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

Supplementary Material

A. Geometry of $\mathbb{S}^3$

In the following we assume that all *points* $\mathbf{x}_i$ are unit vectors in $\mathbb{R}^4$, i.e. $\mathbf{x}_i^\mathsf{T}\mathbf{x}_i = 1, \ldots$, written in column form.

A.1. Simplices

Two points $\mathbf{x}_0, \mathbf{x}_1$ spanning a two-dimensional linear subspace define an edge. We assume the edge connects the points along the shorter of the two arcs connecting the points. Under this assumption, the length of the edge is given as $d(\mathbf{x}_0, \mathbf{x}_1)$. For practical reasons, we flip the sign of $\mathbf{x}_0$ if the scalar product $\mathbf{x}_0^\mathsf{T}\mathbf{x}_1$ is negative. Then the edge is defined as the convex combination of $\mathbf{x}_0$ and $\mathbf{x}_1$ normalized so that the points lie on the unit sphere.

Three points $\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2$ spanning a three-dimensional linear subspace in $\mathbb{R}^4$ define 'planar' spherical faces. From the possible combinations, we pick the face defined by the three shorter arcs defined by the three pairs of points. Similar to the case of edges, for practical purposes we achieve this by flipping the signs of $\mathbf{x}_1$ and/or $\mathbf{x}_2$ based on the signs of the scalar products.

The face has a normal in $\mathbb{S}^3$, which is the direction in $\mathbb{S}^3$ orthogonal to the two directions spanned by the face. Notice that the linear 3-space spanned by $\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2$ contains all normals to the unit sphere inside the face, so the normals to the linear 3-space in $\mathbb{R}^4$ is $\mathbf{n}$. In other words, we can compute $\mathbf{n}$ as the single element in the null space of the 3×4 matrix $(\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2)$.

Four points $(\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3) = \mathbf{X} \in \mathbb{R}^{4\times4}$ with $\mathbf{X}$ having full rank form spherical tetrahedra. Again we select the desired simplex by always taking the shorter among the two possible arcs for pairs of points and make this practical by possibly reflecting the points. We often need the normals of all four faces of the tetrahedron. For this case notice that any normal vector $\mathbf{n}_0$ for the face opposite of $\mathbf{x}_0$ pointing into the direction of $\mathbf{x}_0$ satisfies

$$\mathbf{X}\mathbf{n}_0 = (\lambda, 0, 0, 0)^\mathsf{T}, \quad \lambda > 0. \tag{24}$$

In other words, unit normal vectors pointing into the spherical tetrahedron can be conveniently computed by normalizing the rows of $\mathbf{X}^{-1}$. We will explain how to compute the (3-dimensional) volume of the spherical tetrahedron further below.

A.2. Circumcenters

The circumcenter $\mathbf{c}$ of a simplex is a point with the same distance r to all vertices of the simplex: $d(\mathbf{c}, \mathbf{x}_i) = r$. To define it in case the simplex it is not fully dimensional, i.e. for edges and faces, we also ask that it is contained in the linear span of the vertices in $\mathbb{R}^4$: $\mathbf{c} = \sum_i a_i \mathbf{x}_i$. This means for edges we get a point on the edge, namely the midpoint; and for faces we get a point on the plane through the face. Circumcenters are naturally of interest for geometry on the sphere.

Note that the vertices of the simplex define an affine subspace in $\mathbb{R}^4$. The intersection of the affine subspace with the restriction of $\mathbb{S}^3$ to the linear subspace spanned by the vertices defines a spherical cap. The boundary of this cap is a sphere. It contains all vertices so it is their circumsphere. The circumcenter is the point on the spherical cap furthest from the affine subspace. So the circumcenter is orthogonal to the affine span of the points, i.e. $\mathbf{c}^\mathsf{T}\mathbf{e}_i = 0, i > 0$, where $\mathbf{e}_i = \mathbf{x}_i - \mathbf{x}_0$ is the edge vector given by the vertex $\mathbf{x}_i$ relative to $\mathbf{x}_0$. This leads to a linear system describing the direction of the circumcenter as a point on the affine span of the vertices, i.e. in terms of the coefficients a_i.The the system for a face is:

$$\begin{pmatrix} \mathbf{x}_0^\mathsf{T}\mathbf{e}_1 & \mathbf{x}_1^\mathsf{T}\mathbf{e}_1 & \mathbf{x}_2^\mathsf{T}\mathbf{e}_1 \\ \mathbf{x}_0^\mathsf{T}\mathbf{e}_2 & \mathbf{x}_1^\mathsf{T}\mathbf{e}_2 & \mathbf{x}_2^\mathsf{T}\mathbf{e}_2 \\ 1 & 1 & 1 \end{pmatrix} \begin{pmatrix} a_0 \\ a_1 \\ a_2 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \tag{25}$$

The case for a tetrahedron is analogous, for an edge we have $a_0 = a_1 = 1/2$. The point on the affine span can be normalized to yield the desired circumcenter.

