

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

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

Abstract The problem of generating uniform deterministic samples over the rotation group, $SO(3)$, is fundamental to many fields, such as computational structural biology, robotics, computer graphics, astrophysics. We present the best-known method to date for constructing incremental, deterministic grids on $SO(3)$; it provides the: 1) 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

Discretization of $SO(3)$, the space of 3D rotations, is a difficult problem that arises in numerous engineering and scientific fields. Examples include biological protein docking problems, robot motion planning, 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 if the space is nicely behaved, as in the case of 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 implic-

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

itly defined, three-dimensional surface embedded in \mathbb{R}^9 . One approach is to place a coordinate systems on the surface, causing it to behave like a patch in \mathbb{R}^3 . However, any such coordinates cause metric distortions in comparison to distances on the original surface. Only few representations of $SO(3)$, such as quaternions, 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 . Such curvature prohibits the introduction of distortion-free grids, similarly 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 that only half of S^3 is used, with unusual connectivity in the equatorial three-plane.

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 such criteria as energy functions, discrepancy, dispersion, and mutual distances has been extensively studied in mathematics and statistics [5, 18, 19]. Problems of sampling rotational groups and spheres have been applied in the context of computational structural biology, physics, chemistry, computer graphics, and robotics [4, 12, 14, 17, 20, 21].

In this paper, we introduce the best-known deterministic method to date for $SO(3)$ that provides the: incremental generation, optimal dispersion reduction with each additional sample, explicit neighborhood structure, lowest metric distortion for grid neighbor edges, and equivolumetric partition of $SO(3)$ into grid cells.

The rest of the paper is organized around the presentation of the method. Section 2 defines the topological properties of $SO(3)$ together with different parametrizations and coordinate systems 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 and experimental results in application to several 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*.

- *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 can not 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 covering* 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 group actions. 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 very abstract. 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.

- *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 opposite 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 $x, y \in \mathbb{RP}^3$ is as the length of the shortest arc between x and y on the 3-sphere:

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

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 [7].

Quaternion representation is also useful 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 relationship between the 3-sphere and \mathbb{RP}^3 , hyperspherical coordinates can be used for $SO(3)$. Consider a point $(\theta, \phi, \psi) \in S^3$, in which ψ has a range of $\pi/2$ (to compensate for the identifications, we consider only one hemisphere of S^3), θ has a range of π , and ϕ has a range of 2π . For each ψ , the 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 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 ranges of the spherical coordinates (θ, ϕ) , 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 ranges of the Hopf coordinates (θ, ϕ) , while ψ takes four discrete values over S^1 .

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

The volume element of the $SO(3)$ 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 convenient for integration though, because of the complicated 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 the 3-sphere, and thus for $\mathbb{R}P^3$. They naturally describe the intrinsic structure of both the 3-sphere and $\mathbb{R}P^3$ and provide a natural tool for obtaining uniform distributions on these spaces. The *Hopf fibration* describes $\mathbb{R}P^3$ in terms of a circle S^1 and an ordinary 2-sphere S^2 . Intuitively, $\mathbb{R}P^3$ is composed of non-intersecting fibers, such that each fiber is

a circle corresponding to the 2-sphere. This fiber bundle structure is denoted as $\mathbb{R}P^3 \cong S^1 \tilde{\otimes} S^2$. The Hopf fibration has the important property of locally being a Cartesian product space. The projective space, $\mathbb{R}P^3$, however, is not (globally) a Cartesian product of S^2 and S^1 . Intuitively, $\mathbb{R}P^3$ 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 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 $\mathbb{R}P^3$.

