

A dimension-independent bound on the Wasserstein contraction rate of a geodesic random walk on the sphere*

Philip Schär[†] Thilo D. Stier[‡]

Abstract

We theoretically analyze the properties of a geodesic random walk on the Euclidean d -sphere. Specifically, we prove that the random walk's transition kernel is Wasserstein contractive with a contraction rate which can be bounded from above independently of the dimension d . Our result is of particular interest due to its implications regarding the potential for dimension-independent performance of both geodesic slice sampling on the sphere and Gibbsian polar slice sampling, which are Markov chain Monte Carlo methods for approximate sampling from essentially arbitrary distributions on their respective state spaces.

Keywords: geodesic random walk; Wasserstein contraction; MCMC; slice sampling.

MSC2020 subject classifications: 60G50; 60J05; 65C05.

Submitted to ECP on September 17, 2023, final version accepted on September 17, 2024.

1 Introduction

In this work we investigate quantitative theoretical properties of a geodesic random walk on the Euclidean d -sphere

$$\mathbb{S}^d := \{x \in \mathbb{R}^{d+1} \mid \|x\| = 1\}$$

(where $\|\cdot\|$ denotes the Euclidean norm) for arbitrary dimensions $d \geq 2$. Intuitively, the random walk we consider performs an iteration from its current state $X_{n-1} = x_{n-1}$ to a new state $X_n = x_n$ by two simple steps. First, it selects a closed geodesic g_n of \mathbb{S}^d running through x_{n-1} , uniformly at random among all such geodesics. Then it chooses the new state x_n as a realization of the uniform distribution over all points on g_n .

Our contribution is a novel theoretical result on this geodesic random walk, which establishes bounds on the speed of its convergence (in Wasserstein distance) from an arbitrary initial distribution to its stationary distribution (which is the uniform distribution over the sphere). In particular, we prove the dimension-independent upper bound $\rho = 7/8$ on the Wasserstein contraction rate of the walk's transition kernel (see Corollary 4.2 below).

*PS gratefully acknowledges funding by the Carl Zeiss Stiftung within the program "CZS Stiftungsprofessuren" and the project "Interactive Inference". Moreover, PS is thankful for support by the DFG within project 432680300 – Collaborative Research Center 1456 "Mathematics of Experiment".

[†]Friedrich Schiller University Jena, Germany. E-mail: philip.schaer@uni-jena.de

[‡]Georg August University of Göttingen, Germany. E-mail: thilodaniel.stier@uni-goettingen.de

Let us begin by providing a formal description of the walk as a Markov chain on $(\mathbb{S}^d, \mathcal{B}(\mathbb{S}^d))$, where $\mathcal{B}(\mathbb{S}^d)$ denotes the Borel σ -algebra of \mathbb{S}^d . For this we require some preliminaries. Let σ_d be the natural surface measure on $(\mathbb{S}^d, \mathcal{B}(\mathbb{S}^d))$ and $\mathfrak{s}_d := \sigma_d(\mathbb{S}^d)$ its total mass. For any $x \in \mathbb{S}^d$ let $\mathbb{S}_x^{d-1} \subseteq \mathbb{S}^d$ denote the *great subsphere w.r.t. x*, i.e. the set

$$\mathbb{S}_x^{d-1} := \{z \in \mathbb{S}^d \mid x^\top z = 0\},$$

and $\sigma_{d-1}^{(x)}$ the natural surface measure on \mathbb{S}_x^{d-1} . It is well-known that $\sigma_{d-1}^{(x)}(\mathbb{S}_x^{d-1}) = \mathfrak{s}_{d-1}$ (see e.g. [2, Section 2.2]). Moreover, because \mathbb{S}_x^{d-1} canonically corresponds to the unit sphere in the tangent space to \mathbb{S}^d at x , it parametrizes the geodesics running through x . Specifically, for any $z \in \mathbb{S}_x^{d-1}$, a *closed geodesic* (which on \mathbb{S}^d is the same as a great circle) through x is given by the unit velocity curve

$$g_{x,z}: [0, 2\pi[\rightarrow \mathbb{S}^d, \omega \mapsto \cos(\omega)x + \sin(\omega)z. \tag{1.1}$$

For notational convenience, we denote by $G_{x,z} := g_{x,z}([0, 2\pi[) \subset \mathbb{S}^d$ the image of $g_{x,z}$.

Drawing a reference point $z \in \mathbb{S}_x^{d-1}$ from $\mathcal{U}(\mathbb{S}_x^{d-1}) := \frac{1}{\mathfrak{s}_{d-1}}\sigma_{d-1}^{(x)}$ to determine a geodesic $g_{x,z}$ which runs through x amounts to a uniform distribution over all such geodesics. Furthermore, due to the constant velocity of $g_{x,z}$, drawing $\Omega \sim \mathcal{U}([0, 2\pi[)$ and applying $g_{x,z}$ to the result implements a uniform distribution over $G_{x,z}$. We refer to [2, Section 2] and the references therein for more details on all of these preliminaries.

We may now express the geodesic random walk we described earlier by its transition kernel¹ P given by

$$P(x, A) := \frac{1}{2\pi \mathfrak{s}_{d-1}} \int_{\mathbb{S}_x^{d-1}} \int_0^{2\pi} \mathbb{1}_A(g_{x,z}(\omega)) d\omega \sigma_{d-1}^{(x)}(dz), \tag{1.2}$$

for $x \in \mathbb{S}^d, A \in \mathcal{B}(\mathbb{S}^d)$.

