A Qualitative Calculus for Three-Dimensional Rotations

Azam Asl
Department of Computer Engineering
NYU Polytechnic
aa2821@cs.nyu.edu

Ernest Davis*
Dept. of Computer Science
New York University
New York, NY 10012
davise@cs.nyu.edu

July 31, 2012

Abstract

We have developed a qualitative calculus for three-dimensional directions and rotations. A direction is characterized in terms of the signs of its components relative to an absolute coordinate system. A rotation is characterized in terms of the signs of the components of the associated 3×3 rotation matrix.

A system has been implemented that can solve the following problems:

1. Given the signs of direction $\hat{v}$ and rotation matrix P, find the possible signs of the image of $\hat{v}$ under P. Moreover, for each possible sign vector of $\hat{v} \cdot P$, generate numerical instantiations of $\hat{v}$ and P that yields that result.
2. Given the signs of rotation matrices P and Q, find the possible signs of the composition $P \cdot Q$. Moreover, for each possible sign matrix for the composition, generate numerical instantiations of P and Q that yield that result.

We have also proven some related complexity and expressivity results. The satisfiability problem for a qualitative rotation constraint network is NP-complete in two dimensions and NP-hard in three dimensions. In three dimensions, any two directions are distinguishable by a qualitative rotation constraint network.

1 Introduction

The field of Qualitative Spatial Reasoning (QSR) develops methods for carrying out geometric computations using qualitative information about spatial properties and relations, rather than numerically precise information [3]. The majority of the QSR literature has addressed reasoning about topological constraints between regions; the best known theory of this kind is the RCC-8 system of relations [19]. However, other work in the area has addressed other geometric properties such as convexity, relative position, and relative size [3]. A qualitative calculus is a theory that describes how an inference engine can use a constraint network of qualitative geometric relations to draw conclusions that are implicit but not explicit in the network.

The research described in this paper defines, implements, and analyzes a qualitative calculus for three-dimensional directions and rotations. As every rigid motion is the composition of a rotation and a translation, developing a qualitative theory of rotation is a significant step toward a system

*Thanks to Richard Cole for helpful discussions about the proof of theorem 5.

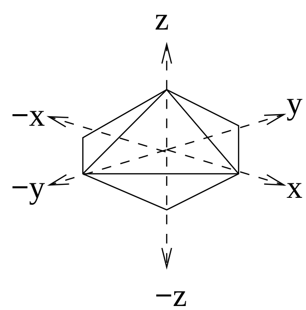

Figure 1: Standard octahedron

that can reason qualitatively about three-dimensional motion generally. This potentially can have a wide range of applications, from planning to robotics to molecular dynamics; these are discussed in section 3.

1.1 Qualitative Characterization of Directions and Rotations

In our theory, we consider a domain consisting of two kinds of spatial entities in three-dimensional space: directions and rotations. There are two operations: applying a rotation Φ to a direction $\hat{v}$, denoted "$\Phi(\hat{v})$" and composing two rotations Φ and Ψ, denoted "$\Phi \circ \Psi$." (We notate composition in the analysts' form: $(\Phi \circ \Psi)(\hat{v}) \equiv \Phi(\Psi(\hat{v})$.) Correspondingly there are two kinds of constraints: "The image of direction $\hat{u}$ under rotation Φ is $\hat{v}$" and "The composition of rotations Φ and Ψ is rotation Θ".

To characterize directions and rotations qualitatively, we proceed as follows. We fix an absolute $\hat{x}, \hat{y}, \hat{z}$ coordinate system. We divide the unit sphere into a JEPD partition of 26 cells corresponding to the 6 vertices, 12 edges, and 8 faces of a standard octahedron, centered at the origin, and with vertices at the points $\hat{x}, -\hat{x}, \hat{y}, -\hat{y}, \hat{z}$ and $-\hat{z}$ (figure 1). A direction $\hat{u}$ is characterized qualitatively by the cell of the partition that contains $\hat{u}$; there are thus 26 qualitative categories of directions. A rotation Φ is characterized qualitatively by the cells of the partitions that contain $\Phi(\hat{x})$, $\Phi(\hat{y})$, $\Phi(\hat{z})$, where $\hat{x}, \hat{y}, \hat{z}$ are the directions along the coordinate axes. As we shall see in section 4.1, there are 336 qualitative categories of rotations.

