

Appendix

This appendix supplements the article “*Exploring Large Action Sets with Hyperspherical Embeddings using von Mises-Fisher Sampling*” and is organized as follows:

- Appendices A, B, C, and D provide detailed proofs and discussions for all theoretical results presented in the paper.
- Appendix E explores the connection between vMF-exp and Thompson Sampling.
- Appendix F explains practical methods for sampling from the vMF distribution.
- Appendix G presents the complete results of our Monte Carlo simulations.
- Appendix H details an additional experimental study using a large-scale, publicly available GloVe dataset.
- Appendix I highlights the successful large-scale deployment of vMF-exp in the private production system of the global music streaming service Deezer for large-scale music recommendation.

A. Asymptotic Behavior of Boltzmann Exploration (Proof of Proposition 4.2)

We begin with the proof of Proposition 4.2 claiming that, in the setting of Section 4.1, we have:

$$P_{\text{B-exp}}(a \mid n, d, V, \kappa) = \underbrace{\frac{f_{\text{vMF}}(A \mid V, \kappa) \mathcal{A}(\mathcal{S}^{d-1})}{n}}_{\text{denoted } P_0(a \mid n, d, V, \kappa)} + o\left(\frac{1}{n\sqrt{n}}\right), \quad (22)$$

with f_{vMF} the probability density function (PDF) of the von Mises-Fisher (vMF) (Fisher, 1953) distribution:

$$\forall A \in \mathcal{S}^{d-1}, f_{\text{vMF}}(A \mid V, \kappa) = C_d(\kappa) e^{\kappa \langle V, A \rangle}, \quad (23)$$

where $\mathcal{A}(\mathcal{S}^{d-1})$ is the surface area of \mathcal{S}^{d-1} , the d -dimensional unit hypersphere, and $C_d(\kappa)$ is the normalizing constant.

Proof. By definition,

$$\begin{aligned} P_{\text{B-exp}}(a \mid n, d, V, \kappa) &= \mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\frac{e^{\kappa \langle V, A \rangle}}{e^{\kappa \langle V, A \rangle} + \sum_{i=1}^n e^{\kappa \langle V, X_i \rangle}} \right] \\ &= \frac{e^{\kappa \langle V, A \rangle}}{n} \mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\frac{1}{\frac{e^{\kappa \langle V, A \rangle}}{n} + \sum_{i=1}^n \frac{e^{\kappa \langle V, X_i \rangle}}{n}} \right] \\ &= \frac{e^{\kappa \langle V, A \rangle}}{n} \mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\frac{1}{D_n} \right]. \end{aligned} \quad (24)$$

We use D_n to denote the denominator of the expression inside the above expectation. D_n is the empirical average of n independent and identically distributed (i.i.d.) random variables (plus a constant). Therefore, by applying the *Central Limit Theorem (CLT)* (Fischer, 2011), we know that as n grows it will be asymptotically distributed according to a Normal distribution with the following expectation:

$$\begin{aligned} \mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} [D_n] &= \mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\frac{e^{\kappa \langle V, A \rangle}}{n} + \sum_{i=1}^n \frac{e^{\kappa \langle V, X_i \rangle}}{n} \right] \\ &= \frac{e^{\kappa \langle V, A \rangle}}{n} + \mathbb{E}_{X \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[e^{\kappa \langle V, X \rangle} \right]. \end{aligned} \quad (25)$$

Moreover, we have:

$$\begin{aligned} \mathbb{E}_{X \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[e^{\kappa \langle V, X \rangle} \right] &= \int_{X \in \mathcal{S}^{d-1}} \frac{e^{\kappa \langle V, X \rangle}}{\mathcal{A}(\mathcal{S}^{d-1})} dX \\ &= \frac{1}{\mathcal{A}(\mathcal{S}^{d-1}) C_d(\kappa)}, \end{aligned} \quad (26)$$

using the fact that $C_d(\kappa)$ is the normalizing constant of a vMF distribution, ensuring that its PDF (Equation (23)) sums to 1 when integrated on the unit hypersphere.

Let us define $\sigma = \text{Var}_{X \sim \mathcal{U}(\mathcal{S}^{d-1})} [e^{\kappa \langle V, X \rangle}]$. Although we do not need an explicit expression for σ , we know it is finite. Additionally, let $g : x \mapsto \frac{1}{x}$ be the inverse function. The CLT ensures that:

$$\sqrt{n} \left[D_n - \frac{1}{\mathcal{A}(\mathcal{S}^{d-1})C_d(\kappa)} \right] \xrightarrow{D} \mathcal{N}(0, \sigma^2), \quad (27)$$

where \xrightarrow{D} denotes convergence in distribution (Jacod & Protter, 2004). Moreover, since g is a differentiable function on \mathbb{R}_+^* , we use the *Delta method* (Oehlert, 1992) to infer that:

$$\sqrt{n} [g(D_n) - g(\frac{1}{\mathcal{A}(\mathcal{S}^{d-1})C_d(\kappa)})] \xrightarrow{D} \mathcal{N}(0, \sigma^2 [g'(\frac{1}{\mathcal{A}(\mathcal{S}^{d-1})C_d(\kappa)})]^2). \quad (28)$$

Replacing g and g' by their respective values, we obtain:

$$\sqrt{n} \left[\frac{1}{D_n} - C_d(\kappa) \mathcal{A}(\mathcal{S}^{d-1}) \right] \xrightarrow{D} \mathcal{N}(0, \sigma^2 (\mathcal{A}(\mathcal{S}^{d-1})C_d(\kappa))^4). \quad (29)$$

Furthermore, recall that if a sequence Z_1, Z_2, \dots of random variables converges in distribution to a random variable Z , then for all bounded continuous function ϕ , $\lim_{n \rightarrow +\infty} \mathbb{E}[\phi(Z_n)] = \mathbb{E}[\phi(Z)]$ (Jacod & Protter, 2004). Since for every n the random variable $Z_n = \sqrt{n}[\frac{1}{D_n} - C_d(\kappa) \mathcal{A}(\mathcal{S}^{d-1})]$ has bounded values, we can simply chose the identity function for ϕ to conclude that :

$$\lim_{n \rightarrow +\infty} \mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\sqrt{n} \left[\frac{1}{D_n} - C_d(\kappa) \mathcal{A}(\mathcal{S}^{d-1}) \right] \right] = 0, \quad (30)$$

which is equivalent to:

$$\mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\frac{1}{D_n} \right] = C_d(\kappa) \mathcal{A}(\mathcal{S}^{d-1}) + o\left(\frac{1}{\sqrt{n}}\right). \quad (31)$$

Finally, by multiplying Equation (31) by $\frac{e^{\kappa \langle V, A \rangle}}{n}$, we obtain Equation (22), concluding the proof. \square

B. Asymptotic Behavior of vMF Exploration in $d = 2$ dimensions (Proof of Proposition 4.3, Part 1)

We now prove Proposition 4.3 when $d = 2$. In 2 dimensions, the vMF distribution takes the special form of the von Mises (vM) distribution (Mardia & Jupp, 2009) which, instead of describing the distribution of the dot product between V and \tilde{V} , describes the distribution of their angle θ . The PDF of a von Mises distribution is defined as follows:

$$\forall \theta \in [-\pi, \pi], f_{\text{vM}}(\theta | \kappa) = \frac{e^{\kappa \cos(\theta)}}{2\pi I_0(\kappa)}. \quad (32)$$

Let us define θ_0 as the angle between V and A . In this section, we prove that:

$$P_{\text{vMF-exp}}(A | n, d = 2, \kappa) = \frac{e^{\kappa \cos(\theta_0)}}{n I_0(\kappa)} + \mathcal{O}\left(\frac{1}{n^2}\right). \quad (33)$$

Proof. By definition,

$$P_{\text{vMF-exp}}(A | n, d = 2, \kappa) = \mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^1)} \left[\mathbb{P}(\tilde{V} \in \mathcal{S}_{\text{Voronoi}}(A | \mathcal{X}_{n+1})) \right], \quad (34)$$

where $\mathcal{S}_{\text{Voronoi}}(X_i | \mathcal{X}_n) = \{\tilde{V} \in \mathcal{S}^{d-1}, \forall j \in \mathcal{I}_n, \langle \tilde{V}, X_i \rangle \geq \langle \tilde{V}, X_j \rangle\}$. Let us call $\mathcal{Y}_n = \{Y_i\}$ the result of the permutation of the indices of \mathcal{X}_n such that the (signed) angles β_i between A and Y_i are sorted in increasing order. Since the $\{X_i\}$ are i.i.d. and uniformly distributed on the circle, then the angles between A and the $\{X_i\}$ are i.i.d. and uniformly distributed on $[0, 2\pi]$. Therefore, the set $\{\beta_i\}$ is the set of the order statistics of n i.i.d. random variables uniformly distributed on $[0, 2\pi]$. Consequently, the set $\{\frac{\beta_i}{2\pi}\}$ is the set of the order statistics of n i.i.d. random variables uniformly distributed on $[0, 1]$, which are known to follow Beta distributions (Gentle, 2009) defined as follows:

$$\forall 1 \leq i \leq n, \frac{\beta_i}{2\pi} \sim \text{Beta}(i, n+1-i). \quad (35)$$

As a consequence, we have:

$$\mathbb{E}[\beta_1] = \frac{2\pi}{n+1}, \quad (36)$$

$$\mathbb{E}[\beta_n] = \frac{2\pi n}{n+1}, \quad (37)$$

$$\text{Var}[\beta_1] = \text{Var}[\beta_n] = \frac{4\pi^2 n}{(n+1)^2(n+2)}. \quad (38)$$

Moreover, for given values of Y_i , we can see from Figure 4 that, in 2 dimensions, Voronoi cells are arcs of the circle and are delimited by perpendicular bisectors of two neighboring points. Specifically, the Voronoi cell of A is delimited by the perpendicular bisector of A and Y_1 on one side, and the perpendicular bisector of A and Y_n on the other side. By denoting θ the (signed) angle between V and \tilde{V} , we have:

$$\begin{aligned} \mathbb{P}\left(\tilde{V} \in \mathcal{S}_{\text{Voronoi}}(A | \mathcal{X}_{n+1})\right) &= \mathbb{P}\left(\theta \in \left[\theta_0 + \frac{\beta_n - 2\pi}{2}, \theta_0 + \frac{\beta_1}{2}\right] \mid \theta \sim \text{vM}(0, \kappa), \beta_1, \beta_n\right) \\ &= \int_{\theta=\theta_0 + \frac{\beta_n - 2\pi}{2}}^{\theta_0 + \frac{\beta_1}{2}} f_{\text{vM}}(\theta | \kappa) d\theta. \end{aligned} \quad (39)$$

Therefore:

$$P_{\text{vMF-exp}}(A | n, d = 2, \kappa) = \mathbb{E}_{\beta_1, \beta_n} \left[\int_{\theta=\theta_0 + \frac{\beta_n - 2\pi}{2}}^{\theta_0 + \frac{\beta_1}{2}} f_{\text{vM}}(\theta | \kappa) d\theta \right]. \quad (40)$$

To get an asymptotic expression of the probability that θ lies between the considered bounds, we can first notice that as n grows, β_1 will approach 0 and β_n will approach 2π . This means that the integral we need to compute will have very narrow