The remainder of this paper is structured as follows. In Section 2, we discuss our motivation for examining the geodesic random walk specified above. Then, in Section 3, we provide some background on Wasserstein contraction, the theoretical framework employed in our analysis. Afterwards, in Section 4, we state and prove our main result. We conclude with a short discussion of that result in Section 5.

2 Motivation

Our examination of the geodesic random walk defined by the transition kernel (1.2) is motivated in several different ways. First of all, the convergence rates of geodesic random walks may be of interest as standalone results, and other authors have previously established results similar to ours on such grounds. In this regard, we wish to highlight [3], in which the authors examined a geodesic random walk on a fairly general Riemannian manifold \mathcal{M} , which could be chosen as \mathbb{S}^d . Like the present work, the main result of [3] relied on the Wasserstein contraction framework to bound the walk's convergence rate by a term which, in the case of the underlying manifold being \mathbb{S}^d , does not depend on the dimension d . On the other hand, the walk they considered for $\mathcal{M} = \mathbb{S}^d$ is (in some sense) simpler than the one we are interested in: It also chooses a random geodesic through the current state (exactly like our walk), but then picks the next state deterministically by taking a step of a fixed size (a hyperparameter) along the selected geodesic (whereas our walk would perform a second sampling step here). Although the aforementioned main result of [3] is a much broader one than ours, in the sense that

¹In our nomenclature, a transition kernel is simply a Markov kernel whose arguments belong to the same measurable space. Basic knowledge of Markov kernels on the reader's part is presumed throughout this work. We refer to [1, Section 1.2] for a sufficiently broad introduction to this topic.

it applies to walks on many different manifolds, whereas we only permit \mathbb{S}^d , our result has the advantage that the walk it considers is also of relevance to a number of methods for (approximate) sampling from essentially arbitrary distributions on either \mathbb{S}^d or \mathbb{R}^d . Before we elaborate on this point, let us emphasize that the proof of our result, despite proceeding within the same theoretical framework as [3], relies on a markedly different strategy, primarily because our restriction to the d -sphere allows us to make extensive use of its unique symmetry properties.

As mentioned above, our motivation for analyzing the transition kernel (1.2) also stems from its relevance to the properties of certain instances of slice sampling, which is a general framework for setting up Markov chain Monte Carlo (MCMC) methods for approximately sampling from arbitrary target distributions given by their (potentially unnormalized²) densities. Specifically, our motivation relates to two particular (hybrid) slice samplers, one based on geodesics and one working in polar coordinates. For each of them, we now explain the sampler’s inner workings in some detail and then point out the connection between its properties and this work’s main result.

Geodesic slice sampling on the sphere (GSSS) [2] is an MCMC method for approximate sampling from target distributions ν on $(\mathbb{S}^d, \mathcal{B}(\mathbb{S}^d))$ (for any $d \geq 1$) given by a density $\varrho: \mathbb{S}^d \rightarrow [0, \infty[$ relative to σ_d , meaning that one’s only access to ν is through evaluations of ϱ and one has

$$\nu(A) = \frac{\int_A \varrho(x) \sigma_d(dx)}{\int_{\mathbb{S}^d} \varrho(x) \sigma_d(dx)}, \quad A \in \mathcal{B}(\mathbb{S}^d).$$

In [2], the authors of GSSS proposed two different variants of the algorithm, *ideal geodesic slice sampling* and *geodesic shrinkage slice sampling*. They perform a transition from the current state $X_{n-1} = x_{n-1}$ to a new state $X_n = x_n$ by the following steps:

1. Draw a threshold $T_n \sim \mathcal{U}(]0, \varrho(x_{n-1})[)$, call the result t_n . This defines the (super-) level set or *slice*

$$L(\varrho, t_n) := \{x \in \mathbb{S}^d \mid \varrho(x) > t_n\},$$

from which the methods then need to draw the new state X_n .

2. Select a random geodesic g_{x_{n-1}, z_n} through x_{n-1} by drawing $Z_n \sim \mathcal{U}(\mathbb{S}_{x_{n-1}}^{d-1})$ and calling the result z_n .
3. In case of ideal geodesic slice sampling, draw

$$\Omega_n \sim \mathcal{U}(\{\omega \in [0, 2\pi[\mid g_{x_{n-1}, z_n}(\omega) \in L(\varrho, t_n)\}), \tag{2.1}$$

call the result ω_n and set $x_n := g_{x_{n-1}, z_n}(\omega_n)$. In case of geodesic shrinkage slice sampling, draw $\Omega_n \sim U_{\text{geo}}(0, \cdot)$, where U_{geo} is an auxiliary transition kernel on $[-2\pi, 2\pi[\times \mathcal{B}([-2\pi, 2\pi])$ which mimics sampling from the distribution in (2.1) (up to shifts by 2π) by means of a shrinkage procedure (a sort of adaptive acceptance/rejection method). Again call the result ω_n and set $x_n := g_{x_{n-1}, z_n}(\omega_n)$.

In the specific case of the target density ϱ being constant (i.e. $\varrho(x) = c > 0, x \in \mathbb{S}^d$), the transition kernel of each of the two GSSS variants is easily seen to coincide with (1.2), and so both variants simplify to the geodesic random walk we are interested in. Naturally, ϱ being constant corresponds to the target distribution ν being the uniform distribution over the entire sphere, $\nu = \mathcal{U}(\mathbb{S}^d)$. While this case was already covered by [2, Theorems 15, 26], those theorems only yielded a dimension-dependent convergence rate bound³ $\rho(d)$ which deteriorates with increasing dimension, in the sense that $\lim_{d \rightarrow \infty} \rho(d) = 1$.