A standard method for representing directions and rotations numerically is in terms of vectors and matrices. (We discuss other representations briefly in section 1.3.) A direction $\hat{v}$ is represented as the three-dimensional vector of its x, y, and z coordinates; a vector represents a direction if its magnitude is 1. A rotation Φ is represented as the 3×3 matrix M whose rows are respectively $\Phi(\hat{x})$, $\Phi(\hat{y})$ and $\Phi(\hat{z})$. For example the matrix

$$M = \begin{bmatrix} 1/3 & 2/3 & 2/3 \\ 2/3 & -2/3 & 1/3 \\ 2/3 & 1/3 & -2/3 \end{bmatrix}$$

is the rotation that maps $\hat{x}$ to the direction $\langle 1/3, 2/3, 2/3 \rangle$; maps $\hat{y}$ to the direction $\langle 2/3, -2/3, 1/3 \rangle$; and maps $\hat{z}$ to the direction $\langle 2/3, 1/3, -2/3 \rangle$.

A numerical matrix M describes a rotation if it satisfies the following two properties.

A. M is orthogonal. That is $M \cdot M^T = I$. Equivalently, every row and every column has unit length; each row is orthogonal to every other row; and every column is orthogonal to every other column. A consequence is that if two elements in a row/column are zero then the third element must be ± 1, and conversely.

B. The determinant of M is equal to 1. A matrix satisfying (A) is either a rotation or a reflection. Rotations have determinant 1; reflections have determinant -1.

In this representation, the application of a rotation Φ to a direction $\hat{v}$ can be computed by multiplying the vector corresponding to $\hat{v}$ by the matrix corresponding to Φ. The composition of two rotations, Φ followed by Ψ is computed as the product of the matrix corresponding to Φ times the matrix corresponding to Ψ. (See [4], chap. 6).

The qualitative characterization of a direction can be computed as the signs of the components of the corresponding vector; the qualitative characterization of a rotation can be computed as the signs of the components of the corresponding matrix. For example, the the direction $\langle 3/5, 0, -4/5 \rangle$ can be qualitatively characterized in terms of the sign vector $\langle +, 0, - \rangle$. The matrix M above can be qualitatively characterized in terms of the sign matrix

$$\begin{bmatrix} + & + & + \\ + & - & + \\ + & + & - \end{bmatrix}$$

A sign vector with two 0 components corresponds to a vertex of the standard octahedron; for example, $\langle 0, 0, - \rangle$ is the $-\hat{z}$ vertex. A sign vector with one 0 component corresponds to an edge; for example $\langle +, 0, - \rangle$ corresponds to the edge connecting $\hat{x}$ to $-\hat{z}$. A sign vector with no 0 components corresponds to a face; for example, the vector $\langle -, +, + \rangle$ corresponds to the face with vertices $-\hat{x}, \hat{y}, \hat{z}$.

The structure of the calculus developed in this paper is structurally more closely akin to the sign calculus of [5] than to relational calculi such as RCC [19]. It is fundamentally a functional qualitative calculus rather than a relational qualitative calculus. In a functional qualitative calculus, the domain consists of one or more spaces of entities. Each space is partitioned into a JEPD collection of categories called "qualitative" categories. There is a fixed (small) collection of functions over the entities. The functional calculus addresses the question, "Given the qualitive categories of entities $x_1 \ldots x_k$, what are the possible qualitative categories of $f(x_1 \ldots x_k)$?" By contrast, in a relational qualitative calculus such as RCC, the domain generally consists of a single space of entities, and the qualitative theory is a JEPD collection of relations over those entities.