A.3. Angles

The angle ϕ enclosed by two edges $\mathbf{x}_0, \mathbf{x}_1$ and $\mathbf{x}_0, \mathbf{x}_2$ can be computed by projecting on to the tangent hyper-plane at $\mathbf{x}_0$ and we get

$$\cos\phi = \mathbf{x}_1^\mathsf{T}\left(\mathbf{I} - \mathbf{x}_0\mathbf{x}_0^\mathsf{T}\right)\mathbf{x}_2. \tag{26}$$

The dihedral angle between two spherical faces (e.g. the dihedral angles in a spherical tetrahedron) can be computed based on the scalar product of their normals. Notice that no projection onto the sphere is necessary (this is similar to computing the angle between two edges in $\mathbb{S}^2$ based on normals to the planes defined by the edges). The six dihedral angles in a spherical tetrahedron can be computed conveniently as the Gram matrix of the normalized inverse of the coordinates.

A.4. Volume

We call the measure of closed regions in $\mathbb{S}^3$ *volume*, because they are 3-dimensional. Another commonly used term is *hyper-area*, because they are a subset of the surface of an embedded manifold.

The total volume of the surface of a unit 3-sphere is $2\pi^2$.

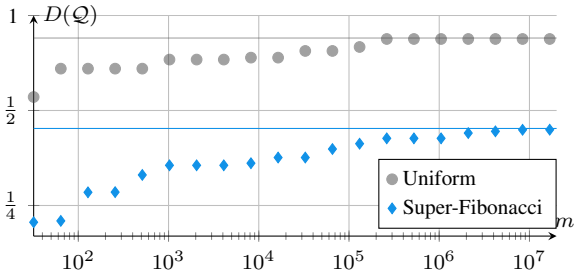

Figure 7. Approx. of discrepancy relative to number of centers.

A hyper-spherical cap can be specified either by the angle θ between a ray through the center of the cap and the perimeter; or by the height h of the center of the cap over the hyper-plane separating the cap from the rest of the sphere. These parameters are related by $h = 1 - \cos\theta$. The boundary of the cap is the intersection of this hyper-plane with the sphere, so it is a Euclidean 3-sphere. Its radius is $r = \sin\theta$, so its surface area (in the plane or in $\mathbb{S}^3$) is

$$A(\theta) = 4\pi \sin^2\theta, \qquad A(h) = 4\pi h(h-2) \tag{27}$$

We can find the volume of the cap by integrating the surface area over the angle:

$$\begin{aligned} V(\theta) = \int_0^\theta 4\pi \sin^2 r \, dr &= 2\pi(\theta - \sin\theta\cos\theta) \\ &= \pi(2\theta - \sin 2\theta). \end{aligned} \tag{28}$$

This form is more attractive in terms of fast computation compared to more general expressions appearing in the literature that involve the incomplete beta function [31].

For computation of the volumes of the Voronoi cells (see below) we also need the volume of a spherical tetrahedron. Expressions for this volume involve hypergeometric series [39, 49]. We use the formulae provided by Murakami [39], as the hypergeometric series appears in the form of the *dilogarithm*, which can be efficiently approximated [36]. Our implementation is based on freely available code [59]. The volume is described in terms of the 6 dihedral angles of the tetrahedron (similar to the area of a spherical triangle in $\mathbb{S}^2$ depending on the three interior angles). The formulae for computing the volume from the dihedral angles are involved – we provide code and refer the reader to the original literature for details.

B. Convergence of Discrepancy

Computing ground truth discrepancy is difficult. There are

$$\binom{n}{2} + \binom{n}{3} + \binom{n}{4} \in O(n^4) \tag{29}$$

different critical spheres defining a cap. For each of them we have to determine the number of samples inside it, which is worst case $O(n)$. So a straightforward implementation is $O(n^5)$ for n samples and can only be computed for very small n. Fig. 7 shows the result for $n = 512$ samples for uniform and Super-Fibonacci sampling. Computing ground truth (horizontal lines) took several hours for each sequence (sequential), while the approximations took a few seconds. The general behavior of increasing discrepancy with the number of centers is similar for larger n.