Each rotation in Hopf coordinates can be written as (θ, ϕ, ψ) , in which ψ parametrizes the circle S^1 and has the range of 2π . The ranges of θ and ϕ are π and 2π respectively, and they represent spherical coordinates for S^2 . The transformation to a quaternion $x = (x_1, x_2, x_3, x_4)$ can be expressed using the formula:

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

Equation (4) represents each rotation from $SO(3)$ as the rotation by an angle $\psi \in S^1$ around an axis z , followed by the rotation, which places z in a position $(\theta, \phi) \in S^2$. Formula (4) is obtained after the *composition* of the two rotations. The volume element on $\mathbb{R}P^3$ which respects the Haar measure is then defined as the surface volume on S^3 :

$$dV = \sin \theta \, d\theta \, d\phi \, d\psi. \quad (5)$$

Note that $\sin \theta \, d\theta \, d\phi$ represents the surface area on the 2-sphere, and $d\psi$ is the length element on the circle. This formula additionally demonstrates that the volumes from the two subspaces, S^2 and S^1 , are simply multiplied to obtain the volume on $SO(3)$. The Hopf coordinates, though, are not convenient for expressing compositions of rotations.

- *Axis-Angle 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 angle and axis representation to quaternions is achieved by:

$$x = (\cos(\theta/2), \sin(\theta/2)n_1, \sin(\theta/2)n_2, \sin(\theta/2)n_3). \quad (6)$$

The angle and axis representation is useful for visualizing the projective space in R^3 . Each rotation is drawn as a vector with direction \mathbf{n} and a magnitude corresponding to θ (a multiple or a function of θ can be used; see Section 5.5, and [1]). Figure 1 shows the visualization of the spherical and Hopf coordinates on $SO(3)$ using the angle and axis representation. From this visualization one can immediately notice the singularities produced 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 non-intersecting; they uniformly cover the $SO(3)$. The fiber structure formed by these circles is also seen on 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 [7]) 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 [15, 22]. 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 [11].

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) [13], 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). We define the *discrepancy* [11] 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.

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 $SO(3)$).

3.2 Problem Formulation

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

- is incremental and 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 the table of Figure 2, several of them are highly related to the problem formulation in Section 3.2. However, none of the methods known to date has all of the desired properties.

Random Sequence of $SO(3)$ Rotations. There are several ways of sampling the space of rotations uniformly at random [16, 23]. The subgroup algorithm [2] for selecting random elements for $SO(3)$ is the most popular method for uniform random sampling of $SO(3)$. Essentially, this method utilizes the Hopf coordinates. Another way of obtaining uniformly distributed random samples over $SO(3)$ is by using spherical symmetry of the multidimensional Gaussian density function [3].

	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	$SO(3)$

Fig. 2: 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 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 [10], 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 can be applied to the problems in which the number of the desired samples is specified in advance. If the sample on S^2 is chosen so that it has a grid structure, the resulting sample on $SO(3)$ has the explicit neighborhood structure. Part of the current work will be in applying this method in a way that provides the incremental quality necessary for our motion planning applications.

Layered Sukharev Grid Sequence. Uniform, deterministic sequences were first designed for the unit cube [8]. 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 $SO(3)$ [22] 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 [9]. Our current paper builds on top of these works by combining the method in [9], with the Successive Orthogonal Images [10] generation of rotations, Hopf coordinates, and the HEALPix spherical sampling method [4] described in the next section.

HEALPix. The HEALPix package [4] was designed for applications in astrophysics. It provides a deterministic, uniform, and multiresolution sampling method for the 2-sphere with additional properties, such as equal-area partitioning, and iso-latitude sampling on the 2-sphere.

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 [9].




Fig. 3: Different sampling methods on S^2 . (a) 200 random samples (b) 192 Sukharev grid samples [22] (c) icosahedron samples (d) 216 HEALPix samples [4]

5 Our Approach

In this section we present an approach which satisfies all of the requirements of Section 3.2, and Figure 2. 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 [9] 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 [10]. It is a much simpler problem to construct nicely behaved grids on the 1-sphere and 2-sphere. Hopf coordinates allow the two grids be lifted to the space of rotations without loss of uniformity. Next we outline the details of this construction.

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 be the grid over the 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 [4] 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.