B.1. Zero-Order Estimate

Let us study the zero-order approximation of $f_{\text{VM}}(\theta \mid \kappa)$ near θ_0 :

$$\begin{aligned}
 \mathbb{E}_{\beta_1, \beta_n} \left[\int_{\theta=\theta_0+\frac{\beta_n-2\pi}{2}}^{\theta_0+\frac{\beta_1}{2}} f_{\text{VM}}(\theta \mid \kappa) d\theta \right] &= \mathbb{E}_{\beta_1, \beta_n} \left[f_{\text{VM}}(\theta_0 \mid \kappa) \left(\theta_0 + \frac{\beta_1}{2} - \left(\theta_0 + \frac{\beta_n - 2\pi}{2} \right) \right) \right] \\
 &= \mathbb{E}_{\beta_1, \beta_n} \left[f_{\text{VM}}(\theta_0 \mid \kappa) \left(\pi - \frac{\beta_n - \beta_1}{2} \right) \right] \\
 &= \pi f_{\text{VM}}(\theta_0 \mid \kappa) \mathbb{E}_{\beta_1, \beta_n} \left[1 - \frac{\beta_n - \beta_1}{2\pi} \right] \\
 &= \frac{e^{\kappa \cos(\theta_0)}}{2I_0(\kappa)} \left(1 - \mathbb{E}_{\beta_1, \beta_n} \left[\frac{\beta_n}{2\pi} \right] + \mathbb{E}_{\beta_1, \beta_n} \left[\frac{\beta_1}{2\pi} \right] \right) \\
 &= \frac{e^{\kappa \cos(\theta_0)}}{2I_0(\kappa)} \frac{n+1-n+1}{n+1} \\
 &= \frac{e^{\kappa \cos(\theta_0)}}{2I_0(\kappa)} \frac{2}{n+1} \\
 &= \frac{e^{\kappa \cos(\theta_0)}}{(n+1)I_0(\kappa)} \\
 &= \frac{e^{\kappa \cos(\theta_0)}}{nI_0(\kappa)} - \frac{e^{\kappa \cos(\theta_0)}}{n(n+1)I_0(\kappa)} \\
 &= \frac{e^{\kappa \cos(\theta_0)}}{nI_0(\kappa)} + \mathcal{O}\left(\frac{1}{n^2}\right).
 \end{aligned} \tag{42}$$

This proves that, asymptotically, the contribution of the zero-order term of f_{VM} to the probability of selecting A is equal to the probability of selecting A using B-exp with the same κ value.

To understand how fast vMF-exp reaches its asymptotic behavior, we now need to study $R_0(\theta)$, the remainder of the Taylor series expansion of f_{VM} around θ_0 .

B.2. Bounding of the Remainder Term

We start by computing the first derivative of f_{VM} :

$$\forall \theta \in [0, 2\pi], |f'_{\text{VM}}(\theta \mid \kappa)| = \frac{|\sin(\theta)| \kappa e^{\kappa \cos(\theta)}}{I_0(\kappa)}, \tag{43}$$

which is bounded⁴ on $[0, 2\pi]$ by $M = \frac{\kappa e^\kappa}{I_0(\kappa)}$. According to the Taylor-Lagrange inequality (Abramowitz & Stegun, 1948), this in turn bounds the remainder term as follows:

$$\forall \theta \in [0, 2\pi], |R_0(\theta)| \leq M|\theta - \theta_0|. \tag{44}$$

⁴We note that a tighter bound could be found by studying the second derivative, but will not be necessary for the purpose of this proof.

In particular, this inequality holds for every $\theta \in [\theta_0 + \frac{\beta_n - 2\pi}{2}, \theta_0 + \frac{\beta_1}{2}]$, and so:

$$\begin{aligned}
 \int_{\theta=\theta_0+\frac{\beta_n-2\pi}{2}}^{\theta_0+\frac{\beta_1}{2}} |R_0(\theta)| \, d\theta &\leq \int_{\theta=\theta_0+\frac{\beta_n-2\pi}{2}}^{\theta_0+\frac{\beta_1}{2}} M|\theta - \theta_0| \, d\theta \\
 &= \int_{\theta=\theta_0}^{\theta_0+\frac{\beta_1}{2}} M(\theta - \theta_0) \, d\theta + \int_{\theta=\theta_0+\frac{\beta_n-2\pi}{2}}^{\theta_0} M(\theta_0 - \theta) \, d\theta \\
 &= \int_{\theta=0}^{\frac{\beta_1}{2}} M\theta \, d\theta - \int_{\theta=\frac{\beta_n-2\pi}{2}}^0 M\theta \, d\theta \\
 &= M \frac{\beta_1^2 + (\beta_n - 2\pi)^2}{8}.
 \end{aligned} \tag{45}$$

The above inequality holds when considering the expected values over uniformly distributed X_i :

$$\begin{aligned}
 \mathbb{E}_{\beta_1, \beta_n} \left[\int_{\theta=\theta_0+\frac{\beta_n-2\pi}{2}}^{\theta_0+\frac{\beta_1}{2}} |R_0(\theta)| \, d\theta \right] &\leq M \frac{\mathbb{E}_{\beta_1, \beta_n} [\beta_1^2] + \mathbb{E}_{\beta_1, \beta_n} [(\beta_n - 2\pi)^2]}{8} \\
 &= M \frac{\text{Var}_{\beta_1, \beta_n} [\beta_1] + (\mathbb{E}_{\beta_1, \beta_n} [\beta_1])^2 + \text{Var}_{\beta_1, \beta_n} [(\beta_n - 2\pi)] + (\mathbb{E}_{\beta_1, \beta_n} [\beta_n - 2\pi])^2}{8} \\
 &= \frac{M}{8} \left(\frac{2 \times 4\pi^2 n}{(n+1)^2(n+2)} + \frac{2 \times 4\pi^2}{(n+1)^2} \right) \\
 &= \frac{M\pi^2}{(n+1)(n+2)} \\
 &= \mathcal{O}\left(\frac{1}{n^2}\right).
 \end{aligned} \tag{46}$$

Since $\left| \mathbb{E}_{\beta_1, \beta_n} \left[\int_{\theta=\theta_0+\frac{\beta_n-2\pi}{2}}^{\theta_0+\frac{\beta_1}{2}} R_0(\theta) \, d\theta \right] \right| \leq \mathbb{E}_{\beta_1, \beta_n} \left[\int_{\theta=\theta_0+\frac{\beta_n-2\pi}{2}}^{\theta_0+\frac{\beta_1}{2}} |R_0(\theta)| \, d\theta \right]$, we have shown:

$$\mathbb{E}_{\beta_1, \beta_n} \left[\int_{\theta=\theta_0+\frac{\beta_n-2\pi}{2}}^{\theta_0+\frac{\beta_1}{2}} R_0(\theta) \, d\theta \right] = \mathcal{O}\left(\frac{1}{n^2}\right). \tag{47}$$

In summary, when combining the asymptotic behavior of the zero-order term and the remainder term, we conclude that when $d = 2$ we have:

$$P_{\text{vMF-exp}}(A \mid n, d = 2, \kappa) = \frac{e^{\kappa \cos(\theta_0)}}{nI_0(\kappa)} + \mathcal{O}\left(\frac{1}{n^2}\right). \tag{48}$$

This proves Proposition 4.3 when $d = 2$. Note that, comparing the asymptotic expressions for $P_{\text{B-exp}}(A \mid n, d = 2, \kappa)$ and $P_{\text{vMF-exp}}(A \mid n, d = 2, \kappa)$, also gives us a proof for Proposition 4.1 when $d = 2$. \square

C. Asymptotic Behavior of vMF Exploration in $d > 2$ dimensions (Proofs of Proposition 4.3, Part 2, and of Proposition 4.4)

We now prove Proposition 4.3 when $d > 2$, starting with a series of intermediary lemmas. We subsequently justify the approximate expression of Proposition 4.4.

C.1. Intermediary Lemmas

We introduce a series of lemmas regarding the properties of the Voronoï cell of A when $\mathcal{X}_n \sim \mathcal{U}^{d-1}$. We recall that, for a given set of embedding vectors \mathcal{X}_n , we use the notation $\mathcal{X}_{n+1} = \mathcal{X}_n \cup \{A\}$.

Lemma C.1. *Let $d \in \mathbb{N}$, $d \geq 2$, $A \in \mathcal{S}^{d-1}$ and $n \in \mathbb{N}^*$. As before, let $\mathcal{A}(\mathcal{S}^{d-1})$ denote the surface area of \mathcal{S}^{d-1} . Then:*

$$\mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\mathcal{A}(\mathcal{S}_{\text{Voronoi}}(A \mid \mathcal{X}_{n+1})) \right] = \frac{\mathcal{A}(\mathcal{S}^{d-1})}{n+1}. \quad (49)$$

Proof. To compute this expectation, one can notice that:

$$\mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\mathcal{A}(\mathcal{S}_{\text{Voronoi}}(A \mid \mathcal{X}_{n+1})) \right] = \mathbb{E}_{\mathcal{X}_{n+1} \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\mathcal{A}(\mathcal{S}_{\text{Voronoi}}(X_{n+1} \mid \mathcal{X}_{n+1}) \mid X_{n+1} = A) \right]. \quad (50)$$

Indeed, considering that A is known is equivalent to considering A as a random vector $X_{n+1} \sim \mathcal{U}(\mathcal{S}^{d-1})$ with the constraint $X_{n+1} = A$. We will now show that the right part of Equation (50) is actually independent of the value of A .

Consider any point $A' \in \mathcal{S}^{d-1}$. One can always define a (not necessarily unique) rotation $R_{A,A'}$ such that $R_{A,A'}(A) = A'$. Since rotations preserve inner products, they also preserve areas of Voronoi cells, which means that for a given set of vectors \mathcal{X}_{n+1} , we have:

$$\mathcal{A}(\mathcal{S}_{\text{Voronoi}}(X_{n+1} \mid \mathcal{X}_{n+1})) = \mathcal{A}(\mathcal{S}_{\text{Voronoi}}(R_{A,A'}(X_{n+1}) \mid R_{A,A'}(\mathcal{X}_{n+1}))). \quad (51)$$

Moreover, the image of the rotation of a random vector uniformly distributed on the hypersphere is also uniformly distributed, which means that:

$$\mathcal{X}_{n+1} \sim \mathcal{U}(\mathcal{S}^{d-1}) \Leftrightarrow R_{A,A'}(\mathcal{X}_{n+1}) \sim \mathcal{U}(\mathcal{S}^{d-1}). \quad (52)$$

Therefore:

$$\begin{aligned} & \mathbb{E}_{\mathcal{X}_{n+1} \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\mathcal{A}(\mathcal{S}_{\text{Voronoi}}(X_{n+1} \mid \mathcal{X}_{n+1}) \mid X_{n+1} = A) \right] \\ &= \mathbb{E}_{\mathcal{X}_{n+1} \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\mathcal{A}(\mathcal{S}_{\text{Voronoi}}(R_{A,A'}(X_{n+1}) \mid R_{A,A'}(\mathcal{X}_{n+1})) \mid X_{n+1} = A) \right] \\ &= \mathbb{E}_{R_{A,A'}(\mathcal{X}_{n+1}) \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\mathcal{A}(\mathcal{S}_{\text{Voronoi}}(R_{A,A'}(X_{n+1}) \mid R_{A,A'}(\mathcal{X}_{n+1})) \mid R_{A,A'}(X_{n+1}) = A') \right] \\ &= \mathbb{E}_{R_{A,A'}(\mathcal{X}_n) \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\mathcal{A}(\mathcal{S}_{\text{Voronoi}}(A' \mid R_{A,A'}(\mathcal{X}_n))) \right] \\ &= \mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\mathcal{A}(\mathcal{S}_{\text{Voronoi}}(A' \mid \mathcal{X}_n)) \right]. \end{aligned} \quad (53)$$

