

Super-Fibonacci Spirals: Fast, Low-Discrepancy Sampling of $SO(3)$

Marc Alexa
 TU Berlin

marc.alex@tu-berlin.de

Abstract

Super-Fibonacci spirals are an extension of Fibonacci spirals, enabling fast generation of an arbitrary but fixed number of 3D orientations. The algorithm is simple and fast. A comprehensive evaluation comparing to other methods shows that the generated sets of orientations have low discrepancy, minimal spurious components in the power spectrum, and almost identical Voronoi volumes. This makes them useful for a variety of applications, in particular Monte Carlo sampling.

1. Introduction

Methods for the generation of orientation or rotation samples are a fundamental building block in science and engineering. They are used in such diverse fields as protein docking [27], crystallography [31], electron cryo-microscopy [8, 48] or additive manufacturing [39], to name but a few. In vision, robotics, and learning they have been used, among others, in path planning [24, 46]. This work has been prompted in particular by the use of orientation samples for estimating pose or establishing rotation equivariance in machine learning methods. For example, convolutional neural networks are linear and have been found to work well with discretized rotations (CNN) [32, 53]. A recent trend in pose estimation is to learn pose distributions, requiring a sample set of orientations [38, 41]. More broadly, Markov-chain Monte Carlo (MCMC) methods work well in a variety of estimation and optimization tasks and are now used heavily in learning. They also apply to problems involving orientations [5] and would profit from using low discrepancy sample sets for driving the Monte Carlo sampler [43].

Computing with orientations requires a representation. In this work, we use the unit 3-sphere $\mathbb{S}^3 \subset \mathbb{R}^4$, which is a *double cover* of $SO(3)$. Every *unit quaternion* $\mathbf{q} \in \mathbb{R}^4$, $\|\mathbf{q}\| = 1$ identifies an orientation, and $-\mathbf{q}$ represents the same orientation. Sampling orientations in this representation means creating pairs of points in \mathbb{R}^4 with unit distance to the origin.

Rather than focusing on a particular application, we try to evaluate the quality of the set using geometric measures on \mathbb{S}^3 . There are various ways to measure ‘well-distribution’ on spheres and sets that are near-optimal in one sense may be far from optimal in another. Classical quality criteria for generating ‘optimal’ *packings* on the sphere [9] or *spherical designs/codes* [52] are based on local measures, i.e. the closest distance between any two orientations or the radius of empty disks – the so-called *dispersion*. Distributions that are optimal in the sense of such local measures often exhibit regularity that leads to aliasing in sampling applications. In Section 4 we discuss several measures that are better suited to measure the quality of an orientation sample and explain how to compute them. In particular, the *discrepancy* of a set is directly related to the error in quasi-Monte Carlo integration by the Koksma–Hlawka inequality [40]. Another classical tool for the analysis of sample sets is the *power spectrum*, describing the spectral properties of the set and revealing unwanted modes in the distribution. We also introduce the idea to project \mathbb{S}^3 onto the Clifford torus for visualizing aliasing artifacts.

Using these methods we compare existing approaches for the task of quickly generating well-distributed orientation samples (reviewed in Section 2) to sampling based on *Super-Fibonacci Spirals* – a new method based on the idea of *Fibonacci* sampling. We provide background and develop and describe the details of this method in Section 3. The resulting implementation is extremely simple, it supports sample refinement, it is faster than all existing methods, and the resulting orientation sets are well suited for sampling.

To summarize, we consider generating sample sets on $SO(3)$ a fundamental building block in vision, robotics, learning, and more generally science and engineering, to which we contribute in the following ways:

- We provide a suite of tools for analyzing the properties of sets of orientations for their use in sampling applications.
- We introduce Super-Fibonacci sampling, a method for generating an arbitrary number of samples on \mathbb{S}^3 . It is simple and very fast and generates high quality orientation samples represented as quaternions.

2. Methods for fast sampling of orientations

For most uses, we believe the following properties are useful: (1) fast sampling (2) an arbitrary number n of samples (3) that can be possibly refined to kn samples for arbitrary k . There are a variety of methods in the literature that fit these desiderata to different degrees.

Uniform. The most straightforward approach is using a pseudo-random uniform distribution on \mathbb{S}^3 . The points can be generated by drawing from any radially symmetric distribution in \mathbb{R}^4 and then normalizing the vector. It is common to use a Gaussian distribution [34]. Generating a sample requires constant time and space, with the constant dominated by the Box-Muller transform [6] necessary for mapping a uniform to a normal distribution. Uniform distributions can be easily refined.

Subdividing 3-polytopes. A common approach in the context of *spherical designs* is the subdivision of regular 3-polytopes [25]. Karney has analyzed the properties of such sets and provides a manual selection with near-optimal dispersion (available on github [26], under MIT license). Computing subdivisions is fast, and provides natural refinement, but severely limits the choice of the number of samples n . We restrict the comparison in Section 4 to the sets selected by Karney, as this is sufficient to illustrate that the regularity introduced by subdivision is unsuitable for sampling.

SOI. Successive Orthogonal Images (SOI) [35] provide a sampling method for $SO(3)$ that boils down to sampling \mathbb{S}^1 regularly and \mathbb{S}^2 as uniformly as possible, with the two sampling densities chosen so that the distance of points on \mathbb{S}^1 and \mathbb{S}^2 is approximately the same. The number of samples n is limited as it has to be the product of two integers that lead to similar spacing on the two sample spaces. This also severely limits refinement. The properties of the approach in terms of computational complexity as well as the resulting distribution depend on the method chosen for sampling \mathbb{S}^2 . We consider the two versions implemented in the `code` made publicly available (individual license terms: “copyrighted but free for commercial use”). The first variant uses subdivisions of an icosahedron. This leads to fast sampling, but introduces some of the unwanted regularity. The second variant uses an optimization process. This makes the method computationally expensive, with the complexity growing super-linearly in the number of samples. The computation time depends on the details of how convergence of the optimization is detected. We use the code as it is provided, which limits the maximal number of iterations, meaning it may favor smaller computation time over quality of the distribution for large sample sets. The results in Section 4 indicate that there is no significant effect on the quality even for large sets.

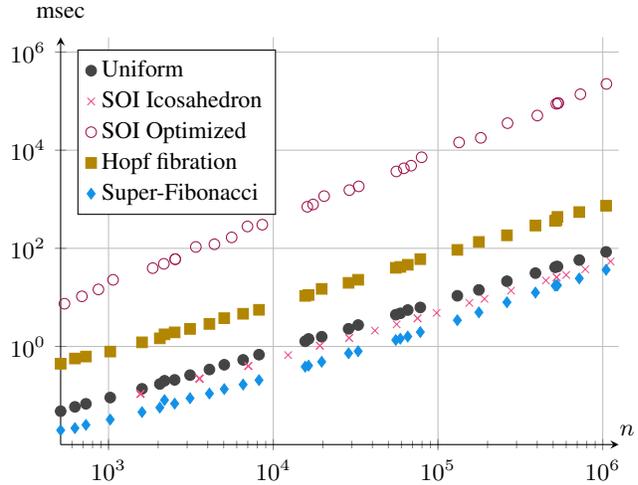


Figure 1. Wall-clock times for generating n orientation samples in quaternion representation. Only SOI with optimization on \mathbb{S}^2 shows super-linear time complexity.

