

Generating Uniform Incremental Grids on $SO(3)$ Using the Hopf Fibration

Anna Yershova, Swati Jain, Steven M. LaValle, and Julie C. Mitchell

Abstract The problem of generating uniform deterministic samples over the rotation group, $SO(3)$, is fundamental to computational biology, chemistry, physics, and numerous branches of computer science. We present the best-known method to date for constructing incremental, deterministic grids on $SO(3)$; it provides: 1) the lowest metric distortion for grid neighbor edges, 2) optimal dispersion-reduction with each additional sample, 3) explicit neighborhood structure, and 4) equivolumetric partition of $SO(3)$ by the grid cells. We also demonstrate the use of the sequence on motion planning problems.

1 Introduction

Numerical computations on continuous spaces often require generation of a representative set of samples. The performance of various methods in engineering and scientific fields, such as numerical optimization and integration as well as collision-free path generation in robot motion planning, rely heavily on the quality of the sampling technique. Hence, it is important that the underlying samples are as good as possible.

A particular problem of discretization of $SO(3)$, the space of 3D rotations, arises in applications, such as biological protein docking problems, robot motion planning,

Anna Yershova
Duke University, Durham, NC 27707, USA, e-mail: yershova@cs.duke.edu

Swati Jain
Duke University, Durham, NC 27707, USA, e-mail: swati.jain@duke.edu

Steven M. LaValle
University of Illinois, Urbana, IL 61801 USA, e-mail: lavalle@cs.uiuc.edu

Julie C. Mitchell
University of Wisconsin-Madison, Madison, WI 53706, e-mail: mitchell@math.wisc.edu

aerospace trajectory design, and quantum computations. Typical operations on this space include numerical optimization, searching, integration, sampling, and path generation. Multiresolution grids are widely used for many of these operations on other spaces which are nicely behaved, such as rectangular subsets of \mathbb{R}^2 or \mathbb{R}^3 .

It would be wonderful to achieve the same for $SO(3)$; however, the space of 3D rotations is substantially more complicated. In its basic form, $SO(3)$ is defined as a set of matrices that satisfy orthogonality and orientation constraints. It is an implicitly defined, three-dimensional surface embedded in \mathbb{R}^9 . One approach is to place a coordinate system on the surface, causing it to behave like a patch in \mathbb{R}^3 . However, many of such coordinates cause metric distortions in comparison to distances on the original surface. Only few representations of $SO(3)$, such as *quaternions* and *Hopf coordinates*, preserve distances and volumes. They treat $SO(3)$ as a unit sphere $S^3 \subset \mathbb{R}^4$ with antipodal points identified. The volumes of surface patches on S^3 correspond to the unique Haar measure for $SO(3)$, which is the only way to obtain distortion-free notions of distance and volume. This implies that if we want to make multiresolution grids on $SO(3)$, we are faced with warping them onto S^3 . It may seem that such curvature prohibits the introduction of distortion-free grids, similar to the problem of making distance-preserving maps of the world (e.g., Greenland usually looks too big on a flat map). In addition, the identification of antipodal points causes a minor complication in which only half of S^3 is used, with unusual connectivity in the equatorial three-plane. However, in this paper we use intrinsic properties unique to $SO(3)$ (first described in [17]) to build almost distortion-free grids and avoid the issue of having to identify the antipodal points on S^3 .

Due to widespread interest in discretizing $SO(3)$ in numerous fields, there have been considerable efforts in the past. The problem of generating point sets on spheres minimizing various criteria, such as energy functions, discrepancy, dispersion, and mutual distances, has been extensively studied in mathematics and statistics [8, 15, 22, 25, 28, 29]. Random sampling methods were also developed in [2, 24, 27, 34]. Problems of sampling rotational groups and spheres have been studied and applied in the context of computational structural biology, physics, chemistry, computer graphics and robotics [4, 7, 16, 19, 21, 26, 30, 31, 32].

In this paper, we introduce the best-known deterministic method to date for $SO(3)$ in terms of providing:

1. incremental generation,
2. optimal dispersion-reduction with each additional sample,
3. explicit neighborhood structure,
4. the lowest metric distortion for grid neighbor edges,
5. equivolumetric partition of $SO(3)$ into grid regions.

The rest of the paper is organized around the presentation of the method. Section 2 defines the topological properties of $SO(3)$ together with its representations that are crucial for presenting our method. Section 3 overviews sampling requirements for the sequence. We discuss the relevant sampling methods that influenced our work in Section 4. Finally, we present our method in Section 5; experimental results and

its application to motion planning problems in Section 6. We conclude our work in Section 7.

2 Properties and Representations of $SO(3)$

The *special orthogonal group*, $SO(3)$, arises from rotations around the origin in \mathbb{R}^3 . Each rotation, by definition, is a linear transformation that preserves the length of vectors and orientation of space. The elements of $SO(3)$ form a group, with the group action being the composition of rotations. $SO(3)$ is not only a group, but also a manifold, which makes it a *Lie group*.

To sample $SO(3)$ uniformly, it is necessary to understand its topology. Any method known to date that produces uniform rotations relies on topology and Haar measure of $SO(3)$ (see Section 4).

Topology of $SO(3)$

$SO(3)$ is diffeomorphic to the *real projective space*, \mathbb{RP}^3 . It is hard to visualize the real projective space, because it cannot be embedded in \mathbb{R}^3 . Fortunately, it can be represented as $\mathbb{RP}^3 = S^3/(x \sim -x)$, the more familiar *3-sphere*, S^3 , embedded in \mathbb{R}^4 , with antipodal points identified. Topologists say that the 3-sphere is a *double cover* of \mathbb{RP}^3 , since one point of the projective space has two corresponding points on the 3-sphere.

Haar Measure on $SO(3)$

Up to a scalar multiple, there exists a unique measure on $SO(3)$ that is invariant with respect to the group action. This is called the *Haar measure*. That is, the Haar measure of a set is equal to the Haar measure of all of the rotations of the set. In our particular situation, we can think of the Haar measure as being invariant under all orthogonal coordinate changes. The Haar measure is an intrinsic property of $SO(3)$ which comes from the group structure, and is independent of its topological structure.

We have not used any coordinate system or parametrization of $SO(3)$ yet, since the notion of Haar measure is abstracted from representations of $SO(3)$. One has to use extreme caution when expressing the measure in terms of any of the representations we describe next. Not all of these naturally preserve the Haar measure.

