (kappa*kappa) - csch*csch;
}

float A3inv(float y, float guess) {
    /* Initial guess */
    float x = guess, residual = 0;

    /* Invert using Newton's method */
    do {
        residual = A3(x)-y, deriv = dA3(x);
        x -= residual/deriv;
    } while (std::abs(residual) > 1e-5f);
    return x;
}

float convolve(float kappa1, float kappa2) {
    return A3inv(A3(kappa1) * A3(kappa2), std::min(kappa1, kappa2));
}
```

References

[1] Fisher, N., Lewis, T., and Embleton, B. *Statistical analysis of spherical data*. Cambridge University Press, 1993.

[2] Han, C., Sun, B., Ramamoorthi, R., and Grinspun, E. Frequency domain normal map filtering. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2007) 26*, 3 (2007), 28:1–28:12.

[3] Jung, S. Generating von mises fisher distribution on the unit sphere (s^2). Tech. rep., University of Pittsburgh, 10 2009.

[4] Mardia, K., and Jupp, P. *Directional statistics*. Wiley, 2000.

[5] Ulrich, G. Computer generation of distributions on the m-sphere. *Applied Statistics* (1984), 158–163.

[6] Wood, A. Simulation of the von mises fisher distribution. *Communications in statistics-simulation and computation 23*, 1 (1994), 157–164.