It is actually possible to construe part of our theory in the form of a relational calculus, as follows. Let the domain be the set of rotations. For each qualitative category S, define the qualitative binary relation $C_S(\Phi, \Psi)$ as holding if there exists a rotation Θ of category S such that $\Psi = \Theta \circ \Phi$. Then, for any two rotations Φ and Ψ, there is a unique relation C_S such that $C_S(\Phi, \Psi)$; namely, where S is the qualitative category of $\Psi \circ \Phi^{-1}$. Therefore the collection of relations C_S does form a JEPD collection of qualitative binary relations over the domain of rotations. Composition, inverse, and identity of these relations then correspond in the natural way to composition, inverse, and identity in the underlying rotation Θ. However, this is not probably the most useful way to conceptualize the domain, since it hides the underlying symmetry between Θ on the one hand and Φ and Ψ on the other.

Moreover, this approach certainly cannot be applied to the qualitative analysis of applying rotations to directions, because two directions are not related by a unique rotation. For any two directions $\hat{u}, \hat{v}$ there are many different rotations Φ, of different qualitative categories, such that $\Phi(\hat{u}) = \hat{v}$. The part of our theory dealing with directions, therefore, can only be construed as a functional calculus and not as a relational calculus.

Having defined this qualitative calculus, we address two fundamental problems:

1. Given the qualitative characterizations of a direction $\hat{u}$ and a rotation Φ, compute the possible qualitative characterizations of $\Phi(\hat{u})$.

2. Given the qualitative characterizations of two rotations Φ and Ψ, compute the qualitative characterization of the composition $\Psi \circ \Phi$.

We have implemented a program that computes these, where the qualitative characterizations are represented in terms of sign vectors and sign matrices.

The most extensively studied form of QSR in general is the problem of determining the consistency or the consequences of a constraint network. In the context of our theory, such a network would have directions as the nodes, labelled by a sign vector, or by a set of possible sign vectors, and rotations as the edges, labelled by a sign matrix or by a set of sign matrices. Our solution to problem (1) above allows label propagation (Waltz propagation) to be carried out over a constraint network. Our solution to problem (2) would allows arc propagation to be carried out; that is, the solution to (2) constitutes a composition table over qualitative rotations. We have not implemented either propagation algorithm, but, given the functionalities implemented in 3DR, this would be straightforward; the super-routines to carry out the propagation are entirely standard ([14]; [21] chap. 6).

The conceptual neighborhood graph over these relations, characterizing continuous change, is easily characterized; the matrix P may instantaneously transition to Q if some of the non-zero signs in Q are changed to 0 in P. In this case P dominates Q, in the terminology of [7].

1.2 Reasoning over three-dimensional rotations as compared to two-dimensional rotations

It should be noted that three-dimensional rotations are intrinsically much more complicated than two-dimensional rotations, in several respects. In two-dimensions the space of directions is isomorphic to the space of rotations; and both directions and rotations can be represented as an angle in the interval $[0, 2\pi)$. Both the application of a rotation to a direction and the composition of two rotations corresponds to addition of angles mod 2π; in particular, two-dimensional rotations are commutative.

The representation of partial specifications of directions and rotations is similarly simple and well behaved. A direction or a rotation can be partially specified by an interval of the unit circle; equivalently, by a subinterval of $[0, 4\pi)$ of length at most 2π. If $\hat{u}$ is a rotation with angle bounded by $[a, b]$ and Γ is a rotation bounded by $[c, d]$, then $\Gamma(\phi)$ is bounded by $[a + c, b + d]$ or some variant, with a small number of easily-characterized special cases. Moreover, any two intervals of the same length represent the same degree of knowledge or ignorance; e.g. bounds of $[0, \pi/6]$ and $[\pi/3, \pi/2]$ are congruent intervals on the unit circle.

