1.2 Optimization and Uniform Point Sets on the Sphere

Holger Heitsch and René Henrion

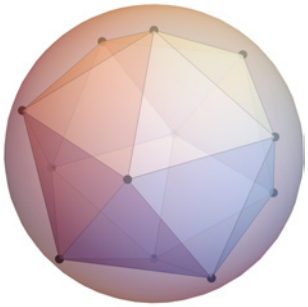

Fig. 1: *Icosahedron inscribed to a sphere such that its 12 vertices yield perfectly uniformly distributed points*

How to distribute points uniformly on a sphere may be considered as a question of mathematical interest even without having in mind any concrete application. A first idea could be related to perfectly regular point constellations as defined by the five platonic solids; see Figure 1. This idea, however, is limited to just a few number of points (4, 6, 8, 12, and 20), moreover, restricted to the classical two-dimensional sphere $\mathbb{S}^2$ in three-dimensional space. On the other hand, large samples of uniformly distributed points on spheres $\mathbb{S}^d$ of arbitrary dimension d (see Figure 2) are important in many problems of physics, chemistry, climate science, engineering and, not to the least, mathematics itself. The starting point for finding uniform point sets on the sphere is to define a criterion measuring uniformity. Such criterion is by no means uniquely defined and so, different criteria would produce different optimal point sets. Two prominent representatives are

$$\varphi_1(X) := \min_{1\le i<j\le N} \|x_i - x_j\|; \quad \varphi_2(X) := \sum_{i,j=1,\dots N, i\neq j} \frac{1}{\|x_i - x_j\|},$$

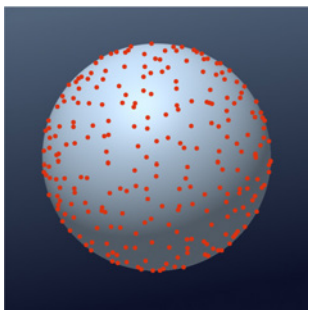

Fig. 2: *Low-discrepancy point set on the sphere*

where $X = \{x_1, \dots, x_N\}$ is a given set of points in $\mathbb{S}^d$. The first criterion measures the minimum pairwise distance of points, which increases with uniformity, while the second criterion is essentially the *Coulomb potential energy*. The latter is minimized by electrons located on the sphere and repulsing each other so that the result is a uniform constellation. In both cases, the relation with optimization is clear, the first criterion has to be maximized, the second one to be minimized over all point sets X on the sphere. The focus here shall be shifted to another measure of uniformity, the so-called *spherical cap discrepancy*, which plays an important role in bounding the integration error of spherical integrals (e.g., approximation of the global mean surface temperature on earth by an average of temperatures measured at finitely many points). Such integrals are important in optimization problems with probabilistic constraints, so that optimization benefits from point sets uniformly distributed on the sphere. On the other hand, contrary to φ_1 and φ_2, the discrepancy criterion is not directly computable by an explicit formula. It has been primarily used so far in the context of theoretical estimates, but not for direct evaluation of concrete samples. It is here that optimization in turn can help to provide an enumerative formula for numerical purposes. Both aspects of connecting optimization with uniform point sets on the sphere will be illustrated in the following.

The spherical cap discrepancy

While the criteria φ_1 and φ_2 mentioned above are based on geometrical or physical intuition about uniformity, the spherical cap discrepancy is related to probability measures by comparing the empirical distribution generated by the given point set with the uniform distribution (normalized surface measure) of $\mathbb{S}^d$. More precisely, for a fairly uniformly distributed collection of points it should hold true that for every hyperplane cutting the sphere, the relative share of points falling

into either of the two arising *spherical caps* is close to the surface measure of the respective cap divided by the surface measure of the whole sphere. In order to specify one out of the two possible caps, one defines a spherical cap as the intersection $\mathbb{S}^d \cap H(w,t)$ of the sphere with the closed half-space

$$H(w,t) := \{x \mid \langle w, x\rangle \geq t\} \quad (w \in \mathbb{S}^d,\ t \in [0,1]),$$

which is generated by a normal vector w and a level t. Figure 3 illustrates some spherical caps for $\mathbb{S}^1$ and $\mathbb{S}^2$.