Orthogonal Matrices

The elements of $SO(3)$ are defined as 3×3 orthogonal matrices with determinant $+1$. The group operation is multiplication of matrices. Because rotation matrices are less efficient and less numerically stable than quaternions, they are generally used less often than quaternions.

Quaternions

One of the most useful representations of the projective space is the set of quaternions. Let $x = (x_1, x_2, x_3, x_4) \in \mathbb{R}^4$ be a unit quaternion, $x_1 + x_2i + x_3j + x_4k$, $\|x\| = 1$, representing a 3D rotation. Because of the topological relationship between the projective space and the 3-sphere, once the identifications of the antipodal points on the 3-sphere are taken into account, metrics similar to those defined for the 3-sphere can be used for the projective space. Moreover, such metrics will respect the Haar measure on $SO(3)$.

The most natural way to define a metric for any two points $x, y \in SO(3)$ is as the length of the shortest arc between x and y on the 3-sphere, which quaternions conveniently allow to do:

$$\rho_{SO(3)}(x, y) = \cos^{-1} |(x \cdot y)|, \quad (1)$$

in which $(x \cdot y)$ denotes the dot product for vectors in \mathbb{R}^4 , and the absolute value, $|\cdot|$, guarantees that the shortest arc is chosen among the identifications of the two quaternions [11].

Quaternion representation is also convenient for calculating the composition of rotations, which is expressed as multiplication of quaternions. Any rotation invariant surface measure on S^3 naturally preserves the Haar measure for $SO(3)$ and can be used for quaternions. However, the surface measure is not straightforwardly expressed using quaternions. Other representations, such as spherical or Hopf coordinates, are more convenient for measuring the volume of surface regions.

Spherical Coordinates for $SO(3)$

Because of the topological relationship between the 3-sphere and $SO(3)$, hyperspherical coordinates can be used for $SO(3)$. Consider a point $(\theta, \phi, \psi) \in S^3$, in which $\psi \in [0, \pi/2]$ (to compensate for the identifications, we consider only one hemisphere of S^3), $\theta \in [0, \pi]$, and $\phi \in [0, 2\pi)$. For each ψ , the full ranges of θ and ϕ define a 2-sphere of radius $\sin(\psi)$. The quaternion $x = (x_1, x_2, x_3, x_4)$ corresponding to the rotation (θ, ϕ, ψ) can be obtained using the formula:




Fig. 1: Visualization of the spherical and Hopf coordinates on $SO(3)$ using angle and axis representation. This representation corresponds to a projection of the S^3 onto the equatorial solid sphere which we draw in \mathbb{R}^3 . (a) The full range of the spherical coordinate $\psi \in [0, \pi/2]$ is shown while the coordinates (θ, ϕ) form a discretization of size 20 over S^2 . (b) The half-spheres show the full ranges of the spherical coordinates $\theta \in [0, \pi]$, and $\phi \in [0, 2\pi)$, while ψ takes four discrete values over $[0, \pi/2]$. (c) The full range of the Hopf coordinate $\psi \in [0, 2\pi)$ is shown while the coordinates (θ, ϕ) form a discretization of size 12 over S^2 . (d) The spheres show the full ranges of the Hopf coordinates $\theta \in [0, \pi]$ and $\phi \in [0, 2\pi)$, while ψ takes four discrete values over $[0, 2\pi)$.

$$\begin{aligned}
 x_1 &= \cos \psi \\
 x_2 &= \sin \psi \cos \theta \\
 x_3 &= \sin \psi \sin \theta \cos \phi \\
 x_4 &= \sin \psi \sin \theta \sin \phi.
 \end{aligned} \tag{2}$$

The volume element on $SO(3)$ defines the Haar measure and has the following expression in spherical coordinates:

$$dV = \sin^2 \psi \sin \theta d\theta d\phi d\psi. \tag{3}$$

This representation is not as convenient for integration as the Hopf coordinates, which have a simpler expression for the Jacobian. Spherical coordinates are also cumbersome for computing compositions of rotations.

Hopf Coordinates for $SO(3)$

As opposed to spherical coordinates for hyperspheres, the *Hopf coordinates* are unique for both $SO(3)$ and the 3-sphere. They naturally describe the intrinsic structure of both of these spaces and provide a natural tool for obtaining uniform distributions on these spaces.

The *Hopf fibration* describes $SO(3)$ in terms of the circle S^1 and the ordinary 2-sphere S^2 . Intuitively, $SO(3)$ is composed of non-intersecting fibers, such that each fiber is a circle S^1 corresponding to a point on the 2-sphere. This fiber bundle structure is denoted as $SO(3) \cong S^1 \tilde{\otimes} S^2$. The Hopf fibration has the important property of locally being a Cartesian product space. The space $SO(3)$, however, is not (globally) the Cartesian product of S^2 and S^1 . Intuitively, $SO(3)$ is the product of S^2 and S^1 similarly to the way the Möbius band is locally the Cartesian product of an interval and a circle S^1 . That is, locally, a sequence of coordinates from each subspace results in a global parametrization of the space, whereas the global embedding into the Euclidean space introduces a twist, and does not have the Cartesian product structure. The Hopf coordinates can also be used for the 3-sphere, because of the topological relationship between the 3-sphere and $SO(3)$.

Each rotation in Hopf coordinates can be written as (θ, ϕ, ψ) , in which $\psi \in [0, 2\pi)$ parametrizes the circle S^1 , and $\theta \in [0, \pi]$ and $\phi \in [0, 2\pi)$ represent spherical coordinates on S^2 . The transformation to a quaternion $x = (x_1, x_2, x_3, x_4)$ can be expressed using the formula:

$$\begin{aligned} x_1 &= \cos \frac{\theta}{2} \cos \frac{\psi}{2} \\ x_2 &= \cos \frac{\theta}{2} \sin \frac{\psi}{2} \\ x_3 &= \sin \frac{\theta}{2} \cos \left(\phi + \frac{\psi}{2} \right) \\ x_4 &= \sin \frac{\theta}{2} \sin \left(\phi + \frac{\psi}{2} \right). \end{aligned} \tag{4}$$

A detailed derivation of the Hopf Coordinates is shown in Appendix 8.1. Briefly, Eq. (4) represents each rotation from $SO(3)$ as a rotation by angle $\psi \in S^1$ around the z axis, *followed* by the rotation, which places z in a position $(\theta, \phi) \in S^2$. Eq. (4) is obtained after the *composition* of these two rotations. The Hopf coordinates define exactly half of S^3 , since the coordinate x_2 never takes negative values. The Hopf coordinates can be extended to the entire S^3 by increasing the range of ψ to be $[0, 4\pi)$.