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 successive 2^3 points define a cube in Hopf coordinates. Each element of the sequence is obtained by combining the corresponding coordinates in the subspaces, $p = (\theta, \phi, \psi)$, using formula (4). If the grid regions are defined on the two subspaces S^1 and S^2 , the corresponding grid regions are also obtained on $SO(3)$ by combining the corresponding coordinates. The dispersion, and discrepancy of the resulting sequence can be easily computed using the representation for the metric and volume element from equations (1), and (5).

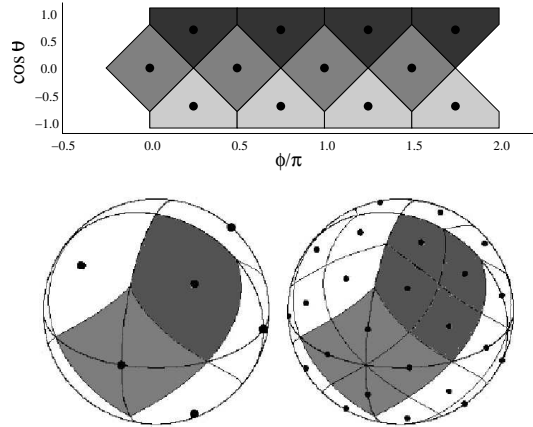


Fig. 4: The base grid of the HEALPix sequence consists of 12 points. The cylindrical projection of the grid cells from S^2 to $(\cos(\theta), \phi)$ coordinates is shown. Each next resolution subdivides each of the spherical squares into 4 squares of equal area[4].

5.1 Choosing the Base Resolution

One of the issues arising when combining the two grids from S^1 and S^2 is the length of a grid cell 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 sides of equal lengths [10]. That is, the following equation should hold for m_1 and m_2 :

$$\frac{2\pi}{m_1} = \sqrt{\frac{4\pi}{m_2}}, \quad (9)$$


in which 2π is the circumference of the circle S^1 , and 4π is the surface area of S^2 .

In our particular case, the base HEALPix grid consists of $m_2 = 12$ cells, and the surface area of each cell is equal to $4\pi/12 = \pi/3$ (Figure 4). Therefore, the length of the side of each grid cell is approximately the square root of that value, that is 1. Then, the number of points in the base resolution of the grid on S^1 needs to be $m_1 = 6$, since it should be close to the length of the circle, which is 2π . Therefore, the base grid of the sequence for $SO(3)$ consists of $m_1 \cdot m_2 = 6 \cdot 12 = 72$ points (the projections of the grid regions on the Hopf coordinates are shown on Figure 5).

5.2 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

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.



resulting sequence, and a method similar to [9] can be used for deciding the ordering of the general base sequences.

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). To do this, one can notice that there are antipodal grid cells in both of the subspaces. Removing antipodal cells from the final sequence can significantly eliminate the number of the base resolution grid cells on $SO(3)$.

In our preliminary experiments in the application to motion planning problems (Section 6) we have manually selected such an ordering. However, it is possible to design a program that would run through the orderings and select the ones that minimize the discrepancy. For the further analysis results we assume that the optimal ordering function $f_{base} : \mathbb{N} \rightarrow [1, \dots, 72]$ is given.

5.3 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 $f_{base}()$ 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. 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 [8].

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 ceiling of the division, $i_{cube} = \lceil i/m \rceil$, to be the index that determines the subregion of the region $f_{base}(i)$ that the i -th element has to be placed within.

3. Call the recursive procedure from [8] 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.

5.4 Analysis

Several claims, similar to those obtained in [8], can be made for the new approach. The most important distinction is that the new sequence provides equal volume partition of the $SO(3)$ which results in strong dispersion guarantees.