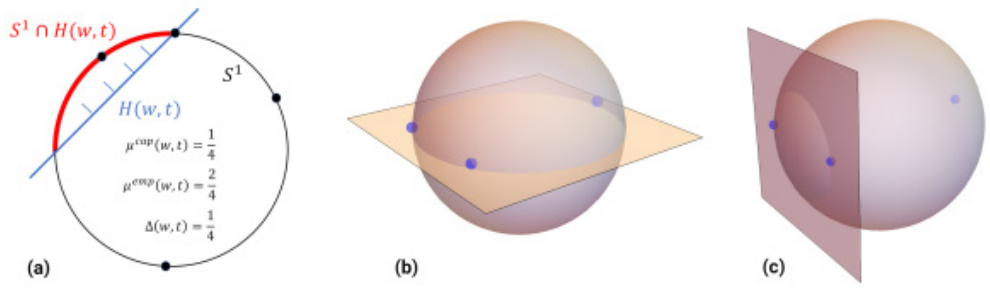

Fig. 3: *Some spherical caps for $\mathbb{S}^1$ and $\mathbb{S}^2$*

In this way, the family of spherical caps can be parameterized by $(w,t) \in \mathbb{S}^d \times [0,1]$, and each such cap can be assigned a *local discrepancy* $\Delta(w,t) := |\mu^{emp}(w,t) - \mu^{cap}(w,t)|$. Here, μ^{emp} is the empirical measure induced by the given point set on the sphere (counting the ratio of points inside the spherical cap or just inside the closed half-space), while μ^{cap} refers to the normalized surface measure providing the ratio between the area of the cap and that of the whole sphere; see Figure 3 (a) for an illustration of these notions. Both quantities, and, hence, $\Delta(w,t)$ are easy to determine. The overall (global) discrepancy between empirical and uniform distribution then results as the largest (supremum) of all the local discrepancies:

$$\Delta := \sup_{w \in \mathbb{S}^d, t \in [0,1]} \Delta(w,t).$$

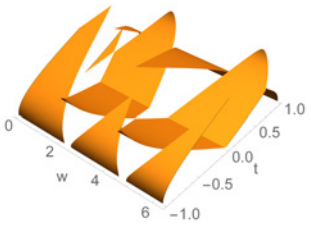

Fig. 4: *Local discrepancy for two points on $\mathbb{S}^1 \sim [0, 2\pi]$*

Unfortunately, $\Delta(w,t)$ is a highly irregular function (see Figure 4) and so, determining its largest value by applying some numerical optimization seems out of hope, much less a direct formula is evident. It is not even clear whether the supremum in the definition of Δ is attained, i.e., whether there exists some concrete spherical cap realizing this value. Indeed, despite the compactness of the set $\mathbb{S}^d \times [0,1]$, the application of the classical Weierstrass theorem fails due to $\Delta(w,t)$ not being upper semicontinuous. Nonetheless, one may show by independent arguments ([4, Prop. 1]) that there does exist a spherical cap realizing the largest value, i.e., $\Delta = \Delta(w^*,t^*)$ for some $(w^*,t^*) \in \mathbb{S}^d \times [0,1]$. Moreover, the boundary of the half-space $H(w^*,t^*)$ associated with this critical cap must contain at least one point of the given set, i.e., $\langle w^*, x^i\rangle = t^*$ for some $i \in \{1,\dots,N\}$. The computation of the global discrepancy drastically simplifies in the special case of the one-dimensional sphere $\mathbb{S}^1$ (circle). We already know that the straight line associated with the critical cap contains at least one point. If it does not contain any other point, then it has to be tangent to the circle because otherwise – after turning the straight line a bit while keeping the point on it – the discrepancy could be locally increased by locally increasing or decreasing μ^{cap} without changing μ^{emp} (see Figure 3 (a)). However, the larger of the two discrepancies defined by