The volume element on $SO(3)$, which is also the surface volume element on S^3 , can be computed from Eq. (4) (see Appendix 8.2 for a detailed derivation), and has the following form:

$$dV = \frac{1}{8} \sin \theta \, d\theta \, d\phi \, d\psi. \tag{5}$$

Note that $\sin \theta \, d\theta \, d\phi$ represents the surface area on the 2-sphere, and $d\psi$ is the length element on S^1 . This formula additionally demonstrates that the length of a portion of S^1 is multiplied by the surface area of the base space, S^2 , to obtain the volume on $SO(3)$. The coefficient $1/8$ results from the fact that neither fibers S^1 nor

the base space S^2 are unit. In fact, in Appendix 8.3 we compute the lengths of the fibers and the surface area of the base space, which is used later for determining the grid cell sizes for our sequence.

As we have shown, the Hopf coordinates preserve the fiber structure of $SO(3)$ and are convenient for integration on $SO(3)$. However, composition of rotations is best expressed using quaternions.

Angle-Axis Representation for $SO(3)$

One of the most intuitive ways to represent rotations is by using *Euler's theorem*, which states that every 3D rotation is a rotation by some angle θ around a unit axis $\mathbf{n} = (n_1, n_2, n_3)$, $\|\mathbf{n}\| = 1$. The transformation from the angle-axis representation to quaternions is achieved by:

$$x = \left(\cos \frac{\theta}{2}, \sin \frac{\theta}{2} n_1, \sin \frac{\theta}{2} n_2, \sin \frac{\theta}{2} n_3 \right). \quad (6)$$

The angle-axis representation is useful for visualizing the projective space in \mathbb{R}^3 . Each rotation is drawn as a vector with direction n and a magnitude corresponding to θ (a multiple or a function of θ can be used; see Section 5.6, and [3]). Figure 1 shows the visualization of the spherical and Hopf coordinates on $SO(3)$ using the angle-axis representation. From this visualization one can immediately notice the singularities introduced by the spherical coordinates. It is also possible to see the advantage of using Hopf coordinates from this visualization. Hopf coordinates do not introduce singularities. The circles represented by the range of the variable ψ are all of equal length (see App. 8.3) and non-intersecting; they uniformly cover $SO(3)$. The fiber structure formed by these circles is also seen in the figure.

Euler Angles Representation

Euler angles are often used in robotics to represent rotations. Each rotation is then a vector (x_1, x_2, x_3) , $x_i \in [-\pi, \pi] / -\pi \sim \pi$. The topology of the resulting space is $S^1 \times S^1 \times S^1$, and, therefore, Euler angles do not correctly capture the structure of $SO(3)$. There are many detrimental consequences of this. Special tricks (see [11]) are needed to implement metric and measure that preserve Haar measure. Moreover, Euler angles are hard to compose, and present problems of singularities and the gimbal lock [23].

Understanding and respecting the global topology of $SO(3)$ is crucial for performing other numerical computations on the space. For example, sampling and interpolation of $SO(3)$ using Euler angles in [33] led to failure in producing motion planning path on a relatively simple problem, which was solved in seconds using the correct parametrizations. In the rest of the paper we use Hopf coordinates and quaternions to represent rotations.

3 Sampling Terminology and Problem Formulation

In applications such as motion planning, the algorithms are often terminated early and the particular order in which samples are chosen becomes crucial. Sampling literature distinguishes between a sample *set* and a sample *sequence*. For a sample set, the number of points, n , is specified in advance, and a set of n points is then chosen to satisfy the requirements of the method. The notion of ordering between points is not defined for a sample set but becomes important for sequences. Successive points in a sequence should be chosen carefully so that the resulting sample sets are all of good quality. Sequences are particularly suitable for motion planning algorithms, in which the number of points needed to solve the problem is not known in advance.

Now that the background definitions for $SO(3)$ have been presented in Section 2, to generate samples over $SO(3)$ we need to formulate the desirable properties for the samples. The first requirement is that samples form a sequence. We also require that samples get arbitrarily close to every point in $SO(3)$, i.e. that the sequence of samples is *dense* in $SO(3)$. Next, we formulate several requirements on the uniformity properties of samples.

3.1 Discrepancy and Dispersion

Additional requirements that the sequence needs to satisfy are described by the uniformity measures, *discrepancy* and *dispersion*.

Intuitively, discrepancy can be thought of as enforcing two criteria: first, that no region of the space is left uncovered; and second, that no region is left too full. Dispersion eliminates the second criterion, requiring only the first. It can be shown that low discrepancy implies low dispersion [18].

To define discrepancy, choose a *range space*, \mathcal{R} , as a collection of subsets of $SO(3)$. Let $R \in \mathcal{R}$ denote one such subset. Range spaces that are usually considered on spheres are the set of spherical caps (intersections of the 3-sphere with half-spaces) or the set of spherical slices (intersections of two 3-hemispheres) [20], which can be used on $SO(3)$ once the identifications of the 3-sphere are taken into account.

Let $\mu(R)$ denote the Haar measure of the subset R . If the samples in the set P are uniform in some ideal sense, then it seems reasonable that the fraction of these samples that lie in any subset R should be roughly $\mu(R)$ divided by $\mu(SO(3))$ (which is simply π^2 , see App. 8.3). We define the *discrepancy* [18] to measure how far from ideal the sample set P is:

$$D(P, \mathcal{R}) = \sup_{R \in \mathcal{R}} \left| \frac{|P \cap R|}{N} - \frac{\mu(R)}{\mu(SO(3))} \right|, \quad (7)$$

in which $|\cdot|$ applied to a finite set denotes its cardinality. Figure 2 (a) demonstrates the notion on the 2-sphere.




Fig. 2: An illustration of the notions of dispersion and discrepancy for a set of points on a 2-sphere. (a) The discrepancy searches for the subset R for which the deviation from the measure of R to the number of samples placed inside R is the largest. (b) The dispersion searches for a point q on the sphere which is the farthest from the sample points.