This result proves that $\mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} [\mathcal{A}(\mathcal{S}_{\text{Voronoi}}(A \mid \mathcal{X}_{n+1}))]$ is independent of A . Then, we use this information along with Equation (50) to obtain:

$$\mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\mathcal{A}(\mathcal{S}_{\text{Voronoi}}(A \mid \mathcal{X}_{n+1})) \right] = \mathbb{E}_{\mathcal{X}_{n+1} \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\mathcal{A}(\mathcal{S}_{\text{Voronoi}}(X_{n+1} \mid \mathcal{X}_{n+1})) \right]. \quad (54)$$

Since $\sum_{i=1}^{n+1} \mathcal{A}(\mathcal{S}_{\text{Voronoi}}(X_i \mid \mathcal{X}_{n+1})) = \mathcal{A}(\mathcal{S}^{d-1})$ (Du et al., 1999; 2010) and the X_i are i.i.d., we derive:

$$\mathbb{E}_{\mathcal{X}_{n+1} \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\mathcal{A}(\mathcal{S}_{\text{Voronoi}}(X_{n+1} \mid \mathcal{X}_{n+1})) \right] = \frac{\mathcal{A}(\mathcal{S}^{d-1})}{n+1}. \quad (55)$$

Combining Equations (50) with Equation (55) leads to Equation (49), concluding the proof. \square

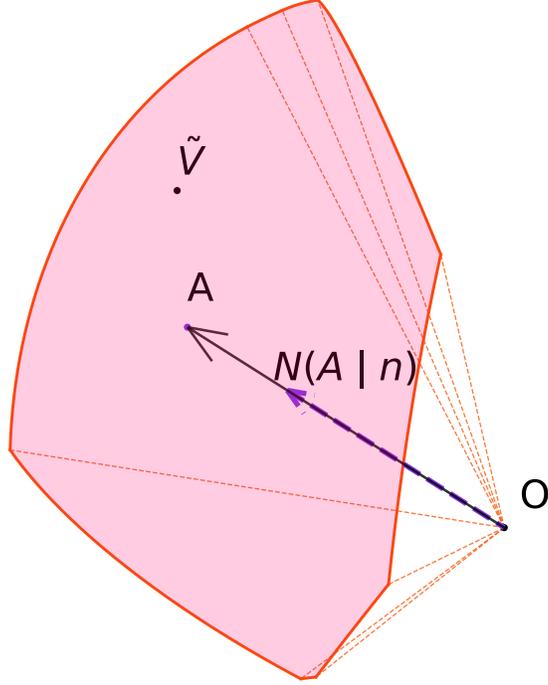


Figure 5. The Voronoi cell of A , $\mathcal{S}_{\text{Voronoi}}(A | \mathcal{X}_{n+1})$, along with the average normal vector of the cell $N(A | n)$. On expectation, $(A | n)$ and A are collinear.

Lemma C.2. Let $d \in \mathbb{N}$, $d \geq 2$, $A \in \mathcal{S}^{d-1}$ and $n \in \mathbb{N}^*$. Then:

$$\exists \lambda \in \mathbb{R}, \mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\int_{\tilde{V} \in \mathcal{S}_{\text{Voronoi}}(A | \mathcal{X}_{n+1})} \tilde{V} \, d\tilde{V} \right] = \lambda A. \quad (56)$$

Proof. We want to prove that the average normal vector of the Voronoi cell of A and A are collinear, as illustrated in Figure 5. To do so, we will show that this average normal vector is invariant to any rotation around A . For every $\theta \in [0, 2\pi]$, we define $R_{A,\theta}$ as the rotation around A of the angle θ . As discussed in the proof of Lemma C.1, $\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1}) \Leftrightarrow R_{A,\theta}(\mathcal{X}_n) \sim \mathcal{U}(\mathcal{S}^{d-1})$. Moreover, $R_{A,\theta}(A) = A$. Let us denote:

$$N(A | n) = \mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\int_{\tilde{V} \in \mathcal{S}_{\text{Voronoi}}(A | \mathcal{X}_{n+1})} \tilde{V} \, d\tilde{V} \right], \quad (57)$$

the expected normal vector of the Voronoi cell of A . Its image by the rotation $R_{A,\theta}$ verifies:

$$\begin{aligned} R_{A,\theta}(N(A | n)) &= R_{A,\theta} \left(\mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\int_{\tilde{V} \in \mathcal{S}_{\text{Voronoi}}(A | \mathcal{X}_{n+1})} \tilde{V} \, d\tilde{V} \right] \right) \\ &= \mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\int_{\tilde{V} \in \mathcal{S}_{\text{Voronoi}}(R_{A,\theta}(A) | R_{A,\theta}(\mathcal{X}_{n+1}))} \tilde{V} \, d\tilde{V} \right] \\ &= \mathbb{E}_{R_{A,\theta}(\mathcal{X}_n) \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\int_{\tilde{V} \in \mathcal{S}_{\text{Voronoi}}(A | R_{A,\theta}(\mathcal{X}_{n+1}))} \tilde{V} \, d\tilde{V} \right] \\ &= N(A | n). \end{aligned} \quad (58)$$

This proves that $N(A | n)$ and A are collinear. \square

Lemma C.3. With the same hypotheses as Lemma C.2:

$$\lambda = \frac{\mathcal{A}(\mathcal{S}^{d-1})}{n+1} \mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1}), \tilde{V} \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\max_i \langle \tilde{V}, X_i \rangle \right]. \quad (59)$$

Proof. λ is defined as follows:

$$\begin{aligned}
 \lambda A &= \mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\int_{\tilde{V} \in \mathcal{S}_{\text{Voronoi}}(\mathcal{A} | \mathcal{X}_{n+1})} \tilde{V} \, d\tilde{V} \right] \\
 \implies \langle \lambda A, A \rangle &= \langle \mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\int_{\tilde{V} \in \mathcal{S}_{\text{Voronoi}}(\mathcal{A} | \mathcal{X}_{n+1})} \tilde{V} \, d\tilde{V} \right], A \rangle \\
 \Leftrightarrow \lambda &= \mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\int_{\tilde{V} \in \mathcal{S}_{\text{Voronoi}}(\mathcal{A} | \mathcal{X}_{n+1})} \langle \tilde{V}, A \rangle \, d\tilde{V} \right] \\
 \Leftrightarrow \lambda &= \mathbb{E}_{\mathcal{X}_{n+1} \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\int_{\tilde{V} \in \mathcal{S}_{\text{Voronoi}}(\mathcal{X}_{n+1} | \mathcal{X}_{n+1})} \langle \tilde{V}, X_{n+1} \rangle \, d\tilde{V} \mid X_{n+1} = A \right] \\
 \Leftrightarrow \lambda &= \mathbb{E}_{\mathcal{X}_{n+1} \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\int_{\tilde{V} \in \mathcal{S}_{\text{Voronoi}}(\mathcal{X}_{n+1} | \mathcal{X}_{n+1})} \max_i \langle \tilde{V}, X_i \rangle \, d\tilde{V} \mid X_{n+1} = A \right].
 \end{aligned} \tag{60}$$

Moreover, as done in the proof of Lemma C.1, we can leverage the invariance by any rotation of the above expression to infer that the conditional expectation is actually independent of A :

$$\lambda = \mathbb{E}_{\mathcal{X}_{n+1} \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\int_{\tilde{V} \in \mathcal{S}_{\text{Voronoi}}(\mathcal{X}_{n+1} | \mathcal{X}_{n+1})} \max_i \langle \tilde{V}, X_i \rangle \, d\tilde{V} \right]. \tag{61}$$

Since, in the above equation, X_{n+1} has the same distribution as every element of \mathcal{X}_{n+1} , a similar expression for λ can be found using each \mathcal{X}_{n+1} element. By summing them together, we obtain:

$$\begin{aligned}
 (n+1)\lambda &= \sum_{j=1}^{n+1} \mathbb{E}_{\mathcal{X}_{n+1} \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\int_{\tilde{V} \in \mathcal{S}_{\text{Voronoi}}(X_j | \mathcal{X}_{n+1})} \max_i \langle \tilde{V}, X_i \rangle \, d\tilde{V} \right] \\
 &= \mathbb{E}_{\mathcal{X}_{n+1} \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\sum_{j=1}^{n+1} \int_{\tilde{V} \in \mathcal{S}_{\text{Voronoi}}(X_j | \mathcal{X}_{n+1})} \max_i \langle \tilde{V}, X_i \rangle \, d\tilde{V} \right] \\
 &= \mathbb{E}_{\mathcal{X}_{n+1} \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\int_{\tilde{V} \in \mathcal{S}^{d-1}} \max_i \langle \tilde{V}, X_i \rangle \, d\tilde{V} \right] \\
 &= \mathbb{E}_{\mathcal{X}_{n+1} \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\int_{\tilde{V} \in \mathcal{S}^{d-1}} \frac{\mathcal{A}(\mathcal{S}^{d-1}) \max_i \langle \tilde{V}, X_i \rangle}{\mathcal{A}(\mathcal{S}^{d-1})} \, d\tilde{V} \right] \\
 &= \mathcal{A}(\mathcal{S}^{d-1}) \mathbb{E}_{\mathcal{X}_{n+1} \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\mathbb{E}_{\tilde{V} \sim \mathcal{U}(\mathcal{S}^{d-1})} [\max_i \langle \tilde{V}, X_i \rangle] \right],
 \end{aligned} \tag{62}$$

which proves the lemma. \square

The last two lemmas are useful to describe the distribution of $\max_i \langle \tilde{V}, X_i \rangle$ when \tilde{V} is fixed, $\mathcal{X}_{n+1} \sim \mathcal{U}(\mathcal{S}^{d-1})$, and n is large.