a tangent through one point equals $1/N$, whereas the larger of the two discrepancies defined by a straight line passing through the two closest points on the circle is easily seen to be at least 1/N. Hence, there always exists a critical cap on $\mathbb{S}^1$ associated with a straight line passing through two distinct points. This fact allows one to enumerate the finite number of local discrepancies defined by couples of points in order to find the global discrepancy as their maximum value. Unfortunately, the idea of finding the critical cap as being defined among the finitely many hyperplanes related to maximal affinely independent subsets of $\{x_1, \dots, x_N\}$ fails for higher-dimensional spheres. This can be seen from Figure 3 (b), where a set of three points on the equator of $\mathbb{S}^2$ is given. The hyperplane passing through these points defines two spherical caps, each of which has local discrepancy $1/2$. On the other hand, the hyperplane passing through the two points on the left-hand side (see Figure 3 (c)) defines, in particular, a small cap to the left whose local discrepancy tends to $2/3$ when the two points converge to the "left pole" along the equator, whereas the previous hyperplane would still have local discrepancy $1/2$. Hence, the maximum discrepancy may not be realized by a hyperplane passing through a subset of $\{x_1, \dots, x_N\}$ with a maximum number (here: three) of affinely independent points.

An enumerative formula for the spherical cap discrepancy

The last observation means that a hyperplane defining a spherical cap with maximum local discrepancy may not be fixed by the set of points x^i it contains. Hence, there may remain a degree of freedom that prevents the global discrepancy Δ from being calculated via straightforward enumeration. This degree of freedom may be removed, however, by an argument from optimization. An essential observation in this direction is that a spherical cap realizing the maximum discrepancy always has an empirical measure not smaller than its normalized surface measure ([4, Cor. 1]). This fact allows one to get rid of the absolute value appearing in the local discrepancies and to write now

$$\Delta := \max_{w\in\mathbb{S}^d, t\in[0,1]} \Delta(w,t) = \max_{w\in\mathbb{S}^d, t\in[0,1]} \mu^{emp}(w,t) - \mu^{cap}(w,t).$$

Here, we already exploited that the supremum in the original definition of Δ is actually a maximum. Now, assume that (w^*, t^*) defines a critical cap ($\Delta = \Delta(w^*, t^*)$) and denote by $I := \{i | \langle w^*, x^i\rangle = t^*\}$ the index set of points x^i located on the hyperplane associated with this cap. Clearly, (w^*, t^*) must be a solution to the maximization problem

$$\max_{w\in\mathbb{S}^d, t\in[0,1]} \mu^{emp}(w,t) - \mu^{cap}(w,t) \quad \text{subject to} \quad \langle w, x^i\rangle = t \quad (i \in I),$$

because (w^*, t^*) itself satisfies the constraints of this problem. Evidently, for (w,t) satisfying these constraints and being close to (w^*, t^*), no additional points may enter the hyperplane, which implies that $\mu^{emp}(w^*, t^*) = \mu^{emp}(w,t)$ (empirical measure locally constant). Therefore, (w^*, t^*) must be a *local* solution to the maximization problem just for the negative spherical measure:

$$\max_{w\in\mathbb{S}^d, t\in[0,1]} -\mu^{cap}(w,t) \quad \text{subject to} \quad \langle w, x^i\rangle = t \quad (i \in I).$$

Observing that the surface measure μ^{cap} just depends on t (not on w) and actually is monotonically decreasing in t, one concludes that (w^*, t^*) must be a *local* solution to the problem

$$\max_{w\in\mathbb{S}^d, t\in[0,1]} t \quad \text{subject to} \quad \langle w, x^i\rangle = t \quad (i \in I). \tag{*}$$