None of these elegant properties hold for three-dimensional rotations. Three-dimensional rotations do not commute. Three-dimensional directions comprise a space with two degrees of freedom (a two-dimensional manifold); whereas three-dimensional rotations comprise a space with three degrees of freedom (a three-dimensional space). There are representations of direction in terms of two angles such as latitude and longitude, and representations of rotations in terms of three angles, such as yaw-pitch-roll or various forms of Eulerian angles. However any such representation necessarily suffers from topological singularities; for instance, longitude is undefined and discontinuous at the north and south poles. (These singularities are related to the "gimbal lock" phenomenon.) There is in general no simple method to express the image of a direction under a rotation, or the composition of two rotations, where these are expressed in angular form; in fact, in general it is necessary to convert the direction to Cartesian coordinates and the rotation(s) to matrix form, multiply, and then convert back to angular form.

Partial specifications of directions and rotations in terms of angles have no simple compositions.

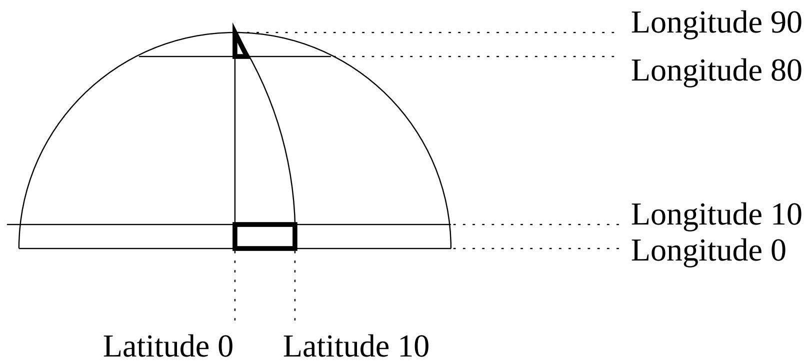

Figure 2: Regions on the sphere bounded by equal angle ranges

Equal angle ranges do not give congruent regions. For example, the region bounded by latitudes 80° and 90° and longitudes 0° and 10° is by no means congruent to the region bounded by latitudes 0° and 10° and longitudes 0° and 10°. The former is a triangle,[1] the latter is a quadrilateral approximately 6 times as large in area (figure 2).

1.3 Alternative representations

Since our qualitative characterization is fundamentally geometric, it is not inherently tied to using a vector/matrix-based representation. There are other frameworks for numerical representations of directions and rotations that are used for a variety of purposes. For instance, in an angular representation, direction can be represented by two angles, latitude and longitude; and rotations are represented by any of a number of possible triples of angles. In the quaternion representation, directions and rotations are represented by a 4-tuple of real numbers; the encoding is too complex to discuss here. Our qualitative categories can certainly be expressed as combinations of numerical inequalities in any of these. For instance the qualitative category of direction $\langle +, -, + \rangle$ can be expressed in the constraints "Latitude is between $-\pi/2$ and 0; longitude is between 0 and $\pi/2$." However, the formulation of the constraints, and the calculation of the qualitative application and composition operators, are substantially more complicated in these alternative representations, and they do not seem to offer any advantage in this application.

2 Related work

A number of earlier studies have considered qualitative directional constraints of one form and another, in two-dimensional geometry. [6] considered relations between points described by the cardinal directions. The STAR calculus of [20] divided the unit circle of directions into m sectors plus m rays, and considered constraints over a domain of points of the form "The direction from **p** to **q** lies in section k." The $\mathcal{OPRA}$ calculus of [16, 18] used a similar qualitative division of the unit circle to characterize the relation between directed points; $\mathcal{OPRA}$, however, used relative frames of reference whereas STAR used an absolute frame of reference. The $\mathcal{LR}$ system of [22] used a system of ternary constraints over points, of the form "Point **r** lies to the left of the directed line going from **p** to **q**." The Cardinal Direction Calculus (CDC) [8, 25] characterized the range of directions between points in two extended objects, relative to absolute coordinate directions.

[1]Not, strictly speaking, a spherical triangle, since the circle of latitude is not a great circle.

