AAAS
Research
Volume 2019, Article ID 1608396, 16 pages
https://doi.org/10.34133/2019/1608396

Research
A SCIENCE PARTNER JOURNAL

Research Article

Quantizing Euclidean Motions via Double-Coset Decomposition

Christian Wülker[1], **Sipu Ruan**[1], **and Gregory S. Chirikjian**[1,2]

[1]*Johns Hopkins University, Baltimore, MD, USA*
[2]*Department of Mechanical Engineering, National University of Singapore, Singapore*

Correspondence should be addressed to Gregory S. Chirikjian; gchirik1@jhu.edu

Received 1 December 2018; Accepted 14 July 2019; Published 15 September 2019

Copyright © 2019 Christian Wülker et al. Exclusive Licensee Science and Technology Review Publishing House. Distributed under a Creative Commons Attribution License (CC BY 4.0).

Concepts from mathematical crystallography and group theory are used here to quantize the group of rigid-body motions, resulting in a "motion alphabet" with which robot motion primitives are expressed. From these primitives it is possible to develop a dictionary of physical actions. Equipped with an alphabet of the sort developed here, intelligent actions of robots in the world can be approximated with finite sequences of characters, thereby forming the foundation of a language in which robot motion is articulated. In particular, we use the discrete handedness-preserving symmetries of macromolecular crystals (known in mathematical crystallography as Sohncke space groups) to form a coarse discretization of the space SE(3) of rigid-body motions. This discretization is made finer by subdividing using the concept of double-coset decomposition. More specifically, a very efficient, equivolumetric quantization of spatial motion can be defined using the group-theoretic concept of a double-coset decomposition of the form $\Gamma\backslash SE(3)/\Delta$, where Γ is a Sohncke space group and Δ is a finite group of rotational symmetries such as those of the icosahedron. The resulting discrete alphabet is based on a very uniform sampling of SE(3) and is a tool for describing the continuous trajectories of robots and humans. An efficient coarse-to-fine search algorithm is presented to round off any motion sampled from the continuous group of motions to the nearest element of our alphabet. It is shown that our alphabet and this efficient rounding algorithm can be used as a geometric data structure to accelerate the performance of other sampling schemes designed for desirable dispersion or discrepancy properties. Moreover, the general "signals to symbols" problem in artificial intelligence is cast in this framework for robots moving continuously in the world.

1. Introduction

The aim of this paper is to develop a "motion alphabet" with which robot motion primitives are expressed. From these primitives one can develop a dictionary of physical actions. Two main themes from the theory of Lie groups are used to construct this alphabet (or quantization) and to efficiently solve the "signals to symbols" problem in this context: (1) the decomposition of a Lie group into cosets, double cosets, and corresponding fundamental domains; (2) the possibility of constructing such fundamental domains as Voronoi or Voronoi-like cells.

Equipped with an alphabet of the sort developed in this paper and the associated algorithms for efficiently rounding off continuous motions to nearby discrete representatives, intelligent actions of robots in the world can be approximated with finite sequences of characters, thereby forming the foundation of a language in which robot motion is articulated.

At the macroscopic scale, the world can be thought of as continuous. Coarse descriptions of this continuous world are used by intelligent systems (*e.g.*, humans and computers) to classify objects, actions, and scenarios. The classical "signals to symbols" problem in artificial intelligence (AI) seeks to bin everything in the continuous world into countable classes characterized by strings of discrete symbols, such as the letters in an alphabet. In a sense, this is the inverse problem of what the genetic code does, since the finite alphabet {A, C, G, T} encodes the morphology and metabolism of every living creature that moves in the continuous world and processes information with an analog brain.

Discrete alphabets (including the Roman alphabet, the radicals that form characters in Asian languages, sign language, etc.) form the basis for all human languages [1]. All discrete characters can be reduced to a binary code, *e.g.*, ASCII or Morse code in the case of western letters. The efficiency of an alphabet (or a code) depends on how much information can be conveyed with a given number of symbols, and how difficult it is to convey those symbols. For example, since the letter "e" is the most often used symbol in the English language, it is represented by a single "·" in

Morse code. Perhaps the simplest and most widely occurring motion of a robot is the "nil motion", or identity element of the motion group, describing the "home" pose (position and orientation), which coincidentally can be denoted simply as "e". All other elements of our motion alphabet will be written as a product of the form $\gamma\delta$ where $\gamma \in \Gamma$, a crystallographic space group, and $\delta \in \Delta$, a finite group of rotational symmetries.

As an example of how discrete symbols classify the world, consider how each of the following sentences describes a class of situations in which there is continuous freedom which becomes more restricted as the number of discrete symbols increases:

(1) The cat is on the table.

(2) The black cat is sitting on the table.

(3) The black cat is sitting on the table and looking at you.

In (1), the cat could be of any color and lying down, sitting, or standing in an infinite variety of postures. The color and posture of the cat get somewhat restricted as the sentences get longer, but we still do not know how the cat is sitting, what its tail is doing, how big it is, its weight, the color of its eyes, etc. Indeed, the sentence would need to be quite long to hone in on what is actually going on, hence the old saying "one picture is worth a thousand words."

The basic problem that must be reconciled by intelligent robots is to approximate, or round off, continuous objects and actions within a discrete descriptive framework and then truncate the discrete description at some finite level given limitations on computational and sensing resources. This is true in both machine-learning (ML) frameworks such as deep learning and classical AI. A specific kind of rounding off (of Euclidean motions) is the main subject of this paper, which is about making precise the round-off statement

$$\text{discrete description} = \lceil \text{continuous description} \rceil \quad (1)$$

in the context of motions of objects and intelligent agents in the world.

A real-world problem that can be addressed using this discrete description of motions is robot motion planning, which is one of the fundamental topics in the field of intelligent robotics. It has a wide range of applications such as autonomous vehicles [2–4], mechanical parts assembly [5, 6], space on-orbit manipulations [7], and protein folding [8]. In general, motion planning seeks to answer the query: "how to plan a path that guides the robot from a given start pose to an end pose subject to some geometric or dynamical constraints". And a popular way that has been applied for decades is to build a "roadmap" [9], which is basically a graph structure consisting of valid vertices and edges. A large number of efficient algorithms have been proposed such as visibility graph [10], Voronoi graph [11], cell decomposition [12], probabilistic roadmap [13], and their variants. For a given environment and constraints (obstacles, nonholonomic constraints, etc.), a roadmap provides information with qualified vertices and connectivity among nearby vertices. And once a query is submitted, graph searching algorithms give a valid optimal sequence of motions from the roadmap. Although a roadmap method is able to answer multiple searching queries, when the environment is changing, those vertices or edges information needs to be updated. To some extent, quantizing a continuous motion of the robot from a predefined motion alphabet is closely related to the roadmap concept. However, the advantage lies in the storage and representation of the motion sequences—once a rich library of alphabets is built a priori, representing different motions in different environment subject to different constraints is just a matter of combinations and ordering of the alphabet indices.

The major contributions of this paper are listed as follows:

(1) A "signals to symbols" framework for Euclidean motions based on double-coset decomposition is proposed and its properties are analyzed.

(2) Concrete formulations of the motion alphabet are constructed through crystallographic symmetry.

(3) Fast decoding algorithms via coarse-to-fine double-coset decomposition are proposed and numerically verified.

(4) Comparisons with existing methods for discretizing the Euclidean group and decoding random motions are performed.

(5) A fast hybrid search method that incorporates some existing sampling methods for rotation group with good dispersion or discrepancy properties has been proposed and verified.

The remainder of this paper is structured as follows: In Section 2, a review of the immense literature on machine intelligence as it applies to intelligent robot action in the world is provided. This is a rapidly changing field and impossible to capture in full detail, but some classical highlights are covered. Section 3 reviews some relevant aspects of abstract group theory. This is made concrete in Section 4 which focuses on the group of rigid-body motions. Section 5 reviews crystallographic symmetry, which is a source of discrete symbols from which a motion alphabet is constructed. In Section 6, we develop a motion alphabet by "dividing up" the group SE(3) of rigid-body motions via fine double-coset decomposition based on a crystallographic Sohncke space group and the subgroup of rotational symmetries of the icosahedron. Section 7 presents other choices for motion alphabets and solves the decoding problem efficiently by introducing a coarse-to-fine search scheme. In Section 8, comparisons with existing methods for sampling rotations and Euclidean motions are performed, and it is shown how our alphabet can be used as a geometric data structure to enhance the speed of these other motion-approximation methods.

2. Literature Review

This section reviews two largely disjoint areas of the literature: (1) the interface between group theory and machine learning and AI; (2) sampling methods and measures of their quality.

2.1. Lie Groups in Machine Learning and AI. The recognition of human (and humanoid-robot) actions has been studied from many different perspectives including [14–19]. Probabilistic graphical models [20], generative models [21],

and recently "SE3-nets" [22] have been used to describe motion uncertainty in the context of learning. Works on vision and reasoning use concepts of quotient operations to mod out irrelevant information [23–25].

Group-theoretic methods (and abstract algebra more generally) can be found sprinkled throughout the AI literature [26–32]. Although group theory has long history [33–35], it is still a useful and popular tool for solving problems related to motions [36–40].

Of particular importance in the current context is the relationship between artificial intelligence and machine learning. AI arose as a branch of cybernetics, focusing on artificial aspects of reasoning and cognition, thereby leading to a redefinition of cybernetics to focus on information and control. Machine learning (and particularly deep neural networks) led by Hinton, LeCun, Bengio, and others can be viewed as an alternative to classical AI [41–45]. Recently, geometric and algebraic methods are being explored in some forms of machine learning [46–48].

The goal of this paper is to develop an alphabet of basic motions from which discrete words that capture the essence of a continuous motion/action are constructed. Throughout the literature, discretization of motions has attracted significant interest, where the Euclidean group is one of the popular ones [49]. In particular, the uniform sampling of rotations, either random or deterministic, has a wide range of applications [50–56]. This is also one of the applications that this work will address. The discretization of the continuous motions will generate part of a dictionary. This dictionary will, in the future, serve as the knowledge base for an expert system that will enable the robot to function at first use "right out of the box."

2.2. Measurements of Sample Quality: Discrepancy, Dispersion, Consistency, and Uniformity. Sampling rotations in an efficient way play important roles in several fields including computer graphics [50, 57], protein crystallography [58], molecular physics [59], materials science [60], and robot motion planning [53, 55].

Several measures of the quality of a finite sampled set of rotations have been proposed in the literature. One is *discrepancy* as defined and used in [53, 61, 62]:

$$D(P,\mathscr{R}) \coloneqq \sup_{R\in\mathscr{R}} \left| \frac{|P\cap R|}{|P|} - \frac{Vol(R)}{Vol(SO(3))} \right|, \tag{2}$$

where $\mathscr{R}$ is a collection of measurable subsets of SO(3), P is a finite set of sample points in SO(3), and $|S|$ is the number of elements in the finite set S. This concept was updated recently to include products of motion groups [63].

Natural measures of distance between rotations, or between rigid-body motions, can be computed in a variety of different ways, as reviewed in detail in [64].

Given any such metric, $\rho : \mathrm{SO}(3)\times\mathrm{SO}(3) \longrightarrow \mathbb{R}_{\geq 0}$, the *dispersion* of the points can be computed as

$$\mathfrak{D}(S,\rho) \coloneqq \max_{R\in SO(3)} \left[\min_{R_S\in S} \rho(R,R_S) \right], \tag{3}$$

where S is a set of the sampled elements in SO(3).

In addition to discrepancy and dispersion, many other measures of sample quality can be defined. For example, a sampling can be called *consistent* if the distribution of the round-off error for any random rotation is concentrated. Therefore, we define the consistency of a set of samples on SO(3) as

$$\mathfrak{C}(S,\rho) \coloneqq \sigma_{R\in S}\{m(R)\}, \tag{4}$$

where

$$m(R) \coloneqq \min_{R_S\in S-\{R\}} \rho(R,R_S) \tag{5}$$

and σ is the standard deviation of the set of distances between each sample point and its nearest neighbor. Here a low value of $\mathfrak{C}(S,\rho)$ indicates high consistency. For example, if $\mathfrak{C}(S,\rho)$ is zero, then every point in the set has nearest neighbors of the same distance. Moreover, if each point in the set has many neighbors that achieve the minimal value of distance, then the set S has a high level of *uniformity*, which can be quantified as follows. For each $R\in S$ compute the subset $A_R \subset S$ as

$$A_R \coloneqq \{Q\in S : \rho(R,Q) = m(R)\}. \tag{6}$$

Then the cardinality of this set measures the number of equally close nearest neighbors to R, and uniformity weights this by the spread of this number:

$$\mathfrak{U}(S,\rho) \coloneqq \frac{\min_{R\in S}|A_R|}{1+\max_{R\in S}|A_R| - \min_{R\in S}|A_R|}, \tag{7}$$

where a high value indicates high uniformity. The numerator reflects the number of nearest neighbors for the worst sample, and the denominator reflects the spread (with 1 included since the difference between max and min can be zero).