While discrepancy is based on measure, a metric-based criterion, *dispersion*, can be introduced:

$$\delta(P, \rho) = \max_{q \in SO(3)} \min_{p \in P} \rho(q, p). \quad (8)$$

Above, ρ denotes any metric on $SO(3)$ that agrees with the Haar measure, such as (1). Intuitively, this corresponds to the spherical radius of the largest empty ball that fits in between the samples (assuming all ball centers lie on $SO(3)$). See Figure 2(b) for an illustration.

3.2 Problem Formulation

In summary, the goal of this paper is to define a sequence of elements from $SO(3)$ which:

- is incremental,
- is deterministic,
- minimizes the dispersion (8) and discrepancy (7) on $SO(3)$,
- has grid structure with respect to the metric (1) on $SO(3)$.

4 Sampling Methods Overview

Our work was influenced by many successful sampling methods developed recently for spheres and $SO(3)$. As demonstrated in Table 1, several of them are highly related to the problem formulated in Section 3.2. However, none of the methods known to date has all of the desired properties.

	Random	Succ. Orth. Images	Layered Sukharev	HEALPix	this work
incremental	yes	no	yes	no	yes
uniform	yes	yes	no	yes	yes
deterministic	no	yes	yes	yes	yes
grid	no	no/yes	yes	yes	yes
spaces	$SO(3)$	$SO(n)$	$SO(3)$ and S^n	S^2	$SO(3)$

Table 1: The comparison of different sampling methods related to the problem of Section 3.2. The rows correspond to the desired properties of these methods.

Random Sequences of Elements from S^d . To generate uniformly distributed random points on a hypersphere S^d , spherical symmetry of the multidimensional Gaussian density function can be exploited [6]. For each of the $i = 0 \dots d + 1$ coordinates use a zero-mean Gaussian distribution with the same variance to generate x_i . This is done approximately by generating k uniformly distributed values from the interval $[-1, 1]$ and adding them following the *Central Limit Theorem*. In practice, any $k \geq 12$ is a reasonable choice. Then the normalized vector $(x_i / \|x_i\|)$ is uniformly distributed over the hypersphere S^d .

Random Sequence of $SO(3)$ Rotations. There are several ways of sampling the space of rotations uniformly at random [2, 24, 27, 34]. The main difficulty in doing so is the choice of a convenient parametrization of $SO(3)$. If a parameter space is sampled uniformly, the resulting samples on $SO(3)$ are not necessarily uniform. As was shown in Section 2, not all of the parametrizations of $SO(3)$ are natural representations of rotations, and some of them lead to measure distortions, and even singularities. Only few parametrizations, such as the Hopf coordinates, result in a local isometry to $SO(3)$.

It is easy to make the mistake of sampling rotations using a wrong parametrization [1]. The subgroup algorithm [5] for selecting random elements for $SO(3)$ is the correct and most popular method for uniform random sampling of $SO(3)$. It uses the fact that any Lie group can be uniformly sampled, by combining elements from a subgroup (in case of $SO(3)$ it is S^1), and the quotient, or coset space (S^2) at random. Essentially, this method utilizes the Hopf coordinates. Random sequences of rotations are used in many applications. However, they lack deterministic uniformity guarantees, and the explicit neighborhood structure.

Successive Orthogonal Images on $SO(n)$. Related to the subgroup method for generating random rotations, is the deterministic method of Successive Orthogonal Images [17], which generates lattice-like sets with a specified length step based on uniform deterministic samples from the subgroup, S^1 , and the coset space, S^2 . The method utilized Hopf coordinates, and is also generalized to arbitrary $SO(n)$.

The deterministic point sets from [17] can be applied to the problems in which the number of the desired samples is specified in advance. If the set of samples on S^2 is chosen so that it has a grid structure, the resulting set of samples on $SO(3)$ has the explicit neighborhood structure. Part of the current work will be in applying




Fig. 3: Different sampling methods on S^2 . (a) 200 random samples (b) 192 Sukharev grid samples [33] (c) icosahedron samples (d) 216 HEALPix samples [7]

this method in a way that provides the incremental quality necessary for our motion planning applications.

Layered Sukharev Grid Sequence for Spheres and $SO(3)$. Uniform, deterministic sequences were first designed for the unit cube $[0, 1]^d$ [13]. To minimize dispersion, the method places one resolution of grid at a time inside of the unit cube. A discrepancy-optimal ordering is then generated for each resolution. The sequence can be extended to spheres and $SO(3)$ [33] using the projection from faces of an inscribed cube. For $SO(3)$, though, the distortions produced by the method result in some grid cells being four times the volume of others.

The general method for designing Layered Sukharev Grid sequences inside Cartesian products was later presented in [14]. Our current paper builds on top of these works by combining the method in [14], with the Successive Orthogonal Images [17] generation of rotations, Hopf coordinates, and the HEALPix spherical sampling method [7] described next.

HEALPix Multiresolution Grids on S^2 . The HEALPix package [7] was designed for efficient and incremental discretization of full-sky maps in application to the satellite missions to measure the cosmic microwave background in astrophysics. It provides a deterministic, uniform, and multiresolution sampling method for the 2-sphere. Moreover, it possesses additional qualities, such as equal area partitioning of the 2-sphere, and isolatitude sampling on the 2-sphere, which make computations of the spherical harmonics integrals even more efficient.

The method takes advantage of the measure preserving cylindrical projection of the 2-sphere. This intrinsic property of the 2-sphere cannot be generalized directly to higher dimensional spheres. However, this work shows that an extremely uniform grid can be constructed on such a non-trivial curvature space as the 2-sphere. It is also not difficult to make this grid incremental using the method from [14]. We have done this as part of our implementation for the current work, and the code can be found at [9].

5 Our Approach

In this section we present our approach of sampling $SO(3)$ which satisfies all of the requirements of Section 3.2, and Table 1. The fiber bundle structure of $SO(3)$ locally behaves similarly to the Cartesian product of two spaces, S^1 and S^2 . Therefore, the method presented in [14] for constructing multiresolution grid sequences for Cartesian products of spaces, can be used for constructing a grid sequence on $SO(3)$. The resulting rotations are computed using the Hopf coordinates, as was first described in [17]. It is a much simpler problem to construct nicely behaved grids on the 1-sphere and 2-sphere. Hopf coordinates allow the two grids to be lifted to the space of rotations without loss of uniformity. Next, we outline the details of this construction.