Hopf fibration. The SOI approach has been modified using coordinates based on the Hopf fibration to allow constructing good orientation samples incrementally [58]. This means that sets of n and $n + 1$ samples only differ by a single additional point. The method has constant computational complexity per sample and works with an arbitrary number of samples. The incremental construction adversely impacts the quality of the distribution. `Code` is available under GNU public license.

Optimization in $SO(3)$. Methods that optimize distributions in $SO(3)$ or \mathbb{S}^3 are not included in the comparison because they are too time consuming: both Gräf and Potts [18] as well as Larsen and Schmidt [29] report *several days* of computation time. This is due to a combination of involved optimization with slow convergence and starting from many different configurations to avoid bad local minima.

Timing. Wall-clock timing for the methods that are provided with source code are shown in Figure 1. As values for $n \lesssim 10^6$ we use: (1) powers of 2, 3, and 5 to sample a sufficient number of values; (2) the *superior highly composite numbers* $n = \{60, 120, 360, 2520, 5040, 55440, 720720\}$, which have a large number of different divisors; (3) the *Fibonacci primes* $n = \{89, 233, 1597, 28657, 514229\}$. The latter two sets of values have been included to see if the properties of the different methods (discussed further below) depend on the divisibility of n . For the SOI methods, we use only the recommended values of n close to this set – this explains the different spacing in the scatter-plot in Figure 1. The analysis reveals that the SOI method using optimization is limited to smaller sample sets because of its super-linear time complexity.

3. Fibonacci sampling on \mathbb{S}^3

Optimizing sets for low discrepancy is difficult. Moreover, optimization in $SO(3)$ is cumbersome because of the underlying geometry. We show how to generate low-discrepancy sets of orientations based on extending the idea of *Fibonacci* sampling.

Fibonacci sampling in the plane. Let t_i be n regularly spaced values in $[0, 1)$, then the samples are given by

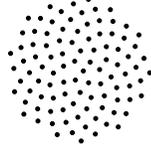
$$x_i = x(t_i), \quad x(t) = \left(\frac{nt}{\phi} - \left\lfloor \frac{nt}{\phi} \right\rfloor \right). \quad (1)$$

The quality of the sampling, i.e. the discrepancy, depends on the constant ϕ . It has been observed that setting $\phi = (1 + \sqrt{5})/2$, the golden ratio, leads to the best results in practice, hence the name. The appearance of the golden ratio may be considered natural, as it is the irrational number whose rational approximations converge slowest. While it is possible to derive upper bounds on the discrepancy based on the continued fraction representation of ϕ [40, Thm. 3.3], tight connections or why the golden ratio is optimal for finite samples remain unclear.

The Fibonacci sampling has been used to generate low-discrepancy sets for two-dimensional domains. Sampling the curve $(x(t), t)$ leads to an aperiodic sampling of the unit square. The area preserving map from $[0, 1]^2$ to the unit disk parameterized in polar coordinates $[0, 2\pi] \times [0, 1]$ given by $(\theta, r) = (2\pi x_0, \sqrt{x_1})$ converts this to a sampling of the unit disk, in cartesian coordinates:

$$\mathbf{y}(t) = \left(\sqrt{t} \sin \frac{2\pi nt}{\phi}, \sqrt{t} \cos \frac{2\pi nt}{\phi} \right). \quad (2)$$

Here, the periodicity of the sine and cosine function cause the floor function to disappear. The inset shows the resulting point distributions for 100 samples. The sampling of the unit disk can be mapped to the 2-sphere, resulting in the omnipresent Fibonacci (or golden) spiral on the sphere [17, 19, 45, 54].



Lifting to \mathbb{S}^3 . The main tool for using Fibonacci sampling on \mathbb{S}^3 is a volume preserving mapping from a solid cylinder in \mathbb{R}^3 to the 3-sphere. Consider the cylinder $\{(h, \mathbf{y} = (y_0, y_1)) \mid -\pi < h \leq \pi, \mathbf{y}^\top \mathbf{y} \leq 1\}$. Then

$$\mathbf{x}(h, \mathbf{y}) = \begin{pmatrix} z \cos h \\ z \sin h \\ y_0 \\ y_1 \end{pmatrix}, \quad z = \sqrt{1 - \mathbf{y}^\top \mathbf{y}} \quad (3)$$

maps points in the cylinder to the unit sphere in \mathbb{R}^4 . The inverse mapping of $\mathbf{x} = (x_0, x_1, x_2, x_3)$ is given by:

$$(h, y_0, y_1) = (\arctan 2(x_1, x_0), x_3, x_4). \quad (4)$$

This shows that the mapping is a bijection between the relative interior of the cylinder and the sphere without the 'equator' $x_0 = x_1 = 0$. The lines $\{-\pi < h \leq \pi, \mathbf{y}^\top \mathbf{y} = 1\}$ on the surface of the cylinder are mapped to the points $(0, 0, y_0, y_1)$ on the equator of the sphere. In the interior of the cylinder $\mathbf{y}^\top \mathbf{y} < 1$, i.e. where the mapping is bijective, we can compute the Jacobian as

$$\mathbf{J}_{h, \mathbf{y}} = \begin{pmatrix} -z \sin h & -\frac{y_0}{z} \cos h & -\frac{y_1}{z} \cos h \\ z \cos h & -\frac{y_0}{z} \sin h & -\frac{y_1}{z} \sin h \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad (5)$$

where we have used

$$\frac{\partial z}{\partial y_0} = -\frac{y_0}{z}, \quad \frac{\partial z}{\partial y_1} = -\frac{y_1}{z}. \quad (6)$$

Using the Jacobian, we can now analyze the change of volume and find the following:

Claim 1 *The mapping*

$$\mathbf{x}(h, \mathbf{y}) : \{(h, \mathbf{y}) \mid -\pi < h \leq \pi, \mathbf{y}^\top \mathbf{y} < 1\} \mapsto \mathbb{S}^3 \subset \mathbb{R}^4$$

is volume preserving.