For example, an integer lattice in 3D Euclidean space has a high level of uniformity with regard to the Euclidean metric, with each point having six nearest neighbors, each with the same minimized value of distance, and so $\mathfrak{C}(S,\rho) = 0$ and $\mathfrak{U}(S,\rho) = 6$. A spherical close-packing can have an even higher value of $\mathfrak{U}(S,\rho)$.

3. Some Relevant Aspects of Group Theory

The alphabets constructed in this paper consist of carefully chosen elements of the group of rigid-body motions, drawn from fundamental domains of double-coset spaces. Since this terminology might not be familiar to some readers with an interest in the topic, the relevant concepts from group theory are reviewed here. The concept of a group itself is assumed to be known.

3.1. Definitions and Properties. Let a group G and a subgroup $H < G$ be given. For an element $g\in G$, the corresponding *left coset* is defined as

$$gH \coloneqq \{gh : h\in H\}. \tag{8}$$

Analogously, the respective *right coset* is defined as

$$Hg \coloneqq \{hg : h\in H\}. \tag{9}$$

A subgroup $N < G$ is called a *normal* subgroup of G if it is conjugated to itself; *i.e.*, $gNg^{-1} = N$ for all $g \in G$. If N is both a normal and a proper subgroup of G, then it is denoted as $N \lhd G$.

A group can be divided into cosets (left or right), all with the same number of elements, and the set of all these cosets is called a *coset space*. In general, the left-coset space

$$G/H \coloneqq \{gH : g \in G\} \tag{10}$$

and the right-coset space

$$H\backslash G \coloneqq \{Hg : g \in G\} \tag{11}$$

are different from each other, except when H is normal in G ($H \lhd G$), in which case $gH = Hg$ for all $g \in G$. In this special case $G/H = H\backslash G =: H/G$ is a *factor* (or *quotient*) *group*.

Given two subgroups $H, K < G$, it is also possible to define *double cosets*

$$HgK \coloneqq \{hgk : h \in H,\ k \in K\}, \quad g \in G, \tag{12}$$

and the corresponding *double-coset space*,

$$H\backslash G/K \coloneqq \{HgK : g \in G\}. \tag{13}$$

Given a left-coset decomposition with respect to a subgroup H, it is possible to define a (nonunique) *fundamental domain*

$$F_{G/H} \subset G \tag{14}$$

consisting of exactly one element per left coset. Since a group can be partitioned into disjoint cosets, it is

$$G = \bigcup_{g \in F_{G/H}} gH \tag{15}$$

and

$$G = \bigcup_{h \in H} F_{G/H} h \tag{16}$$

(and analogously in the right-coset case [65]). It is important to note that we are generally dealing with unions of disjoint sets in this paper. When $H \lhd G$, such fundamental domain is a group with respect to the original group operation modulo H, and this group is isomorphic to G/H.

Given two subgroups $H, K < G$, the corresponding double-coset decomposition is

$$G = \bigcup_{g \in F_{H\backslash G/K}} HgK \tag{17}$$

and we further have that

$$G = \bigcup_{h \in H} \bigcup_{k \in K} h F_{H\backslash G/K} k, \tag{18}$$

where $F_{H\backslash G/K}$ is a fundamental domain for the double-coset space $H\backslash G/K$ consisting of exactly one element per double coset. When G is a *Lie group* (*i.e.*, when G has the structure of a differentiable manifold, and when further the group operation and inversion are compatible with this smooth structure) and H, K are *discrete* subgroups, then such fundamental domains $F_{G/H}$ and $F_{H\backslash G/K}$ will have the same dimensionality as G, but smaller volume.

When G is a Lie group, one way to construct fundamental domains is as Voronoi-like cells: Since G is a smooth manifold, a proper distance function (metric) ρ exists, and we can define

$$\begin{aligned} &F^{\circ}_{G/H} \\ &\quad \coloneqq \{g \in G : \rho(e, g) < \rho(e, gh) \ \text{ for all } h \in H\backslash\{e\}\}, \\ &F^{\circ}_{H\backslash G} \\ &\quad \coloneqq \{g \in G : \rho(e, g) < \rho(e, hg) \ \text{ for all } h \in H\backslash\{e\}\}, \end{aligned} \tag{19}$$

where e is the identity element of G, and when $H \cap K = \{e\}$,

$$\begin{aligned} F^{\circ}_{H\backslash G/K} \coloneqq \{&g \in G : \rho(e, g) < \rho(e, hgk) \ \text{ for all } (h, k) \\ &\in H \times K\backslash\{(e, e)\}\}. \end{aligned} \tag{20}$$

Above we have defined the *interior* of fundamental domains. The union of the corresponding shifts of these open sets will not in general completely exhaust G (cf. (16) and (18)). However, a set of measure zero will at most be missing, which is not relevant for our practical purposes.

A *unimodular* Lie group G is one for which there exists an integration measure, μ, or equivalently a volume element dg, such that for every integrable function $f : G \longrightarrow \mathbb{R}$,

$$\begin{aligned} \mu(f) &\coloneqq \int_G f(g)\, dg = \int_G f(g_0 g)\, dg \\ &= \int_G f(g g_0)\, dg = \int_G f\left(g^{-1}\right) dg. \end{aligned} \tag{21}$$

In the case of a compact Lie group such as $G = SO(3)$, $\mu(1) = Vol(G)$.

In terms of ZXZ Euler angles (α, β, γ), it is well known that the volume element for SO(3) is [64]

$$dg = \sin\beta\, d\alpha\, d\beta\, d\gamma. \tag{22}$$

Consequently, sampling uniformly in Euler angles leads to clumping of samples around $\beta = 0$ and π, and under sampling near $\beta = \pi/2$.

Some authors therefore sample β nonuniformly, by making a change of coordinates as $\beta = \cos^{-1}(x)$. Then $\sin\beta = \sqrt{1 - x^2}$, and $|d\beta/dx| = 1/\sqrt{1 - x^2}$. This results in equivolumetric portioning of $SO(3)$ in the coordinates (α, x, γ) with volume element $dg = d\alpha\ dx\ d\gamma$.

However, equipartitioning into units of equal volume is not the same as equipartitioning into units of equal shape. One attempt to partition based on shape arose nearly 50 years ago in the context of protein crystallography, where Lattman [58] realized that the metric tensor for SO(3) as expressed in Euler angles (α, β, γ) could be diagonalized by changing to $(\alpha', \beta', \gamma') \coloneqq (\alpha+\gamma, \beta, \alpha-\gamma)$. This diagonalization

does not change the volume element, which remains $dg = \sin\beta'\, d\alpha'\, d\beta'\, d\gamma'$.

Group theory has been a cornerstone in all areas of the physical sciences including crystallography, classical mechanics, quantum mechanics, chemistry, and particle physics. Moreover, in classical works on finite automata theory, attempts were made to incorporate the theory of finite groups [26–28]. Many roboticists and computer vision researchers know about the special Euclidean group SE(3), which is a Lie group that describes rigid-body motions. This will be discussed later, after first reviewing the group of pure rotations.

3.2. A Concrete Example: The Rotation Group. In the following, the abstract definitions are illustrated with a concrete example. The set of 3×3 rotation matrices is called the *special orthogonal group* and is denoted as SO(3). That is,

$$\mathrm{SO}(3) \coloneqq \left\{R \in \mathbb{R}^{3\times 3} : R^{\mathrm{T}}R = \mathbb{I} \text{ and } \det R = 1\right\}, \tag{23}$$

where $\mathbb{I}$ is the identity matrix. Here the group operation is simply matrix multiplication. It is not difficult to show that $\mathbb{I}$ is the group identity, given any two $R_1, R_2 \in \mathrm{SO}(3)$ that the matrix product $R_1R_2 \in \mathrm{SO}(3)$, and that $R^{-1} = R^{\mathrm{T}}$ is the inverse of $R \in \mathrm{SO}(3)$.

Explicitly for SO(3), elements of the associated Lie algebra $\mathfrak{so}(3)$, which correspond to infinitesimal rotations, are *skew-symmetric matrices*

$$X = \begin{bmatrix} 0 & -x_3 & x_2 \\ x_3 & 0 & -x_1 \\ -x_2 & x_1 & 0 \end{bmatrix}, \tag{24}$$

and the *matrix exponential* (or *exponential*) *map* gives

$$R(\boldsymbol{x}) = \exp(X) = \mathbb{I} + \frac{\sin\|\boldsymbol{x}\|_2}{\|\boldsymbol{x}\|_2} X + \frac{1-\cos\|\boldsymbol{x}\|_2}{\|\boldsymbol{x}\|_2^2} X^2, \tag{25}$$

where $\boldsymbol{x} = (x_1, x_2, x_3)^{\mathrm{T}} =: X^{\vee}$ is the dual vector corresponding to X. The opposite operation gives $\widehat{\boldsymbol{x}} \coloneqq X$. The parameters x_1, x_2, and x_3 can be thought of as Cartesian coordinates in the Lie algebra $\mathfrak{so}(3)$, and when these coordinates are restricted to the range $\|\boldsymbol{x}\|_2 \leq \pi$, they can be used to parameterize all of SO(3) through the exponential map. When $\|\boldsymbol{x}\|_2 = \pi$ the point is at the boundary. In such a case $\boldsymbol{x}$ and $-\boldsymbol{x}$ describe the same rotation, and so one model for SO(3) is that of a solid ball of radius π in Euclidean space, with antipodal points identified as being equivalent, or "glued."

The inverse map for each is the *matrix logarithm*. This degenerates when the rotation angle $\theta \coloneqq \|\boldsymbol{x}\|_2$ is π. By restricting the discussion to the case when $\theta < \pi$, the logarithm is uniquely defined on a subset of SO(3) depleted by the set of measure zero defined by $\theta = \pi$. This depletion will have no effect on our formulation. Indeed, we can define the metric

$$\rho(R_1, R_2) \coloneqq \left\|\log^{\vee}\left(R_1^{\mathrm{T}}R_2\right)\right\|_2 \tag{26}$$

when $R_1^{\mathrm{T}}R_2$ is not a rotation by π, and otherwise $\rho(R_1, R_2) \coloneqq \pi$.

It is interesting to note that the above distance function ρ for SO(3) is *bi-invariant; i.e.*,

$$\rho(R_1, R_2) = \rho(RR_1, RR_2) = \rho(R_1R, R_2R), \quad R \in \mathrm{SO}(3). \tag{27}$$

Using this particular metric ρ it is possible to construct *Voronoi cells* (in the classical sense) in SO(3) for fundamental domains $F_{H\backslash\mathrm{SO}(3)}$ and $F_{H\backslash\mathrm{SO}(3)/K}$, because then (19) and (20) become

$$\begin{aligned} F^{\circ}_{\mathrm{SO}(3)/H} &= F^{\circ}_{H\backslash\mathrm{SO}(3)} \\ &= \left\{R \in \mathrm{SO}(3) : \rho(R, \mathbb{I}) < \rho(R, h) \text{ for all } h \in H\backslash\{\mathbb{I}\}\right\} \end{aligned} \tag{28}$$

and

$$\begin{aligned} F^{\circ}_{H\backslash\mathrm{SO}(3)/K} = \big\{&R \in \mathrm{SO}(3) : \rho(R, \mathbb{I}) \\ &< \rho(R, hk) \text{ for all } (h,k) \in H \times K\backslash\{(\mathbb{I}, \mathbb{I})\}\big\}, \end{aligned} \tag{29}$$

respectively.

Of particular interest to us are the cases where H is one of the finite groups of rotational symmetries of the Platonic solids. This is shown in Figure 1 (see also [54]). In this figure the fundamental domains $F_{H\backslash\mathrm{SO}(3)}$ are depicted in exponential coordinates in $\mathfrak{so}(3)$ (identified with a ball of radius π, as explained above). Note that this is a conceptual plot, since actually the edges and faces of these Voronoi cells are slightly bent. The number $|H|$ of elements in H is 12 for the group of tetrahedral, 24 for the group of octahedral, and 60 for the group of icosahedral rotational symmetries. By the left (or right) action of H on the respective fundamental domain, it is possible to (almost completely) cover SO(3); cf. (16).