5.1 Description of the Grid Structure on $SO(3)$

Let ψ be the angle parametrizing the circle, S^1 , and (θ, ϕ) be the spherical coordinates parametrizing the sphere, S^2 . Using these coordinates, define T_1 to be the multiresolution grid over the circle and T_2 to be the multiresolution grid over the 2-sphere. Let m_1 and m_2 be the number of points at the base resolution 0 of the grids T_1 and T_2 respectively.

There are numerous grids that can be defined on S^2 (see Figure 3 for an illustration of some). In this work we have selected the HEALPix grid [7] on S^2 , and the ordinary grid for S^1 . Both of these grids are uniform, have simple neighborhood structure, and can have multiple resolutions. Moreover, HEALPix divides the surface of the 2-sphere into subregions of equal area. After multiplying these by equal length fibers of S^1 , this results in equivolumetric partition of $SO(3)$ into the grid regions.

Next, consider the space $S^2 \tilde{\otimes} S^1$. The multiresolution grid sequence that we define for $SO(3)$ has $m_1 \cdot m_2 \cdot 2^{3l}$ points at the resolution level l , in which every 2^3 points falling into a single grid cell comprise a cube in Hopf coordinates. Each element of the sequence is obtained by combining the corresponding coordinates in the subspaces, S^1 and S^2 , using the Eq. (4). The dispersion and discrepancy of the resulting sequence can then be computed using the representation for the metric and volume element from Eq. (1) and (5).

5.2 Choosing the Base Resolution on S^1 and S^2

One of the issues arising when combining the two grids from S^1 and S^2 is the length of a grid cell edge along each of the coordinates. For this, we have to match the number of cells in each base grids on both of the subspaces, so that they have cell

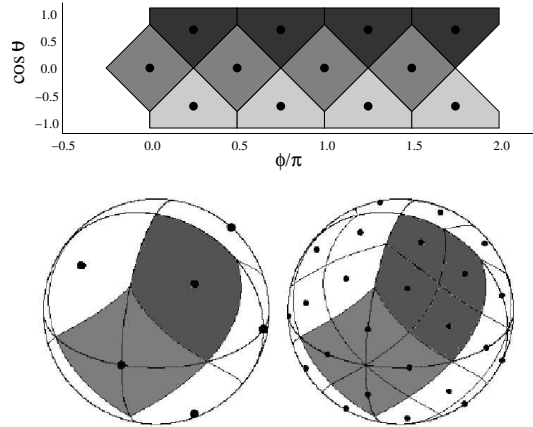


Fig. 4: The base grid of the HEALPix sequence consists of 12 points. The cylindrical projection of the grid cells from S^2 to $(\cos(\theta), \phi)$ coordinates is shown. Each next resolution subdivides each of the spherical squares into 4 squares of equal area [7].

sides of equal lengths [17]. That is, the following equation should hold for m_1 and m_2 :

$$\frac{\mu(S^1)}{m_1} = \sqrt{\frac{\mu(S^2)}{m_2}}, \quad (9)$$

in which $\mu(S^1)$ is the circumference of the circle S^1 and $\mu(S^2)$ is the surface area of S^2 . In Appendix 8.3 we explicitly show that both of these values are equal to π .


In our particular case, the base HEALPix grid consists of $m_2 = 12$ cells (Figure 4). Therefore, the number of points in the base resolution of the grid on S^1 is $m_1 = 6$. The base grid of the sequence for $SO(3)$ then consists of $m_1 \cdot m_2 = 6 \cdot 12 = 72$ points (the projections of the grid regions on the Hopf coordinates are shown in Figure 5).

5.3 Choosing the Base Ordering

The next step is to choose the ordering of the $m = m_1 m_2$ points within the base resolution on $SO(3)$. In general, the initial ordering will influence the quality of the resulting sequence, and a method similar to [14] can be used for deciding the ordering of the base sequence.

In our case, we have to define the ordering on the first 72 points of the sequence (see Figure 5 for the illustration of the associated grid regions). In our implementation [9] we have manually selected such an ordering. However, it is possible to

Fig. 5 The base grid of the proposed $SO(3)$ sequence consists of 72 points. For the Hopf coordinates (θ, ϕ, ψ) the projections of the grid cells on each of the coordinates are shown. Grid cells for ψ are chosen according to the ordinary grid on S^1 . The grid cells for $(\cos(\theta), \phi)$ are obtained using HEALPix. The manually selected ordering of the cells is shown for both of the projections, S^1 and S^2 .



design a program that would run through the orderings and select the one that minimizes the discrepancy or any other desired property. For the purpose of further analysis we assume that such an optimal ordering function $f_{base} : \mathbb{N} \rightarrow [1, \dots, 72]$ is available.

In our implementation [9] we first have selected an ordering of the 12 base points on S^2 and 6 base points on S^1 (these orderings are shown in Figure 5). For each point on S^2 we then generated the 6 points on S^1 according to the S^1 ordering. The points on S^2 are chosen according to the S^2 ordering.

5.4 The Sequence

The sequence for $SO(3)$ is constructed one resolution level at a time. The order in which the points from each resolution level are placed in the sequence can be described as follows. The ordering $f_{base} : \mathbb{N} \rightarrow [1, \dots, m]$ of the first m points in the base resolution determines the order of the grid regions within $SO(3)$ and is taken from the previous section. The points in other resolutions fall into the base resolution grids according to the function $f_{base}(i) = f_{base}(i \bmod m)$. Every successive m points in the sequence should be placed in these grid regions in the same order. Each of the grid regions is isomorphic to the $[0, 1]^3$, and is subdivided into 8 grid regions in each successive resolution. Where exactly each point should be placed within each of the grid regions is determined by the ordering $f_{cube} : \mathbb{N} \rightarrow [1, \dots, 8]$ and recursion procedure defined for the cube $[0, 1]^3$ in [13].

The resulting procedure for obtaining the coordinates of the i th element in the sequence is the following:

1. Assign $f_{base}(i)$ to be the index of the base grid region that the i -th element has to be placed within.
2. Assign the floor of the division, $i_{cube} = \lfloor i/m \rfloor$, to be the number of subregions already generated in the base grid region. This index then determines the subregion of the region $f_{base}(i)$ that the i -th element has to be placed within.