The volume element can be computed as the product of the singular values of the Jacobian. Noting that all three columns of $\mathbf{J}_{h, \mathbf{y}}$ are tangents on the unit sphere in $\mathbf{x}(h, \mathbf{y})$, we can compute this as $\det(\mathbf{x}(h, \mathbf{y}), \mathbf{J}_{h, \mathbf{y}})$, because $\mathbf{x}(h, \mathbf{y})$ is orthogonal to the tangent plane and has unit length. Developing by the first column we find

$$\begin{aligned} & + z \cos h(z \cos h) - z \sin h(z \sin h) \\ & + y_0 (y_0 \sin^2 h + y_0 \cos^2 h) \\ & - y_1 (-y_1 \cos^2 h - y_1 \sin^2 h) \\ & = (1 - \mathbf{y}^\top \mathbf{y}) (\cos^2 h + \sin^2 h) + \mathbf{y}^\top \mathbf{y} = 1, \end{aligned} \quad (7)$$

as claimed. \square

Given the mapping, the main idea is to use Fibonacci sampling twice to generate points in the cylinder: (1) along the main axis h of the cylinder and (2) on the unit disk (y_0, y_1) orthogonal to the main axis. Using two different constants ϕ and ψ for the two samplings we get:

$$\mathbf{z}(t) = \left(\frac{nt}{\psi} - \left\lfloor \frac{nt}{\psi} \right\rfloor, \sqrt{t} \sin \frac{2\pi nt}{\phi}, \sqrt{t} \cos \frac{2\pi nt}{\phi} \right)^\top. \quad (8)$$

Plugging this sampling pattern into the mapping to the 3-sphere we find the following simple curve, exhibiting the expected symmetry:

$$\mathbf{w}(t) = \begin{pmatrix} \sqrt{t} \sin \frac{2\pi nt}{\phi} \\ \sqrt{t} \cos \frac{2\pi nt}{\phi} \\ \sqrt{1-t} \sin \frac{2\pi nt}{\psi} \\ \sqrt{1-t} \cos \frac{2\pi nt}{\psi} \end{pmatrix}. \quad (9)$$

Sampling this curve at regular values t_i is a natural method for generating orientation samples. For easy reproduction, Algorithm 1 provides the details for implementing this approach.

Algorithm 1: Generating n samples on $SO(3)$ as unit quaternions

Function Super-Fibonacci (n, ϕ, ψ)

```

for  $i \in \{0, \dots, n-1\}$  do
   $s \leftarrow i + \frac{1}{2}$ 
   $t \leftarrow \frac{s}{n}, d \leftarrow 2\pi s$ 
   $r \leftarrow \sqrt{t}, R \leftarrow \sqrt{1-t}$ 
   $\alpha \leftarrow \frac{d}{\phi}, \beta \leftarrow \frac{d}{\psi}$ 
   $\mathbf{q}_i \leftarrow (r \sin \alpha, r \cos \alpha, R \sin \beta, R \cos \beta,)$ 

```

This algorithm clearly has constant complexity per sample. In addition, it is faster in practice than all alternatives we have tested (c.f. Figure 1). The algorithm also directly shows that the a set of kn samples contains the set generated for n samples or, more generally, sets with m and n samples share every k -th sample, where k is the gcd of m and n , i.e. it is easy refine sets that have already been used for computation.

Parameters. An important factor for the quality of the sample set are the constants ϕ and ψ . Not only do we need ψ and ϕ to be irrational, also their relation is important: if they are rational multiples of each other the resulting sets are sub-optimal. Since mathematical theory connecting the properties of these numbers and the resulting sample sets is lacking, heuristic search and experimental exploration are the only remaining options. For this we have used the quality criteria that can be computed deterministically from the Voronoi diagram (see Section 4.3 and the supplementary material).

It may seem natural to set one of the constants to the golden ratio. In this scenario, a choice for the other constant that provides good results is the *super-golden ratio*, the only real solution to $\psi^3 = \psi^2 + 1$. This observation led to the name 'Super-Fibonacci'. A more exhaustive exploration of the roots of small degree polynomials, however, reveals slightly better options. The suggested choice are the positive real roots of:

$$\boxed{\phi^2 = 2, \quad \psi^4 = \psi + 4, \quad \phi, \psi \in \mathbb{R}^+} \quad (10)$$

Unlike $\phi = \sqrt{2}$, the solution for ψ has no simple expression, although as a depressed quartic it could be represented in terms of (nested) square and cube roots. For reference and easy reproduction, its decimal expansion sufficient for double precision is:

$$\psi = 1.533751168755204288118041 \dots \quad (11)$$

Interestingly, ψ is not a *Pisot* number [4], because all roots of $\psi^4 - \psi - 4$ have modulus larger 1. Pisot numbers are connected to aperiodic tilings [36] and specific instances such as the golden ratio, the silver ratio, the super golden ratio and, specifically for two-dimensional grids, the plastic number [21, 50], have been suggested in the context of (Fibonacci) sampling. The values of ψ and ϕ as defined above are used for all evaluations.

4. Analysis and Comparison

Measuring the distance of two orientations by the natural metric induced on the sphere

$$d(\mathbf{p}, \mathbf{q}) = \arccos |\langle \mathbf{p}, \mathbf{q} \rangle|. \quad (12)$$

is equivalent to the geodesic distance in the natural Riemannian metric of $SO(3)$ [22] or, intuitively, the smallest possible angle of rotation turning one orientation into the other (taking the absolute value accounts for the fact that \mathbf{q} and $-\mathbf{q}$ represent the same orientation). We will more generally use \mathbb{S}^3 to measure geometric quantities. In particular, we need the volume of a spherical cap with radius θ (in radians)

$$V_o(\theta) = \pi(2\theta - \sin 2\theta). \quad (13)$$

and the area of the sphere bounding this cap:

$$A_o(\theta) = 4\pi \sin^2 \theta. \quad (14)$$

Derivations and other geometric tools for computing on \mathbb{S}^3 as well as data and code for the methods described below are provided in the supplementary material.

4.1. Discrepancy

Discrepancy measures the error being made when estimating the volume of a (convex) region by considering only the number of samples in the region [28]. Each sample represents the expected measure, i.e. the total measure of the domain divided by the number of samples. For the quaternion representation, this means assigning the volume $\frac{2\pi^2}{2n} = \frac{\pi^2}{n}$ to each sample of $SO(3)$, as \mathbb{S}^3 has total volume $2\pi^2$ and each sample is represented by the *pair* $\pm \mathbf{q}$.

For Euclidean domains it is common to consider *star* discrepancy, which is based on the measure of a box [59]. The appropriate analogy are hyper-spherical caps on \mathbb{S}^3 , ensuring that the strong connection to Monte-Carlo techniques remains [12, 9.1.5]. Consequently, the quaternionic spherical cap discrepancy for a set $\mathcal{Q} = \{\mathbf{q}_i \in \mathbb{R}^4, \|\mathbf{q}_i\| = 1\}$ of size $n = |\mathcal{Q}|$ is

$$D(\mathcal{Q}) = \sup_{\mathbf{p} \in \mathbb{S}^3, r \in [0, \pi[} \left| V_o(r) - \frac{\pi^2}{n} \#\{d(\mathbf{p}, \mathbf{q}_i) \leq r\} \right|. \quad (15)$$

Only caps smaller than the semi-sphere are considered because the unit quaternions are a double cover.