The coset spaces resulting from quotienting the rotation group by discrete subgroups have been studied in the pure mathematics literature under the names "spherical space forms" [66] and Poincaré homology 3-spheres [67]. The geometric and topological properties of these manifolds are related to how the opposing faces of our tiles can be glued together.

If H is the group of rotational symmetry operations of the icosahedron, then $|H| = 60$ and $F_{H\backslash\mathrm{SO}(3)}$ can be viewed as a dodecahedral cell centered at the origin of the Lie algebra $\mathfrak{so}(3)$ (see Figure 1, right). If we choose the second subgroup K to be a conjugated group of the tetrahedral, octahedral, or icosahedral symmetries (*i.e.*, $K \coloneqq gPg^{-1}$, where P is the group of the rotational symmetries of the respective Platonic solid and $g \in \mathrm{SO}(3)$ is chosen such that $H \cap K = \{\mathbb{I}\}$), then the Voronoi cell $F_{H\backslash\mathrm{SO}(3)/K}$ takes a shape as exemplarily shown in Figure 2. On the other hand, if we choose $K \coloneqq H$, then $F_{H\backslash\mathrm{SO}(3)/H}$ cannot be constructed as a Voronoi cell, but it can be chosen as an irregular tetrahedron (the ruby region in Figure 3), yielding a subdivision of the dodecahedral cell $F_{H\backslash\mathrm{SO}(3)}$ by conjugation of the tetrahedron with the elements in H. Note that such conjugation has the effect of rotation in $\mathfrak{so}(3)$ since $\log^{\vee}(gRg^{-1}) = g\log^{\vee}R$. Similarly, if $K < H$, then a $|K|$-fold division of $F_{H\backslash\mathrm{SO}(3)}$ can be computed to

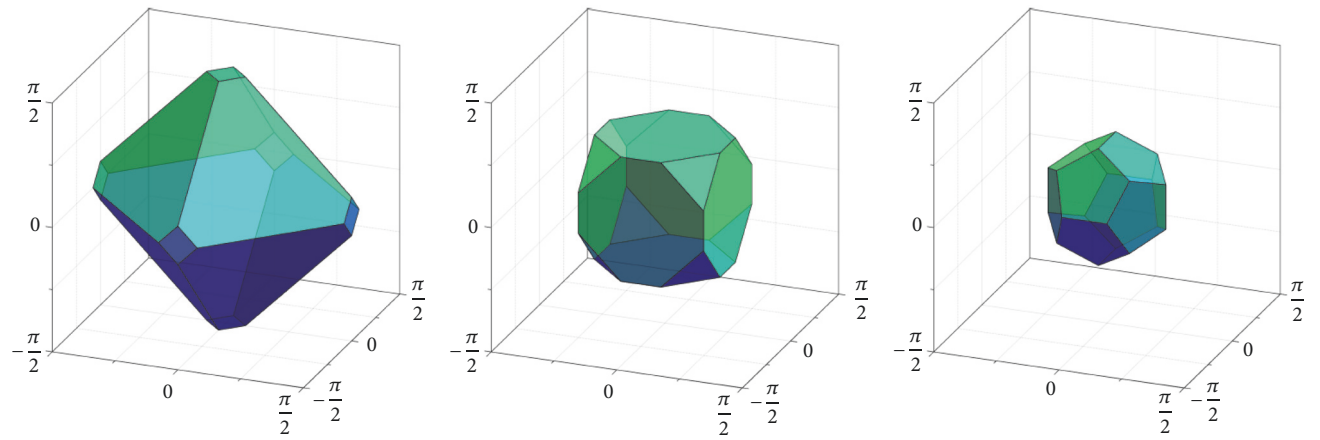


FIGURE 1: Fundamental domain $F_{H\backslash SO(3)}$ with H as the group of (from left to right) tetrahedral, octahedral, and icosahedral symmetries, constructed as Voronoi cells, viewed in exponential coordinates.

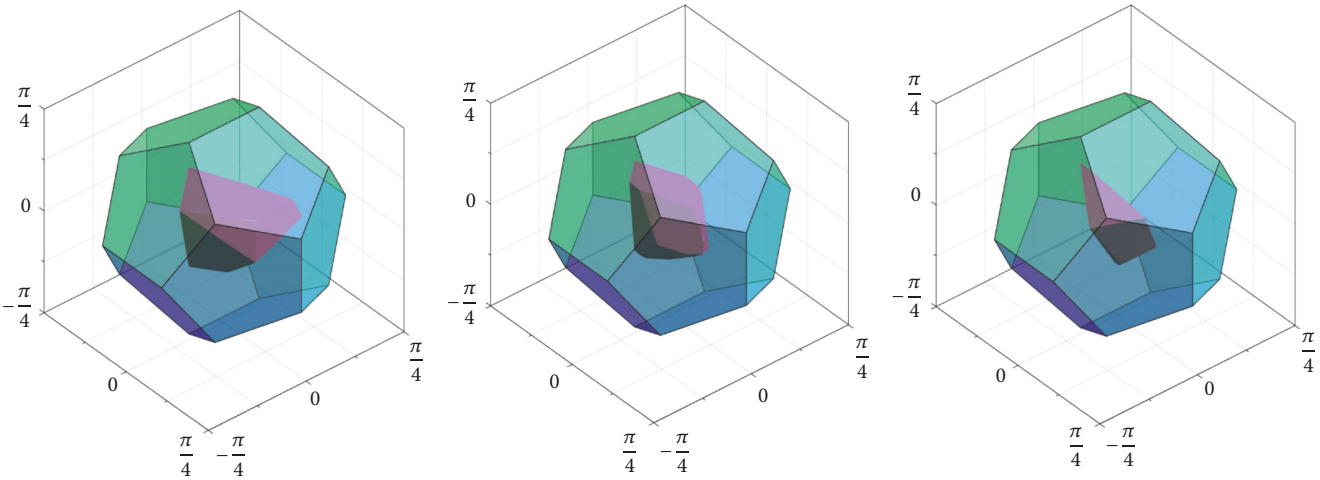


FIGURE 2: Center Voronoi cell $F_{H\backslash SO(3)}$ in coset space (emerald region) with H as the group of icosahedral symmetries, and center Voronoi cell $F_{H\backslash SO(3)/K}$ in double-coset space (ruby region) with H as the group of icosahedral symmetries for all cases and K as a conjugated group of (from left to right) tetrahedral, octahedral, and icosahedral symmetries, respectively.

represent $F_{H\backslash SO(3)/K}$, and $F_{H\backslash SO(3)}$ can be reconstructed from these pieces, again by conjugation.

When choosing H to be icosahedral and $K = H$ or $K = gHg^{-1}$, this means that we can divide SO(3) into 3600 pieces of equal size. The 3600 respective barycentric or Voronoi centers of these cells can be taken as a discretization of SO(3), and any of these centers can be written in a unique way as $R_{ij} = h_i k_j$ where $(h_i, k_j) \in H \times K$ with $i, j \in \{1, 2, \ldots, 60\}$. This means that any one of the 3600 points R_{ij} corresponds to a two-letter word (h_i, k_j).

A natural question to ask is then, for $R \in$ SO(3), how do we find the closest word (h_i, k_j) to approximate it? This decoding or "signals to symbols" problem is addressed in Section 7.1.

4. Rigid-Body Motions as a Group Used for Framing Robots

Given any rigid object or multi-rigid-body actor (such as a human, robot, self-driving car, or smart house), the behavior and use of this item involve the time evolution of its conformation and its overall position and orientation, or "pose". A common descriptive framework for both the internal (conformational) degrees of freedom and their relative pose (configuration) is to attach reference frames on each rigid component.

Let $g = (R, \boldsymbol{t})$ denote a rigid-body motion relative to a reference frame fixed in space, where $R \in$ SO(3) is a rotation matrix and $\boldsymbol{t} \in \mathbb{R}^3$ is a translation vector. The set of all such motions forms a six-dimensional Lie group, the *special Euclidean group* SE(3). This is a group because the composition operation

$$g_1 g_2 = (R_1, \boldsymbol{t}_1)(R_2, \boldsymbol{t}_2) \coloneqq (R_1 R_2, R_1 \boldsymbol{t}_2 + \boldsymbol{t}_1) \tag{30}$$

satisfies the properties of closure and associativity, the identity exists and is simply $e = (\mathbb{I}, \mathbf{0})$ (with zero translation), and the inverse of g is $g^{-1} = (R^{\mathrm{T}}, -R^{\mathrm{T}}\boldsymbol{t})$. Note that SE(3) is not

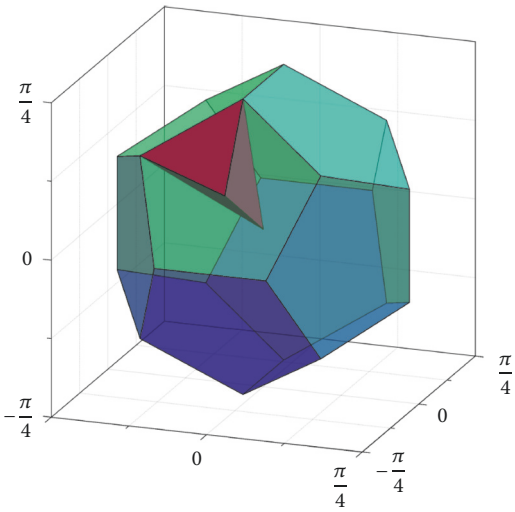


FIGURE 3: Dodecahedral cell $F_{H\backslash SO(3)}$ (emerald region) and tetrahedral wedge $F_{H\backslash SO(3)/H}$ (ruby region), with H as the group of icosahedral symmetry. The dodecahedral cell can be decomposed into 60 identical tetrahedral wedges, with five packed on each pentagonal face.

commutative (Abelian). The group operation is the same as multiplying homogeneous transformation matrices; *i.e.*,

$$H(g_1 g_2) = H(g_1) H(g_2),$$
$$\text{where } H(g) \coloneqq \begin{bmatrix} R & \boldsymbol{t} \\ \mathbf{0}^{\mathrm{T}} & 1 \end{bmatrix}. \tag{31}$$

In the case of *planar* motions, we deal with the special Euclidean group SE(2), where $R \in \mathrm{SO}(2)$ is parameterized by a single angle θ, and $\boldsymbol{t} \in \mathbb{R}^2$ has components (x, y), totalling three degrees of freedom. SO(2) is the group of rotations in $\mathbb{R}^2$, defined analogously to (23). In general, SE(d) is an example of an *outer* (or *external*) *semidirect product* which combines the two groups $(\mathbb{R}^d, +)$ and SO(d) into the new group

$$\mathrm{SE}(d) \coloneqq \mathrm{SO}(d) \ltimes \mathbb{R}^d. \tag{32}$$

The underlying set of this group is the Cartesian product $\mathrm{SO}(d) \times \mathbb{R}^d$, but the symbol $\ltimes$ reflects the fact that the group operation is not simply $(R_1, \boldsymbol{t}_1)(R_2, \boldsymbol{t}_2) = (R_1 R_2, \boldsymbol{t}_1 + \boldsymbol{t}_2)$, which is also a group (called the *direct product*), but does not reflect the way that rigid-body motions work.

The next section reviews the handedness-preserving (Sohncke) crystallographic space groups, which are discrete subgroups of SE(3) and which form an important component of the motion alphabets developed in this paper.

5. Crystallographic Groups

A *crystallographic group* is a discrete (and also cocompact) subgroup of the *Euclidean group* $\mathrm{E}(d) \coloneqq \mathrm{O}(d) \ltimes \mathbb{R}^d$, where O($d$) is the *orthogonal group* consisting of all orthogonal real-valued $d \times d$ matrices (defined as in (23) for $d = 3$, but also allowing $\det R = -1$ there). In addition to rotations, the group O(d) also contains *reflections* and *rotoreflections* (*improper rotations*). If $d = 3$, a crystallographic group is commonly called a *space group*, for $d = 2$, it is referred to as a *wallpaper group*. The literature on mathematical crystallography is immense and spans many decades. Some notable introductions include [68–71]. The relationship between torsion-free crystallographic groups (*i.e.*, Bieberbach groups; see below) and flat manifolds has been studied extensively [72–76].