²As slice sampling approaches generally do not require the target density to be normalized, we will omit the “potentially unnormalized” in the following.

³In the two referenced results, this $\rho(d)$ plays the role that the total variation distance between the distribution of the n -th Markov chain iterate X_n and the target distribution ν is upper-bounded by $\rho(d)^n$.

In contrast, our result proves that GSSS for the uniform target distribution $\nu = \mathcal{U}(\mathbb{S}^d)$ works well in any dimension $d \geq 2$.

As one may be inclined to disregard this result on the grounds that exact sampling from $\mathcal{U}(\mathbb{S}^d)$ is computationally tractable (so that there is no point in applying MCMC to sample from this distribution in the first place), we should point out that the Wasserstein contraction property we establish is known to be robust under small perturbations of the transition kernel⁴, cf. [5, Remark 54]. Moreover, since the behavior of GSSS is easily seen to be fairly robust under small perturbations of the target distribution ν (unlike e.g. gradient-based methods), we expect such changes to the target would only cause small perturbations to its transition kernel, so that the Wasserstein contraction rate from the uniform case $\nu = \mathcal{U}(\mathbb{S}^d)$ should, to some extent, be preserved. Under this premise, our result on GSSS for the uniform target should intuitively transfer to GSSS for “near-uniform” target distributions, i.e. ones whose densities ϱ are not quite constant, but “flat” in the sense that their range $\sup \varrho - \inf \varrho$ is relatively small compared to their minimum value $\inf \varrho$. That is, the true convergence rate of GSSS for such near-uniform targets should be much closer to the dimension-independent rate bound we provide for the case $\nu = \mathcal{U}(\mathbb{S}^d)$ than to the dimension-dependent rate bounds one obtains from the results in [2] for such near-uniform targets.

This in turn suggests GSSS could be useful as a base sampler for parallel tempering [4] on high-dimensional spheres: In such schemes, one needs to apply the base sampler not only to a given target density ϱ on the sphere, but also to “heated-up” versions of ϱ , which become flatter and flatter with increasing (simulated) temperature and thereby eventually enter the aforementioned near-uniform regime. Whereas the performances of typical choices of the base sampler, e.g. the random walk Metropolis algorithm, are known to slowly deteriorate with increasing dimension, GSSS should, according to our above considerations, be able to maintain an excellent performance even in high dimensions (at least for the near-uniform high-temperature versions of the target). We therefore surmise that the use of GSSS as the base sampler in parallel tempering schemes on the sphere has the potential to substantially improve the overall performance of these schemes, compared to the use of a more conventional base sampler.

Gibbsian polar slice sampling (GPSS) [10] is an MCMC method that can, in principle, be applied to any target distribution ζ on $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$ (for any $d \geq 2$) given by a (potentially unnormalized) Lebesgue density $\eta: \mathbb{R}^d \rightarrow [0, \infty[$, meaning

$$\zeta(A) = \frac{\int_A \eta(x) dx}{\int_{\mathbb{R}^d} \eta(x) dx}, \quad A \in \mathcal{B}(\mathbb{R}^d).$$

GPSS replaces a given target density η by the transformation $\tilde{\eta}(x) := \|x\|^{d-1} \eta(x)$, which allows it to generate its samples in polar coordinates, i.e. as the product of a $]0, \infty[$ -valued random variable (the *radius*) and an \mathbb{S}^{d-1} -valued random variable (the *direction*). Specifically, GPSS performs a transition from the current state $X_{n-1} = x_{n-1}$ to a new state $X_n = x_n$ by the following steps:

1. Draw a threshold $T_n \sim \mathcal{U}(]0, \tilde{\eta}(x_{n-1})[)$, call the result t_n and define the slice

$$L(\tilde{\eta}, t_n) := \{x \in \mathbb{R}^d \mid \tilde{\eta}(x) > t_n\},$$

from which the new state X_n must then be drawn.

2. Let $r_{n-1} := \|x_{n-1}\| \in]0, \infty[$ and $\theta_{n-1} := x_{n-1}/r_{n-1} \in \mathbb{S}^{d-1}$.

⁴More concretely, a sufficiently small perturbation of a Wasserstein contractive kernel is still Wasserstein contractive, albeit likely with a slightly worse contraction rate.

3. Direction update: Draw $\Theta_n \sim Q_n(\theta_{n-1}, \cdot)$, where Q_n denotes the transition kernel of GSSS for the target distribution

$$\nu_n := \mathcal{U}(\{\theta \in \mathbb{S}^{d-1} \mid r_{n-1}\theta \in L(\tilde{\eta}, t_n)\}),$$

given by its density

$$\varrho_n(\theta) := \mathbb{1}_{L(\tilde{\eta}, t_n)}(r_{n-1}\theta), \quad \theta \in \mathbb{S}^{d-1}.$$

Call the result θ_n .

4. Radius update: Draw $R_n \sim U_n(r_{n-1}, \cdot)$, where U_n denotes an auxiliary transition kernel on a subset of \mathbb{R} , designed to mimic sampling from

$$\mathcal{U}(\{r \in]0, \infty[\mid r\theta_n \in L(\tilde{\eta}, t_n)\}).$$

Call the result r_n . Finally, set $x_n := r_n\theta_n$.