Lemma C.4. Let $B : (z_1, z_2) \mapsto \int_0^1 t^{z_1-1} (1-t)^{z_2-1} dt$ denote the Beta function. Let $d \geq 3$, $\tilde{V} \in \mathcal{S}^{d-1}$ and X be a random vector with $X \sim \mathcal{U}(\mathcal{S}^{d-1})$. Let F_{radial} be the cumulative distribution function (CDF) of $\langle \tilde{V}, X \rangle$. The Taylor series expansion of F_{radial} near 1 is:

$$F_{\text{radial}}(t) = 1 - \frac{2^{\frac{d-1}{2}}}{(d-1)B(\frac{1}{2}, \frac{d-1}{2})} (1-t)^{\frac{d-1}{2}} + o((1-t)^{\frac{d-1}{2}}). \tag{63}$$

Proof. The distribution of $\langle \tilde{V}, X \rangle$ has been studied in directional statistics (Mardia & Jupp, 2009). Its PDF is known to be:

$$\begin{aligned}
 f_{\text{radial}}(t) &= \frac{(1-t^2)^{\frac{d-1}{2}-1}}{B(\frac{1}{2}, \frac{d-1}{2})} \\
 &= \frac{(1-t)^{\frac{d-1}{2}-1}(1+t)^{\frac{d-1}{2}-1}}{B(\frac{1}{2}, \frac{d-1}{2})} \\
 &= \frac{(1-t)^{\frac{d-1}{2}-1}(2-(1-t))^{\frac{d-1}{2}-1}}{B(\frac{1}{2}, \frac{d-1}{2})} \\
 &= \frac{2^{\frac{d-1}{2}-1}(1-t)^{\frac{d-1}{2}-1}(1-\frac{(1-t)}{2})^{\frac{d-1}{2}-1}}{B(\frac{1}{2}, \frac{d-1}{2})} \\
 &= \frac{2^{\frac{d-1}{2}-1}(1-t)^{\frac{d-1}{2}-1}}{B(\frac{1}{2}, \frac{d-1}{2})} \left(\sum_{i=0}^{\infty} \binom{\frac{d-1}{2}-1}{i} \left(\frac{1-t}{2}\right)^i \right).
 \end{aligned} \tag{64}$$

The last line above was obtained using Newton's generalized binomial theorem for real exponent (Coolidge, 1949). It involves the term $\binom{\frac{d-1}{2}-1}{i} = \frac{(\frac{d-1}{2}-1)_i}{i!}$ with $(\cdot)_i$ the Pochhammer symbol used to designate a falling factorial (Abramowitz & Stegun, 1948). We have obtained an expression of f_{radial} involving an infinite weighted sum of powers of $(1-t)$ with exponents greater or equal to 0 since $d \geq 3$. Therefore, by uniqueness of the Taylor polynomial, we derive that the Taylor series expansion of f_{radial} near 1 is:

$$f_{\text{radial}}(t) = \frac{2^{\frac{d-1}{2}-1}(1-t)^{\frac{d-1}{2}-1}}{B(\frac{1}{2}, \frac{d-1}{2})} + o((1-t)^{\frac{d-1}{2}-1}). \tag{65}$$

Since by definition F_{radial} is the primitive of f_{radial} on $[-1, 1]$ and that $F_{\text{radial}}(1) = 1$, we can integrate the above equation to get:

$$\begin{aligned}
 F_{\text{radial}}(t) &= 1 - \frac{2}{d-1} \frac{2^{\frac{d-1}{2}-1}(1-t)^{\frac{d-1}{2}}}{B(\frac{1}{2}, \frac{d-1}{2})} + o((1-t)^{\frac{d-1}{2}}) \\
 &= 1 - \frac{2^{\frac{d-1}{2}}(1-t)^{\frac{d-1}{2}}}{(d-1)B(\frac{1}{2}, \frac{d-1}{2})} + o((1-t)^{\frac{d-1}{2}}).
 \end{aligned} \tag{66}$$

Since this is exactly the Equation (63), this completes the proof. \square

Lemma C.5. Let $d \geq 3$, $\tilde{V} \in \mathcal{S}^{d-1}$ and let F_{radial} be defined as in Lemma C.4. For $n \in \mathbb{N}^*$, let $\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})$ be a set of n i.i.d. random vectors uniformly distributed on \mathcal{S}^{d-1} , and let F_n be the CDF of $\max_i \langle \tilde{V}, X_i \rangle$. Then, for $u \in [-1, 1]$:

$$\lim_{n \rightarrow +\infty} F_n(a_n u + b_n) = e^{-(1+\gamma u)^{\frac{-1}{\gamma}}}, \tag{67}$$

where $\gamma = -\frac{2}{d-1}$, $a_n = \frac{1}{2} \left(\frac{(d-1)B(\frac{1}{2}, \frac{d-1}{2})}{n} \right)^{\frac{2}{d-1}}$ with B the Beta function, and $b_n = 1 - \frac{2}{d-1} a_n$.

Proof. The proof relies on the Fisher-Tippett-Gnedenko theorem (Gnedenko, 1943) which states that if there exists a couple of sequences a_n and b_n such that the left term of Equation (67) converges, then its limit should be the CDF of a Generalized Extreme Value distribution (GEV) with shape parameter γ , which is the right term of Equation (67). Theorem 5 of Gnedenko (1943) provides a necessary and sufficient convergence condition for a random variable with maximal value x_{\max} and CDF F , provided that $\gamma < 0$:

$$\lim_{t \rightarrow 0^+} \frac{1 - F(x_{\max} - u t)}{1 - F(x_{\max} - t)} = u^{\left(\frac{-1}{\gamma}\right)} \text{ for all } u > 0. \tag{68}$$

Recall that Lemma C.4 gives us the Taylor expansion of F_{radial} near 1 : $F_{\text{radial}}(t) = 1 - K(1-t)^{\frac{d-1}{2}} + o((1-t)^{\frac{d-1}{2}})$ with $K = \frac{2^{\frac{d-1}{2}}}{(d-1)B(\frac{1}{2}, \frac{d-1}{2})}$. Knowing that $x_{\max} = 1$, we obtain that, $\forall u > 0$:

$$\begin{aligned} \lim_{t \rightarrow 0^+} \frac{1 - F_{\text{radial}}(1-ut)}{1 - F_{\text{radial}}(1-t)} &= \lim_{t \rightarrow 0^+} \frac{1 - (1 - K(ut)^{\frac{d-1}{2}}) + o((t)^{\frac{d-1}{2}})}{1 - (1 - K(t)^{\frac{d-1}{2}}) + o((t)^{\frac{d-1}{2}})} \\ &= \lim_{t \rightarrow 0^+} \frac{K(ut)^{\frac{d-1}{2}} + o((t)^{\frac{d-1}{2}})}{K(t)^{\frac{d-1}{2}} + o((t)^{\frac{d-1}{2}})} \\ &= u^{\left(\frac{d-1}{2}\right)}, \end{aligned} \quad (69)$$

which guarantees convergence and in the same time gives the value of $\gamma = -\frac{2}{d-1}$.

To find suitable sequences a_n and b_n , we can use the fact that $F_n(t) = F_{\text{radial}}(t)^n$ and study the behavior of $\ln F_n(t)$ near $t = 1$:

$$\begin{aligned} \ln F_n(t) &= \ln (F_{\text{radial}}(t)^n) \\ &= n \ln (F_{\text{radial}}(t)) \\ &= n(\ln(1 - K(1-t)^{\frac{-1}{\gamma}} + o((1-t)^{\frac{-1}{\gamma}}))) \text{ as } t \rightarrow 1^- \\ &= -nK((1-t)^{\frac{-1}{\gamma}} + o((1-t)^{\frac{-1}{\gamma}})) \text{ as } t \rightarrow 1^-. \end{aligned} \quad (70)$$

By defining $a_n = -\gamma(Kn)^\gamma$, $b_n = 1 - (Kn)^\gamma$ and doing the change of variable $u = \frac{t-b_n}{a_n}$, we see that:

$$\begin{aligned} t &= a_n u + b_n \\ &= 1 - (1 + \gamma u)(Kn)^\gamma. \end{aligned} \quad (71)$$

Since for every u , $\lim_{n \rightarrow +\infty} (1 + \gamma u)(Kn)^\gamma = 0$ (recall that $\gamma < 0$), the term $o((1-x)^{\frac{-1}{\gamma}})$ as $x \rightarrow 1^-$ is equivalent to $o(\frac{1}{n})$ as $n \rightarrow +\infty$. this means that:

$$\begin{aligned} \ln (F_n(a_n u + b_n)) &= -nK \left(((1 + \gamma u)(Kn)^\gamma)^{\frac{-1}{\gamma}} + o\left(\frac{1}{n}\right) \right) \text{ as } n \rightarrow +\infty. \\ &= -(1 + \gamma u)^{\frac{-1}{\gamma}} + o(1) \text{ as } n \rightarrow +\infty. \end{aligned} \quad (72)$$

We can now consider the exponential of the above expression to get our asymptotic maximum distribution:

$$\lim_{n \rightarrow +\infty} F_n(a_n u + b_n) = e^{-(1+\gamma u)^{\frac{-1}{\gamma}}}, \quad (73)$$

which concludes the proof. \square

Corollary C.6. With $\Gamma : z \mapsto \int_0^\infty t^{z-1} e^{-t} dt$ the Gamma function (Abramowitz & Stegun, 1948), we have:

$$\mathbb{E} x_n \sim \mathcal{U}(S^{d-1}) \left[\max_i \langle \tilde{V}, X_i \rangle \right] = 1 - \frac{\Gamma(\frac{d+1}{d-1})}{2} \left(\frac{(d-1)B(\frac{1}{2}, \frac{d-1}{2})}{n} \right)^{\frac{2}{d-1}} + o\left(\frac{1}{n^{\frac{2}{d-1}}}\right). \quad (74)$$

Proof. According to the Portmanteau theorem (Billingsley, 2013), Lemma C.5 is equivalent to:

$$\frac{\max_i \langle \tilde{V}, X_i \rangle - b_n}{a_n} \xrightarrow{D} \text{GEV}(\gamma), \quad (75)$$

where $\text{GEV}(\gamma)$ is a generalized extreme value distribution with shape parameter γ (Gnedenko, 1943). Recall that if a sequence Z_1, Z_2, \dots of random variables converges in distribution to a random variable Z , then for all bounded continuous

function ϕ , $\lim_{n \rightarrow +\infty} \mathbb{E}[\phi(Z_n)] = \mathbb{E}[\phi(Z)]$. Since $\frac{\max_i \langle \tilde{V}, X_i \rangle - b_n}{a_n}$ is bounded for every n , we can consider the identity function for ϕ and obtain:

$$\lim_{n \rightarrow +\infty} \mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\frac{\max_i \langle \tilde{V}, X_i \rangle - b_n}{a_n} \right] = \mathbb{E}[\text{GEV}(\gamma)] = \frac{\Gamma(1-\gamma) - 1}{\gamma}. \quad (76)$$

Replacing γ , a_n and b_n by their respective expressions, it implies that:

$$\begin{aligned} \lim_{n \rightarrow +\infty} \frac{\mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\max_i \langle \tilde{V}, X_i \rangle \right] - 1 + (Kn)^{-\frac{2}{d-1}}}{(Kn)^{-\frac{2}{d-1}}} + \Gamma\left(\frac{d-1}{d-1}\right) - 1 &= 0 \\ \Rightarrow \lim_{n \rightarrow +\infty} \frac{\mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\max_i \langle \tilde{V}, X_i \rangle \right] - 1 + K^{-\frac{2}{d-1}} \Gamma\left(\frac{d-1}{d-1}\right)}{n^{-\frac{2}{d-1}}} &= 0. \end{aligned} \quad (77)$$

Since $K^{-\frac{2}{d-1}} = \frac{1}{2} \left(\frac{(d-1)\text{B}(\frac{1}{2}, \frac{d-1}{2})}{n} \right)^{\frac{2}{d-1}}$, this is equivalent to writing:

$$\begin{aligned} \mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\max_i \langle \tilde{V}, X_i \rangle \right] - 1 + \frac{1}{2} \left(\frac{(d-1)\text{B}(\frac{1}{2}, \frac{d-1}{2})}{n} \right)^{\frac{2}{d-1}} \Gamma\left(\frac{d-1}{d-1}\right) &= o\left(\frac{1}{n^{\frac{2}{d-1}}}\right) \\ \Leftrightarrow \mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\max_i \langle \tilde{V}, X_i \rangle \right] &= 1 - \Gamma\left(\frac{d-1}{d-1}\right) \frac{1}{2} \left(\frac{(d-1)\text{B}(\frac{1}{2}, \frac{d-1}{2})}{n} \right)^{\frac{2}{d-1}} + o\left(\frac{1}{n^{\frac{2}{d-1}}}\right). \end{aligned} \quad (78)$$