Elements of a crystallographic group Γ can be expressed as pairs

$$\gamma = \left(R_\gamma, \boldsymbol{t}_\gamma + \boldsymbol{v}\left(R_\gamma\right)\right), \tag{33}$$

where $R_\gamma \in \mathbb{P}$ (a discrete *point group*, *i.e.*, a subgroup of O(d)), $\boldsymbol{t}_\gamma \in \mathbb{L}$ (a lattice in $\mathbb{R}^d$), and $\boldsymbol{v} : \mathbb{P} \longrightarrow \mathbb{R}^d$. In particular, $\boldsymbol{v}$ is the translational part of a glide-reflection or screw-displacement lattice motion. In general $\boldsymbol{v}$ will satisfy the *cocycle identities*

$$\begin{aligned} \boldsymbol{v}(\mathbb{I}) &= \mathbf{0}, \\ \boldsymbol{v}\left(R_{\gamma_1} R_{\gamma_2}\right) &= \left(R_{\gamma_1} \boldsymbol{v}\left(R_{\gamma_2}\right) + \boldsymbol{v}\left(R_{\gamma_1}\right)\right) \bmod T, \end{aligned} \tag{34}$$

where $T \coloneqq \{\mathbb{I}\} \ltimes \mathbb{L}$ is the subgroup of pure (or *primitive*) translations in Γ, which is always normal ($T \triangleleft \Gamma$). The "mod T" removes components in the sum that are in T, in analogy with, for example, $(1 + 5) \bmod 4 = 2$ in modulo-4 arithmetic.

If an element $\gamma \in \Gamma \backslash \{e\}$ is of *finite order* (*i.e.*, if there exists an $n \in \mathbb{N}$ such that $\gamma^n = e$), it is called a *torsion element*. The group Γ is called *torsion-free* (or a *Bieberbach group*) if it is free of torsion elements. This is equivalent to the property that no element $\gamma \in \Gamma$ other than the identity e has a fixed point (*i.e.*, a point $\boldsymbol{p} \in \mathbb{R}^d$ with $\gamma \boldsymbol{p} = \boldsymbol{p}$). If $\boldsymbol{v} \equiv \mathbf{0}$ in (33), then Γ can be written as the semidirect product $\Gamma = \mathbb{P} \ltimes \mathbb{L}$ and the group is called *symmorphic*. Of the 230 possible types of space groups, 73 can be decomposed in this way. These are the *symmorphic space groups*. Bieberbach groups are not symmorphic.

In the theory of crystallographic groups, a well-known isomorphism is due to the above-mentioned fact that the translation subgroup T of Γ is normal:

$$\frac{\Gamma}{T} \cong \mathbb{P}. \tag{35}$$

Moreover, it can be shown that $\mathbb{P}$ is not only discrete, but must be *finite*. In the case $d = 3$ the point group $\mathbb{P}$ will belong to one of 32 discrete *crystallographic point groups* that constitute the so-called *crystal classes*.

In the context of space groups, we can distinguish between Bieberbach (*i.e.*, torsion-free) groups as one extreme, and groups which contain only the identity and torsion elements (rotations, reflections, and improper rotations) as the other extreme, with all other space groups lying somewhat "in between." For a symmorphic space group Γ, a Bieberbach subgroup $\Gamma_{\mathrm{B}} < \Gamma$ with minimum index in Γ is the subgroup T of primitive translations; for many nonsymmorphic space groups, on the other hand, there is a Bieberbach subgroup Γ_{B} with index $[\Gamma : \Gamma_{\mathrm{B}}] < [\Gamma : T]$

allowing for a decomposition of Γ as a group product [77, p. 719]

$$\Gamma = \Gamma_B S \coloneqq \{\gamma_B s : \gamma_B \in \Gamma_B, s \in S\}, \tag{36}$$

where $S < \Gamma$ is a proper subgroup of $\mathbb{P} \ltimes \{\mathbf{0}\} < \mathrm{E}(3)$, and thus $\Gamma_B \cap S = \{e\}$. If $\Gamma_B \lhd \Gamma$, then (36) is an *inner* (or *internal*) semidirect product ($\Gamma = \Gamma_B \rtimes S$). In this case, the decomposition $\gamma = \gamma_B s$ is unique for each $\gamma \in \Gamma$. Another useful property is that each of the 65 *Sohncke space groups* (*i.e.*, handedness-preserving space groups $\Gamma < \mathrm{SE}(3)$) can be written as a product

$$\Gamma = \Gamma_B \Gamma_S \coloneqq \{\gamma_B \gamma_S : \gamma_B \in \Gamma_B,\ \gamma_S \in \Gamma_S\}, \tag{37}$$

where Γ_B and Γ_S are, respectively, Bieberbach and symmorphic subgroups (see [77, Thm. 3]). If $\Gamma_B \cap \Gamma_S = \{e\}$ and Γ_B or Γ_S is normal, then (37) is an inner semidirect product, and for every $\gamma \in \Gamma$ there exist unique $\gamma_B \in \Gamma_B$ and $\gamma_S \in \Gamma_S$ such that $\gamma = \gamma_B \gamma_S$. If both Γ_B and Γ_S are normal, then Γ is decomposed into the inner direct product $\Gamma = \Gamma_B \times \Gamma_S$.

6. A Motion Alphabet Based on a Fine Double-Coset Decomposition

The essence of language is communicating information about the continuous world using a finite number of symbols, characters, or patterns. In the same spirit, in classical AI, a fundamental goal is to convert perceptual information from the continuous world to strings of symbols drawn from a finite alphabet, so that the machinery of finite-state automata reasoning can be applied. Every action of a robot in the physical world can be described as a sequence of paths in SE(3). Such paths are continuous mathematical trajectories which can be viewed as "signals".

A fundamental problem in AI and sensory perception is the so-called "signals to symbols" problem [78–80], in which observations in the continuous world are converted to coarsified representations (*i.e.*, symbols), on which AI systems can execute logical reasoning algorithms. A natural way to discretize the space of paths is to first discretize time, thereby reducing the infinite dimensions of the path space to a finite number of dimensions, and then to replace each sampled pose on the path with a rounded-off version in a discrete set (the alphabet of maneuvers). Because SE(3) does not have dense subgroups, a clever discretization needs to be constructed.

In the three-dimensional case, one of the finest space groups is $\Gamma = \mathrm{P432}$ which has a total of 24 rotational elements corresponding to the rotational symmetry operations of a cube (cf. [69, p. 634f]). Therefore, if we want to quantize rigid-body motions with a resolution that is reasonable for real-world tasks, something finer than rounding off to the nearest element of Γ is necessary.

Building on ideas introduced in the context of protein-packing models in X-ray crystallography [55, 65, 77, 81, 82], we can augment Γ by an auxiliary discrete rotation group. For example, let Δ denote the group of rotational symmetries of the icosahedron (as a discrete subgroup of SE(3)). Δ has 60 elements, and there is no lattice of discrete translations that corresponds to it. In particular, $\Gamma, \Delta < \mathrm{SE}(3)$ and $\Gamma \cap \Delta = \{e\}$, the identity element. This means that the double-coset space $\Gamma\backslash\mathrm{SE}(3)/\Delta$ is a compact Riemannian manifold. It is possible to define a compact fundamental domain $F_{\Gamma\backslash\mathrm{SE}(3)/\Delta} \subset \mathrm{SE}(3)$ using (20). Restricting the quotient map from SE(3) to $\Gamma\backslash\mathrm{SE}(3)/\Delta$ to the fundamental domain $F_{\Gamma\backslash\mathrm{SE}(3)/\Delta}$ then gives a bijective (*i.e.*, one-to-one) mapping from $F_{\Gamma\backslash\mathrm{SE}(3)/\Delta}$ to $\Gamma\backslash\mathrm{SE}(3)/\Delta$. Moreover, the action of Γ (from the left) and Δ (from the right) gives a way to tile SE(3) with disjoint shifts of $F_{\Gamma\backslash\mathrm{SE}(3)/\Delta}$, because (cf. (18))

$$\mathrm{SE}(3) = \bigcup_{\gamma\in\Gamma} \bigcup_{\delta\in\Delta} \gamma F_{\Gamma\backslash\mathrm{SE}(3)/\Delta} \delta. \tag{38}$$

Intensive study of the fundamental domains $F_{\Gamma\backslash\mathrm{SE}(3)}$ and $F_{\Gamma\backslash\mathrm{SE}(3)/\Delta}$ has been conducted (*ibid.*).

When the fundamental domain $F_{\Gamma\backslash\mathrm{SE}(3)/\Delta}$ is constructed using (20), it has the identity element e at its center, and so the tiling in (38) has the effect of sampling each center point by moving from e to $\gamma\delta$ where $\gamma \in \Gamma$ and $\delta \in \Delta$. In other words, the product $\Gamma \times \Delta < \mathrm{SE}(3)^2$ can be used as a quantized version of SE(3). The number of rotational elements will be $24 \times 60 = 1440$ which is sufficiently fine to capture the essence of any frame along a trajectory during a robot task. The alphabet defined by $\Gamma \times \Delta$ is infinite, but by limiting the extent of translations to be contained in a bounded region, it becomes finite. This means that continuous trajectories can be translated into a finite string of alphabet characters (Figure 4). This opens up the possibility of mapping these quantized trajectories into words expressed in a natural language.

An important advantage of the quantization scheme (38) is that the shifted fundamental domains $\gamma F_{\Gamma\backslash\mathrm{SE}(3)/\Delta}\delta$ will all have the same volume; *i.e.*,

$$\mu\left(\gamma F_{\Gamma\backslash\mathrm{SE}(3)/\Delta}\delta\right) = \mu\left(F_{\Gamma\backslash\mathrm{SE}(3)/\Delta}\right) \quad \text{for all } (\gamma, \delta) \in \Gamma \times \Delta, \tag{39}$$

where μ is the (left- and right-invariant) Haar measure on the (unimodular) Lie group SE(3). This means that the centers $\gamma\delta$ of these shifted fundamental domains used for quantization at the same time also allow for a very uniform sampling of the group SE(3).

7. More Alphabets and Coarse-to-Fine Decoding Algorithms

As explained above, after wisely designing a pair of discrete subgroups $\Gamma, \Delta < \mathrm{SE}(3)$ with $\Gamma \cap \Delta = \{e\}$, we can construct a fundamental domain as in (20) and decompose SE(3) as in (38). There are many other ways to choose $F_{\Gamma\backslash\mathrm{SE}(3)}$ and $F_{\Gamma\backslash\mathrm{SE}(3)/\Delta}$, as explained in [65]. A particularly simple choice is the Cartesian product

$$F_{\Gamma\backslash\mathrm{SE}(3)/\Delta} \coloneqq F_{\mathbb{P}\backslash\mathrm{SO}(3)/\Delta} \times F_{\mathbb{L}\backslash\mathbb{R}^3}, \tag{40}$$

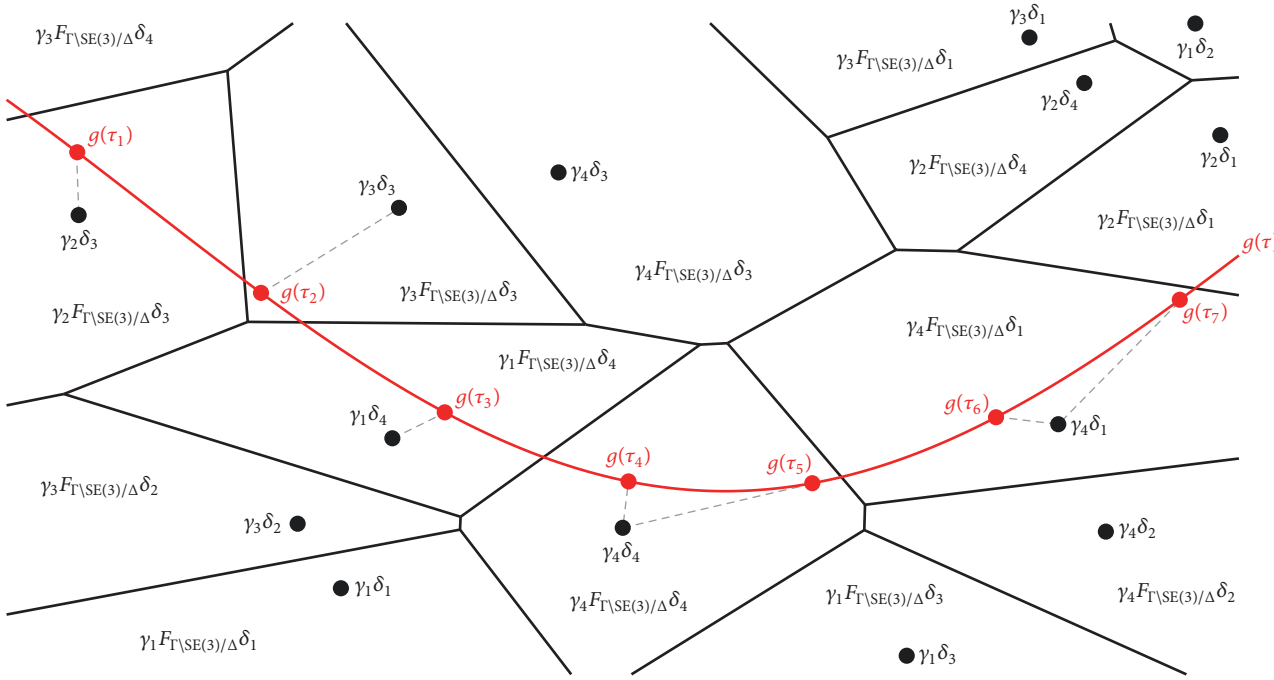


FIGURE 4: Discretizing a continuous motion trajectory g at times $\tau_1, \ldots, \tau_7$ using the alphabet $\Gamma \times \Delta$ (conceptual plot). After discretization the continuous motion can be expressed as the sentence $(\gamma_2, \delta_3), (\gamma_3, \delta_3), (\gamma_1, \delta_4), (\gamma_4, \delta_4), (\gamma_4, \delta_4), (\gamma_4, \delta_1), (\gamma_4, \delta_1)$.