Fig. 6: Different sets of samples on $SO(3)$ (a) 2000 random samples (b) 2048 Sukharev grid samples (c) 1944 icosahedral samples (d) 1944 HEALPix samples

3. Call the recursive procedure from [13] to determine the coordinates of the subregion of the cube $[0, 1]^3$ determined by the index i_{cube} and the ordering f_{cube} . The i -th element is then placed within this subregion of the $f_{base}(i)$ region.

5.5 Analysis

Several claims, similar to those obtained in [13], can be made for the new approach. The most important distinction is that the new sequence provides equal volume partition of $SO(3)$ which results in a strong dispersion guarantee.

Proposition 0.1. *The dispersion of the sequence T at the resolution level l satisfies:*

$$\delta(T) \leq 2 \sin^{-1} \left(\frac{1}{2} \sqrt{\delta^2(T_2) + \left(\frac{\pi}{m_1 2^l} \right)^2} \right),$$

in which $\delta(T_2)$ is the dispersion of the sequence T_2 defined over S^2 .

Proof: It was shown in [17] that the fibers S^1 are locally orthogonal to the base space S^2 in the sense that an equivalent of the Pythagorean theorem holds for the Hopf coordinates. The bound follows directly from the Pythagorean theorem, and the dispersion bound on the ordinary grid on S^1 at the resolution level l . \square

Proposition 0.2. *The sequence T has the following properties:*

- *It is discrepancy-optimal with respect to the set of axis-aligned grid regions defined over S^1 and S^2 .*
- *The position of the i -th element of T can be generated in $O(\log i)$ time.*
- *For any i -th sample, any of the $2d$ nearest grid neighbors from the same layer can be found in $O((\log i)/d)$ time.*

Proof: The proof closely follows similar considerations in [13]. \square

5.6 Visualization of the Results

To visualize our sequence and compare it with other sequences designed for $SO(3)$, we use the angle-axis, (θ, n) , representation from Section 2. It can be shown that if the rotations are uniformly distributed, then the distribution of the angle θ is $(\sin(\theta) - \theta)/\pi$. This allows us to draw the elements of $SO(3)$ as the points inside a ball in such a way that every radial line has uniform distribution of elements. This provides a more intuitive visualization, which partially preserves the uniformity. See Figure 6 for visualization of several of the methods of sampling over $SO(3)$, compared to the proposed approach. Specifically, the images show points in the direction of the axis of rotation and with distance to the origin equal to $(\sin(\theta) - \theta)/\pi$. Using this representation, the distribution of points increases linearly as a function of distance from the origin. In comparison, a set of points that was uniform with respect to the measure on \mathbb{R}^3 would have a distribution that varies as the cube of distance from the origin.

6 Implementation and Application to Motion Planning

We have implemented our sampling algorithm in C++ as part of the new library publicly available at [9]. The experiments reported here were performed on a 2.2 GHz Pentium IV running Linux and compiled under GNU C++.

We first compared the uniformity of the new sequence with the Layered Sukharev sequence and random sequence. To demonstrate the importance of understanding the topology of $SO(3)$, we have included the evaluation of uniformity for random Euler angles in this experiment. For each of the deterministic sequences, we generated a fixed set of points, in the range 50 to 100,000, which is shown on the x axis of the graph in Figure 7. We then calculated the distance from a randomly generated point on $SO(3)$ to the nearest neighbor in each of the sets, and selected the largest such distance among 10,000 random points. We have averaged the same computations over 10 runs for each of the random sequences. The obtained value approximates the dispersion. The results are shown for each of the sequences as a separate curve in Figure 7. In all of the cases, the smallest obtained value was the one generated with our new method, which demonstrates that the resulting samples are more uniformly distributed compared to other sequences known for $SO(3)$.




Fig. 7: (a) For each of the deterministic sequences, we generated a fixed set of points, in the range 50 to 100,000, which is shown on the x axis of the graph. We then calculated the distance from a randomly generated point on $SO(3)$ to the nearest neighbor in each of the sets, and selected the largest such distance among 10,000 random points. The obtained value approximates the dispersion and is shown on the y axis for each of the sequences. We have averaged the same computations over 10 runs for each of the random sequences. A portion of the graph is magnified in (b).

Even though it might appear that the actual difference in dispersion is not significant for data sets of a particular size, there is another interpretation of the results on the graph. Consider a particular value of dispersion on the graph, for example, $d = 0.06$. If a sample set has this dispersion then no ball of radius $r > 0.06$ can be placed between the samples. To achieve such dispersion, the Hopf sequence required around 50,000 samples, whereas the sample set generated using random Euler angles with twice as many points does not reach the same resolution (see Fig. 7(b)).

We also used our library as the sampling method in the implementation of PRM-based planner [10] in the Motion Strategy Library [12]. It is important to note that the experiments we present here are just one of possible applications of the developed sequences to motion planning problems. Alternate applications may exist in other areas of computer science, or related fields.

In our experimental setup we consider the rotation-only models for which the configuration space is $SO(3)$. For the two problems shown in Figure 8 we have compared the number of nodes generated by the basic PRM planner using the pseudo-random sequence (with quaternion components [24]), the layered Sukharev grid sequence, and the new sequence. For the first problem the results are: 258, 250, and 248 nodes, respectively. To solve the second problem the PRM planner needed 429, 446 and 410 nodes, respectively. In each trial a fixed, random quaternion rotation was premultiplied to each deterministic sample, to displace the entire sequence. The results obtained were averaged over 50 trials.

Based on our results we have observed that the performance of our method is equivalent or better than the performance of the previously known sequences for the basic PRM-based planner. This makes our approach an alternative approach for use in motion planning. It is important to note, however, that for some applications, such as verification problems, only strong resolution guarantees are acceptable.




Fig. 8: Motion planning problems involving: a) moving a robot (black) from the north pole to the south pole. Multiple views of the geometry of the problem are shown (obstacles are drawn in lighter shades); and b) moving a robot along the corridor.

7 Conclusions

In conclusion, we have developed and implemented a deterministic incremental grid sequence on $SO(3)$ that is highly uniform, can be efficiently generated, and divides the surface of $SO(3)$ into regions of equal volume. Sequences that minimize uniformity criteria, such as dispersion and discrepancy, at each step of generation are especially useful in applications in which the required number of samples is not known in advance.