We have thus obtained Equation (74), concluding the proof of the corollary. \square

C.2. Proof of Proposition 4.3

We now return to Proposition 4.3. In this section we consider the case of vMF-exp when $d > 2$ and X_i embeddings are uniformly distributed on \mathcal{S}^{d-1} . Under those assumptions:

$$P_{\text{vMF-exp}}(a \mid n, d, V, \kappa) = \frac{f_{\text{vMF}}(A \mid V, \kappa) \mathcal{A}(\mathcal{S}^{d-1})}{n} + \mathcal{O}\left(\frac{1}{n^{1+\frac{2}{d-1}}}\right). \quad (79)$$

Proof. Similarly to the 2 dimensional case, the definition of $P_{\text{vMF-exp}}(a \mid n, d, V, \kappa)$ is:

$$P_{\text{vMF-exp}}(A \mid n, d, V, \kappa) = \mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\mathbb{P}(\tilde{V} \in \mathcal{S}_{\text{Voronoi}}(A \mid \mathcal{X}_{n+1}) \mid \tilde{V} \sim \text{vMF}(V, \kappa)) \right], \quad (80)$$

which can be written using the PDF of the vMF distribution:

$$P_{\text{vMF-exp}}(A \mid n, d, V, \kappa) = \mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\int_{\tilde{V} \in \mathcal{S}_{\text{Voronoi}}(A \mid \mathcal{X}_{n+1})} f_{\text{vMF}}(\tilde{V} \mid V, \kappa) d\tilde{V} \right]. \quad (81)$$

As done in the 2D case, we study the Taylor expansion of f_{vMF} near A :

$$\begin{aligned} \forall \tilde{V} \in \mathcal{S}_{\text{Voronoi}}(A \mid \mathcal{X}_{n+1}), f_{\text{vMF}}(\tilde{V} \mid \kappa, V) &= C_d(\kappa) e^{\kappa \langle V, \tilde{V} \rangle} \\ &= C_d(\kappa) e^{\kappa \langle V, A \rangle} e^{\kappa \langle V, \tilde{V} - A \rangle} \\ &= f_{\text{vMF}}(A \mid V, \kappa) \sum_{i=0}^{\infty} \frac{(\kappa \langle V, \tilde{V} - A \rangle)^i}{i!} \\ &= f_{\text{vMF}}(A \mid V, \kappa) (1 + \kappa \langle V, \tilde{V} - A \rangle + R_1(\tilde{V})). \end{aligned} \quad (82)$$

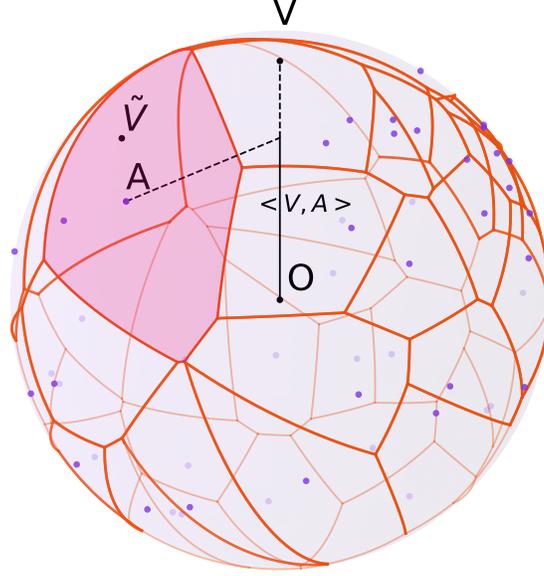


Figure 6. For $d = 3$: vMF-exp explores the action A when \tilde{V} lies in its Voronoi cell, shown in red.

with $R_1(\tilde{V}) = \sum_{i=2}^{\infty} \frac{(\kappa \langle V, \tilde{V} - A \rangle)^i}{i!}$. Leveraging the linearity property of both integration and expectation (Jacod & Protter, 2004), we can study $P_{\text{vMF-exp}}(A | n, d, V, \kappa)$ by assessing separately the contribution of the different terms of the expansion of f_{vMF} in:

$$P_{\text{vMF-exp}}(A | n, d, V, \kappa) = \mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\int_{\tilde{V} \in \mathcal{S}_{\text{Voronoi}}(A | \mathcal{X}_{n+1})} f_{\text{vMF}}(A | V, \kappa) (1 + \kappa \langle V, \tilde{V} - A \rangle + R_1(\tilde{V})) d\tilde{V} \right]. \quad (83)$$

However, contrary to the 2D case where $\mathcal{S}_{\text{Voronoi}}(A | \mathcal{X}_{n+1})$ is always defined as the arc between 2 angles on the circle, for $d > 2$ the shape of $\mathcal{S}_{\text{Voronoi}}(A | \mathcal{X}_{n+1})$ is highly dependent of the layout of the elements of \mathcal{X}_n that share a frontier with A . Figure 6 provides an illustration of the complexity and diversity of the shapes of Voronoi cells for uniformly sampled points on the 3D sphere.

As a consequence, expliciting the bounds of integration, as we did in the 2D case, can be somewhat tedious. Instead, we will leverage the geometrical properties of the problem at hand to estimate $P_{\text{vMF-exp}}(A | n, d, V, \kappa)$. We start with the zero-order term.

C.2.1. ZERO-ORDER TERM

Since the zero-order term is constant, its integral over $\mathcal{S}_{\text{Voronoi}}(A | \mathcal{X}_{n+1})$ can be expressed as:

$$\int_{\tilde{V} \in \mathcal{S}_{\text{Voronoi}}(A | \mathcal{X}_{n+1})} f_{\text{vMF}}(A | V, \kappa) d\tilde{V} = f_{\text{vMF}}(A | V, \kappa) \mathcal{A}(\mathcal{S}_{\text{Voronoi}}(A | \mathcal{X}_{n+1})), \quad (84)$$

where $\mathcal{A}(\mathcal{S}_{\text{Voronoi}}(A | \mathcal{X}_{n+1}))$ is the value of the surface area of $\mathcal{S}_{\text{Voronoi}}(A | \mathcal{X}_{n+1})$. To assess the expected value of the above equation for uniformly distributed \mathcal{X}_n , we use Lemma C.1 and obtain:

$$\begin{aligned} \mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\int_{\tilde{V} \in \mathcal{S}_{\text{Voronoi}}(A | \mathcal{X}_{n+1})} f_{\text{vMF}}(A | V, \kappa) d\tilde{V} \right] &= \frac{f_{\text{vMF}}(A | V, \kappa) \mathcal{A}(\mathcal{S}^{d-1})}{n+1} \\ &= \frac{f_{\text{vMF}}(A | V, \kappa) \mathcal{A}(\mathcal{S}^{d-1})}{n} + \mathcal{O}\left(\frac{1}{n^2}\right). \end{aligned} \quad (85)$$

C.2.2. FIRST-ORDER TERM

We want to estimate the value of:

$$\begin{aligned} & \mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\int_{\tilde{V} \in \mathcal{S}_{\text{Voronoi}}(A | \mathcal{X}_{n+1})} f_{\text{vMF}}(A | V, \kappa) \kappa \langle V, \tilde{V} - A \rangle d\tilde{V} \right] \\ &= f_{\text{vMF}}(A | V, \kappa) \kappa \left(\langle V, \mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\int_{\tilde{V} \in \mathcal{S}_{\text{Voronoi}}(A | \mathcal{X}_{n+1})} \tilde{V} d\tilde{V} \right] \rangle - \frac{\langle V, A \rangle \mathcal{A}(\mathcal{S}^{d-1})}{n} \right). \end{aligned} \quad (86)$$

Using Lemmas C.2 and C.3 as well as Corollary C.6, the left term inside the parentheses is:

$$\begin{aligned} & \langle V, \mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\int_{\tilde{V} \in \mathcal{S}_{\text{Voronoi}}(A | \mathcal{X}_{n+1})} \tilde{V} d\tilde{V} \right] \rangle \\ &= \langle V, A \rangle \frac{\mathcal{A}(\mathcal{S}^{d-1})}{n+1} \mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1}), \tilde{V} \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\max_i \langle \tilde{V}, X_i \rangle \right] \\ &= \langle V, A \rangle \frac{\mathcal{A}(\mathcal{S}^{d-1})}{n+1} \left(1 - \frac{\Gamma(\frac{d+1}{d-1})}{2} \left(\frac{(d-1)B(\frac{1}{2}, \frac{d-1}{2})}{n} \right)^{\frac{2}{d-1}} + o\left(\frac{1}{n^{\frac{2}{d-1}}}\right) \right). \end{aligned} \quad (87)$$

Re-injecting this expression into Equation (86) gives the following expression for the contribution of the first-order term to the probability of sampling A :

$$\begin{aligned} & \mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\int_{\tilde{V} \in \mathcal{S}_{\text{Voronoi}}(A | \mathcal{X}_{n+1})} f_{\text{vMF}}(A | V, \kappa) \kappa \langle V, \tilde{V} - A \rangle d\tilde{V} \right] \\ &= -f_{\text{vMF}}(A | V, \kappa) \frac{\mathcal{A}(\mathcal{S}^{d-1})}{n+1} \kappa \langle V, A \rangle \left(\frac{\Gamma(\frac{d+1}{d-1})}{2} \left(\frac{(d-1)B(\frac{1}{2}, \frac{d-1}{2})}{n} \right)^{\frac{2}{d-1}} + o\left(\frac{1}{n^{\frac{2}{d-1}}}\right) \right). \end{aligned} \quad (88)$$

C.2.3. REMAINDER TERM

As done in the 2D proof, we leverage the Taylor-Lagrange inequality (Abramowitz & Stegun, 1948). The second derivative of the function $f(x) = C_d(\kappa)e^{\kappa x}$ is $f(x)^{(2)} = \kappa^2 f(x)$, which is bounded on $x \in [-1, 1]$ by $M = \kappa^2 C_d(\kappa)e^{\kappa x}$. This implies that:

$$\begin{aligned} |R_1(\tilde{V})| &\leq \frac{M \langle V, \tilde{V} - A \rangle^2}{2} \\ &\leq \frac{M \|\tilde{V} - A\|_2^2}{2} \text{ (according to the Cauchy-Schwarz inequality (Jacod & Protter, 2004))} \\ &= M(1 - \langle \tilde{V}, A \rangle). \end{aligned} \quad (89)$$

This inequality holds for every $\tilde{V} \in \mathcal{S}_{\text{Voronoi}}(A | \mathcal{X}_{n+1})$ when $\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})$, which means that:

$$\begin{aligned} & \mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\int_{\tilde{V} \in \mathcal{S}_{\text{Voronoi}}(A | \mathcal{X}_{n+1})} f_{\text{vMF}}(A | V, \kappa) |R_1(\tilde{V})| d\tilde{V} \right] \\ &\leq f_{\text{vMF}}(A | V, \kappa) \mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\int_{\tilde{V} \in \mathcal{S}_{\text{Voronoi}}(A | \mathcal{X}_{n+1})} M(1 - \langle \tilde{V}, A \rangle) d\tilde{V} \right] \\ &= f_{\text{vMF}}(A | V, \kappa) \frac{\mathcal{A}(\mathcal{S}^{d-1})}{n+1} M(1 - \mathbb{E}_{\mathcal{X}_{n+1} \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\int_{\tilde{V} \in \mathcal{S}^{d-1}} \max_i \langle \tilde{V}, X_i \rangle d\tilde{V} \right]) \\ &= f_{\text{vMF}}(A | V, \kappa) M \frac{\mathcal{A}(\mathcal{S}^{d-1})}{n+1} \left(\frac{\Gamma(\frac{d+1}{d-1})}{2} \left(\frac{(d-1)B(\frac{1}{2}, \frac{d-1}{2})}{n} \right)^{\frac{2}{d-1}} + o\left(\frac{1}{n^{\frac{2}{d-1}}}\right) \right) \\ &= O\left(\frac{1}{n^{1+\frac{2}{d-1}}}\right). \end{aligned} \quad (90)$$