C. Measures from Delaunay triangulation

Two measures commonly used for assessing the quality of spherical designs can be conveniently computed from the Delaunay triangulation of the samples:

- The *covering radius* or *dispersion* is the radius of the largest empty sphere. It can be computed as the maximal radius of the circumspheres of the Delaunay tetrahedra.
- The *packing distance* is the largest radius of non-overlapping spheres placed at the samples. It is half of the shortest Delaunay edge.

Optimal values for these measures are achieved by regular tetrahedral subdivisions, of which there are only 3 instances that can be realized: the 3 regular 3-polytopes with tetrahedral cells. By making the assumption that all tetrahedra are regular, unattainable lower bounds can be derived. For any tetrahedral mesh on the sphere, by Euler-Poincaré we have

$$n - e + t = 0, \tag{30}$$

where e is the number of edges, t is the number of tetrahedra, and n, as before, is the number of vertices. Here we have already exploited that the number of triangles is twice the number of tetrahedra. Assuming a regular spherical tetrahedron with dihedral angle θ the average edge degree is $\frac{2\pi}{\theta}$. Plugging this in we find the relations

$$n = t\left(\frac{3\theta}{\pi} - 1\right) = e\left(1 - \frac{\pi}{3\theta}\right). \tag{31}$$

This shows that optimal values depend on the size of the set. We can compute the volume $V_\Delta(\theta)$ of a regular spherical tetrahedron with dihedral angle θ using the above mentioned computation of the volume for given dihedral angles. This allows to express bounds on the dihedral angle and other properties for given sample count:

$$n(\theta) = \frac{2\pi^2}{V_\Delta(\theta)}\left(\frac{3\theta}{\pi} - 1\right). \tag{32}$$

Computing θ for given n can be done by bisection as $V_\Delta(\theta)$ is quite involved. Given θ we can construct the spherical tetrahedron explicitly and then compute the edge length

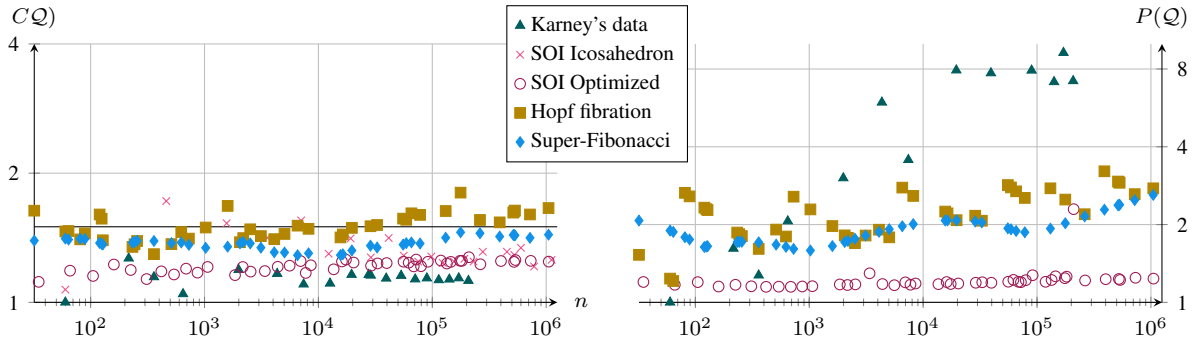

Figure 8. Dispersion (or covering radius, left) and smallest distance between points (packing distance, right) relative to lower bounds for these values, derived based on assuming the sphere could be tiled with congruent regular spherical tetrahedra.

or circumradius of the spherical tetrahedron to get lower bounds on packing distance or dispersion. The numerical values for the bounds on dispersion derived in this way appear to be similar to the ones by Larsen and Schmidt [29], although our derivation is based on other principles and perhaps simpler.

Figure 8 shows scatter plots of the dispersion and packing distance relative to the (generally unattainable) lower, resp. upper bounds. We have excluded the uniform distribution, as the random placement of samples may lead to arbitrarily bad relative values. Likewise, the close samples in the SOI method based on subdividing an icosahedron has led to unusable data for packing distance and the data is not shown. Karney's data has been constructed and selected to minimize dispersion and yields to lowest dispersion, but has rather bad packing distance. The SOI method based on optimization leads to good dispersion and very low packing distance. Super-Fibonacci sampling is generally better than Hopf fibration in both cases.