Given how GPSS relies on GSSS within its transition mechanism, it is perhaps not too surprising that there are strong links between their respective properties. Here we are particularly interested in the properties of GPSS for target densities η which are *rotationally invariant*, meaning $\eta(x) = h_\eta(\|x\|)$ for some $h_\eta: [0, \infty[\rightarrow [0, \infty[$ and all $x \in \mathbb{R}^d$. Let r_{n-1} and t_n be as in the above enumeration. Since, by construction,

$$\tilde{\eta}(r_{n-1}\theta) = (r_{n-1})^{d-1}h_\eta(r_{n-1}) = \tilde{\eta}(x_{n-1}) > t_n, \quad \theta \in \mathbb{S}^{d-1},$$

one then has $r_{n-1}\theta \in L(\tilde{\eta}, t_n)$ for any $\theta \in \mathbb{S}^{d-1}$, and thus $\nu_n = \mathcal{U}(\mathbb{S}^{d-1})$. In other words, GPSS applied to a rotationally invariant target density performs its direction update by means of GSSS applied to $\mathcal{U}(\mathbb{S}^{d-1})$. Thus our main result in this work not only permits an alternative view as a result on GSSS for the uniform target $\nu = \mathcal{U}(\mathbb{S}^d)$, but also another one as a result on the direction update of GPSS for rotationally invariant targets.

Moreover, the examination of GPSS for such targets is well motivated by both theoretical considerations and practical observations from previous works. The motivation from the theoretical side arises as follows: In [10], GPSS is constructed as a computationally efficient approximation to *polar slice sampling* (PSS) [7], which transitions similarly to GPSS but, instead of drawing R_n and Θ_n separately, relies on the joint update

$$(R_n, \Theta_n) \sim \mathcal{U}(\{(r, \theta) \in]0, \infty[\times \mathbb{S}^{d-1} \mid r\theta \in L(\tilde{\eta}, t_n)\}).$$

The recent works [8, 9] showed that PSS performs dimension-independently well (in terms of its spectral gap) on two classes of rotationally invariant target densities, one light-tailed and one heavy-tailed. Naturally, these results raise the question whether or not the dimension-independence of PSS is retained by GPSS. According to the recent work [6], the key quantity to examine in order to answer this question is the spectral gap of the *on-slice transition kernel* H_t , which in our notation is determined by the current threshold $t := t_n$, takes the current state x_{n-1} as input, and outputs the new state x_n . Admittedly, for a complete proof that GPSS does indeed have dimension-independent spectral gaps in the same cases as PSS, we would still need to conduct a considerable amount of further analysis, which we leave for future work. Nevertheless, the result we establish here should be an important ingredient for such a proof (besides the results on PSS from [8, 9]), particularly because the Wasserstein contraction property we prove is well-known to imply an explicit lower bound on the spectral gap [5, Proposition 30].

From the practical side, GPSS was already empirically observed to exhibit *seemingly* dimension-independent performance for (certain) rotationally invariant targets by its authors in [10]. However, the empirical analysis of an MCMC method's output is

generally insufficient for definitively confirming such a conjecture, so to be sure one must undertake a theoretical examination instead.

A strong additional motivation for analyzing GPSS for rotationally invariant targets stems from the recent work [11]. There the authors propose *parallel affine transformation tuning* (PATT), a sampling scheme based on multiple interacting chains which collaboratively learn an affine transformation of the target density, aimed at bringing it into isotropic position (i.e. mean zero, covariance is identity matrix). Essentially, each of these chains proceeds by running a *base sampler* for the transformed target. While the base sampler may be an arbitrary MCMC method, the authors in [11] advocate designating GPSS the default choice for this role. Furthermore, for sufficiently regular target densities, their affine transformations into isotropic position either attain rotational invariance outright or at least come very close to it. Hence the performance of PATT with its default base sampler GPSS, denoted PATT-GPSS in [11], is strongly dependent on the performance of GPSS for targets which are rotationally invariant or close to it. Like for GSSS and uniform targets, we surmise the performance of GPSS for targets which are nearly rotationally invariant should be close to its performance for rotationally invariant targets, but for now it remains an open problem to theoretically confirm this.

3 Background on Wasserstein contraction

Before we can move on to our main result, we need to briefly introduce the relevant notions concerning Wasserstein contraction, for which we loosely follow the corresponding parts of [9]. For two probability measures ξ_1, ξ_2 on $(\mathbb{S}^d, \mathcal{B}(\mathbb{S}^d))$, denote by $\Gamma(\xi_1, \xi_2)$ the set of couplings between them, i.e. of probability measures γ on $(\mathbb{S}^d \times \mathbb{S}^d, \mathcal{B}(\mathbb{S}^d \times \mathbb{S}^d))$ that satisfy

$$\begin{aligned} \gamma(A \times \mathbb{S}^d) &= \xi_1(A), \\ \gamma(\mathbb{S}^d \times A) &= \xi_2(A) \end{aligned}$$

for any $A \in \mathcal{B}(\mathbb{S}^d)$. The restriction of the Euclidean distance on \mathbb{R}^{d+1} to $\mathbb{S}^d \subseteq \mathbb{R}^{d+1}$ is easily seen to constitute a metric on \mathbb{S}^d . We can therefore define the Wasserstein distance (of order 1) between ξ_1 and ξ_2 w.r.t. the Euclidean distance,

$$\mathcal{W}_1(\xi_1, \xi_2) := \inf_{\gamma \in \Gamma(\xi_1, \xi_2)} \int_{\mathbb{S}^d \times \mathbb{S}^d} \|v_1 - v_2\| \gamma(\mathrm{d}v_1 \times \mathrm{d}v_2).$$