We used Lemmas C.2 and C.3 to go from line 2 to 3, and Corollary C.6 to go from line 3 to 4. In essence, we have bounded the contribution of $R_1(\tilde{V})$ to the probability of sampling A as follows:

$$\mathbb{E}_{\mathcal{X}_n \sim \mathcal{U}(\mathcal{S}^{d-1})} \left[\int_{\tilde{V} \in \mathcal{S}_{\text{Voronoi}}(A | \mathcal{X}_{n+1})} f_{\text{vMF}}(A | V, \kappa) |R_1(\tilde{V})| d\tilde{V} \right] = O\left(\frac{1}{n^{1+\frac{2}{d-1}}}\right) \quad (91)$$

Finally, adding up Equations (85), (89), and (91), we conclude the proof of Proposition 4.3 for $d \geq 3$ and (via the first-order term) simultaneously justify the approximate probability $P_1(a | n, V, \kappa)$ introduced in Proposition 4.4. \square

D. Similar Asymptotic Behavior of B-exp and vMF-exp for Large Action Sets (Proof of Proposition 4.1)

Finally, Propositions 4.2 and 4.3 allow us to derive Proposition 4.1, i.e., that in the setting of Section 4.1, we have:

$$\lim_{n \rightarrow +\infty} \frac{P_{\text{B-exp}}(a \mid n, d, V, \kappa)}{P_{\text{vMF-exp}}(a \mid n, d, V, \kappa)} = 1. \quad (92)$$

Proof. According to Proposition 4.2, we have:

$$P_{\text{B-exp}}(a \mid n, d, V, \kappa) = \frac{f_{\text{vMF}}(A \mid V, \kappa) \mathcal{A}(\mathcal{S}^{d-1})}{n} + o\left(\frac{1}{n\sqrt{n}}\right). \quad (93)$$

Moreover, according to Proposition 4.3, we have:

$$P_{\text{vMF-exp}}(a \mid n, d, V, \kappa) = \frac{f_{\text{vMF}}(A \mid V, \kappa) \mathcal{A}(\mathcal{S}^{d-1})}{n} + \begin{cases} \mathcal{O}\left(\frac{1}{n^2}\right) & \text{if } d = 2, \\ \mathcal{O}\left(\frac{1}{n^{1+\frac{d}{2}-1}}\right) & \text{if } d > 2. \end{cases} \quad (94)$$

Therefore:

$$\begin{aligned} \lim_{n \rightarrow +\infty} \frac{P_{\text{B-exp}}(a \mid n, d, V, \kappa)}{P_{\text{vMF-exp}}(a \mid n, d, V, \kappa)} &= \lim_{n \rightarrow +\infty} \frac{n P_{\text{B-exp}}(a \mid n, d, V, \kappa)}{n P_{\text{vMF-exp}}(a \mid n, d, V, \kappa)} \\ &= \frac{f_{\text{vMF}}(A \mid V, \kappa) \mathcal{A}(\mathcal{S}^{d-1}) + 0}{f_{\text{vMF}}(A \mid V, \kappa) \mathcal{A}(\mathcal{S}^{d-1}) + 0} \\ &= 1. \end{aligned} \quad (95)$$

□

E. Link with Thompson Sampling

At first glance, one might draw some similarities between vMF-exp and Thompson Sampling (TS) with Gaussian prior for contextual bandits (Chapelle & Li, 2011). Admittedly, vMF-exp shares a common spirit with TS, where action selection is preceded by sampling individual weights according to a Normal distribution centered on an observed context/state vector. However, vMF-exp also presents two major differences:

- Firstly, in vMF-exp, vector sampling is performed according to a vMF hyperspherical distribution, centered on the state embedding vector V . This choice of distribution ensures that vectors with the same inner product with the state vector have the same probability of being sampled, as illustrated in Figure 1(a). This aligns better with the similarity used to retrieve nearest neighbors and, as emphasized in this paper, leads to probabilities of exploring actions asymptotically comparable to Boltzmann Exploration (with better scalability) under the theoretical assumptions of Section 4.1.
- Secondly, vMF-exp is not designed to maximize the expected reward of a policy in an RL or contextual bandit environment and does not impose any parameter update strategy. Instead, it serves as an action selection tool for any scenario where policy updates cannot be performed regularly (as in the batch RL setting commonly found in industrial applications), yet broad exploration must still be guaranteed between consecutive updates.

F. Sampling from the von Mises-Fisher Distribution

F.1. Radial-tangent decomposition

Given a vector $\tilde{V} \in \mathcal{S}^{d-1}$ and a concentration $\kappa \in \mathbb{R}^+$, the algorithm described in (Pinzón & Jung, 2023) sample from a vMF(V, κ) by leveraging the radial-tangent decomposition of the elements of \mathcal{S}^{d-1} . For any $\tilde{V} \in \mathcal{S}^{d-1}$, let us call $\tilde{t} = \langle V, \tilde{V} \rangle$. Then we have:

$$\tilde{V} = \tilde{t}V + \sqrt{1 - \tilde{t}^2}\tilde{V}_O \quad (96)$$

where the vector \tilde{V}_O has a unit norm and is orthogonal to V .

F.2. vMF distribution

If, $\tilde{V} \sim \text{vMF}(V, \kappa)$, then :

- \tilde{t} is a scalar valued random variable.
- \tilde{V}_O is a random vector uniformly distributed on the (d-2) dimensional sub-sphere that is centered at and perpendicular to V . For instance, for $d = 3$, this would mean a circle centered around V .
- \tilde{t} and \tilde{V}_O are independent.

Since the reciprocal is also true, t and \tilde{V}_O can thus be separately sampled to obtain \tilde{V} .

F.3. Sampling $\tilde{t} = \langle V, \tilde{V} \rangle$

The PDF of \tilde{t} is known (Fisher, 1953) and follows:

$$f_{\text{radial}}(t; \kappa, d) = \frac{(\kappa/2)^{\frac{d}{2}-1}}{\Gamma(\frac{1}{2})\Gamma(\frac{d-1}{2})I_{\frac{d}{2}-1}(\kappa)} e^{t\kappa} (1 - t^2)^{\frac{d-3}{2}} \quad (97)$$

This PDF can be used to sample r through rejection sampling (Gentle, 2009).

F.4. Sampling \tilde{V}_O

\tilde{V}_O can be obtained by following the steps of algorithm 1.

Algorithm 1 Sample \tilde{V}_O

- 1 - Sample vector U uniformly from \mathcal{S}^{d-1} ;
 - 2 - Compute projection of U on V : $W = \langle U, V \rangle V$;
 - 3 - Subtract projection and normalize: $\tilde{V}_O = \frac{U - W}{\|U - W\|}$;
 - 4 - return \tilde{V}_O
-

Note that a simple way to sample U uniformly on \mathcal{S}^{d-1} is to sample d standard Gaussians independently (one for each dimension) and then normalize the resulting vector (Gentle, 2009).

F.5. Wrapping up

The vector \tilde{V} can now be computed by summing the right term of equation 97. Overall, we see that sampling \tilde{V} from vMF(V, κ) is **data-independent**, hence the scalability of the approach.