In the paper we report the performance of the deterministic sequence on motion planning examples for demonstration purposes only. We believe that the value of this method is in providing strong provable guarantees (the bound on dispersion, neighborhood structure, and deterministic generation). These guarantees are not always required in motion planning. Therefore, we do not perform extensive experimental evaluation. However, for the motion planning examples in Section 6, the provable guarantees of the method come at no additional cost. Moreover, it is consistent with the performance of other deterministic sampling methods on Euclidean spaces we have observed on numerous motion planning examples in our previous works. Therefore, we conclude that the sequence can be applied in the context of motion planning, in case deterministic guarantees are required, or in any other applications with such guarantees.

There are a number of ways to improve the current work which we consider as future directions. It is an interesting problem to determine the criteria for an optimal selection of the base sequence on $SO(3)$ to improve the performance of the sequence. It is also tempting to assess the general rate of convergence for motion planning solutions using different sampling sequences.

There are many general open problems related to the presented work. Nicely distributed grids are not yet developed for general n -spheres, $n > 3$. Implicitly defined manifolds, such as the ones arising from motion planning for closed linkages, are very hard to efficiently and uniformly sample. Such manifolds also arise as the conformation spaces of protein loops. In such cases, efficient parametrization is the bottleneck for developing sampling schemes.

8 Appendix

8.1 Derivation of the Hopf Coordinates

The group of all 3D rotations, $SO(3)$, has a subgroup of all planar rotations, S^1 . Therefore, $SO(3)$ can be represented as a disjoint union of all of the cosets of this subgroup. Consider a subgroup S^1 of all rotations around z axis by angle $\psi \in [0, 2\pi)$. Such rotations can be represented by a unit quaternion of the form:

$$h_1 = \left(\cos \frac{\psi}{2}, 0, 0, \sin \frac{\psi}{2} \right).$$

To obtain the left coset of this subgroup, the z axis should be placed in an arbitrary position on the sphere S^2 . If the sphere is parametrized using the spherical coordinates $\theta \in [0, \pi]$ and $\phi \in [0, 2\pi)$, then the rotation of the z axis to the (θ, ϕ) position on the sphere corresponds to the quaternion of the form:

$$h_2 = \left(\cos \frac{\theta}{2}, \sin \frac{\theta}{2} \sin \phi, \sin \frac{\theta}{2} \cos \phi, 0 \right).$$

To obtain the Hopf coordinates for an element $x \in SO(3)$, we compose the two rotations $x = h_1 * h_2$ and obtain the expression in Eq. (4):

$$x = \left(\cos \frac{\theta}{2} \cos \frac{\psi}{2}, \sin \frac{\theta}{2} \sin \left(\phi + \frac{\psi}{2} \right), \cos \left(\phi + \frac{\psi}{2} \right) \sin \frac{\theta}{2}, \cos \frac{\theta}{2} \sin \frac{\psi}{2} \right).$$

8.2 Derivation of the Volume Element Using the Hopf Coordinates

The Jacobian, $J = (\partial x / \partial \theta, \partial x / \partial \psi, \partial x / \partial \phi)$, of the coordinate transformation in Eq. (4) is the following:

$$J = \begin{pmatrix} -\frac{1}{2} \sin \frac{\theta}{2} \cos \frac{\psi}{2} & -\frac{1}{2} \cos \frac{\theta}{2} \sin \frac{\psi}{2} & 0 \\ -\frac{1}{2} \sin \frac{\theta}{2} \sin \frac{\psi}{2} & \frac{1}{2} \cos \frac{\theta}{2} \cos \frac{\psi}{2} & 0 \\ \frac{1}{2} \cos \frac{\theta}{2} \sin \left(\phi + \frac{\psi}{2} \right) & \frac{1}{2} \sin \frac{\theta}{2} \cos \left(\phi + \frac{\psi}{2} \right) & \sin \frac{\theta}{2} \cos \left(\phi + \frac{\psi}{2} \right) \\ \frac{1}{2} \cos \frac{\theta}{2} \cos \left(\phi + \frac{\psi}{2} \right) & -\frac{1}{2} \sin \left(\phi + \frac{\psi}{2} \right) \sin \frac{\theta}{2} & -\sin \frac{\theta}{2} \sin \left(\phi + \frac{\psi}{2} \right) \end{pmatrix}.$$

We next compute the volume element of the transformation by taking a square root of the following determinant:

$$dV = \sqrt{\det(J^T J)} d\theta d\phi d\psi = \frac{1}{8} \sin \theta d\theta d\phi d\psi.$$

8.3 Lengths of the S^1 Fibers and the Area of the S^2 Base Space for the Hopf Coordinates

By fixing θ and ϕ in Eq. (4) we obtain half circles on the hypersphere S^3 (which are S^1 fibers on $SO(3)$, after the identifications of the antipodal points are taken into account on S^3). To compute the length of a fiber, we follow similar derivations to the Appendix 8.2. For the fixed values of θ and ϕ , the Jacobian of the resulting transformation is the matrix:

$$J = (\partial x / \partial \psi) = \begin{pmatrix} -\frac{1}{2} \cos \frac{\theta}{2} \sin \frac{\psi}{2} \\ \frac{1}{2} \cos \frac{\theta}{2} \cos \frac{\psi}{2} \\ \frac{1}{2} \sin \frac{\theta}{2} \cos(\phi + \frac{\psi}{2}) \\ -\frac{1}{2} \sin(\phi + \frac{\psi}{2}) \sin \frac{\theta}{2} \end{pmatrix}.$$

Then the length element for S^1 is:

$$\sqrt{\det(J^T J)} d\psi = \frac{1}{2} d\psi.$$

Therefore, the length of each fiber is:

$$\frac{1}{2} \int_0^{2\pi} d\psi = \pi.$$

The area of the corresponding base sphere, S^2 , is then obtained from the volume $V(SO(3)) = \pi^2$ as the remaining contribution. Therefore, it is π . Note, that these derivations demonstrate that the base space S^2 and the fibers S^1 are not unit for the Hopf fibration.

Acknowledgements Primary funding for this project was provided by the National Science Foundation (NSF CISE-0535007) in support of Steve LaValle and Anna Yershova at UIUC, along with the Alfred P Sloan Foundation and US Department of Energy (DE-FG02-04ER25627) in support of Julie Mitchell. In addition, Anna Yershova's work while at Duke University was supported by the following grants to Bruce R. Donald: NIH grants R01 GM-65982 and R01 GM-078031.