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: The bound follows directly from the Pythagorean theorem, and the dispersion bound on the ordinary grid T_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 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 [8]. \square

5.5 Visualization of the Results

To visualize our sequence and compare it with other sequences designed for $SO(3)$, we use the angle and axis, (θ, n) , representation from Section 2. It can be shown that if the rotations are uniformly distributed, then the distribution of an angle θ is $(\sin(\theta) - \theta)/\pi$ [1]. 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




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

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 Application: Motion Planning

We have implemented our algorithm in C++ and applied to implementations of PRM-based planner [6] in the Motion Strategy Library. The experiments were performed on a 2.2 Ghz Pentium IV running Linux and compiled under GNU C++.

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 7 we have compared the number of nodes generated by the basic PRM planner using the pseudo-random sequence (with quaternion components [16]), the layered Sukharev grid sequence, and the new sequence derived from the use of Successive Orthogonal Images, the HEALPix method and incremental Sukharev ordering. 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.




Fig. 7: 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.

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.

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. One such application is robotic motion planning. We have demonstrated the use of our method on several motion planning problems, showing that resolution completeness guarantees can be achieved at small computational cost.

There is a number of ways to improve current work, which we plan to address in the future. We plan to complete a more rigorous analysis, as well as comparison of different base sequences for S^2 to improve our understanding of the benefits of our method. A more extensive experimental evaluation of the sequences is also a part of our future work. It is yet inconclusive, but tempting, to assess the general rate of convergence for motion planning solution using different sampling sequences.

There are many general open problems arising from this research. 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.

Acknowledgements This work was partially supported by NSF CISE-0535007.

References

1. G. S. Chirikjian and A. B. Kyatkin. *Engineering Applications of Noncommutative Harmonic Analysis*. CRC Press, Boca Raton, 2001.
2. P. Diaconis and M. Shahshahani. The subgroup algorithm for generating uniform random variables. *Prob. in Eng. and Info. Sci.*, 1:15–32, 1987.
3. G. F. Fishman. *Monte Carlo: Concepts, Algorithms, and Applications*. Springer-Verlag, Berlin, 1996.
4. 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.
5. D. P. Hardin and E. B. Saff. Discretizing manifolds via minimum energy points. *Notices of the American Mathematical Society*, 51(10):1186–1194, 2004.
6. L. E. Kavradi, 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.
7. J. Kuffner. Effective sampling and distance metrics for 3D rigid body path planning. In *IEEE Int. Conf. Robot. & Autom.* IEEE, May 2004.
8. 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.
9. S. R. Lindemann, A. Yershova, and S. M. LaValle. Incremental grid sampling strategies in robotics. In *Workshop on the Algorithmic Foundations of Robotics*, 2004.
10. J. C. Mitchell. Sampling rotation groups by successive orthogonal images. *SIAM J. Sci. Comput.*, 30(1):525–547, 2007.
11. H. Niederreiter. *Random Number Generation and Quasi-Monte-Carlo Methods*. Society for Industrial and Applied Mathematics, Philadelphia, USA, 1992.
12. 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.
13. 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.
14. 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.
15. K. Shoemake. Animating rotation with quaternion curves. In *Proc. of the 12th Annual Conference on Computer Graphics and Interactive Techniques*. ACM Press, 1985.
16. K. Shoemake. Uniform random rotations. In D. Kirk, editor, *Graphics Gems III*, pages 124–132. Academic Press, 1992.
17. 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.
18. X. Sun and Z. Chen. Spherical basis functions and uniform distribution of points on spheres. *J. Approx. Theory*, 151(2):186–207, 2008.
19. G. Wagner. On a new method for constructing good point sets on spheres. *Journal of Discrete and Computational Geometry*, 9(1):119–129, 1993.
20. 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.
21. 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.
22. A. Yershova and S. M. LaValle. Deterministic sampling methods for spheres and $SO(3)$. In *Proc. IEEE International Conference on Robotics and Automation*, 2004.
23. K. Zyczkowski and M. Kus. Random unitary matrices. *J. Phys. A*, 27(12):4235–4245, 1994.