where $\mathbb{P} \cong \Gamma/T$ is the point group of Γ and $\mathbb{L}$ the lattice of primitive translations. Alternatively, given decomposition (36), another natural choice is

$$F_{\Gamma\backslash SE(3)/\Delta} \coloneqq F_{S\backslash SO(3)/\Delta} \times F_{\Gamma_B\backslash\mathbb{R}^3}. \tag{41}$$

Here $S\backslash SO(3)/\Delta$ and $\Gamma_B\backslash\mathbb{R}^3$ are not (double-)coset spaces—because $S \not\leq SO(3)$ and $\Gamma_B \not\leq \mathbb{R}^3$—but are *orbit spaces* consisting, respectively, of *orbits* $SR\Delta \coloneqq \{sR\delta : s \in S, \delta \in \Delta\}$ and $\Gamma_B \boldsymbol{x} \coloneqq \{\gamma_B \boldsymbol{x} : \gamma_B \in \Gamma_B\}$ ($R \in SO(3)$, $\boldsymbol{x} \in \mathbb{R}^3$). The fundamental domains $F_{S\backslash SO(3)/\Delta}$ and $F_{\Gamma_B\backslash\mathbb{R}^3}$ above can be constructed by choosing exactly one point per orbit, which is consistent with the definitions in Section 3. Different fundamental domains $F_{\Gamma\backslash SE(3)/\Delta}$ such as those above can be used to express different quantizations via (38).

7.1. The Purely Rotational Case. The choice for $F_{\Gamma\backslash SE(3)/\Delta}$ in (40) allows us to bootstrap off of the fundamental domains for double-coset spaces for SO(3) discussed earlier. In fact, we can go even further and describe an SE(3) motion trajectory $g(\tau) = (R(\tau), \boldsymbol{t}(\tau))$ in $SO(3) \times \mathbb{R}^3$ (as a direct product rather than a semidirect product). This is not merely to make things easier—viewing pose change trajectories in this way has some advantages, as described in [40], where the direct product $SO(3) \times \mathbb{R}^3$ is called the *pose change group* and is denoted as PCG(3). Therefore, below we describe in some detail how the "signals to symbols" problem can be solved efficiently in the purely rotational case.

Since (40) is a set rather than a group, we can view it as a subset of either SE(3) or PCG(3). Either way, the general decoding problem reduces to this: Given $H, K < SO(3)$, and $R \in SO(3)$, how can we efficiently find the unique pair $(h_i, k_j) \in H \times K$ such that

$$R = h_i Q k_j \tag{42}$$

with $Q \in F_{H\backslash SO(3)/K}$? In particular, if the Voronoi choice is made for $F_{H\backslash SO(3)/K}$, solving (42) allows for simply rounding off R to $h_i k_j$, as indicated in Section 3.2.

The question then becomes how to do this. With the crystallographic constraint, in SE(3) it is possible to define H such that $|H| = 24$ (octahedral symmetry) and $|K| = 60$ (icosahedral symmetry), leading to $24 \times 60 = 1440$ combinations. In PCG(3), on the other hand, subgroups need not be restricted to the crystallographic constraint, and we can have more rotational elements (see below). The question is, is there a better way to test for h_i and k_j than two nested for loops over i and j resulting in a large number of evaluations to find where $\rho(R, h_i k_j)$ is minimized, which is equivalent to solving (42) when the Voronoi choice is made for $F_{H\backslash SO(3)/K}$?

The answer is positive, and we shall now present a technique to achieve this. Consider the double-coset space $H\backslash SO(3)/K$, where H is the group of rotational symmetries of the icosahedron and $K = gHg^{\mathrm{T}}$ is a conjugated group, with g being chosen so that $H \cap K = \{\mathbb{I}\}$. It is thus $|H\backslash SO(3)/K| = |H| \times |K| = 60^2 = 3600$. We can construct a fundamental domain $F_{H\backslash SO(3)}$ for the coset space $H\backslash SO(3)$ as a dodecahedral Voronoi cell (cf. (28) and Figure 1, right). Due to the Voronoi property, we can find the shifted tile $h_i F_{H\backslash SO(3)}$ containing the rotation R of interest by computing the distance $\rho(R, h_i)$ of R to the 60 tile centers $h_i \in H$. We then know that $h_i^{\mathrm{T}} R$ lies in the identity-centered tile $F_{H\backslash SO(3)}$. We can construct the fundamental domain $F_{H\backslash SO(3)/K}$ for the

double-coset space as a Voronoi cell, too (cf. (29) and Figure 2, right). Since the shifts $h_i F_{H\backslash \mathrm{SO}(3)/K} k_j$ of this identity-centered fundamental domain will cover SO(3), they will also cover $F_{H\backslash \mathrm{SO}(3)}$. However, the number of required shifts will be much smaller than 3600. Indeed, we found that approximately 180 shifted fundamental domains $h_i F_{H\backslash \mathrm{SO}(3)/K} k_j$ are sufficient when the conjugation element g is empirically chosen, for example, to minimize the extent

$$\sup_{R \in F_{H\backslash \mathrm{SO}(3)/K}} \rho(R, \mathbb{I}) \tag{43}$$

of $F_{H\backslash \mathrm{SO}(3)/K}$. We chose g in this particular way so as to minimize the round-off error in the quantization scheme. We can now quickly find the shifted fundamental domain $h_{i'} F_{H\backslash \mathrm{SO}(3)/K} k_j$ containing the above $h_i^{\mathrm{T}} R$ by exploiting the Voronoi property of these shifted domains. We then have that $R \in (h_i h_{i'}) F_{H\backslash \mathrm{SO}(3)/K} k_j$, and so the decomposition (42) is found. Instead of the 3600 distances on SO(3), in the above technique, we have to compute only about $60 + 180 = 240$ distances, resulting in a potential speedup of approximately $3600/240 = 15$. We further exploited the fact that $\rho(R_1, R_2) = \arccos((1/2)[\mathrm{tr}(R_1^{\mathrm{T}} R_2) - 1])$, the trace wherein can be computed efficiently as

$$\mathrm{tr}\left(R_1^{\mathrm{T}} R_2\right) = \sum_{i=1}^{3} \sum_{j=1}^{3} [R_1]_{ij} [R_2]_{ij}, \tag{44}$$

as well as the fact that $\arccos x < \arccos y$ if and only if $x > y$.

An even more efficient decoding algorithm can be obtained when considering the double-coset space $H\backslash \mathrm{SO}(3)/H$, where H is the group of rotational icosahedral symmetry. Here again we can use the dodecahedral Voronoi fundamental domain $F_{H\backslash \mathrm{SO}(3)}$ for the coset space $H\backslash \mathrm{SO}(3)$ and find the shifted tile $h_i F_{H\backslash \mathrm{SO}(3)}$ containing the rotation R easily as described above. As a fundamental domain $F_{H\backslash \mathrm{SO}(3)/H}$ for the double-coset space, we can choose the tetrahedral wedge shown in Figure 3. We can find the conjugated wedge $h_j F_{H\backslash \mathrm{SO}(3)/H} h_j^{\mathrm{T}}$ containing the "pulled-back" rotation $h_i^{\mathrm{T}} R$ by using one of the standard query methods. We then know that $R \in (h_i h_j) F_{H\backslash \mathrm{SO}(3)/H} h_j^{\mathrm{T}}$, and so (42) is solved.

We note that as is the case in Section 6 (cf. (39)), the shifts of the fundamental domains $F_{H\backslash \mathrm{SO}(3)/K}$ and $F_{H\backslash \mathrm{SO}(3)/H}$ above all have the same volume, which is an important advantage of the double-coset approach presented in this paper.

7.2. Planar-Motion Alphabets Based on Wallpaper Groups. As mentioned in Section 5, crystallographic groups are called wallpaper groups in the two-dimensional setting. Figure 5 shows fundamental domains $F_{\mathrm{p}_i\backslash \mathrm{SE}(2)}$ constructed using (19) for instances of the well-known wallpaper groups (cf. [69, Chap. 6]) p_1, p_2, p_4, p_3, and p_6 (see also [55]), all of which are symmorphic. Here SE(2) is identified with $\mathbb{R}^2 \times (-\pi, \pi) \subset \mathbb{R}^3$, with the x and y axes representing translations in x and y direction and the z axis representing the rotation angle θ. These fundamental domains are generated using the Euclidean metric $\| \cdot - \cdot \|_2$ on $\mathbb{R}^3$, adapted so as to take into account the 2π-periodicity in the rotation angle θ. It is important to note that this metric is left- but not right-invariant (there is no bi-invariant metric on SE(2)). Therefore, the fundamental domains shown in Figure 5 are actually Voronoi rather than Voronoi-like cells (as is the case for SO(3), cf. (28)).

The group p_1 consists solely of translations, constituting a parallelogrammatic lattice in the translational x-y plane. This results in a box with (irregular) hexagonal shape in the x-y plane and height 2π as the fundamental domain $F_{\mathrm{p}_1\backslash \mathrm{SE}(2)}$. In addition to the translations in p_1, the wallpaper group p_2 also contains a rotation of order two (*i.e.*, with angle π). Therefore, the fundamental domain $F_{\mathrm{p}_2\backslash \mathrm{SE}(2)}$ also has a hexagonal shape in the translational plane, but the height is only π (from $-\pi/2$ to $\pi/2$) instead of 2π. The group p_4 is a group with rotations of order four (*i.e.*, with angles $\pi/2$, π, and $3\pi/2$), as well as translations in a square lattice. Thus the fundamental domain $F_{\mathrm{p}_4\backslash \mathrm{SE}(2)}$ has the shape of a square in the translational plane, with its height being $\pi/2$ (from $-\pi/4$ to $\pi/4$). The groups p_3 and p_6 both have a hexagonal translation lattice. In addition to these translations, p_3 contains rotations of order three (rotation angles $2\pi/3$ and $4\pi/3$), while p_6 contains rotations of order six (rotation angles $\pi/3$, $2\pi/3$, π, $4\pi/3$, and $5\pi/3$). Both fundamental domains $F_{\mathrm{p}_3\backslash \mathrm{SE}(2)}$ and $F_{\mathrm{p}_6\backslash \mathrm{SE}(2)}$ have a regular hexagonal shape in the translational plane, with a height of $2\pi/3$ (from $-\pi/3$ to $\pi/3$) and $\pi/3$ (from $-\pi/6$ to $\pi/6$), respectively.