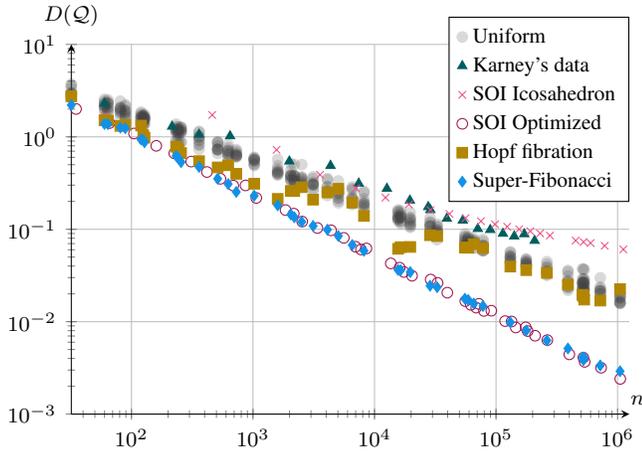


Figure 2. Approximation of spherical cap discrepancy.

For computing this value, note that changing the volume of the cap without changing the number of samples in the cap also changes the discrepancy. This implies that critical points of $D(Q)$ are achieved if the boundary sphere of the cap contains one or more of the samples. If this is the case, an infinitesimal change of the cap includes or excludes the points on the boundary. Alas, exploring all $\binom{n}{4} \in O(n^4)$ configurations with points on the boundary of the cap is prohibitive even for moderate n . For Euclidean spaces, sweeping techniques and specific rectangular decompositions reduce the complexity to $O(n^3)$ [13, 42]. It is not clear how to lift these techniques to the sphere and the cubic complexity is still unmanageable. We have therefore opted for a sampling technique that still exploits the observation that critical caps have samples on their boundary.

The idea is to fix m centers $\mathbf{c}_j \in \mathbb{S}^3$ uniformly random. For each center \mathbf{c}_j , we consider the caps defined by the n radii $d(\mathbf{c}_j, \mathbf{q}_i)$. Notice that these caps have \mathbf{q}_i on their boundary. The caps are sorted by radii, meaning the \mathbf{q}_i are sorted based on distances $d(\mathbf{c}, \mathbf{q}_i)$. For reasonably distributed samples this can be done in linear time [11]. Walking through the sorted list provides the number of points inside the cap without any overhead. The overall complexity of this procedure is linear in m and n and can be easily performed in parallel for different centers \mathbf{c}_j . The supplemental material shows convergence behavior for small n . When comparing different distributions we use the same set of centers.

Figure 2 show the result of this comparison as a scatter plot of discrepancy against sample size. We find that SOI using optimization on \mathbb{S}^2 and Super-Fibonacci sampling have at least an order of magnitude lower discrepancy than all other methods across all sample sizes. The very regular subdivisions of polyhedra (Karney's data or SOI based on subdivision) perform worse than uniform sampling.

4.2. Radial distribution function

Stationary processes are commonly analyzed based on first order statistics. The *radial distribution* or *pair correlation* function describes the density of the samples dependent on the relative position of pairs of samples [23]. The relative position may be simplified to the distance between the samples, if the process generating the samples is assumed to be isotropic, or if there are no preferred directions. A density function of pairwise distance is computed by convolving with a kernel κ , typically a Gaussian.

For each sample at distance r (measured in radians), we need to normalize by the *expected* number of samples at that distance, which is proportional to the *surface area* of the sphere at the boundary of a cap with radius r . This leads to the following density function, up to constant factors

$$g_Q(r) = \frac{1}{nA_Q(r)} \sum_{\mathbf{q}_i, \mathbf{q}_j \in Q, i \neq j} \kappa(r - d(\mathbf{q}_i, \mathbf{q}_j)). \quad (16)$$

For κ being the Dirac δ -function, $g(r)$ is closely related to the *power spectrum* [20], describing the distribution of power over spatial frequencies, in this case spherical harmonics. Both functions contain the exact same information. While the radial distribution function is often used in statistical mechanics, the power spectrum is more common in engineering [56]. Distributions with good properties for sampling are commonly associated with *blue noise* [30].

Computing the radial distribution function amounts to adding the kernel $\kappa(r - d(\mathbf{q}_i, \mathbf{q}_j))$ for each pair of different samples $i \neq j$ to the function $g_Q(r)$. Since the quadratic complexity leads to unmanageable run-times we take a random subset of m samples and then compute the distances to all other n samples, resulting in a complexity of $O(mn)$.

The density function $g(r)$ can be represented discretely using equally spaced r_i in $[0, \pi/2]$ so that $g(r_i)$ forms a vector. For the kernel we use a Gaussian. We adapt the variance to the radius of the sphere, whose volume is $1/n$ -th of the total volume of \mathbb{S}^3 :

$$\frac{2\pi^2}{n} = \frac{4\pi r^3}{3} \iff r = \left(\frac{3\pi}{2n}\right)^{1/3} \quad (17)$$

Then we set the variance of the Gaussian to half of the radius, i.e. $\sigma^2 = r^2/4$ (other factors would be possible as well). In order to avoid iterating over all r_i for each pair we compute the width of this Gaussian for which its contribution is larger than the machine epsilon in single precision.

Figure 3 shows the the result of this computation for sample sets with $n \approx 2 \cdot 10^4$ points – curves for other values of n are qualitatively similar. The uniform distribution results in a nearly flat radial distribution, as expected, because all pair-wise distance appear with equal probability. The methods using subdivided polyhedra, such as Karney's

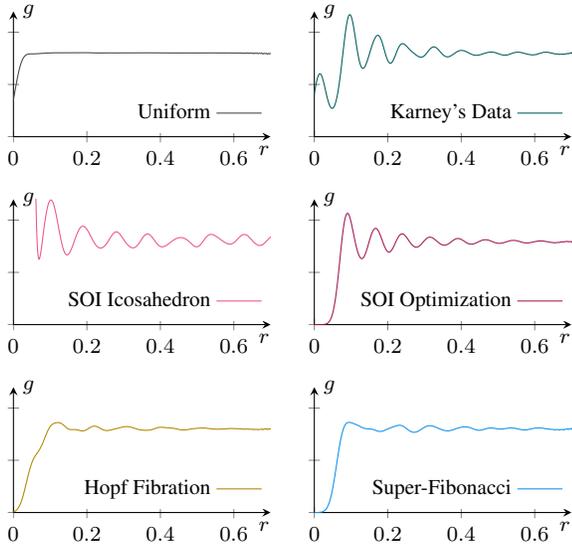


Figure 3. Radial distribution functions for $n \approx 2 \cdot 10^4$.

data and the SOI method based on subdividing an icosahedron exhibit unwanted peaks for small and large pairwise distances. The SOI method using optimization on \mathbb{S}^2 has no pairs with small distance but the typical ringing resulting from the avoidance of such pairs. The Hopf fibration and the Super-Fibonacci sampling show the desired behavior, with no pairs at small distances and quickly decaying variation for pairs at higher distances – see Singh et al. [51] for examples and discussion on the relation between the properties of a sample set and its radial distribution function.