G. Additional Monte Carlo Simulations

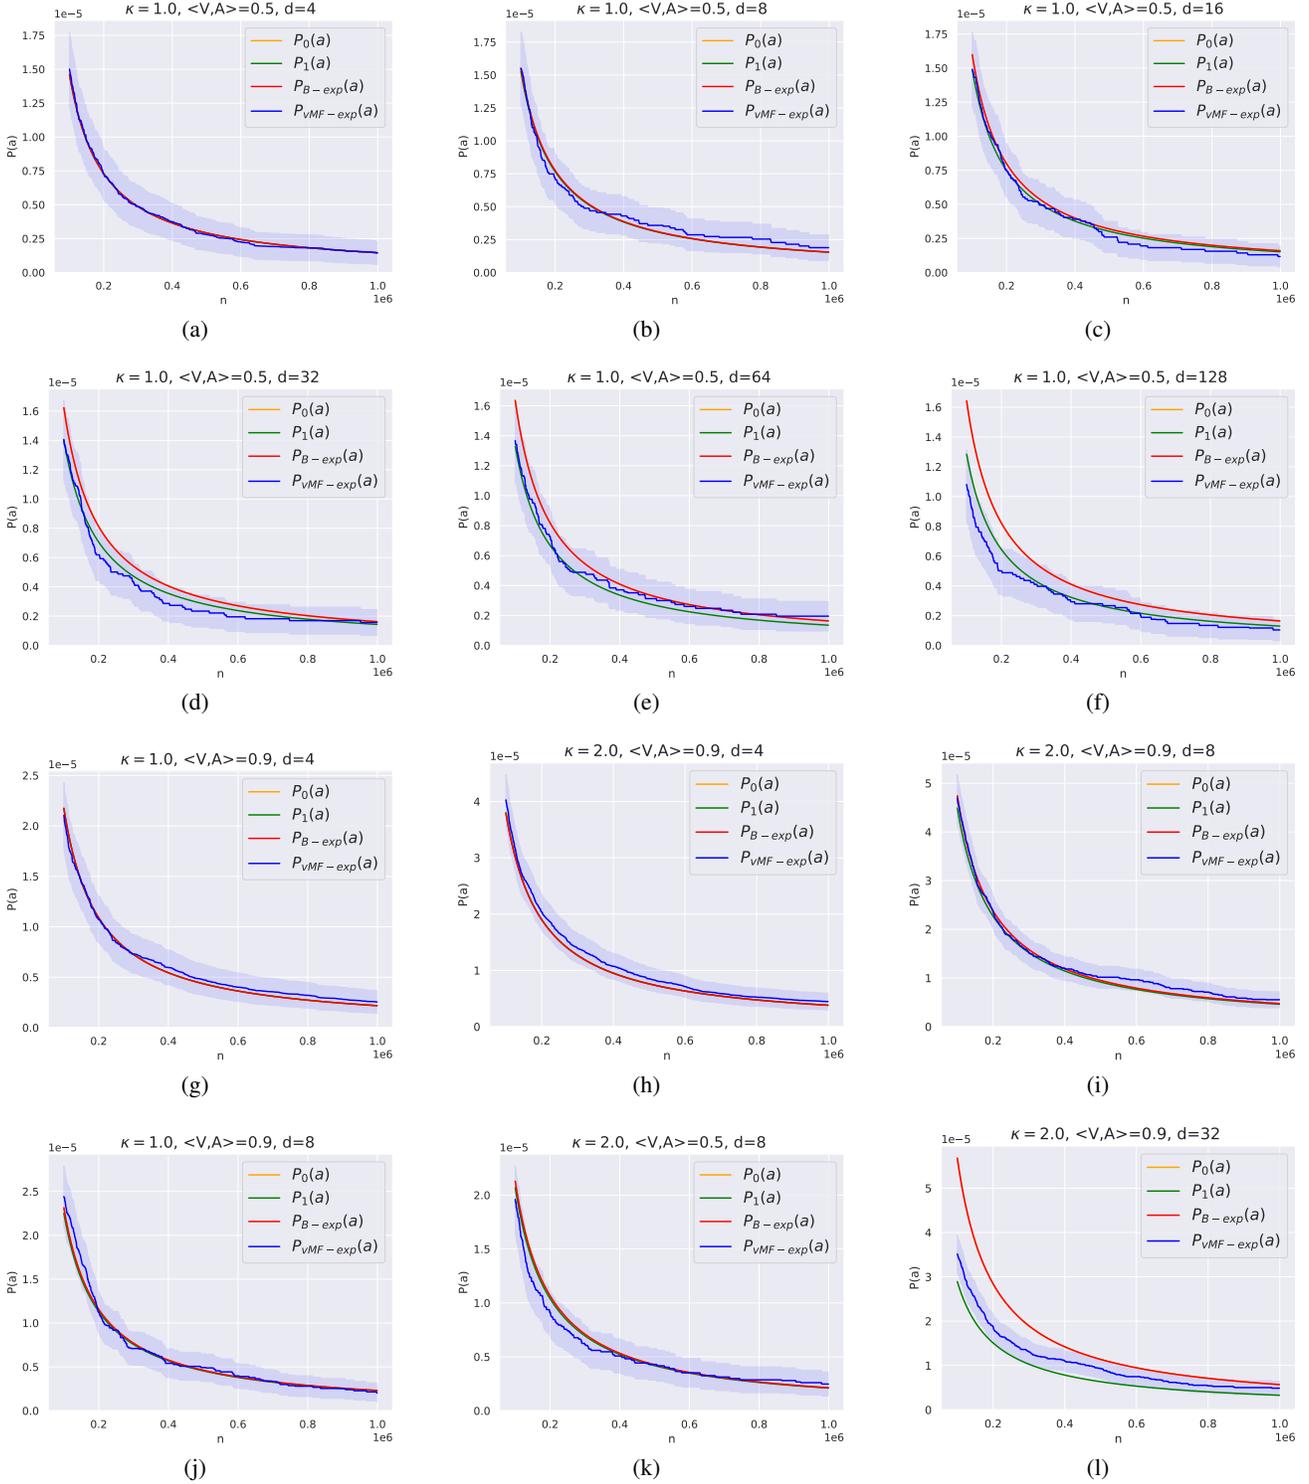


Figure 7. We report complete results for the Monte Carlo simulations presented and discussed in Section 4.3, involving more combinations of d , κ , and $\langle V, A \rangle$. We recall that $P_{B-\text{exp}}(a)$ and $P_0(a)$ are indistinguishable for this range of n values. We emphasize that the y-axis is on a 10^{-5} scale; hence, all probabilities are extremely close.

H. Additional Experiments on a Real-World Dataset of GloVe Word Embedding Vectors

While our main contributions in this work are theoretical, we aimed in the main paper to validate our key findings with Monte Carlo simulations, which involved synthetic data. Understanding that some readers may wish to further explore our topic through reproducible experiments on real-world data, we present an additional study in this Appendix H. This study experimentally validates the main properties of vMF-exp on a large-scale, publicly available real-world dataset.

H.1. Experimental setting

We present vMF-exp experiments on real-world, publicly available data. Specifically, we compare the behaviors of B-exp and vMF-exp on the GloVe-25 dataset of 1 million GloVe word embedding vectors with dimension $d = 25$ (Pennington et al., 2014). Each vector, learned using word2vec (Mikolov et al., 2013) from 2 billion tweets, represents a word token. We subtract the set’s average from each vector and divide them by their norms. We obtain a vector set, denoted \mathcal{G} , with all vectors lying on the unit hypersphere, making GloVe-25 a relevant large-scale dataset for our study.

Our experiments follow the protocol outlined in Section 4.3 of the main paper. In this section, we compared the empirical probabilities $P_{\text{B-exp}}(a)$ and $P_{\text{vMF-exp}}(a)$ of sampling an action a represented by a vector A given a state vector V , for varying action numbers n and inner products $\langle V, A \rangle$. While we relied on Monte Carlo simulations with uniformly drawn vectors $\mathcal{X}_n \sim \mathcal{U}(S^{d-1})$, in this Appendix H, vectors are sampled from \mathcal{G} , with V and A also drawn from \mathcal{G} such that $\langle V, A \rangle$ matches the pre-selected values. Our goal is to empirically compare $P_{\text{B-exp}}(a)$ and $P_{\text{vMF-exp}}(a)$, while verifying the claims that **P1**, **P2**, and **P3** simultaneously hold for vMF-exp.

Finally, in our Sections 4.2 and 4.3, we provided analytical approximations of $P_{\text{B-exp}}(a)$ and $P_{\text{vMF-exp}}(a)$ in the presence of independent and identically distributed (i.i.d.) uniform embedding vectors. We will assess the usefulness of these approximations on these GloVe vectors, which do not strictly satisfy these strong assumptions.

Finally, regarding these experiments on GloVe-25, we note that:

- The GloVe-25 dataset is available for download at: <https://nlp.stanford.edu/projects/glove/>.
- In our experiments, we use the Python vMF sampler from Pinzón & Jung (2023) to efficiently explore large action sets.
- All results are reproducible using our source code: <https://github.com/deezer/vMF-exploration>.

H.2. Results and Discussion

On P1 We now discuss our results. We first focus on **P1**. While B-exp requires computing $\langle V, X_i \rangle$ and softmax values for all n vectors $X_i \in \mathcal{X}_n$, vMF-exp only involves sampling a d -dimensional vector (in constant time with respect to n) and finding its approximate nearest neighbor (ANN) in \mathcal{X}_n .

Table 1 compares the performance of four popular ANN algorithms on GloVe-25. Following standard ANN literature (Simhadri et al., 2024), our performance metric is the maximum throughput, measured in queries per second (QPS), for which the average recall of the exact top-10 neighbors exceeds 90%. We also report the throughput of exhaustive search, as an indicator of B-exp’s inefficiency.

Table 1 shows that exhaustive search yields throughput 2 to 3 orders of magnitude lower than ANN methods. This confirms the significantly better scalability (**P1**) of vMF-exp compared to B-exp.

Table 1. Performance of popular ANN algorithms on GloVe-25, extracted from the benchmark of Aumüller et al. (2017). Following Simhadri et al. (2024), our performance metric is the maximum throughput, measured in Queries Per Second (QPS), for which the average recall of the exact top-10 neighbors exceeds 90%. The evaluated algorithms include two implementations of HNSW (Malkov & Yashunin, 2018), one from the Faiss library (Douze et al., 2024) and the other from NMSLIB (Boysov & Naidan, 2013), as well as ScaNN (Guo et al., 2020) and NGT-QG (Iwasaki & Miyazaki, 2018). Exhaustive search is 2 to 3 orders of magnitude slower than ANN methods.

Algorithm	Exhaustive Search	HNSW (Faiss)	HNSW (NMSLIB)	ScaNN	NGT-QG
Maximum Throughput in Queries Per Second (QPS)	34	6197	14080	23436	22733

On P2 and P3 Figure 8 compares $P_{B\text{-exp}}(a)$ and $P_{vMF\text{-exp}}(a)$ for increasing values of n and $\langle V, A \rangle$. Figure 8(a) highlights the two properties that make B-exp popular in RL: the ability to sample actions with unrestricted radius (**P2**) and the ordering of sampling probabilities based on action similarity to V (**P3**).

Importantly, Figure 8(b) confirms that vMF-exp also satisfies both properties. In our tests, A always has a positive sampling probability, which strictly increases with $\langle V, A \rangle$. Thus, vMF-exp also satisfies **P2** and **P3** on GloVe-25.

On Theoretical Approximations Finally, Figure 9 shows that, although GloVe vectors are not i.i.d. and uniform, the analytical approximations of the main paper for $P_{B\text{-exp}}(a)$ and $P_{vMF\text{-exp}}(a)$ often remain accurate, particularly for B-exp. Also, vMF-exp closely matches B-exp’s probabilities for low absolute values of $\langle V, A \rangle$.

However, as $|\langle V, A \rangle|$ increases, the gap between $P_{vMF\text{-exp}}(a)$ and $P_{B\text{-exp}}(a)$ grows more rapidly than predicted by approximations, highlighting the limitations of the i.i.d. and uniform assumptions and opening the way for future research on more general expressions.

Conclusion This additional study confirmed the key theoretical and scalability properties of vMF-exp on a large-scale and publicly available real-world dataset. Our results highlight its potential as a practical solution for exploring large action sets when hyperspherical embedding vectors represent these actions.

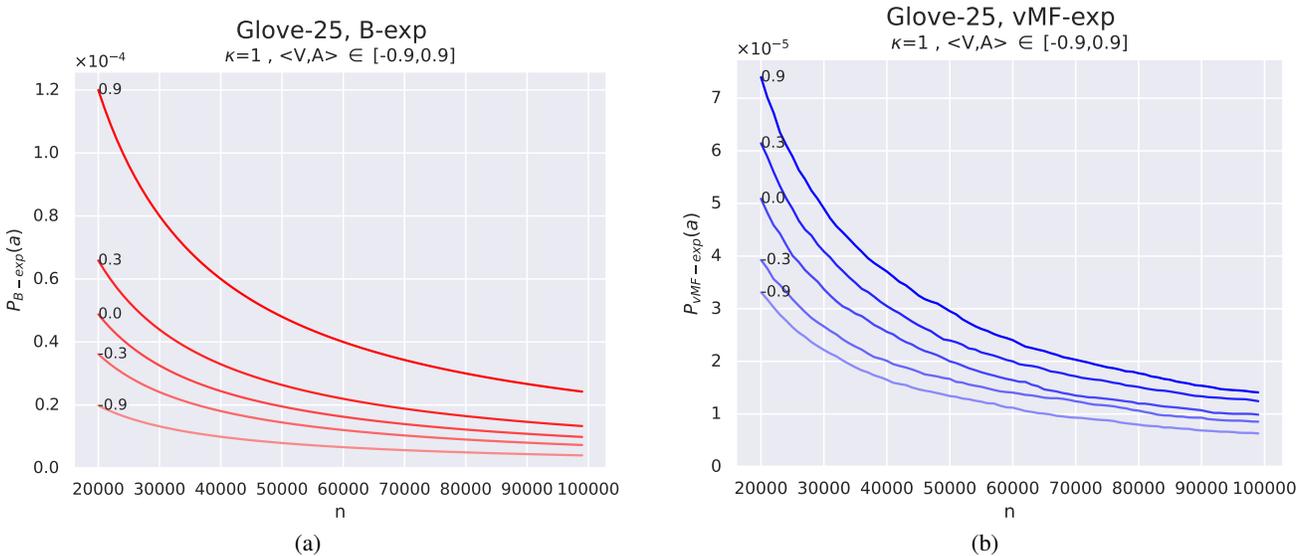


Figure 8. We report the empirical probabilities $P_{B\text{-exp}}(a)$ and $P_{vMF\text{-exp}}(a)$ of sampling an action a represented by a GloVe embedding vector A , using B-exp and vMF-exp, respectively, given a state vector V , with $20000 \leq n \leq 100000$ and $\langle V, A \rangle \in \{0.9, 0.3, 0.0, -0.3, -0.9\}$, and with $d = 25$ and $\kappa = 1$. Sampling is repeated 30 million times and averaged to obtain precise estimates. For both methods, the probability of sampling a for exploration is strictly positive (**P2**) and is a strictly increasing function of the inner product similarity $\langle V, A \rangle$ (**P3**). Results remain consistent when κ is modified.