As an example of a planar-motion alphabet similar to the one described in Section 6, let us consider the double-coset space $\Gamma\backslash \mathrm{SE}(2)/\Delta$ with $\Gamma := \mathrm{p}_4$ and $\Delta := \mathrm{C}_{2n-1} \ltimes \{\mathbf{0}\} < \mathrm{SE}(2)$, where

$$\mathrm{C}_{2n-1} := \left\{ \begin{bmatrix} \cos\theta_j & -\sin\theta_j \\ \sin\theta_j & \cos\theta_j \end{bmatrix} : \theta_j = \frac{2\pi j}{2n-1},\ j = 0, \ldots, 2(n-1) \right\} < \mathrm{SO}(2) \tag{45}$$

is the group of rotations of order $2n-1$ ($n \in \mathbb{N}$). It is $\Gamma \cap \Delta = \{e\}$ and so we can construct a fundamental domain $F_{\Gamma\backslash \mathrm{SE}(2)/\Delta}$ as a Voronoi-like cell using (20). This is shown in Figure 6 for $n = 3$. In fact, because the left-invariant metric used here is also invariant under purely rotational actions from the right, the fundamental domain $F_{\Gamma\backslash \mathrm{SE}(2)/\Delta}$ is a classical Voronoi cell here, too (again as in the case of SO(3), cf. (29)). Analogously as in Section 6, we can use the alphabet $\Gamma \times \Delta < \mathrm{SE}(2)^2$ for an equivolumetric quantization of SE(2), which at the same time allows for a very uniform sampling of the group. Because the scaling of the translational lattice in the wallpaper groups is arbitrary, we can make the alphabet $\Gamma \times \Delta$ arbitrarily fine by reducing the translational scaling in p_4 and increasing the parameter n above.

To illustrate the use of the planar-motion alphabets constructed above, let us consider the SE(2) trajectory $g(\tau) := (R(\tau), \boldsymbol{t}(\tau))$, $\tau \in [0, 2\pi)$, where $R(\tau)$ is a rotation by an angle of τ and

$$\boldsymbol{t}(\tau) := \left[4\cos\tau, 6\left(\frac{\tau}{2\pi}\right) - 3\right]^{\mathrm{T}}. \tag{46}$$

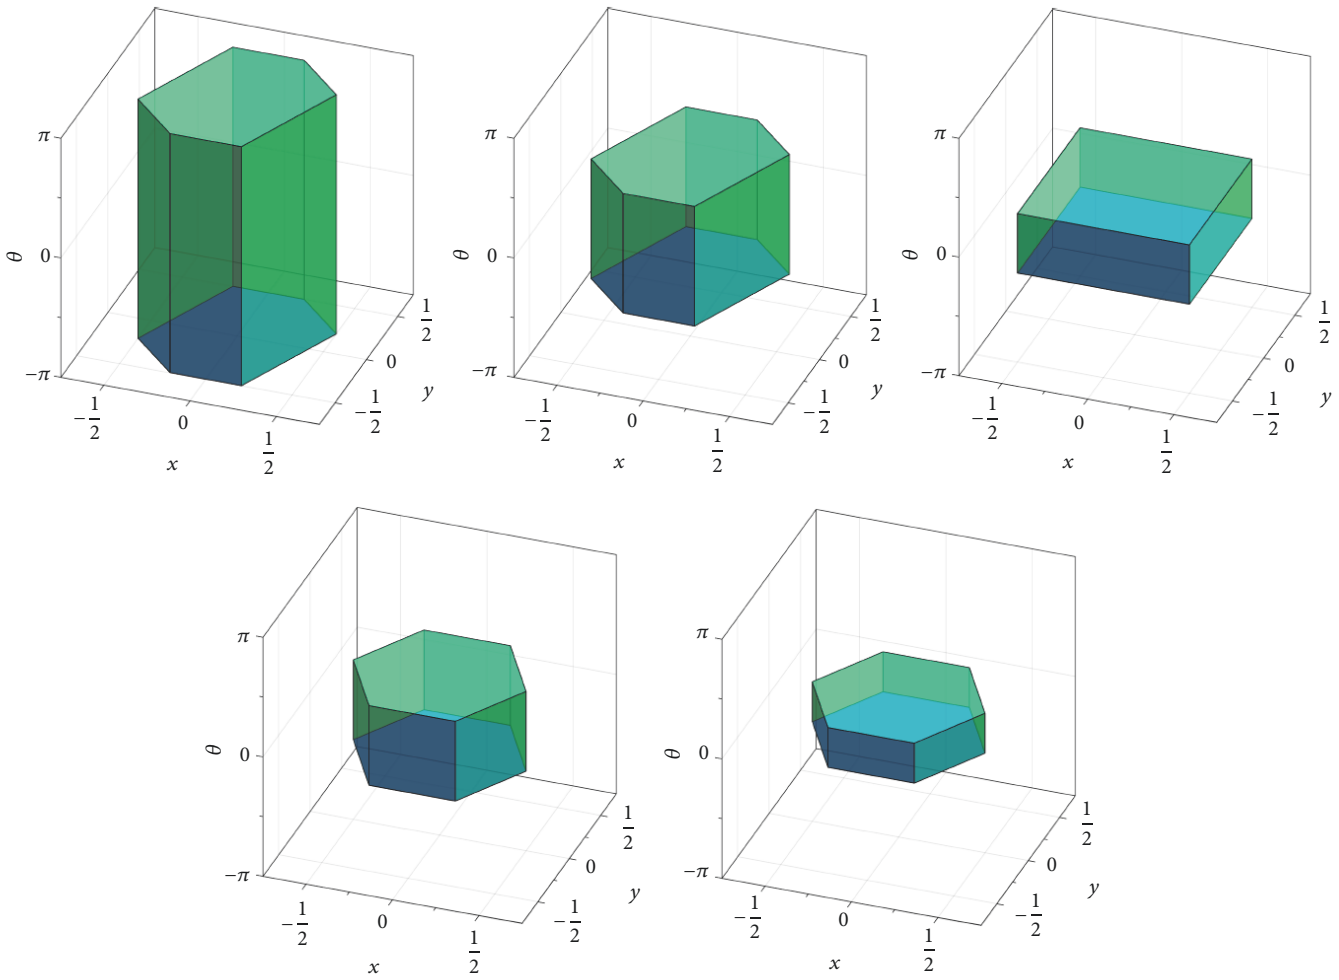


FIGURE 5: Center Voronoi cell $F_{p_i\backslash SE(2)}$ for certain instances of the wallpaper groups (left to right, top to bottom) p_1, p_2, p_4, p_3, and p_6.

We may discretize g at the five equidistant time points $\pi(1 + 2k/5)$, $k = -2, \ldots, 2$. As a motion alphabet for SE(2), let us use $\Gamma \times \Delta = p_4 \times C_5$. We can denote the elements of C_5 by δ_j with the index j as in (45). Let us denote the elements of p_4 as γ_{lmn}, where $m \in \mathbb{Z}$ and $n \in \mathbb{Z}$ denote the translation in x and y direction, respectively, while $l \in \{0, 1, 2, 3\}$ indicates a rotation by an angle of $l\pi/2$. The continuous motion trajectory g can now be expressed as the sentence $(\gamma_{2,3,-2}, \delta_3)$, $(\gamma_{2,-1,-1}, \delta_4)$, $(\gamma_{2,-4,0}, \delta_0)$, $(\gamma_{2,-1,1}, \delta_1)$, $(\gamma_{2,3,2}, \delta_2)$.

To close this section, we shall discuss how such decoding problem can be solved efficiently. We can use the same coarse-to-fine search scheme that we used in the purely rotational case in the previous section: In a first step, for a given element $g \in SE(2)$, we find $\gamma \in p_4$ such that $g \in \gamma F_{\Gamma\backslash SE(2)}$. This step is particularly easy in the case of p_4 when compared with, *e.g.*, the group p_1 (with anisotropic translational lattice). In fact, as implied by Figure 7, in the case of p_4 the above step can be realized by appropriately rounding off (in the decimal sense) the translational components of g, as well as the rotation angle. In the case of p_1, on the other hand, we would have to compute the distances of g to the Voronoi centers. In a second step, we search for the shifted fundamental domain $\gamma' F_{\Gamma\backslash SE(2)/\Delta}\delta$ containing the pulled-back element $\gamma^{-1}g$ by a purely rotational search. The decomposition of g then reads $(\gamma\gamma')Q\delta$ with $Q = (\gamma\gamma')^{-1}g\delta^{-1} \in F_{\Gamma\backslash SE(2)/\Delta}$. Of course it is also possible to first treat the translational part of g and then, say, the rotational part. With appropriate modifications, the above coarse-to-fine decoding scheme can also be used with the fine alphabet for SE(3) developed in Section 6.

8. Comparisons and Applications

To clearly demonstrate the advantageous potential of our proposed discretization and decoding algorithms, we provide comparisons of performance with existing methods and introduce a hybrid sampling method to take advantage of the speed and low dispersion properties.

8.1. The Accuracy and Speed of Rounding Off Motions

8.1.1. Uniformity of Sampling on SO(3). Discretization on SO(3) is an important application of this work, which gives equivolumetric decomposition of the group in the sense

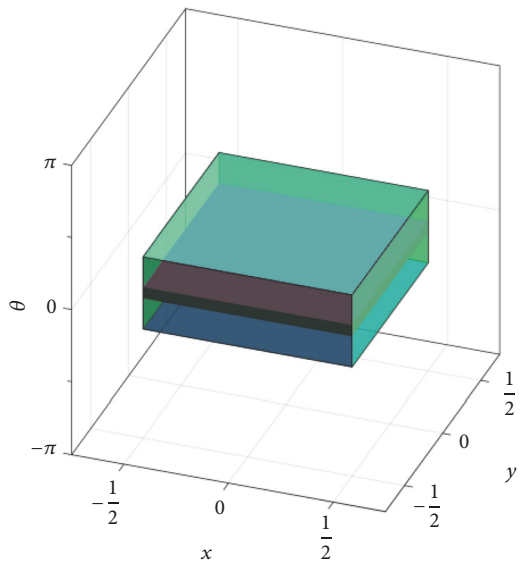


FIGURE 6: Center Voronoi cell $F_{\Gamma\backslash SE(2)}$ in single-coset space (emerald region) based on an instance of the wallpaper group $\Gamma = p_4$, and center Voronoi cell $F_{\Gamma\backslash SE(2)/\Delta}$ in double-coset space (ruby region) with $\Delta = C_5 \ltimes \{\mathbf{0}\}$ as the group of rotations of order five.

that any rotation is located inside an identical Voronoi cell. The uniformity is essential to evaluate how good a sampling method is, and can be measured by *dispersion*, *discrepancy*, *consistency*, etc., as introduced in Section 2.2. For the computations of dispersion and consistency in this work, $\rho(R, R_s) := \|\log^\vee(R^T R_S)\|_2$ is the distance metric defined in Section 3.2.

We compare the dispersion (3), consistency (4), and uniformity (7) of our method in sampling from SO(3) with uniform Euler angles and uniform sampling using Hopf fibration [53]. For the dispersion comparison, we randomly generate 10000 rotations and compute the distance with the nearest sample. The maximum value of these 10000 resulting distances approximates the dispersion. For the consistency, on the other hand, for each sample, we compute the distance to its nearest sample and take the standard deviation. And for the uniformity measure, we compute the number of nearest neighbors of each sample.

For our method, we decompose SO(3) using the double-coset space $H\backslash SO(3)/K$, where H is the group of rotational symmetries of the icosahedron and $K = gHg^{-1}$, where g is chosen such that

$$\log^\vee g := [0.4359, -0.07692, -0.1282]^T. \quad (47)$$

We found that the 181 shifted fundamental domains $h_i F_{H\backslash SO(3)/K} k_j$ with center $h_i k_j$ closest to the identity are sufficient to cover $F_{H\backslash SO(3)}$. And for the Euler angle, we choose ZYZ parameterization and uniformly generate α and γ within their range $[-\pi, \pi]$ and β such that $\cos\beta \in [-1, 1]$.

In terms of the dispersion, our method is higher than the other two methods: the dispersion of ours is 0.4291, Hopf fibration method is 0.2690, and Euler angle method is 0.3401. However, for the consistency, ours can achieve 0 deviation, while Hopf is 0.0264 and Euler angle is 0.0865. This consistency result is a significant advantage of our method in the sense that the sampling grid always has equi-distance edges. Also since the number of nearest neighbors with minimum distance for our method is always 2, the uniformity is 2. This also outperforms the other two methods, where Hopf fibration has a uniformity of 0.25, and Euler angle is only 0.0303. The results show that our method can be used in uniformly sampling rotations from SO(3).