4.3. Spherical Voronoi diagram and areas

The *Voronoi cell* Ω_i associated to a *site* is the set of points in a given domain closest to the site, i.e. for quaternion sites $\{\mathbf{q}_i\}$ on \mathbb{S}^3

$$\Omega_i = \{\mathbf{x} \in \mathbb{S}^3 : d(\mathbf{x}, \mathbf{q}_i) \leq d(\mathbf{x}, \mathbf{q}_j), i \neq j\}. \quad (18)$$

Voronoi cells consists of two congruent antipodal regions. The boundaries of Voronoi cells are composed of bisectors, which are planes through the origin (the locus of points with equal distance to two sites in \mathbb{R}^4) intersected with the unit sphere \mathbb{S}^3 . In other words, Voronoi cells are (pairs of) spherical polyhedra. One may consider the quality of a sample set reflected by the properties of the Voronoi cells – commonly considered is the volume [3], which one wants to be identical across samples.

The Voronoi diagram is dual to the Delaunay triangulation [2]. The combinatorics of the Delaunay triangulation on a sphere are identical to the convex hull of the points under mild conditions, namely that no hemisphere is empty. In order to use standard software for computation of the con-

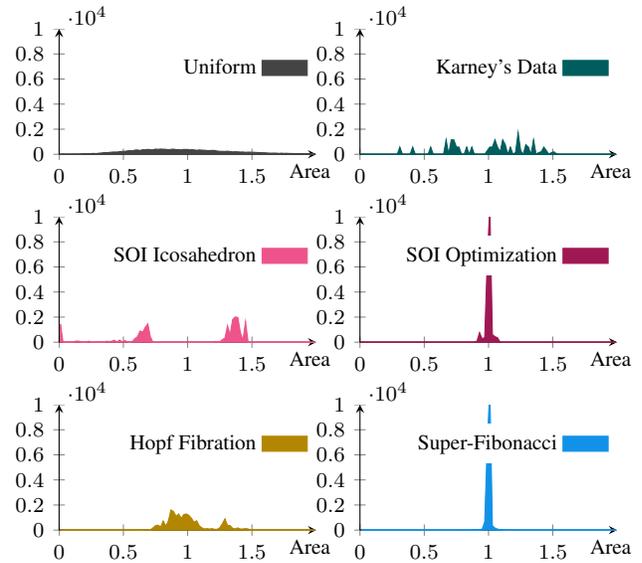


Figure 4. Histograms of Voronoi areas relative to the equal area division $2\pi^2/n$ for $n \approx 2 \cdot 10^4$.

vex hull, we add for each sample \mathbf{q}_i its antipodal point $-\mathbf{q}_i$, which also guarantees that no hemisphere is empty. For computation of the convex hull we use CGAL's *dD* triangulation code [10] (while it is possible to compute the spherical Voronoi diagram directly [7], there is no publicly available stable implementation). An important trick is to add one point in the interior of the sphere. This has no effect on the convex hull, but avoids the computationally costly handling of the degenerate situation that all simplices share the same circumsphere. If the samples are well-distributed, the Delaunay triangulation can be computed in $O(n)$ [1].

We find the Voronoi vertices as the circumcenters of the (spherical) Delaunay tetrahedra (see supplementary material). For a given sample \mathbf{q}_i , the Voronoi vertices arising from the Delaunay tetrahedra incident on \mathbf{q}_i form a spherical polyhedron, which is the Voronoi cell Ω_i . The faces of Ω_i are spherical polygons orthogonal to Delaunay edges emanating from \mathbf{q}_i . For computation of the volume of the Voronoi cell, we decompose it into spherical tetrahedra, whose volume can be computed using hypergeometric series [37, 47].

Figure 4 shows a histogram of volumes relative to the equal area volume tessellation $2\pi^2/n$ for $n \approx 2 \cdot 10^4$. The uniform distribution has the expected smooth and large variation in volumes. Distributions based on subdivided polyhedra such as Karney's data or the SOI method based on an icosahedron exhibit a discrete set of volumes, with a large variation. Hopf fibration, interestingly, shows two modes, both with moderate variation in volume. The SOI method for points optimized on \mathbb{S}^2 and the Super-Fibonacci sampling both have a remarkably narrow peak.

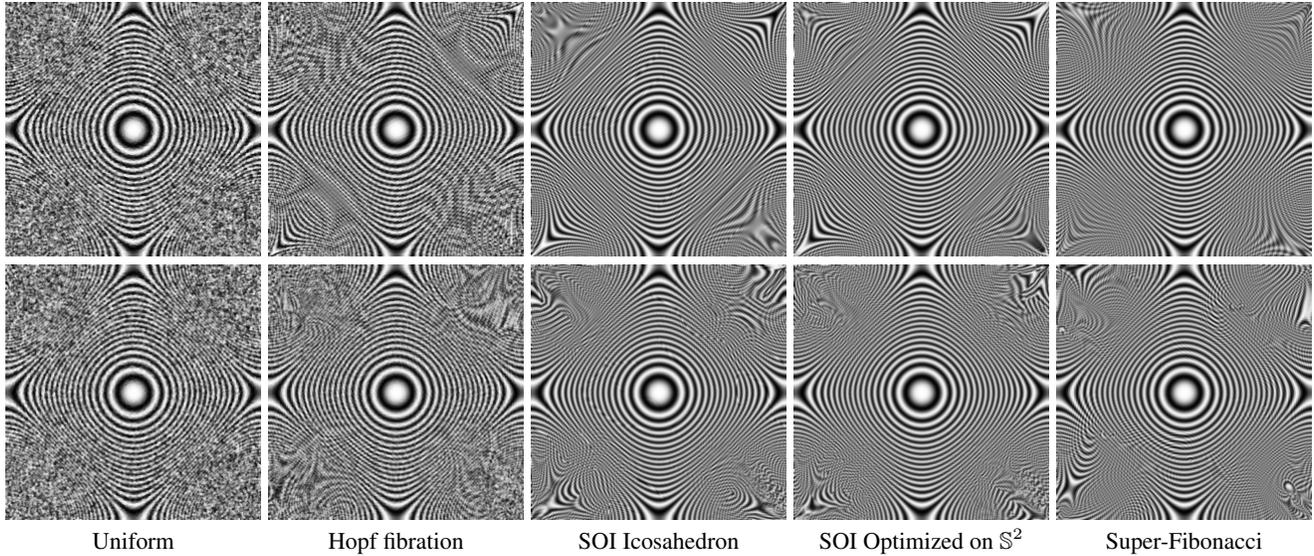


Figure 5. Zone plates on the Clifford torus reconstructed with Gaussian kernels centered in $n \approx 5 \cdot 10^5$ samples. Upper row is based on canonical orientation of torus and samples, lower row based on (the same) random orientation of the samples relative to the torus.