Rodney Brooks' ACRONYM computer vision system [2] used a 3-angle system to characterize the three-dimensional orientation of objects and their parts, and used angular constraints like $-\pi/2 < TILT < -\pi/6$ to specify partial information about the orientations of both objects and camera. Special purpose inference modules, essentially carrying out interval propagation, were used to carry out inferences from these bounds. In themselves, constraints of this kind are of course a much more expressive language for partial information than that considered in this paper. However, it would appear that ACRONYM only considered absolute orientation and not relative orientation, and that inferences only involved two objects at a time (the camera and an object in the image).

The FROB spatial inference system of McDermott and Davis [15] likewise used interval bounds on Eulerian angles. We do not know of any more recent system that did any kind of reasoning with partial specifications of three-dimensional orientation. In vision, at least, reasoning with hard constraints of this kind has largely been superceded by reasoning with soft (probabilistic) constraints. The complexities of three-dimensional rotations are still significant in reasoning with soft constraints, but they are easier to ignore.

The sign calculus has been used extensively in qualitative reasoning, particularly qualitative physical reasoning; a systematic discussion may be found in [5]. Kim's qualitative theory of linkages [11] used sign vectors to characterize orientations of segments in a two-dimensional linkage.

3 Potential Applications

Qualitative reasoning is potentially useful in situations where an object undergo rotations in three-dimensions, and only partial information about these rotations is available, or where the relative three-dimensional orientation of objects or object parts is only partially known or specified. This partial information may arise in a number of different ways:

1. Sensors provide limited information. ([9] is an interesting study of the difficulties of planning motion of a three-dimensional robotic arm with realistic limitations on sensors.)
2. Actions are carried out with limited accuracy. Hence, if one is projecting the consequence of a sequence of actions, or using dead reckoning to estimate current orientation after a sequence of actions, then the uncertainties of the rotations are combined.
3. The information comes from a plan that is only partially specified.
4. The information comes from inferences. For instance, if it is observed that a bottle holds liquid, then one can infer that its orientation has a positive upward component. If one observes the shadow of an object of unknown height on the ground, then its projection on the x-y plane is known but not its orientation in the z-direction.
5. Only certain relative orientations are physically possible. For example, two segments of an articulated object meet at a joint with a limited range.

As an example of (5) above, consider the following model of a robotic arm.[2]. The arm consists of n links of length $L_1 \ldots L_n$ connected in sequence by pin joints. The angle at the connection between link $i-1$ and link i is θ_i and when all the θ_i are zero, then the links lie stretched out along the positive x-axis. However, the direction of the pins alternates. Specifically, the pin at the origin is always vertical, so that the first link rotates in the $x-y$ plane. When $\theta_1 = 0$ the second pin is parallel to the y axis, so the second link rotates in the $x-z$ plane. When all the θ's are 0 the third pin is again vertical, the fourth pin is again parallel to the y axis, the fifth is vertical, the sixth is parallel to the y axis and so on. (Figure 3)

[2]We do not know of any actual arm built like this, but it is a convenient model