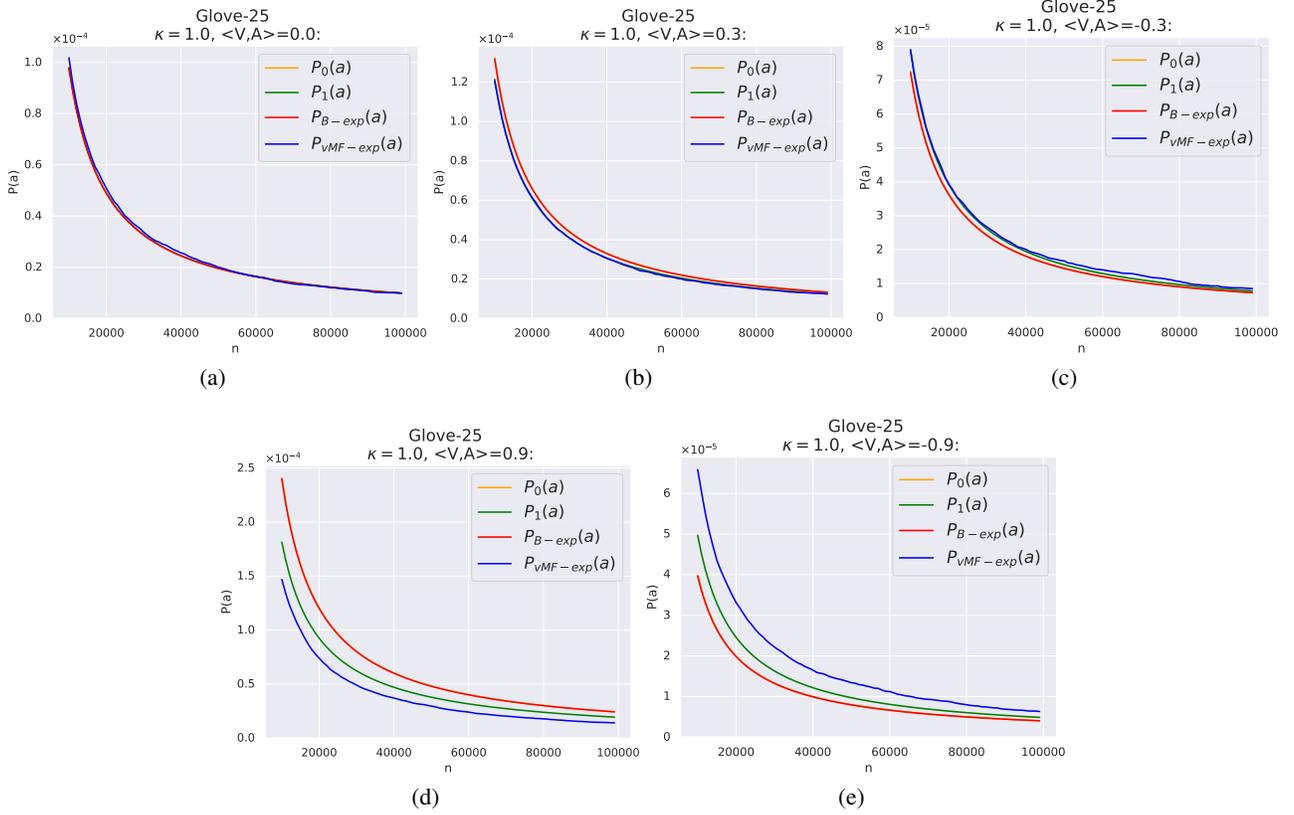


Figure 9. We compare $P_{B-\text{exp}}(a)$ and $P_{vMF-\text{exp}}(a)$ on GloVe-25 to the analytical approximations $P_0(a)$ and $P_1(a)$ stated in Propositions 4.1 to 4.4 of the main paper under the assumption of i.i.d. and uniformly distributed vectors. Our tests confirm the usefulness of these approximations on GloVe-25. The yellow curve ($P_0(a)$) is indistinguishable from the red curve ($P_{B-\text{exp}}(a)$), indicating that Proposition 4.2 holds across all configurations. Furthermore, $P_{vMF-\text{exp}}(a)$ (blue) remains close to $P_{B-\text{exp}}(a)$ (red) for low absolute values of $\langle V, A \rangle$ (Figures 9(a), 9(b), and 9(c)), as anticipated by Propositions 4.1 and 4.3. In Figures 9(b) and 9(c), the small difference between $P_{vMF-\text{exp}}(a)$ and $P_{B-\text{exp}}(a)$ aligns with the alternative expression $P_1(a)$ (green) derived in Proposition 4.4. However, as $|\langle V, A \rangle|$ increases (Figures 9(d) and 9(e)), the difference between vMF-exp and B-exp grows more rapidly than predicted by Proposition 4.4, highlighting the limitations of the i.i.d. and uniform assumptions.

I. Application to Large-Scale Music Recommendation

Our analysis of vMF-exp in this paper was intentionally general, as the method can be applied to various problem settings. In this Appendix I, we showcase a real-world application of vMF-exp.

I.1. Experimental Setting

We consider the “Mixes inspired by” feature of the global music streaming service Deezer⁵. This recommender system is deployed at scale and available on the homepage of this service (Bendada et al., 2023a). As shown in Figure 10, it displays a personalized shortlist of songs, selected from those previously liked by each user. A click on a song generates a playlist of 40 songs “inspired by” the initial one, with the aim of helping users discover new music within a catalog including several millions of recommendable songs.

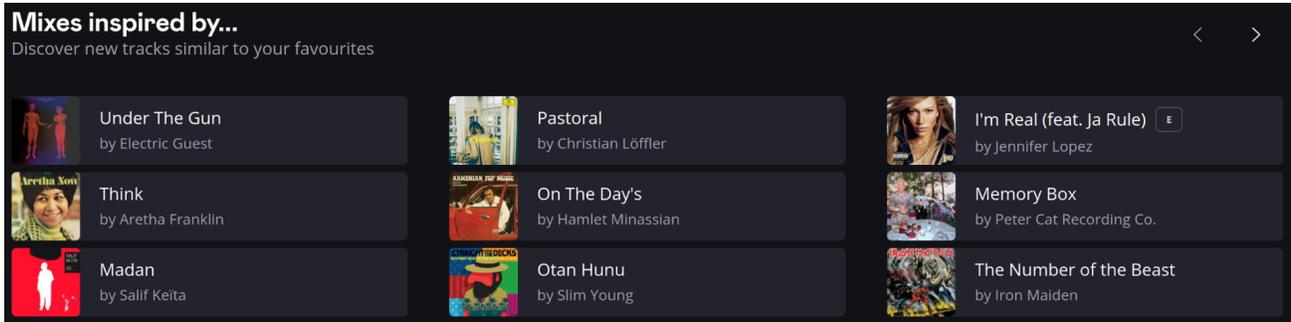


Figure 10. Interface of the “Mixes inspired by” recommender system on the music streaming service Deezer.

To generate playlists, Deezer leverages a collaborative filtering model (Koren & Bell, 2015). This model learns unit norm song embedding representations of dimension $d = 128$ by factorizing a mutual information matrix based on song co-occurrences in various listening contexts, using singular value decomposition (SVD) (Banerjee & Roy, 2014; Briand et al., 2021; 2024). Inner product proximity in the resulting embedding space aims to reflect user preferences. When a user selects an initial song, the model retrieves its embedding, then (approximately) identifies its neighbors in the embedding space using the efficient Faiss library (Johnson et al., 2019) for ANN. Currently, Deezer generates the entire playlist at once in production.

The service is considering RL approaches to, instead, recommend songs one by one while adapting to user feedback on previous songs of the playlist (likes, skips, etc.). However, as explained in Section 1, adopting such approaches would require exploring millions of possible actions/songs, significantly increasing the complexity of this task.

In this Appendix I, we continue generating “Mixes inspired by” playlists all at once, but take a step towards RL by comparing three methods for exploring large action sets of millions of songs:

- vMF-exp: we use the embedding of the user’s selected song as the initial state V . We sample a random state embedding \tilde{V} according to the vMF distribution, using the estimator of Banerjee et al. (2005) to tune κ (see Equation (4) of Sra (2012)). Finally, we recommend the 40 nearest neighbors of \tilde{V} in the embedding space according to the ANN engine.
- TB-exp: comparing vMF-exp to full B-exp is practically intractable at this scale. We compare vMF-exp to TB-exp with a similar κ . We first retrieve the $m = 500$ nearest neighbors of the initial song in the embedding space, according to the ANN engine. Then, we generate the playlist by sampling 40 songs from these 500 using a truncated Boltzmann distribution.
- Reference: we also compare vMF-exp to a baseline that retrieves the 500 nearest neighbors of the initial song using ANN, then shuffles them randomly to generate a playlist of 40 songs.

In early 2024, we conducted an industrial-scale online A/B test on the music streaming service Deezer to compare these exploration strategies in real conditions. The test involved millions of users worldwide, randomly split and unaware of the test.

⁵<https://www.deezer.com/en/>

I.2. Results and Discussion

Firstly, it is important to highlight that we were able to successfully deploy vMF-exp in Deezer’s production environment, achieving a sampling latency of just a few milliseconds, comparable to the other methods. This industrial deployment on a service used by millions of users on a daily basis confirms the claimed scalability of vMF-exp and its practical relevance for large-scale applications.

Using vMF-exp or TB-exp for exploration improved the daily number of recommended songs “liked” by users through “Mixes inspired by” (liking a song adds it to their list of favorites), compared to the reference baseline. For confidentiality, we do not report exact numbers of likes or users in each cohort, but present relative rates with respect to the reference. On average, users exposed to vMF-exp or TB-exp added 11% more recommended songs to their playlists than the reference cohort. These differences were statistically significant at the 1% level (p -value < 0.01). No apparent differences were observed between vMF-exp and TB-exp, showing that vMF-exp is competitive with TB-exp.

In addition, vMF-exp, which does not suffer from the restricted radius of TB-exp, recommended more diverse playlists. We measured the average Jaccard similarity (Tan et al., 2016) of playlists generated from the same initial selection, to assess how similar the songs sampled from the same state embedding were, for each method. Results reveal that TB-exp had an average Jaccard similarity 35% higher (less diverse playlists) than vMF-exp, a statistically significant difference at the 1% level (p -value < 0.01). Therefore, vMF-exp allowed for a more substantial exploration, without compromising performance.

At the time of writing, Deezer continues to use vMF-exp for “Mixes inspired by” recommendations. Playlists are still generated at once, but our work equips this service with an effective strategy to explore their large and embedded action set of millions of songs. This opens interesting avenues for further investigation of RL for recommendation. In the near future, Deezer will launch tests involving actor-critic RL models (Konda & Tsitsiklis, 1999; Sutton & Barto, 2018) to explore and generate songs sequentially based on user feedback.