D. Parameter search

We can use the deterministic computation of dispersion relative to the upper bound to search for suitable parameters ϕ and ψ in Super-Fibonacci sampling. We consider roots of small degree polynomials with integer coefficients:

$$p(x) = \sum_i = 0^d c_i x^i. \tag{33}$$

Inspired by the use of *Pisot* numbers [4] in aperiodic tilings [38] we limit the coefficients for the non-constant parts to $c_i \in \{-1, 0, 1\}$ and further the constant part to be negative $c_0 \in \{-1, \ldots, -\hat{c}_0\}$. For each such polynomial we use Newton's method with starting value $x_0 = 1$ and store the result as candidate r_i if Newton's method converged. Early experiments showed that the values for ϕ and ψ should not be too different from each other so we limit the candidates to those satisfying $1 < r_i < 2$.

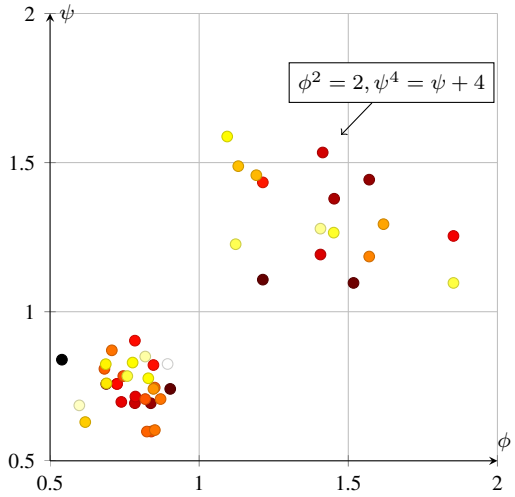

Figure 9. The pairs (ϕ, ψ) whose dispersion realtive to the lower bound is smaller than 1.5 for all $n = 2^k, k \in \{6, \ldots, 20\}$.

Based on suggestion for two dimensional Fibonacci sample [21, 52], we use the candidates to define the following pairs

$$(\phi, \psi) = \begin{cases} (r_i, r_j), i \neq j \\ \left(r_i^{-1}, r_j^{-1}\right), i \neq j \\ \left(r_i, r_i^2\right) \\ \left(r_i^{-1}, r_i^{-2}\right) \end{cases}. \tag{34}$$

With this pool of pairs, we compute the dispersion relative to the upper bound and retain only the pairs that are below a threshold. We use 1.5 (compare Fig. 8, left). Our strategy is to start with small samples size n, quickly eliminating a large number of pairs. Then we continue with the remaining pairs for increasing n, repeatedly removing pairs whose dispersion is more than 1.5 the (unattainable) lower bound. This leaves us with a set of pairs that have comparably low dispersion for a range of sample sizes.

Limiting the polynomial degree to 5, and the constant

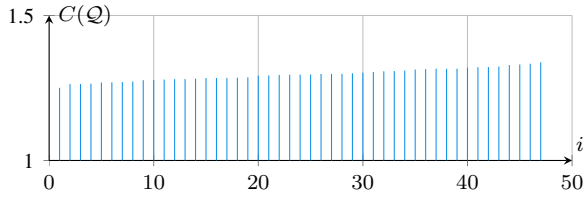

Figure 10. Average dispersion relative to lower bound over $n = 2^k, k \in \{6, \ldots, 20\}$ for different pairs of (ϕ, ψ).

to -4 generates 126 candidates r_i, resulting in roughly 16K initial pairs. The pairs remaining when rejecting pairs based on $n = 2^k, k \in \{6, \ldots, 20\}$ results in 47 pairs, shown in Figure 9. All pairs result from either taking $(\phi, \psi) = (r_i, r_j)$ or $(\phi, \psi) = (r_i^{-1}, r_j^{-1})$. There is no discernible relation between location in the plane and average dispersion. Figure 10 shows that the average dispersion is very similar among the selected pairs. Our particular choice based on $\phi^2 = 2, \psi^4 = \psi + 4$ is at index 10. This selection is based mostly on picking simple polynomials with few non-zero coefficients but also checking the resulting discrepancies (which is more time consuming than computing dispersion). However, the visualizations suggests that other choices for (ϕ, ψ) would lead to alternative Super-Fibonacci samples with similar properties.