Let P be a Markov kernel on $\mathbb{S}^d \times \mathcal{B}(\mathbb{S}^d)$, then its *Dobrushin coefficient* (w.r.t. the Euclidean distance) is given by

$$\text{Dob}(P) := \sup_{x, y \in \mathbb{S}^d, x \neq y} \frac{\mathcal{W}_1(P(x, \cdot), P(y, \cdot))}{\|x - y\|}, \tag{3.1}$$

see [1, Definition 20.3.1, Lemma 20.3.2]. We now say that P is *Wasserstein contractive with rate bound* ρ if $\text{Dob}(P) \leq \rho < 1$. This type of contraction is known to be an extremely powerful property: If P is Wasserstein contractive with rate bound ρ and admits ν as its invariant distribution⁵, then for any distribution μ on $(\mathbb{S}^d, \mathcal{B}(\mathbb{S}^d))$ one has

$$\mathcal{W}_1(\mu P^n, \nu) \leq \rho^n \mathcal{W}_1(\mu, \nu), \quad n \in \mathbb{N}_0 \tag{3.2}$$

(see e.g. [1, Theorem 20.3.4]), where μP^n corresponds to the distribution of X_n when $(X_n)_{n \in \mathbb{N}_0}$ is a Markov chain with initial distribution μ and transition kernel P . Moreover,

⁵Note that we know the invariant distribution of the transition kernel P defined in (1.2) to be $\nu := \mathcal{U}(\mathbb{S}^d)$, since the kernel coincides with that of GSSS for target distribution $\mathcal{U}(\mathbb{S}^d)$ (cf. Section 2) and GSSS always leaves its target distribution invariant [2, Lemma 14].

it is well-known (see e.g. [5, Proposition 30]) that, under the same assumptions, one gets

$$\text{gap}_\nu(P) \geq 1 - \rho, \tag{3.3}$$

where $\text{gap}_\nu(P)$ denotes the $(L_2(\nu))$ -spectral gap of P . The spectral gap is known to be a key property of Markov chains, as it quantifies both the convergence speed in terms of the total variation distance and the asymptotic sample quality in terms of Monte Carlo integration errors. For details, we refer to [8, Sections 1-2] and the references therein. Simply put, an explicit upper bound on the *Wasserstein contraction rate* (we use this term synonymously to Dobrushin coefficient) of a kernel P yields not only the geometric convergence (3.2) but also the explicit bound (3.3) on the spectral gap and by extension the various additional results such a bound entails.

4 Main result

Using the concepts presented in Section 3, we can now formulate our main result:

Theorem 4.1. *For any fixed $d \geq 2$, the transition kernel P defined in (1.2), which describes the geodesic random walk on \mathbb{S}^d , is Wasserstein contractive with rate bound*

$$\rho := \frac{1}{2\pi} \int_0^{2\pi} \sqrt{\cos^2(\omega) + \frac{\sin^2(\omega)}{d}} \, d\omega. \tag{4.1}$$

Clearly, this bound is strictly decreasing in d . Thus, by computing its value for the smallest permissible dimension, $d = 2$, we obtain a dimension-independent bound on the rate that applies in all dimensions $d \geq 2$. Numerical integration shows that for $d = 2$ the rate bound (4.1) is approximately 0.86. Rounding this up to a nicer value, we end up with the following corollary.

Corollary 4.2. *For any fixed $d \geq 2$, the transition kernel P defined in (1.2) is Wasserstein contractive with rate bound $\rho := 7/8$.*

We postpone any further discussion of Theorem 4.1 and Corollary 4.2 to the next section and focus here on proving the theorem. Fix P to be as in (1.2) for the remainder of the section. Before we delve into any computations, we observe that the inherent symmetry of \mathbb{S}^d and the random walk kernel (1.2) on it allow us to make some important simplifications to the definition (3.1) of $\text{Dob}(P)$ in our particular case. Namely, thanks to the rotational symmetry⁶, $\mathcal{W}_1(P(x, \cdot), P(y, \cdot))$ for any $x, y \in \mathbb{S}^d$ can only depend on the (Euclidean) distance between x and y , not on their absolute positions on the sphere. It is therefore sufficient to arbitrarily fix $x \in \mathbb{S}^d$ and only consider points y from the image of a geodesic⁷ $g_{x,z}$ through x from which x itself is removed, i.e. $G_{x,z} \setminus \{x\}$. As we may freely choose both x and z , we choose them as points whose Euclidean coordinates are convenient to our proof, specifically

$$\begin{aligned} x &:= (1, 0, \dots, 0)^\top \in \mathbb{S}^d, \\ z &:= (0, 1, 0, \dots, 0)^\top \in \mathbb{S}_x^{d-1}. \end{aligned} \tag{4.2}$$

Note that this choice of z yields

$$G_{x,z} \setminus \{x\} = \{y_\alpha := (\cos(\alpha), \sin(\alpha), 0, \dots, 0)^\top \in \mathbb{S}^d \mid \alpha \in]0, 2\pi[\}. \tag{4.3}$$

Using the definitions in (4.2), (4.3), we may now rewrite $\text{Dob}(P)$ as

$$\text{Dob}(P) = \sup_{\alpha \in]0, 2\pi[} \frac{\mathcal{W}_1(P(x, \cdot), P(y_\alpha, \cdot))}{\|x - y_\alpha\|}. \tag{4.4}$$

⁶Specifically, due to the symmetry of \mathbb{S}^d as a Riemannian manifold, the transition probabilities of P are unaffected by rotations around the origin (applied to both its arguments).