This optimization problem (see Figure 5) is simple enough to identify its solutions from the necessary optimality conditions. More precisely, denote by X_I the matrix whose columns are generated by x^i for $i \in I$ and by $\tilde{X}_I := \begin{pmatrix} X_I \\ -\mathbf{1}^T \end{pmatrix}$ the extended matrix with $\mathbf{1} := (1, \dots, 1)^T$. Without loss of generality, we may assume that $\operatorname{rank} \tilde{X}_I = \#I$. Then, the critical cap is characterized explicitly by ([4, Lemma 3])

$$t^* = \left(\frac{1-\gamma_I}{\gamma_I}\right)^{1/2}, \quad w^* = \frac{1+(t^*)^2}{t^*} X_I \left(\tilde{X}_I^T \tilde{X}_I\right)^{-1} \mathbf{1} \quad (\gamma_I := \mathbf{1}^T \left(\tilde{X}_I^T \tilde{X}_I\right)^{-1} \mathbf{1}).$$

One may show that $0 < \gamma_I \le 1$. Clearly, the solution above is not defined for $\gamma_I = 1$, which is equivalent to $t^* = 0$. This corresponds to a hyperplane passing through the origin or a cap that is a hemisphere. This case is technically more delicate to treat. For randomly generated point sets as in Figure 6, one will have $t^* \neq 0$ with probability one, but for constructed sets as in Figure 2, this degenerate case may be relevant as well. Altogether, we arrive at the following explicit enumerative formula for the spherical cap discrepancy Δ ([4, Th. 1]) for which a MATLAB implementation is available at https://www.wias-berlin.de/people/heitsch/capdiscrepancy.

Keeping the notation introduced above, define

$$\begin{aligned} \Phi_1 &:= \left\{ I \subseteq \{1, \dots, N\} \,\middle|\, 1 \le \operatorname{rank} \tilde{X}_I = \#I \le \min\left\{n, \operatorname{rank} \tilde{X}\right\};\ \gamma_I < 1 \right\}, \\ \Phi_0 &:= \left\{ I \subseteq \{1, \dots, N\} \,\middle|\, \operatorname{rank} \tilde{X}_I = \#I = \min\left\{n, \operatorname{rank} \tilde{X}\right\};\ \gamma_I = 1 \right\}, \end{aligned}$$

$$t_I := \begin{cases} \left(\frac{1-\gamma_I}{\gamma_I}\right)^{1/2} & I \in \Phi_1 \\ 0 & I \in \Phi_0 \end{cases}, \quad w_I := \begin{cases} \frac{1+t_I^2}{t_I} X_I \left(\tilde{X}_I^T \tilde{X}_I\right)^{-1} \mathbf{1} & I \in \Phi_1 \\ \in \operatorname{Ker} X_I^T \cap \mathbb{S}^{n-1} & I \in \Phi_0 \end{cases},$$

where the selection of w_I in case of $I \in \Phi_0$ is arbitrary. Then, $\Delta = \max\{\Delta_1, \Delta_0\}$, where

$$\Delta_1 := \max_{I\in\Phi_1} \Delta(w_I, t_I), \quad \Delta_0 := \max_{I\in\Phi_0} \max\{\Delta(w_I, 0), \Delta(-w_I, 0)\}.$$

The formula may be exploited now in order to evaluate and compare point sets on spheres with respect to their uniformity. Figure 7 shows normalized Monte Carlo and quasi-Monte Carlo samples of a Gaussian distribution with independent components (yielding a uniform distribution on the sphere) as well as a low-discrepancy sequence using the Lambert transformation [2] on $\mathbb{S}^2$. Figure 8 plots the true discrepancies and a simple lower bound used in [2, p. 1005] for randomly generated samples on $\mathbb{S}^4$.

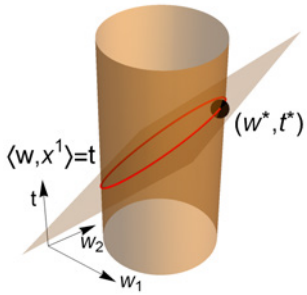

Fig. 5: *Optimization problem (*) in a trivial constellation ($\mathbb{S}^1$, $I = \{1\}$, $x^1 = (1, 0)$). The red curve represents the feasible set.*