8.1.2. Computational Time of SO(3) Nearest Neighbor Search. Another key factor for performance evaluation is the running time when searching for the nearest sample for a random rotation. A common and efficient way is using the Euler angle parameterization. Suppose we are given a set of sample points constructed using Euler angles (either with or without $\cos^{-1}$ sampling, or Lattman's diagonalization of the metric tensor, as discussed in the Introduction). Then given an arbitrary rotation, R, one can compute its Euler angles $(\alpha_R, \beta_R, \gamma_R)$ and attempt to round off to the nearest Euler angles in the sample set as $([\alpha_R], [\beta_R], [\gamma_R])$. This is simple and fast to compute, but because Euler angles are not an equi-metric spacing, the resulting rounded rotation matrix $[R] = R_3([\alpha_R])R_1([\beta_R])R_3([\gamma_R])$ may not be the closest of the sampled rotations to R. In contrast, our method is both fast and accurate in the sense of metric round-off.

To illustrate this, we perform a comparison here at an Intel Core i7-4790 CPU at 3.60 GHz × 8 and with Matlab R2018b. For our method, we use the double-coset space $H\backslash SO(3)/H$, where we have a faster version of the decoding algorithm. Comparisons are performed with the Euler angle searching method (described above). 1000 random rotations are generated and localized to the nearest sample from the sampling list. The accuracy of computing the nearest neighbor is evaluated using the brute-force nearest neighbor search (*i.e.*, minimization of the distance $\rho(R, h_i h_j)$ with respect to both i and j).

The result shows that the proposed decoding algorithm yields an average runtime for each search at around 55.9μs in our method, while the Euler angle method runs 53.6μs per rotation. Both methods are at the same level of efficiency, but ours outperforms in terms of round-off accuracy. Figure 8 shows the minimum distance between the queried rotations (50 of all the testing rotations are shown for a clearer plot) to the set of samples. Ours can always find the true nearest neighbor, while Euler angle sometimes returns the sample which is not the nearest to the queried rotation.

8.2. Combining the Benefits of Speed and Good Dispersion Properties. Sampling methods such as those based on the Hopf fibration were designed for good performance in terms of minimal dispersion and consequently outperform both Euler angles and our sampling scheme in terms of minimizing dispersion. Ours was designed for rapid query. However, if one wants both, it is possible to simply combine them in a natural way as follows:

(1) Partition any given set of sample rotations with desirable properties (such as dispersion and discrepancy) by determining in which shard each sample point belongs.

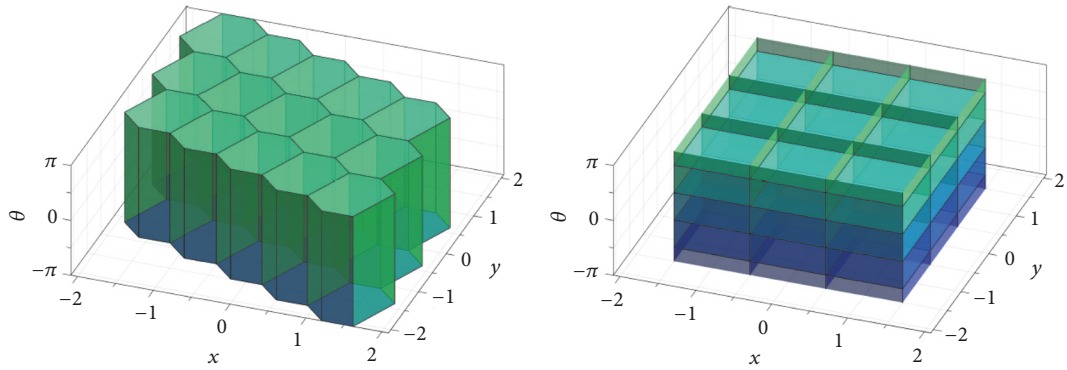


FIGURE 7: Tessellation of SE(2) illustrated in $\mathbb{R}^2 \times (-\pi, \pi)$, based on the fundamental domain (left) $F_{P1\backslash SE(2)}$ and (right) $F_{P4\backslash SE(2)}$ (conceptual plot).

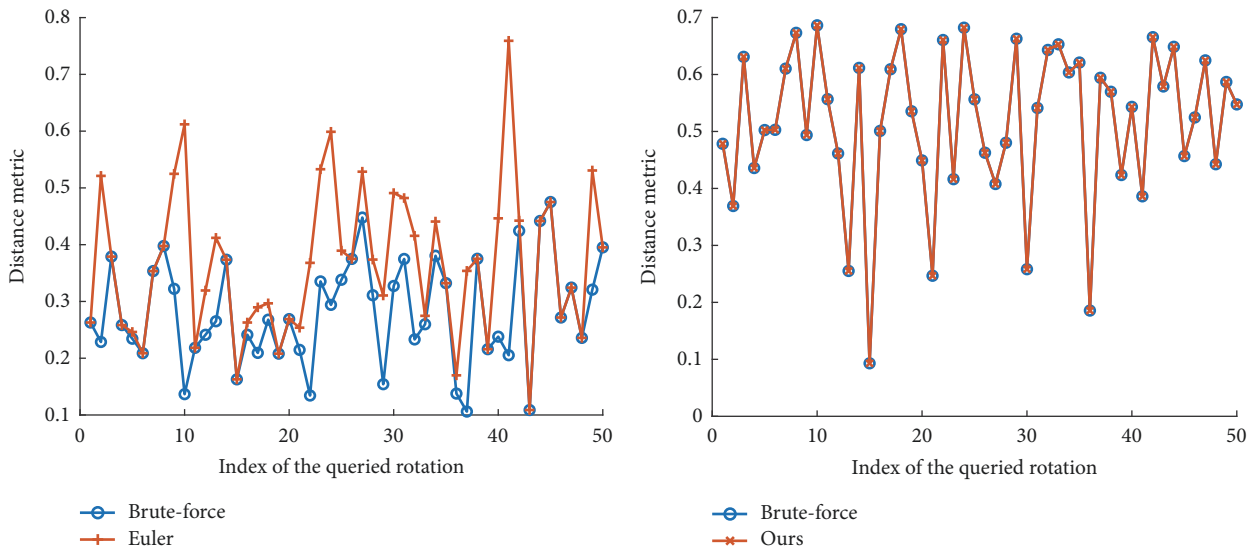


FIGURE 8: Comparisons of the minimum distance between the queried rotations to the set of samples. The true values are computed using the brute-force nearest neighbor search, which is shown in blue curve. The Euler angle search (left figure) sometimes returns higher values of distance, but ours (right figure) can always give the correct answer.

(2) Given an arbitrary rotation, determine to which double-coset fundamental domain it belongs.

(3) Compute the distance between the arbitrary given rotation and all sample points in the same shard, and those in the nearest surrounding shards.

In this scheme the number of sample points could be in the same order or even much higher than the number of shards.

We perform a numerical experiment to verify this proposed hybrid searching method using the $H\backslash SO(3)/H$ decomposition. As a preprocessing step, we first compute the nearest neighbors for each of the 60 elements in H, and we find that each rotation is surrounded by 12 neighbors. This step is to construct a connectivity map for the single-coset space or, in other words, for each icosahedron cell. Then we uniformly sample 10000 rotations using Hopf fibration, which is known to have low dispersion and discrepancy. For each sample, we decompose the rotation group via our proposed fast decoding algorithm and precompute a list of index pair, which locates the cell and shard of those samples. Afterwards, we randomly generate 1000 rotations and the goal is to find their closest sampling rotations. For each of the random rotations, we decompose it, locate the cell (determined by the first index), and then calculate the minimum distance with all the samples located in the same cell and the 12 neighboring cells. By using this hybrid method, the running time is around 5 times faster than the brute-force searching method, and the resulting minimum distance is verified to be 100% correct.

Another experiment has also been performed using the $H\backslash SO(3)/K$ decomposition. The difference with the previous test is the preprocess, where here we precompute the nearest neighbors for the 3600 elements, *i.e.*, $h_i k_j \in H \times K$ where $i, j = 1, 2, \ldots, 60$. This is equivalent to finding the neighbors of each shard in the double-coset space. The same sampling

and testing sets are input into the hybrid algorithm and we achieve a speedup of around 20 times than the brute-force method. Also, the resulting minimum distance is verified as 100% accurate.

9. Conclusion

Robot tasks involve continuous motions in space. The quantization of these motions by introducing a class of motion alphabets has been established in this paper. With such an alphabet, continuous motion trajectories can be captured with finite words/sentences. It was demonstrated in some examples how the possibility of constructing fundamental domains for coset and double-coset spaces as Voronoi or Voronoi-like cells can be used to solve this decoding or "signals to symbols" problem efficiently via a coarse-to-fine search scheme. The performance, such as uniformity, of the proposed group discretization method and the fast decoding algorithms are compared with other existing methods. The alphabets developed here will be used in the future to facilitate the connection between advances in artificial intelligence (such as the use of artificial neural networks) and physical robots acting in the world.

Disclosure

The findings and opinions expressed here are only those of the authors, and not of the funding agencies.

Conflicts of Interest

The authors declare that there are no conflicts of interest regarding the publication of this article.

Acknowledgments

This work was performed under Office of Naval Research Award N00014-17-1-2142 and National Science Foundation grant CCF-1640970. The authors gratefully acknowledge the supporting agencies.

References

[1] M. Studdert-Kennedy, "How did language go discrete," *Evolutionary Prerequisites for Language*, 2005.

[2] A. Sgorbissa, "Integrated robot planning, path following, and obstacle avoidance in two and three dimensions: Wheeled robots, underwater vehicles, and multicopters," *International Journal of Robotics Research*, vol. 38, no. 7, pp. 853–876, 2019.

[3] Z. Shiller and Y.-R. Gwo, "Dynamic motion planning of autonomous vehicles," *IEEE Transactions on Robotics and Automation*, vol. 7, no. 2, pp. 241–249, 1991.

[4] D. Dolgov, S. Thrun, M. Montemerlo, and J. Diebel, "Path planning for autonomous vehicles in unknown semi-structured environments," *International Journal of Robotics Research*, vol. 29, no. 5, pp. 485–501, 2010.

[5] A. Gandia, S. Parascho, R. Rust, G. Casas, F. Gramazio, and M. Kohler, "Towards automatic path planning for robotically assembled spatial structures," in *Robotic Fabrication in Architecture, Art and Design*, pp. 59–73, Springer, 2018.

[6] M. Dogar, A. Spielberg, S. Baker, and D. Rus, "Multi-robot grasp planning for sequential assembly operations," *Autonomous Robots*, vol. 43, no. 3, pp. 649–664, 2019.

[7] F. Capolupo and P. Labourdette, "Receding-horizon trajectory planning algorithm for passively safe on-orbit inspection missions," *Journal of Guidance, Control, and Dynamics*, vol. 42, no. 5, pp. 1–10, 2019.

[8] N. M. Amato and G. Song, "Using motion planning to study protein folding pathways," *Journal of Computational Biology*, vol. 9, no. 2, pp. 149–168, 2002.

[9] L. Jean-Claude, *Robot Motion Planning*, vol. 124, Springer Science & Business Media, 2012.

[10] H. Han-Pang and C. Shu-Yun, "Dynamic visibility graph for path planning," in *Proceedings of the 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (IEEE Cat. No. 04CH37566)*, vol. 3, pp. 2813–2818, IEEE, Sendai, Japan, 2004.

[11] P. Bhattacharya, M. L. Gavrilova, and L. Marina, "Voronoi diagram in optimal path planning," in *Proceedings of the 4th International Symposium on Voronoi Diagrams in Science and Engineering (ISVD '07)*, pp. 38–47, IEEE, July 2007.

[12] F. Lingelbach, "Path planning using probabilistic cell decomposition," in *Proceedings of the IEEE International Conference on Robotics and Automation, 2004 (ICRA'04)*, vol. 1, pp. 467–472, IEEE, 2004.

[13] L. E. Kavraki and L. Jean-Claude, "Probabilistic roadmaps for robot path planning," 1998.

[14] A. Yilmaz and M. Shah, "Actions sketch: a novel action representation," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR '05)*, vol. 1, pp. 984–989, June 2005.

[15] A. Yilmaz and M. Shah, "A differential geometric approach to representing the human actions," *Computer Vision and Image Understanding*, vol. 109, no. 3, pp. 335–351, 2008.

[16] P. Turaga, R. Chellappa, V. S. Subrahmanian, and O. Udrea, "Machine recognition of human activities: a survey," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 18, no. 11, pp. 1473–1488, 2008.

[17] J. Liu, B. Kuipers, and S. Savarese, "Recognizing human actions by attributes," in *Proceedings of the Conference on Computer Vision and Pattern Recognition*, pp. 3337–3344, 2011.