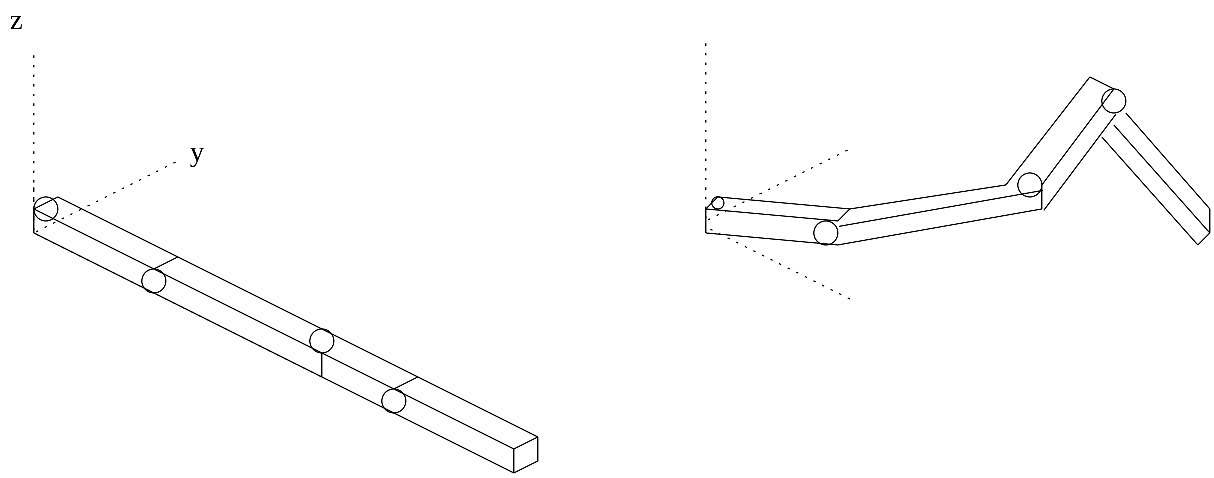

Figure 3: Three dimensional robotic arm

In that case, the orientation of the final segment of the arm is the composition of the rotation matrix from each segment to the next; that is, the product of matrices of the form

$$\begin{bmatrix} \cos(\theta_i) & -\sin(\theta_i) & 0 \\ \sin(\theta_i) & \cos(\theta_i) & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} \cos(\theta_i) & 0 & -\sin(\theta_i) \\ 0 & 1 & 0 \\ \sin(\theta_i) & 0 & \cos(\theta_i) \end{bmatrix}$$

There may additionally be constraints on the values of θ_i; for example, it may not be possible to bend any of the joints more than 90°. A question such as "Is it possible for the final segment to be facing in the positive $x - y$ direction?" or "Given two such arms can they be positioned so that the hands face in opposite directions?" can then be cast as a satisfiability problem over the qualitative calculus that we develop here. (Reasoning about the *positions* of the objects involved would require extensions to the theory developed here.)

The use of other forms of qualitative spatial reasoning for planning for robotic manipulation and navigation has been considered in [23] and [17].

Other domains where this kind of reasoning could be useful include interpreting the motion of a three-dimensional articulated object (e.g. a video of an acrobat); reasoning about the motions of satellite or other object in outer space; or reasoning about molecular motions. We are not making any strong claims as to the immediate applicability of the theory presented here to any particular practical problems. However, the general problem of qualitative reasoning about three-dimensional rotations is certainly potentially important in a wide range of applications; and the theory developed in this paper is a substantial advance in developing techniques for that kind of reasoning.

4 Qualitative Rotations

In this section, we describe the implementation of 3DR. We describe first how the symmetries of the geometry can be used to dramatically simplify the case analysis, and then how the reduced case analysis is carried out.

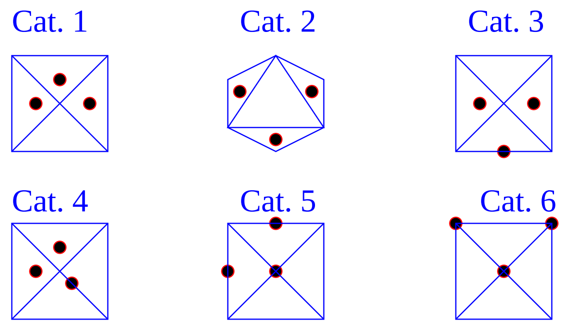

Figure 4: Categories of rotations. The lines are the standard octahedron; the red dots are the images of $\hat{x}, \hat{y}, \hat{z}$

4.1 Octahedral Rotations

The *octahedral* rotations map the standard octahedron into itself; equivalently, they map $\hat{x}$, $\hat{y}$, $\hat{z}$ into some combination of $\pm\hat{x}$, $\pm\hat{y}$, $\pm\hat{z}$. There are 24 such rotations. Proof: You can choose the image $\Gamma(\hat{x})$ to be any of the 6 vertices of the octahedron; then you can choose $\Gamma(\hat{y})$ to be any of the 4 orthogonal vertices; then $\Gamma(\hat{z})$ is fixed. Alternate proof: Pick a face F of the octahedron. An octahedral rotation maps F to any of the 8 faces of the octahedron; and for each of these there are 3 ways that the two triangles can be aligned.