References

1. J. Arvo. Random rotation matrices. In *Graphics Gems II*, J. Arvo (editor), pages 355–356. Academic Press, Boston, 1991.
2. J. Arvo. *Fast random rotation matrices*, pages 117–120. Academic Press, Boston, 1992.
3. G. S. Chirikjian and A. B. Kyatkin. *Engineering Applications of Noncommutative Harmonic Analysis*. CRC Press, Boca Raton, 2001.
4. R. A. Crowther. The fast rotation function in the molecular replacement method. *Int. Sci. Rev. Ser.*, 12:173–178, 1972.
5. P. Diaconis and M. Shahshahani. The subgroup algorithm for generating uniform random variables. *Prob. in Eng. and Info. Sci.*, 1:15–32, 1987.

6. G. F. Fishman. *Monte Carlo: Concepts, Algorithms, and Applications*. Springer-Verlag, Berlin, 1996.
7. K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann. HEALPix: a framework for high-resolution discretization and fast analysis of data distributed on the sphere. *arXiv:astro-ph/0409513*, 622:759–771, Apr. 2005.
8. D. P. Hardin and E. B. Saff. Discretizing manifolds via minimum energy points. *Notices of the American Mathematical Society*, 51(10):1186–1194, 2004.
9. S. Jain. Library for uniform deterministic sequences/sets of samples over 2-sphere and $SO(3)$. 2009. <http://www.cs.duke.edu/~yershova/so3sampling/so3sampling.htm>.
10. L. E. Kavraki, P. Svestka, J. Latombe, and M. H. Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Trans. Robot. & Autom.*, 12(4):566–580, June 1996.
11. J. Kuffner. Effective sampling and distance metrics for 3D rigid body path planning. In *IEEE Int. Conf. Robot. & Autom.* IEEE, May 2004.
12. S. M. LaValle. MSL: Motion strategy library. <http://msl.cs.uiuc.edu/msl/>.
13. S. R. Lindemann and S. M. LaValle. Incremental Low-Discrepancy lattice methods for motion planning. In *IEEE Int'l Conf. on Robotics and Automation*, pages 2920–2927, 2003.
14. S. R. Lindemann, A. Yershova, and S. M. LaValle. Incremental grid sampling strategies in robotics. In *Workshop on the Algorithmic Foundations of Robotics*, 2004.
15. A. Lubotsky, R. Phillips, and P. Sarnak. Hecke operators and distributing points on the sphere I. *Comm. Pure Appl. Math.*, XXXIX:S149–S138, 1986.
16. J. G. Mandell, V. A. Roberts, M. E. Pique, V. Kotlovyi, J. C. Mitchell, E. Nelson, I. Tsigelny, and L. F. Ten Eyck. Protein docking using continuum electrostatics and geometric fit. *Protein Engineering*, 14(2):105–113, 2001.
17. J. C. Mitchell. Sampling rotation groups by successive orthogonal images. *SIAM J. Sci. Comput.*, 30(1):525–547, 2007.
18. H. Niederreiter. *Random Number Generation and Quasi-Monte-Carlo Methods*. Society for Industrial and Applied Mathematics, Philadelphia, USA, 1992.
19. S. Ramamoorthy, R. Rajagopal, Q. Ruan, and L. Wenzel. Low-discrepancy curves and efficient coverage of space. In *Proc. of the Workshop on Algorithmic Foundations of Robotics*, 2006.
20. G. Rote and R. F. Tichy. Spherical dispersion with applications to polygonal approximation of the curves. *Anz. Österreich. Akad. Wiss. Math.-Natur, Kl. Abt. II*, 132:3–10, 1995.
21. J. Rovira, P. Wonka, F. Castro, and M. Sbert. Point sampling with uniformly distributed lines. *Point-Based Graphics, 2005. Eurographics/IEEE VGTC Symposium Proc.*, June 2005.
22. E. B. Saff and A. B. J. Kuijlaars. Distributing many points on a sphere. *Mathematical Intelligencer*, 19(1):5–14, 1997.
23. K. Shoemake. Animating rotation with quaternion curves. In *Proc. of the 12th Annual Conference on Computer Graphics and Interactive Techniques*. ACM Press, 1985.
24. K. Shoemake. Uniform random rotations. In D. Kirk, editor, *Graphics Gems III*, pages 124–132. Academic Press, 1992.
25. N. J. A. Sloane, T. S. D. R. H. Hardin, and J. H. Conway. Minimal-energy clusters of hard spheres. *Disc Comp Geom*, 14:237–259, 1995.
26. M. J. E. Sternberg and G. Moont. Modelling protein-protein and protein-DNA docking. In *Bioinformatics – From Genomes to Drugs, T. Lengauer (editor)*, pages 361–404. Wiley-VCH, Weinheim, 2002.
27. G. W. Stewart. The efficient generation of random orthogonal matrices with an application to condition estimation. *SIAM J Numer Anal*, 17(3):403409, 1980.
28. X. Sun and Z. Chen. Spherical basis functions and uniform distribution of points on spheres. *J. Approx. Theory*, 151(2):186–207, 2008.
29. G. Wagner. On a new method for constructing good point sets on spheres. *Journal of Discrete and Computational Geometry*, 9(1):119–129, 1993.
30. D. Wales and J. Doye. Global optimization by Basin-Hopping and the lowest energy structures of Lennard-Jones clusters containing up to 110 atoms. *J Phys Chem A*, 101, 1997.

31. A. K. Yan, C. J. Langmead, and B. R. Donald. A Probability-Based similarity measure for saupe alignment tensors with applications to residual dipolar couplings in NMR structural biology. *Int. J. Robot. Res.*, 24(2-3):162–182, 2005.
32. G. Yang and I.-M. Chen. Equivolumetric partition of solid spheres with applications to orientation workspace analysis of robot manipulators. *IEEE Transactions on Robotics*, 22(5):869–879, Oct. 2006.
33. A. Yershova and S. M. LaValle. Deterministic sampling methods for spheres and $SO(3)$. In *Proc. IEEE International Conference on Robotics and Automation*, 2004.
34. K. Zyczkowski and M. Kus. Random unitary matrices. *J. Phys. A*, 27(12):4235–4245, 1994.