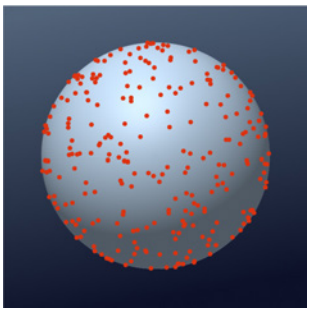

Fig. 6: *Set of randomly generated (Monte-Carlo) points on the sphere*

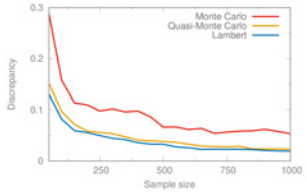

Fig. 7: *Discrepancies as functions of the sample size for three sampling methods on $\mathbb{S}^2$*

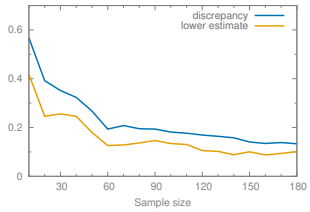

Fig. 8: *True discrepancies and simple lower estimates for randomly generated samples as functions of the sample size on* $\mathbb{S}^4$

Optimization under probabilistic constraints

Efficient samples of the uniform distribution on the sphere are in turn essential for the solution of optimization problems subject to probabilistic constraints

$$\text{minimize } f(x) \text{ subject to } \varphi(x) := \mathbb{P}(g_i(x,\xi) \geq 0 \ (i \in I)) \geq p, \qquad (**)$$

where x is a finite- or infinite-dimensional decision, f is some cost function, ξ is a random vector, and g represents some finite or infinite inequality system. Under such constraint, a decision x is declared to be feasible if the random inequality system is satisfied at least with a probability p. The essential ingredient of such optimization problems is the probability function φ whose values and gradients – needed in any numerical solution approach – are not given explicitly, but have to be approximated. In the case that ξ has an elliptically symmetric distribution, the probability can be represented as a spherical integral that promises a significant reduction of variance for the resulting estimate when compared to sampling "in space." For instance, if $\xi \sim \mathcal{N}(\mu, \Sigma)$ has a $d-$ dimensional normal distribution, then

$$\varphi(x) = \int_{\mathbb{S}^{d-1}} \alpha(z) d\nu^U(z); \quad \alpha(z) := \nu^\chi\{r \geq 0 \mid g_i(x, \mu + rLz) \geq 0 \ (i \in I)\},$$

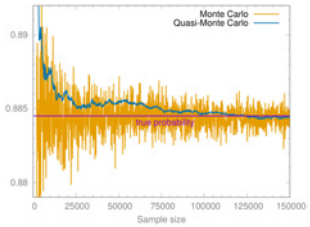

Fig. 9: *Probability estimated according to (***) (d=36) by Monte Carlo (Mersenne Twister) and quasi-Monte Carlo*

where $LL^T = \Sigma$, ν^U is the uniform distribution on $\mathbb{S}^{d-1}$, and ν^χ is the one-dimensional Chi-distribution with d degrees of freedom. Under suitable conditions on g (growth conditions, constraint qualifications), α can be shown to be differentiable [1], so that $\nabla\phi$ is obtained as a spherical integral as well by differentiating under the integral. For numerical purposes the spherical integral is approximated by the average over a finite number of points uniformly sampled over the sphere:

$$\varphi(x) \approx N^{-1} \sum_{i=1}^{N} \alpha(z^i) \quad (z^i \in \mathbb{S}^{d-1}). \qquad (***)$$

It is here that the mentioned relation between the spherical cap discrepancy and the integration error for spherical integrals comes into play (in particular, the latter tends to zero if the former does so). As a consequence, the difference in discrepancy between Monte Carlo and quasi-Monte Carlo sampling, which is visible in Figure 7, reflects also in the goodness of estimations of probabilities (and their gradients) as supported by Figure 9. Optimization problems with probabilistic constraints like (**) find a lot of applications in engineering, economics, power management, telecommunications, and other fields. Figures 10 and 11 illustrate two instances from optimal control under uncertainty. In the first application, an optimal Neumann boundary control of the vibrating string is considered under random initial conditions (random Fourier coefficients for initial position of the string) [3]. The aim is to find a cost-minimal control driving the terminal energy of the string close to zero (smaller than a given tolerance) with given probability.