The Delaunay triangulation directly yields the *covering radius* or *dispersion* (the radius of the largest empty sphere among the points) and the *packing distance* (the largest radius of non-overlapping spheres placed at the samples). These measures are less relevant in the context of sampling applications and discussed in the supplementary material.

4.4. Clifford torus visualization

The quality of sample distributions is often demonstrated using visualization, exploiting that the human visual system is sensitive for (ir)regularities [57]. Apart from showing the samples, another method is by reconstructing a function with varying frequency content from the values in the point samples, exhibiting aliasing artifacts [44]. To accommodate these approaches for samples on \mathbb{S}^3 we consider the *Clifford torus*, which is a slice of \mathbb{S}^3 given in canonical orientation as

$$x_0^2 + x_1^2 = x_2^2 + x_3^2 = \frac{1}{2}. \quad (19)$$

The Clifford torus is intrinsically flat, meaning it can be mapped to the (standard Euclidean) plane isometrically. The common parameterization is

$$\mathbf{x}(\theta, \phi) = \frac{1}{\sqrt{2}} (\cos \theta, \sin \theta, \cos \phi, \sin \phi). \quad (20)$$

The orthogonal projection of a point $\mathbf{x} \in \mathbb{S}^3 \setminus \{x_0 = x_1 = 0, x_2 = x_3 = 0\}$ is given by

$$\mathbf{x}' = \frac{1}{\sqrt{2}} \left(\frac{x_0}{\sqrt{x_0^2 + x_1^2}}, \frac{x_1}{\sqrt{x_0^2 + x_1^2}}, \frac{x_2}{\sqrt{x_2^2 + x_3^2}}, \frac{x_3}{\sqrt{x_2^2 + x_3^2}} \right) \quad (21)$$

showing that the image of orthogonal projections in the plane is

$$(\arctan2(x_0, x_1), \arctan2(x_2, x_3)). \quad (22)$$

For visualizing aliasing artifacts we use the reconstruction of a zone plate function, where the distance to a fixed point \mathbf{q} on the Clifford is measured only along the torus to avoid aliasing introduced by slicing:

$$z(\mathbf{x}) = \frac{1 + \cos(kd(\mathbf{x}', \mathbf{q})^2)}{2}. \quad (23)$$

This function is sampled on \mathbb{S}^3 and reconstructed using a Gaussian kernel, restricted to the Clifford torus. The image in $[-\pi, \pi]^2$ is represented discretely on a $m \times m$ grid. The variance of the Gaussian kernel is set to the diagonal of the grid cells, i.e. $\sigma = m^{-1}\sqrt{2}\pi$. The isometry between the flat image and the embedding in \mathbb{R}^4 allows using the same variance on \mathbb{S}^3 .

Figure 5 shows reconstructions of the zone plate. The number of samples $n \approx 5 \cdot 10^5$ has been chosen to make sure that each grid point in the torus may be sufficiently close to a sample (for an ideal distribution of samples over \mathbb{S}^3). The upper row images are based on the canonical orientation of the sample set. As there are potential directional preferences in the samples sets, the lower row reconstructions are based on rotating the sample sets against the torus, using the same (random) rotation for all methods. The reconstructions based on the SOI methods and Super-Fibonacci sampling clearly outperform Hopf fibration and uniform sampling. Aliasing artifacts are less severe for the canonical orientation and least conspicuous for Super-Fibonacci sampling.

5. Discussion

Super-Fibonacci sampling provides a simple and fast approach to computing an arbitrary number of well-distributed orientations. The samples have significantly lower discrepancy than similarly fast methods and are on par with optimization-based techniques. For illustrating the quality of the distribution relative to the time necessary to generate the samples, we compare the reconstruction of the zone plate based on the number of samples generated in $25ms$. The result is shown in Figure 6). In this comparison, Super-Fibonacci samplings clearly outperforms all other methods.

The analysis of Voronoi cells reveals that the volume of the cells is almost identical - the equal volume property has been observed to avoid unwanted regularity in optimization algorithms [3, 16]. The reason for the good volume distribution is possibly the volume preserving mapping from the full cylinder, carrying the good properties of Fibonacci sampling in the interval and unit disk over to \mathbb{S}^3 .

The fact that Super-Fibonacci sampling is faster than the simple uniform random sampling of the unit sphere is due to fewer and more implementation-friendly use of non-elementary functions. Overall there are only two square roots, and the two pairs of sine and cosine functions with the same arguments are optimized by modern compilers to exploit parallel `sincos` hardware implementations. The fact that for two sets of size n and m , every k -th point will be the same, where k is the gcd of n and m makes it easy to refine a sequence of n points, simply by generating a multiple of the n points (and then skipping the original points).

Limitations. The dispersion of sets generated with Super-Fibonacci is worse than those generated with the SOI method based on optimization – see the supplementary material for details. This may be relevant in some applications [15]. For standard spherical Fibonacci sampling, it has been observed that the smallest empty spheres appear at the beginning and end of the sequence. This problem can be alleviated quite easily [49]. We have not tried such modifications for Super-Fibonacci sampling in order to improve dispersion. It should also be noted that we consider the SOI methods only for those values of n that admit a factorization into a sampling on \mathbb{S}^1 and \mathbb{S}^2 with similar spacing. Enforcing any composite value n would be possible, but result in very uneven sampling on \mathbb{S}^3 .

The constants ϕ and ψ do affect the properties of the distributions. They have been determined experimentally for sample sizes up to $n \approx 10^5$ considering different quality measures. Theoretical analysis connecting these constants to the resulting distribution is lacking and would be welcome. Moreover, for larger sample sizes and/or specific use cases different constants may be more suitable.

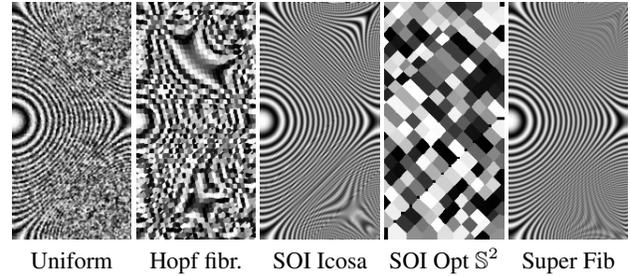


Figure 6. Right half of the zone plates on the Clifford torus reconstructed based on the number of samples that had been generated by the sampling methods in approximately $25ms$ total running time.