[18] R. Vemulapalli, F. Arrate, and R. Chellappa, "Human action recognition by representing 3D skeletons as points in a lie group," in *Proceedings of the 27th IEEE Conference on Computer Vision and Pattern Recognition (CVPR '14)*, pp. 588–595, Columbus, Ohio, USA, June 2014.

[19] K. Hausman, S. Niekum, S. Osentoski, and G. S. Sukhatme, "Active articulation model estimation through interactive perception," in *Proceedings of the 2015 IEEE International Conference on Robotics and Automation, ICRA 2015*, pp. 3305–3312, USA, May 2015.

[20] D. Koller and N. Friedman, *Probabilistic Graphical Models: Principles and Techniques*, MIT Press, Cambridge, MA, USA, 2009.

[21] M. Leyton, *A Generative Theory of Shape*, Springer-Verlag, Berlin, Heidelberg, Germany, 2003.

[22] A. Byravan and D. Fox, "SE3-nets: learning rigid body motion using deep neural networks," in *Proceedings of the 2017 IEEE International Conference on Robotics and Automation, ICRA 2017*, pp. 173–180, Singapore, June 2017.

[23] S. Soatto, "Actionable information in vision," in *Proceedings of the 12th International Conference on Computer Vision, ICCV 2009*, pp. 2138–2145, Japan, October 2009.

[24] Y. Li, C. Fermuller, Y. Aloimonos, and H. Ji, "Learning shift-invariant sparse representation of actions," in *Proceedings of the 2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, CVPR 2010*, pp. 2630–2637, USA, June 2010.

[25] L. Zhang and B. Zhang, *Quotient Space Based Problem Solving: A Theoretical Foundation of Granular Computing*, Morgan Kaufmann Publishers Inc, Burlington, MA, USA, 2014.

[26] J. Hartmanis, *Algebraic Structure Theory of Sequential Machines*, Prentice-Hall, Englewood Cliffs, NJ, USA, 1966.

[27] A. Ginzburg, *Algebraic Theory of Automata*, Academic Press, New York, NY, USA, 1968.

[28] M. A. Arbib, *Theories of Abstract Automata*, Prentice-Hall, Englewood Cliffs, NJ, USA, 1969.

[29] M. Minsky and S. Papert, *Perceptrons: An Introduction to Computational Geometry*, The MIT Press, Cambridge, MA, USA, 1987.

[30] Y. Liu and R. T. Collins, "Computational model for repeated pattern perception using Frieze and wallpaper groups," in *Proceedings of the CVPR '2000: IEEE Conference on Computer Vision and Pattern Recognition*, vol. 1, pp. 537–544, June 2000.

[31] T. C. Henderson, X. Fan, S. Devnani, S. Kumar, E. Cohen, and E. Grant, "Symmetry as an organizational principle in cognitive sensor networks," UUCS-09-005, University of Utah, 2009.

[32] C. Kiddon and P. Domingos, "Symmetry-based semantic parsing," in *Proceedings of the 2014 Workshop on Learning Semantics*, 2015.

[33] H. F. Baker, "Alternants and continuous groups," *Proceedings of the London Mathematical Society*, vol. 2, no. 1, pp. 24–47, 1905.

[34] D. A. Suprunenko, *Matrix Groups*, vol. 45, American Mathematical Society, 1976.

[35] G. Birkhoff and S. M. Lane, *A Survey of Modern Algebra*, AK Peters/CRC Press, 1998.

[36] J. M. Selig, "Lie groups and Lie algebras in robotics," in *Computational Noncommutative Algebra and Applications*, pp. 101–125, Springer, 2004.

[37] R. M. Murray, S. S. Sastry, and Z. Li, *A Mathematical Introduction to Robotic Manipulation*, CRC Press, Boca Raton, FL, USA, 1st edition, 1994.

[38] Q. Ma, Z. Goh, S. Ruan, and G. S. Chirikjian, "Probabilistic approaches to the AXB= YCZ calibration problem in multi-robot systems," *Autonomous Robots*, vol. 42, no. 7, pp. 1497–1520, 2018.

[39] A. Müller, "Screw and Lie group theory in multibody dynamics: recursive algorithms and equations of motion of tree-topology systems," *Multibody System Dynamics*, vol. 42, no. 2, pp. 219–248, 2018.

[40] G. S. Chirikjian, R. Mahony, S. Ruan, and J. Trumpf, "Pose changes from a different point of view," *Journal of Mechanisms and Robotics*, vol. 10, no. 2, Article ID 021008, 2018.

[41] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2323, 1998.

[42] G. E. Hinton, S. Osindero, and Y.-W. Teh, "A fast learning algorithm for deep belief nets," *Neural Computation*, vol. 18, no. 7, pp. 1527–1554, 2006.

[43] Y. Bengio, P. Lamblin, D. Popovici, and H. Larochelle, "Greedy layer-wise training of deep networks," in *Proceedings of the 20th Annual Conference on Neural Information Processing Systems (NIPS '06)*, pp. 153–160, Cambridge, Mass, USA, December 2006.

[44] Y. Bengio, "Learning deep architectures for AI," *Foundations and Trends in Machine Learning*, vol. 2, no. 1, pp. 1–127, 2009.

[45] J. Tompson, A. Jain, Y. LeCun, and C. Bregler, "Joint training of a convolutional network and a graphical model for human pose estimation," in *Proceedings of the 28th Annual Conference on Neural Information Processing Systems 2014, NIPS 2014*, vol. 1, pp. 1799–1807, Canada, December 2014.

[46] M. Vejdemo-Johansson, F. T. Pokorny, P. Skraba, and D. Kragic, "Cohomological learning of periodic motion," *Applicable Algebra in Engineering, Communication and Computing*, vol. 26, no. 1-2, pp. 5–26, 2015.

[47] T. S. Cohen and M. Welling, "Group equivariant convolutional networks," in *Proceedings of the 33rd International Conference on Machine Learning*, vol. 48, pp. 2990–2999, 2016.

[48] S. Dieleman, J. De Fauw, and K. Kavukcuoglu, "Exploiting cyclic symmetry in convolutional neural networks," in *Proceedings of the 33rd International Conference on Machine Learning*, vol. 48, pp. 1889–1898, 2016.

[49] J. J. Kuffner, "Effective sampling and distance metrics for 3D rigid body path planning," in *Proceedings of the International Conference on Robotics and Automation (ICRA)*, Citeseer, pp. 3993–3998, 2004.

[50] A. James, "Fast random rotation matrices," in *Graphics Gems III (IBM Version)*, pp. 117–120, Elsevier, 1992.

[51] A. Yershova and S. M. LaValle, "Deterministic sampling methods for spheres and SO(3)," in *Proceedings of the IEEE International Conference on Robotics and Automation, 2004 (ICRA'04)*, vol. 4, pp. 3974–3980, IEEE, USA, May 2004.

[52] J. C. Mitchell, "Sampling rotation groups by successive orthogonal images," *SIAM Journal on Scientific Computing*, vol. 30, no. 1, pp. 525–547, 2008.

[53] A. Yershova, S. Jain, S. M. Lavalle, and J. C. Mitchell, "Generating uniform incremental grids on SO_3 using the hopf fibration," *International Journal of Robotics Research*, vol. 29, no. 7, pp. 801–812, 2010.

[54] Y. Yan and G. S. Chirikjian, "Almost-uniform sampling of rotations for conformational searches in robotics and structural biology," in *Proceedings of the 2012 IEEE International Conference on Robotics and Automation, ICRA 2012*, pp. 4254–4259, USA, May 2012.

[55] Y. Yan and G. S. Chirikjian, "Voronoi cells in lie groups and coset decompositions: implications for optimization, integration," in *Proceedings of the 52nd Conference on Decision and Control*, pp. 1137–1143, 2013.

[56] R. Blaser and P. Fryzlewicz, "Random rotation ensembles," *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 126–151, 2016.

[57] K. Shoemake, "Uniform random rotations," in *Graphics Gems III (IBM Version)*, pp. 124–132, Elsevier, 1992.

[58] E. E. Lattman, "Optimal sampling of the rotation function," *Acta Crystallographica Section B: Structural Crystallography and Crystal Chemistry*, vol. 28, no. 4, pp. 1065–1068, 1972.

[59] D. J. Evans, "On the representatation of orientation space," *Molecular Physics*, vol. 34, no. 2, pp. 317–325, 1977.

[60] A. Morawiec and D. P. Field, "Rodrigues parameterization for orientation and misorientation distributions," *Philosophical Magazine A*, vol. 73, no. 4, pp. 1113–1130, 1996.

[61] B. Chazelle, *The Discrepancy Method: Randomness and Complexity*, Cambridge University Press, 2001.

[62] J. Matousek, *Geometric Discrepancy: An Illustrated Guide*, vol. 18, Springer Science & Business Media, 2009.

[63] B. Chandrajit, B. Abhishek, E. Chattopadhyay, and D. Zuckerman, "On low discrepancy samplings in product spaces of motion groups," 2014, https://arxiv.org/abs/1411.7753.

[64] G. S. Chirikjian and A. B. Kyatkin, *Harmonic Analysis for Engineers and Applied Scientists*, Dover Publications, New York, NY, USA, 2nd edition, 2016.

[65] G. S. Chirikjian, S. Sajjadi, B. Shiffman, and S. M. Zucker, "Mathematical aspects of molecular replacement: IV. Measure-theoretic properties of motion spaces," *Acta Crystallographica*, vol. 73, no. 5, pp. 387–402, 2017.

[66] M. M. Postnikov, "Three-dimensional spherical forms," *Akademiya Nauk SSSR. Trudy Matematicheskogo Instituta Imeni V. A. Steklova*, vol. 196, pp. 114–146, 1991.

[67] R. C. Kirby and M. G. Scharlemann, "Eight faces of the Poincaré homology 3-sphere," in *Geometric Topology*, pp. 113–146, Elsevier, 1979.

[68] T. Janssen, *Crystallographic Groups*, Elsevier, New Yory, NY, USA, 1973.

[69] T. Hahn, Ed., *International Tables for Crystallography, Volume A: Space Group Symmetry*, International Union for Crystallography (IUCr), Chester, UK, 2002.

[70] M. I. Aroyo, J. M. Perez-Mato, C. Capillas et al., "Bilbao crystallographic server: I. databases and crystallographic computing programs," *Zeitschrift für Kristallographie*, vol. 221, no. 1, pp. 15–27, 2006.

[71] H. Wondratschek and U. Müller, Eds., *International Tables for Crystallography, Volume A1: Symmetry Relations Between Space Groups*, International Union for Crystallography (IUCr), Chester, UK, 2008.

[72] W. Hantzsche and H. Wendt, "Dreidimensionale euklidische raumformen," *Mathematische Annalen*, vol. 110, no. 1, pp. 593–611, 1935.

[73] L. S. Charlap, *Bieberbach Groups and Flat Manifolds*, Springer-Verlag, New York, NY, USA, 1986.

[74] J. M. Montesinos, *Classical Tessellations and Three-Manifolds*, Springer-Verlag, Berlin, Heidelberg, Germany, 1987.

[75] J. A. Wolf, *Spaces of Constant Curvature*, AMS Chelsea Publishing, New York, NY, USA, 2010.

[76] A. Szczepanski, *Geometry of Crystallographic Groups*, World Scientific, Singapore, 2012.

[77] G. S. Chirikjian, K. Ratnayake, and S. Sajjadi, "Decomposition of Sohncke space groups into products of Bieberbach and symmorphic parts," *Zeitschrift fur Kristallographie*, vol. 230, no. 12, pp. 719–741, 2015.

[78] R. S. Jackendoff, *Semantics and Cognition*, MIT Press, Cambridge, MA, USA, 1985.

[79] T. Winograd and F. Flores, *Understanding Computers and Cognition: A New Foundation for Design*, Ablex Publishing Corporation, Norwood, NJ, USA, 1986.

[80] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, Pearson Education Inc, Upper Saddle River, NJ, USA, 3rd edition, 2009.

[81] G. S. Chirikjian, "Mathematical aspects of molecular replacement. I. algebraic properties of motion spaces," *Acta Crystallographica Section A Foundations of Crystallography*, vol. 67, no. 5, pp. 435–446, 2011.

[82] G. S. Chirikjian, "Kinematics meets crystallography: the concept of a motion space," *Journal of Computing and Information Science in Engineering*, vol. 15, no. 1, Article ID 011012, 2015.