⁷A suitably chosen half of a geodesic would also suffice.

It remains to show that (4.4) is upper-bounded by (4.1). For the sake of a simpler proof structure, we first establish the following auxiliary result.

Lemma 4.3. *For any $\alpha \in \mathbb{R}$, the rotation matrix*

$$R_\alpha = \begin{pmatrix} \cos(\alpha) & -\sin(\alpha) & & & \\ \sin(\alpha) & \cos(\alpha) & & & \\ & & 1 & & \\ & & & \ddots & \\ & & & & 1 \end{pmatrix} \in \mathbb{R}^{(d+1) \times (d+1)}$$

defines a transport map between $P(x, \cdot)$ and $P(y_\alpha, \cdot)$ via $v \mapsto R_\alpha v$. That is, the probability measure γ_α on $\mathbb{S}^d \times \mathbb{S}^d$ determined by

$$\gamma_\alpha(A \times B) := \int_{\mathbb{S}^d} \mathbb{1}_A(v) \mathbb{1}_B(R_\alpha v) P(x, dv), \quad A, B \in \mathcal{B}(\mathbb{S}^d)$$

is a coupling between $P(x, \cdot)$ and $P(y_\alpha, \cdot)$.

Proof. We immediately get

$$\gamma_\alpha(A \times \mathbb{S}^d) = \int_{\mathbb{S}^d} \mathbb{1}_A(v) P(x, dv) = P(x, A), \quad A \in \mathcal{B}(\mathbb{S}^d),$$

which is already half of what we need to show. The second half turns out to be slightly more complicated but still straightforward. For the sake of brevity, we denote by $c_d := (2\pi \mathfrak{s}_{d-1})^{-1}$ the normalization constant of P . We then get

$$\begin{aligned} \gamma_\alpha(\mathbb{S}^d \times B) &= \int_{\mathbb{S}^d} \mathbb{1}_B(R_\alpha v) P(x, dv) = \int_{\mathbb{S}^d} \mathbb{1}_{R_{-\alpha}(B)}(v) P(x, dv) \\ &= P(x, R_{-\alpha}(B)) = c_d \int_{\mathbb{S}_x^{d-1}} \int_0^{2\pi} \mathbb{1}_{R_{-\alpha}(B)}(g_{x,\tilde{z}}(\omega)) d\omega \sigma_{d-1}^{(x)}(d\tilde{z}) \\ &= c_d \int_{\mathbb{S}_x^{d-1}} \int_0^{2\pi} \mathbb{1}_{R_{-\alpha}(B)}(\cos(\omega)x + \sin(\omega)\tilde{z}) d\omega \sigma_{d-1}^{(x)}(d\tilde{z}) \\ &= c_d \int_{\mathbb{S}_x^{d-1}} \int_0^{2\pi} \mathbb{1}_B(\cos(\omega)R_\alpha x + \sin(\omega)R_\alpha \tilde{z}) d\omega \sigma_{d-1}^{(x)}(d\tilde{z}) \\ &= c_d \int_{\mathbb{S}_{y_\alpha}^{d-1}} \int_0^{2\pi} \mathbb{1}_B(\cos(\omega)y_\alpha + \sin(\omega)\tilde{z}) d\omega ((R_\alpha)_\# \sigma_{d-1}^{(x)})(d\tilde{z}) \\ &= c_d \int_{\mathbb{S}_{y_\alpha}^{d-1}} \int_0^{2\pi} \mathbb{1}_B(\cos(\omega)y_\alpha + \sin(\omega)\tilde{z}) d\omega \sigma_{d-1}^{(y_\alpha)}(d\tilde{z}) \\ &= P(y_\alpha, B) \end{aligned}$$

for any $B \in \mathcal{B}(\mathbb{S}^d)$. Note that we make repeated use of $R_\alpha x = y_\alpha$ (which trivially follows from their respective definitions) and that the second to last equation relies on the fact that R_α is an isometry from \mathbb{S}_x^{d-1} to $\mathbb{S}_{y_\alpha}^{d-1}$. \square

With the identity (4.4) and Lemma 4.3 at hand, we are now well-equipped to prove the theorem.