Relation to other work and use in other context. The high speed generation of samples may be useful for quasi Monte Carlo techniques, as the sample set can be generated on the fly. Optimization techniques may profit from a deterministic starting point, much closer to the optimum than random uniform samples.

Analytic techniques typically aim at optimizing local measures, such as packing distance or dispersion. Optimizing discrepancy directly is difficult, but optimizing sample sets towards specified radial distribution functions may be possible [55]. A general strategy is to minimize the sum of squared distance to the samples, known as Lloyd’s method [33], which for continuous domains leads to *centroidal Voronoi diagrams* [14]. In this context, it has been observed that forcing the sample set to have equal Voronoi volumes avoids the tendency of such optimization approaches to introduce unwanted regularity [16]. Recall that Super-Fibonacci sampling generates Voronoi cells with almost identical volume, which is exactly what these optimization methods are aiming for.

Balzer et al. [3] suggest the optimization of the distribution under a discrete capacity constraint, which makes the optimization oblivious to the underlying geometry. This requires a ‘background’ sample set to enforce the equal capacities, which needs to be a large multiple of the number of desired samples. Super-Fibonacci sampling is ideal for generating this background set.

Because of its natural refinement property, a set of kn samples generated with Super-Fibonacci sampling contains m different but individually well-distributed sets of k orientations. This can be used for rejection-sampling of density functions: the density function is quantized into k levels and then compared against the index of a sample modulo k .

The samples in Super-Fibonacci sampling can be computed independently based on the index i and the total number n . This enables use in distributed computing without the necessity of message passing.

References

- [1] Nina Amenta, Dominique Attali, and Olivier Devillers. Complexity of delaunay triangulation for points on lower-dimensional polyhedra. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '07, page 1106–1113, USA, 2007. Society for Industrial and Applied Mathematics. 6
- [2] Franz Aurenhammer, Rolf Klein, and Der-Tsai Lee. *Voronoi Diagrams and Delaunay Triangulations*. WORLD SCIENTIFIC, 2013. 6
- [3] Michael Balzer, Thomas Schlömer, and Oliver Deussen. Capacity-constrained point distributions: A variant of lloyd's method. *ACM Trans. Graph.*, 28(3), July 2009. 6, 8
- [4] Marie J. Bertin, Annette Decomps-Guilloux, Marthe Grandet-Hugot, Martine Pathiaux-Delefosse, and Jean Schreiber. *Pisot and Salem Numbers*. Birkhäuser, Basel, 1992. 4
- [5] Tolga Birdal, Umut Şimşekli, M. Onur Eken, and Slobodan Ilic. Bayesian pose graph optimization via bingham distributions and tempered geodesic mcmc. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS'18, page 306–317, Red Hook, NY, USA, 2018. Curran Associates Inc. 1
- [6] G. E. P. Box and Mervin E. Muller. A Note on the Generation of Random Normal Deviates. *The Annals of Mathematical Statistics*, 29(2):610 – 611, 1958. 2
- [7] Manuel Caroli, Pedro M. M. de Castro, Sébastien Lorient, Olivier Rouiller, Monique Teillaud, and Camille Wormser. Robust and efficient delaunay triangulations of points on or close to a sphere. In Paola Festa, editor, *Experimental Algorithms*, pages 462–473, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg. 6
- [8] Yu H Chen, Se Un Park, Dennis Wei, Greg Newstadt, Michael A Jackson, Jeff P Simmons, Marc De Graef, and Alfred O Hero. A dictionary approach to electron backscatter diffraction indexing. *Microscopy and Microanalysis*, 21(3):739, 2015. 1
- [9] John Horton Conway and Neil James Alexander Sloane. *Sphere packings, lattices and groups*, volume 290. Springer Science & Business Media, 2013. 1
- [10] Olivier Devillers, Samuel Hornus, and Clément Jamin. dD triangulations. In *CGAL User and Reference Manual*. CGAL Editorial Board, 5.2.1 edition, 2021. 6
- [11] L. Devroye and T. Klincsek. Average time behavior of distributive sorting algorithms. *Computing*, 26(1):1–7, 1981. 5
- [12] Josef Dick and Friedrich Pillichshammer. Discrepancy theory and quasi-monte carlo integration. In William Chen, Anand Srivastav, and Giancarlo Travaglini, editors, *A Panorama of Discrepancy Theory*, pages 539–619. Springer International Publishing, Cham, 2014. 4
- [13] Carola Doerr, Michael Gnewuch, and Magnus Wahlström. Calculation of discrepancy measures and applications. In William Chen, Anand Srivastav, and Giancarlo Travaglini, editors, *A Panorama of Discrepancy Theory*, pages 621–678. Springer International Publishing, Cham, 2014. 5
- [14] Qiang Du, Vance Faber, and Max Gunzburger. Centroidal voronoi tessellations: Applications and algorithms. *SIAM Review*, 41(4):637–676, 1999. 8
- [15] Benoît Gandar, Gaëlle Loosli, and Guillaume Deffuant. Sample dispersion is better than sample discrepancy for classification. Technical report, Oct. 2010. 8
- [16] Fernando de Goes, Katherine Breeden, Victor Ostromoukhov, and Mathieu Desbrun. Blue noise through optimal transport. *ACM Trans. Graph.*, 31(6), Nov. 2012. 8
- [17] Alvaro Gonzáles. Measurement of areas on a sphere using fibonacci and latitude–longitude lattices. *Mathematical Geosciences*, 42(1):49–64, Nov 2009. 3
- [18] Manuel Gräf and Daniel Potts. On the computation of spherical designs by a new optimization approach based on fast spherical fourier transforms. *Numerische Mathematik*, 119:699–724, 2011. 2
- [19] J H Hannay and J F Nye. Fibonacci numerical integration on a sphere. *Journal of Physics A: Mathematical and General*, 37(48):11591–11601, nov 2004. 3
- [20] Daniel Heck, Thomas Schlömer, and Oliver Deussen. Blue noise sampling with controlled aliasing. *ACM Trans. Graph.*, 32(3), July 2013. 5
- [21] Doug Hensley and Francis Edward Su. Random walks with badly approximable numbers. *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, 64:95–102, 2004. 4
- [22] Du Q. Huynh. Metrics for 3d rotations: Comparison and analysis. *Journal of Mathematical Imaging and Vision*, 35:155–164, 2009. 4
- [23] Janine Illian, Antti Penttinen, Helga Stoyan, and Dietrich Stoyan. *Statistical analysis and modelling of spatial point patterns*, volume 70. John Wiley & Sons, 2008. 5
- [24] Lucas Janson, Brian Ichter, and Marco Pavone. Deterministic sampling-based motion planning: Optimality, complexity, and performance. *The International Journal of Robotics Research*, 37(1):46–61, 2018. 1
- [25] Charles F.F. Karney. Quaternions in molecular modeling. *Journal of Molecular Graphics and Modelling*, 25(5):595–604, Jan 2007. 2
- [26] Charles F. F. Karney. Nearly optimal coverings of orientation space. <https://github.com/cffk/orientation>, 2006. [License: MIT]. 2
- [27] Dima Kozakov, David R Hall, Bing Xia, Kathryn A. Porter, Dzmityr Padhorny, Christine Yueh, Dmitri Beglov, and Sándor Vajda. The cluspro web server for protein-protein docking. *Nature Protocols*, 12:255–278, 2017. 1
- [28] L. Kuipers and H. Niederreiter. *Uniform Distribution of Sequences*. A Wiley-Interscience publication. Wiley, 1974. 4
- [29] Peter Mahler Larsen and Søren Schmidt. Improved orientation sampling for indexing diffraction patterns of polycrystalline materials. *Journal of Applied Crystallography*, 50(6):1571–1582, Dec 2017. 2
- [30] D. L. Lau, R. Ulichney, and G. R. Arce. Blue and green noise halftoning models. *IEEE Signal Processing Magazine*, 20(4):28–38, 2003. 5
- [31] S. F. Li and R. M. Suter. Adaptive reconstruction method for three-dimensional orientation imaging. *Journal of Applied Crystallography*, 46(2):512–524, Apr 2013. 1