The matrix for an octahedral rotation has one non-zero element, either 1 or -1, in each row and column. For example, the rotation Γ such that $\Gamma(\hat{x}) = -\hat{z}, \Gamma(\hat{y}) = \hat{y}, \Gamma(\hat{z}) = \hat{x}$ corresponds to the matrix

$$\begin{bmatrix} 0 & 0 & -1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix}$$

There are 24 such matrices with determinant 1.

If $\hat{u}$ and $\hat{v}$ are directions on the same piece of the octahedron, and Γ is an octahedral rotation then clearly $\Gamma(\hat{u})$ and $\Gamma(\hat{v})$ are still in the same piece of the octahedron. Therefore, if A and B are two numerical rotational matrices with the same signs, and P is an octahedral rotation matrix, then $A \cdot P$ and $B \cdot P$ have the same sign. Therefore, we can define an equivalence relation over the sign rotation matrices: Q is equivalent to R if $Q = R \cdot P$ for some octahedral rotation P. Every equivalence class of sign rotations has 24 elements.

Since the octahedral rotations preserve the geometry of the octahedron, we can be sure that two sign rotations A and B are from different equivalence classes if the two triples $\langle \hat{x} \cdot A, \hat{y} \cdot A, \hat{z} \cdot A \rangle$ and $\langle \hat{x} \cdot B, \hat{y} \cdot B, \hat{z} \cdot B \rangle$ are geometrically different, in terms of features of the octahedron. Note that these are the rows of A and B respectively.

Below is a list of 14 representatives, with geometric features that guarantee that they are distinct. These fall into 6 categories (figure 4).

Category 1 includes three representatives. In these, $\hat{x} \cdot R$, $\hat{y} \cdot R$, and $\hat{z} \cdot R$ are all in faces that share the vertex $\langle 1, 0, 0 \rangle$.

$$R1 = \begin{bmatrix} + & + & + \\ + & - & + \\ + & + & - \end{bmatrix} \qquad R2 = \begin{bmatrix} + & + & + \\ + & - & + \\ + & - & - \end{bmatrix} \qquad R3 = \begin{bmatrix} + & + & + \\ + & - & - \\ + & + & - \end{bmatrix}$$

Note that the face corresponding to $\hat{x} \cdot R1$ shares an edge with both $\hat{y} \cdot R1$ and with $\hat{z} \cdot R1$, but that $\hat{y} \cdot R1$ and $\hat{z} \cdot R1$ do not share an edge; that $\hat{y} \cdot R2$ shares an edge with both $\hat{x} \cdot R2$ and with $\hat{z} \cdot R2$; $\hat{z} \cdot R3$ shares an edge with both $\hat{x} \cdot R3$ and with $\hat{y} \cdot R3$. (Two faces share an edge if the sign vectors are equal in two places; they share a vertex if the sign vectors are equal in one place; they are antipodal if the signs are opposite throughout.)

Category 2 includes a single representative. In this $\hat{x} \cdot R$, $\hat{y} \cdot R$, and $\hat{z} \cdot R$ are all inside faces. Each of the faces connects to each of the other at a vertex.

$$R4 = \begin{bmatrix} + & + & + \\ + & - & - \\ - & + & - \end{bmatrix}$$

Category 3 includes 3 representatives. In these, two of the row vectors lie in faces F1, F2 and one lies in an edge E. Faces F1, F2 have a common vertex V; edge E connects a vertex of F1 not equal to V with a vertex of F2 not equal to V. For instance in R5 below, F1 (the second row) is the face with vertices $-\hat{x}$, $\hat{y}, \hat{z}$; F2 (the third row) is the face with vertices $\hat{x}$, $-\hat{y}, \hat{z}$; these meet at vertex V=$\hat{z}$. Edge E (the first row) connects $\hat{x}$ with $\hat{y}$.