Proof of Theorem 4.1. Fix $\alpha \in]0, 2\pi[$ for now. Observe that for $v = (v_1, \dots, v_{d+1})^\top \in \mathbb{S}^d$, one has

$$R_\alpha v = (\cos(\alpha)v_1 - \sin(\alpha)v_2, \sin(\alpha)v_1 + \cos(\alpha)v_2, v_3, \dots, v_{d+1})^\top,$$

so that

$$\begin{aligned} \|v - R_\alpha v\|^2 &= ((1 - \cos(\alpha))v_1 + \sin(\alpha)v_2)^2 + ((1 - \cos(\alpha))v_2 - \sin(\alpha)v_1)^2 \\ &= (1 - \cos(\alpha))^2 v_1^2 + 2(1 - \cos(\alpha)) \sin(\alpha) v_1 v_2 + \sin^2(\alpha) v_2^2 \\ &\quad + (1 - \cos(\alpha))^2 v_2^2 - 2(1 - \cos(\alpha)) \sin(\alpha) v_1 v_2 + \sin^2(\alpha) v_1^2 \\ &= ((1 - \cos(\alpha))^2 + \sin^2(\alpha))(v_1^2 + v_2^2) \\ &= (1 - 2 \cos(\alpha) + \cos^2(\alpha) + \sin^2(\alpha))(v_1^2 + v_2^2) \\ &= 2(1 - \cos(\alpha))(v_1^2 + v_2^2). \end{aligned}$$

Moreover, since $R_\alpha x = y_\alpha$, plugging $v := x$ (still defined as in (4.2)) into the above computation yields

$$\|x - y_\alpha\|^2 = 2(1 - \cos(\alpha)).$$

Applying Lemma 4.3 and then plugging in the above, we obtain

$$\begin{aligned} \frac{\mathcal{W}_1(P(x, \cdot), P(y_\alpha, \cdot))}{\|x - y_\alpha\|} &\leq \frac{1}{\|x - y_\alpha\|} \int_{\mathbb{S}^d \times \mathbb{S}^d} \|v - \tilde{v}\| \gamma_\alpha(\mathbf{d}v \times \mathbf{d}\tilde{v}) \\ &= \frac{1}{\|x - y_\alpha\|} \int_{\mathbb{S}^d} \|v - R_\alpha v\| P(x, \mathbf{d}v) \\ &= \int_{\mathbb{S}^d} \sqrt{v_1^2 + v_2^2} P(x, \mathbf{d}v). \end{aligned}$$

Remarkably, the final expression above does not depend on α anymore. Therefore, combining this result with (4.4), we immediately get

$$\text{Dob}(P) \leq \int_{\mathbb{S}^d} \sqrt{v_1^2 + v_2^2} P(x, \mathbf{d}v). \tag{4.5}$$

In the remainder of the proof, we further simplify the above integral. To increase legibility, we rely on the following notation: For a probability space (H, \mathcal{H}, ξ) and a measurable function $f : (H, \mathcal{H}) \rightarrow (\mathbb{R}, \mathcal{B}(\mathbb{R}))$, we denote by

$$\mathbb{E}_{W \sim \xi}(f(W)) := \int_H f(w) \xi(\mathbf{d}w)$$

the expectation of the transformation $f(W)$ of a random variable $W \sim \xi$.

By definition, points in \mathbb{S}_x^{d-1} are perpendicular to x , so that

$$u_1 = x^\top u = 0 \quad \forall u = (u_1, \dots, u_{d+1})^\top \in \mathbb{S}_x^{d-1}. \tag{4.6}$$

Moreover, $\mathcal{U}(\mathbb{S}_x^{d-1})$ is invariant under permutations of the vector components at indices 2 to $d + 1$, so that we must have

$$1 = \mathbb{E}_{U \sim \mathcal{U}(\mathbb{S}_x^{d-1})}(\|U\|^2) \stackrel{(4.6)}{=} \sum_{i=2}^{d+1} \mathbb{E}_{U \sim \mathcal{U}(\mathbb{S}_x^{d-1})}(U_i^2) = d \cdot \mathbb{E}_{U \sim \mathcal{U}(\mathbb{S}_x^{d-1})}(U_2^2),$$

meaning

$$\mathbb{E}_{U \sim \mathcal{U}(\mathbb{S}_x^{d-1})}(U_2^2) = \frac{1}{d}. \tag{4.7}$$

Finally, we bound $\text{Dob}(P)$ by applying, in this order (one bullet point per (in)equality),

- inequality (4.5)
- the definitions (1.2), (1.1) of P and $g_{x,u}$
- the definition (4.2) of x and equation (4.6)
- Fubini's theorem and the reformulation of an integral as an expectation
- the fact that $r \mapsto \sqrt{\cos^2(\omega) + \sin^2(\omega)r}$ is a concave function on $[0, \infty[$ (for any ω), Jensen's inequality and monotonicity of the integral
- equation (4.7)

and obtain

$$\begin{aligned}
 \text{Dob}(P) &\leq \int_{\mathbb{S}^d} \sqrt{v_1^2 + v_2^2} P(x, dv) \\
 &= \frac{1}{2\pi \mathfrak{s}_{d-1}} \int_{\mathbb{S}_x^{d-1}} \int_0^{2\pi} \sqrt{\sum_{i=1}^2 (\cos(\omega)x_i + \sin(\omega)u_i)^2} d\omega \sigma_{d-1}^{(x)}(du) \\
 &= \frac{1}{2\pi \mathfrak{s}_{d-1}} \int_{\mathbb{S}_x^{d-1}} \int_0^{2\pi} \sqrt{\cos^2(\omega) + \sin^2(\omega)u_2^2} d\omega \sigma_{d-1}^{(x)}(du) \\
 &= \frac{1}{2\pi} \int_0^{2\pi} \mathbb{E}_{U \sim \mathcal{U}(\mathbb{S}_x^{d-1})} \left(\sqrt{\cos^2(\omega) + \sin^2(\omega)U_2^2} \right) d\omega \\
 &\leq \frac{1}{2\pi} \int_0^{2\pi} \sqrt{\cos^2(\omega) + \sin^2(\omega)\mathbb{E}_{U \sim \mathcal{U}(\mathbb{S}_x^{d-1})}(U_2^2)} d\omega \\
 &= \frac{1}{2\pi} \int_0^{2\pi} \sqrt{\cos^2(\omega) + \frac{\sin^2(\omega)}{d}} d\omega. \quad \square
 \end{aligned}$$