- [32] Kyaw Zaw Lin, Weipeng Xu, Qianru Sun, Christian Theobalt, and Tat-Seng Chua. Learning a disentangled embedding for monocular 3d shape retrieval and pose estimation, 2019. [1](#)
- [33] S. Lloyd. Least squares quantization in pcm. *IEEE Trans. Inf. Theor.*, 28(2):129–137, Sept. 2006. [8](#)
- [34] George Marsaglia. Choosing a Point from the Surface of a Sphere. *The Annals of Mathematical Statistics*, 43(2):645 – 646, 1972. [2](#)
- [35] Julie C. Mitchell. Sampling rotation groups by successive orthogonal images. *SIAM Journal on Scientific Computing*, 30(1):525–547, 2008. [2](#)
- [36] R.V. Moody. *The Mathematics of Long-Range Aperiodic Order*. Nato Science Series C:. Springer Netherlands, 1997. [4](#)
- [37] Jun Murakami. Volume formulas for a spherical tetrahedron. *Proceedings of the American Mathematical Society*, 140(9):3289–3295, 2012. [6](#)
- [38] Kieran A Murphy, Carlos Esteves, Varun Jampani, Srikanth Ramalingam, and Ameesh Makadia. Implicit-pdf: Non-parametric representation of probability distributions on the rotation manifold. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 7882–7893. PMLR, 18–24 Jul 2021. [1](#)
- [39] Saigopal Nelaturi, Morad Behandish, Amir M. Mirzende-hdel, and Johan de Kleer. Automatic support removal for additive manufacturing post processing. *Computer-Aided Design*, 115:135–146, 2019. [1](#)
- [40] Harald Niederreiter. *Random Number Generation and Quasi-Monte Carlo Methods*. Society for Industrial and Applied Mathematics, 1992. [1](#), [3](#)
- [41] Brian Okorn, Mengyun Xu, Martial Hebert, and David Held. Learning orientation distributions for object pose estimation, 2020. [1](#)
- [42] Mark H. Overmars and Chee-Keng Yap. New upper bounds in klee’s measure problem. *SIAM Journal on Computing*, 20(6):1034–1045, 1991. [5](#)
- [43] Art B. Owen and Seth D. Tribble. A quasi-monte carlo metropolis algorithm. *Proceedings of the National Academy of Sciences*, 102(25):8844–8849, 2005. [1](#)
- [44] A. Cengiz Öztireli. A Comprehensive Theory and Variational Framework for Anti-aliasing Sampling Patterns. *Computer Graphics Forum*, 2020. [7](#)
- [45] Davide Piccini, Arne Littmann, Sonia Nielles-Vallespin, and Michael O. Zenge. Spiral phyllotaxis: The natural way to construct a 3d radial trajectory in mri. *Magnetic Resonance in Medicine*, 66(4):1049–1056, 2011. [3](#)
- [46] Barak Raveh, Angela Enosh, and Dan Halperin. A little more, a lot better: Improving path quality by a path-merging algorithm. *IEEE Transactions on Robotics*, 27(2):365–371, 2011. [1](#)
- [47] Jason M. Ribando. Measuring solid angles beyond dimension three. *Discrete & Computational Geometry*, 36(3):479–487, 2006. [6](#)
- [48] J Peter Rickgauer, Nikolaus Grigorieff, and Winfried Denk. Single-protein detection in crowded molecular environments in cryo-em images. *eLife*, 6:e25648, may 2017. [1](#)
- [49] Martin Roberts. How to evenly distribute points on a sphere more effectively than the canonical fibonacci lattice. Blog post, available online. [8](#)
- [50] Martin Roberts. The unreasonable effectiveness of quasirandom sequences. Blog post, available online. [4](#)
- [51] Gurprit Singh, Cengiz Öztireli, Abdalla G.M. Ahmed, David Coeurjolly, Kartic Subr, Oliver Deussen, Victor Ostromoukhov, Ravi Ramamoorthi, and Wojciech Jarosz. Analysis of sample correlations for monte carlo rendering. *Computer Graphics Forum*, 38(2):473–491, 2019. [6](#)
- [52] Neil Sloane et al. Spherical codes. <http://neilsloane.com/packings/index.html>. [1](#)
- [53] Yifan Sun, Jiacheng Zhuo, Arnav Mohan, and Qixing Huang. K-best transformation synchronization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. [1](#)
- [54] Richard Swinbank and R. James Purser. Fibonacci grids: A novel approach to global modelling. *Quarterly Journal of the Royal Meteorological Society*, 132(619):1769–1793, 2006. [3](#)
- [55] O.U. Uche, F.H. Stillinger, and S. Torquato. On the realizability of pair correlation functions. *Physica A: Statistical Mechanics and its Applications*, 360(1):21–36, 2006. [8](#)
- [56] Robert Ulichney. *Digital Halftoning*. MIT Press, Cambridge, MA, USA, 1987. [5](#)
- [57] John I. Yellott. Spectral consequences of photoreceptor sampling in the rhesus retina. *Science*, 221(4608):382–385, 1983. [7](#)
- [58] Anna Yershova, Swati Jain, Steven M. LaValle, and Julie C. Mitchell. Generating uniform incremental grids on so(3) using the hopf fibration. *The International Journal of Robotics Research*, 29(7):801–812, 2010. [2](#)
- [59] Stanisław K. Zaremba. The mathematical basis of monte carlo and quasi-monte carlo methods. *SIAM Review*, 10(3):303–314, 1968. [4](#)