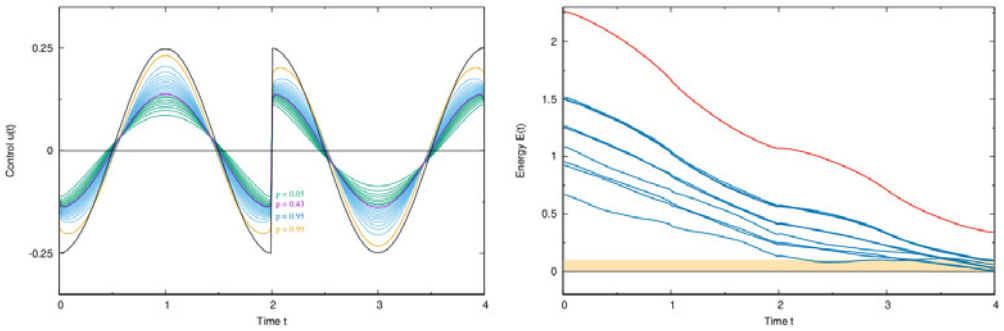

Fig. 10: Optimal control of the vibrating string under random initial condition and probabilistic constraint for terminal energy. The figure shows the optimal control functions (left) for different required probabilities $p \in (0,1)$ and the energy as function of time for $p = 0.9$ for ten generated scenarios of the initial condition (one instance not reaching the required small terminal energy).

The second application is the capacity maximization problem in gas networks under random loads. The network owner aims at documenting additional free capacity of the network for serving potential new clients under the constraint that the loads of new clients (arbitrary up to free capacity) plus the random loads of current clients (distribution estimated from historical data) can be served physically (here: respecting a lower pressure bound in pipes) with given high probability.

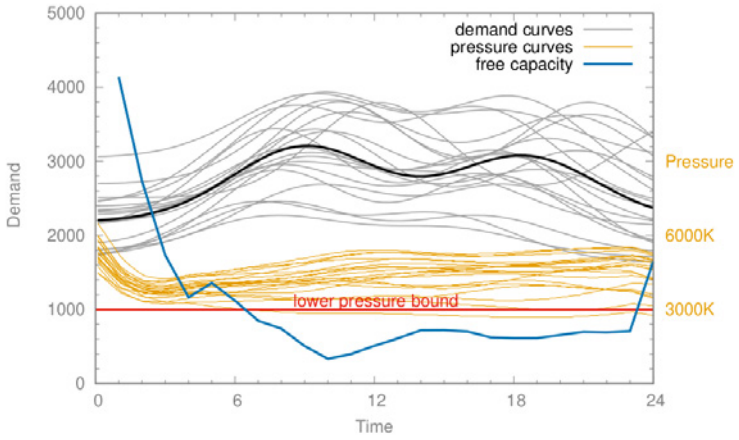

Fig. 11: Maximization of free capacity in gas networks. Mean (black) and simulated scenarios (gray) for load profiles of current clients. Resulting scenarios for pressure (yellow), most of which satisfy the lower bound.

References

[1] W. van Ackooij, R. Henrion, P. Pérez-Aros, *Generalized gradients for probabilistic/robust (probust) constraints*, Optimization, **69** (2020), pp. 1451–1479.

[2] C. Aistleitner, J. Brauchart, J. Dick, *Point sets on the sphere S2 with small spherical cap discrepancy*, Discrete Comput. Geom., **48** (2012), pp. 990–1024.

[3] M.H. Farshbaf Shaker, M. Gugat, H. Heitsch, R. Henrion, *Optimal Neumann boundary control of a vibrating string with uncertain initial data and probabilistic terminal constraints*, SIAM J. Control Optim., **58** (2020), pp. 2288–2311.

[4] H. Heitsch, R. Henrion, *An enumerative formula for the spherical cap discrepancy*, J. Comput. Appl. Math., **390** (2021), pp. 113409/1–113409/14.