5 Discussion

Let us conclude the paper with some brief comments on two aspects of our main result, Theorem 4.1. Firstly, in the limit $d \rightarrow \infty$, the rate bound (4.1) simplifies to

$$\rho = \frac{1}{2\pi} \int_0^{2\pi} |\cos(\omega)| d\omega = \frac{2}{\pi} \approx 0.64.$$

As can be seen in Figure 1, the value of (4.1) takes quite some time to converge to this asymptotic rate bound, only becoming “visually indistinguishable” from it when the dimension is in the order of hundreds or thousands.




Figure 1: Convergence of the contraction rate bound (4.1) to its asymptotic value. The blue line represents the value of (4.1) at each integer dimension $d = 2, 3, \dots, 10^3$. The dashed black line marks the asymptotic rate bound $2/\pi$.

Secondly, the rate bound (4.1) being strictly decreasing means that Theorem 4.1 actually yields *stronger* guarantees as the dimension d increases. This suggests that we may be dealing with a case of the blessing of dimensionality, where increases in the dimension lead to improvements in a method’s performance. However, we believe this effect to merely be a peculiar consequence of our proof strategy, and conjecture that the true contraction rate of the random walk (i.e. its Dobrushin coefficient) is in fact either constant in the dimension d or increasing in it.

References

- [1] Randal Douc, Eric Moulines, Pierre Priouret, and Philippe Soulier, *Markov chains*, Springer Series in Operations Research and Financial Engineering, Springer, Cham, 2018. MR3889011
- [2] Michael Habeck, Mareike Hasenpflug, Shantanu Kodgirwar, and Daniel Rudolf, *Geodesic slice sampling on the sphere*, arXiv preprint arXiv:2301.08056v2, 2023.
- [3] Oren Mangoubi and Aaron Smith, *Rapid mixing of geodesic walks on manifolds with positive curvature*, *The Annals of Applied Probability* **28** (2018), no. 4, 2501–2543. MR3843835
- [4] Radford M. Neal, *Sampling from multimodal distributions using tempered transitions*, *Statistics and Computing* **6** (1996), 353–363.
- [5] Yann Ollivier, *Ricci curvature of Markov chains on metric spaces*, *Journal of Functional Analysis* **256** (2009), no. 3, 810–864. MR2484937
- [6] Sam Power, Daniel Rudolf, Björn Sprungk, and Andi Q. Wang, *Weak poincaré inequality comparisons for ideal and hybrid slice sampling*, arXiv preprint arXiv:2402.13678, 2024.
- [7] Gareth O. Roberts and Jeffrey S. Rosenthal, *The polar slice sampler*, *Stochastic Models* **18** (2002), no. 2, 257–280. MR1904830
- [8] Daniel Rudolf and Philip Schär, *Dimension-independent spectral gap of polar slice sampling*, *Statistics and Computing* **34** (2024), no. 1. MR4662468
- [9] Philip Schär, *Wasserstein contraction and spectral gap of slice sampling revisited*, *Electronic Journal of Probability* **28** (2023), 1–28, article 136. MR4664458
- [10] Philip Schär, Michael Habeck, and Daniel Rudolf, *Gibbsian polar slice sampling*, *Proceedings of the 40th International Conference on Machine Learning*, *Proceedings of Machine Learning Research*, vol. 202, PMLR, 2023, pp. 30204–30223.
- [11] Philip Schär, Michael Habeck, and Daniel Rudolf, *Parallel affine transformation tuning of Markov chain Monte Carlo*, *Proceedings of the 41st International Conference on Machine Learning*, *Proceedings of Machine Learning Research*, vol. 235, PMLR, 2024, pp. 43571–43607.

Acknowledgments. We are thankful for the helpful feedback of an anonymous referee. We thank Björn Sprungk for the discussion that led to our investigation, in particular for drawing our attention to the question of dimension dependence of GSSS for a constant target, in the context of its relevance for GPSS. We thank Mareike Hasenpflug for discussions on the topic and helpful comments on a preliminary version of the manuscript.

Electronic Journal of Probability

Electronic Communications in Probability

Advantages of publishing in EJP-ECP

- Very high standards
- Free for authors, free for readers
- Quick publication (no backlog)
- Secure publication (LOCKSS¹)
- Easy interface (EJMS²)

Economical model of EJP-ECP

- Non profit, sponsored by IMS³, BS⁴, ProjectEuclid⁵
- Purely electronic

Help keep the journal free and vigorous

- Donate to the IMS open access fund⁶ (click here to donate!)
- Submit your best articles to EJP-ECP
- Choose EJP-ECP over for-profit journals

¹LOCKSS: Lots of Copies Keep Stuff Safe <http://www.lockss.org/>

²EJMS: Electronic Journal Management System: <https://vtex.lt/services/ejms-peer-review/>

³IMS: Institute of Mathematical Statistics <http://www.imstat.org/>

⁴BS: Bernoulli Society <http://www.bernoulli-society.org/>

⁵Project Euclid: <https://projecteuclid.org/>

⁶IMS Open Access Fund: <https://imstat.org/shop/donation/>