$$R5 = \begin{bmatrix} + & + & 0 \\ - & + & + \\ + & - & + \end{bmatrix} \qquad R6 = \begin{bmatrix} + & + & + \\ + & 0 & - \\ - & + & - \end{bmatrix} \qquad R7 = \begin{bmatrix} + & + & + \\ - & - & + \\ + & - & 0 \end{bmatrix}$$

Category 4 includes 3 representatives. In these, two of the row vectors lie in a face and one lies in an edge. The two faces have a common edge; one vertex of the edge meets one of the shared vertices of the faces.

$$R8 = \begin{bmatrix} + & + & + \\ + & + & - \\ - & + & 0 \end{bmatrix} \qquad R9 = \begin{bmatrix} + & + & + \\ - & 0 & + \\ + & - & + \end{bmatrix} \qquad R10 = \begin{bmatrix} + & + & 0 \\ - & + & - \\ - & + & + \end{bmatrix}$$

Category 5 includes 3 representatives. In these, one vector is mapped to a vertex and the other two are mapped to edges. The representatives are rotations around the coordinate axes.

$$R11 = \begin{bmatrix} + & 0 & 0 \\ 0 & + & - \\ 0 & + & + \end{bmatrix} \qquad R12 = \begin{bmatrix} + & 0 & - \\ 0 & + & 0 \\ + & 0 & + \end{bmatrix} \qquad R13 = \begin{bmatrix} + & - & 0 \\ + & + & 0 \\ 0 & 0 & + \end{bmatrix}$$

Category 6 is the class of octahedral rotation. The single representative is the identity. All the vectors are mapped to vertices.

$$R14 = \begin{bmatrix} + & 0 & 0 \\ 0 & + & 0 \\ 0 & 0 & + \end{bmatrix}$$

We can show, conversely, that this list of classes is complete. Let Φ be a rotation matrix; then $\Phi(\hat{x}), \Phi(\hat{y})$, and $\Phi(\hat{z})$ must be orthogonal. We now carry out the following case analysis:

Case 1. At least one of $\Phi(\hat{x})$, $\Phi(\hat{y})$, $\Phi(\hat{z})$ lie on an octahedral vertex . Consider for example the case where $\Phi(\hat{z}) = \hat{z}$. Then either $\Phi(\hat{x})$ and $\Phi(\hat{y})$ lie on the positive or negative $\hat{x}$ and $\hat{y}$ axes (category 6) or $\Phi(\hat{x})$ and $\Phi(\hat{y})$ lie on equatorial edges $\pi/2$ apart (category 5). By the same token if any one of $\Phi(\hat{x})$, $\Phi(\hat{y})$, $\Phi(\hat{z})$ lie on an octahedral vertex, then Φ is either in category 5 or 6.

Case 2. Case (1) does not apply, but at least one of $\Phi(\hat{x})$, $\Phi(\hat{y})$, $\Phi(\hat{z})$ lies on an octahedral edge. If two of these lie on octahedral edges, then the third axis would have to lie on the orthogonal vertex, so we would be back in case 1. So the other two must both lie in the interior of the octants. Suppose, for example that $\Phi(\hat{x})$ lies on the equatorial edge between $\hat{x}$ and $\hat{y}$; that is $\text{Sign}(\hat{x}) = \langle +, +, 0 \rangle$. Since $\Phi(\hat{y})$ and $\Phi(\hat{z})$ are orthogonal to $\hat{x}$ and have non-zero z components, they must each have either the form $\langle +, -, ? \rangle$ or $\langle -, +, ? \rangle$. If they have the same x and y signs but opposite z signs, then this is category 4; if they have the same z sign and opposite x and y signs, then this is category 3. The analogous analysis applies whenever one of the coordinate axes is mapped to a point with one 0 coordinate.

Case 3. All three of $\Phi(\hat{x})$, $\Phi(\hat{y})$, $\Phi(\hat{z})$ lies in the interior of octants. No two can lie in the same octant, or they would be closer than $\pi/2$ and no two can lie in antipodal octants, or they would be farther than $\pi/2$. Therefore either there is one octant that meets each of the other two in an edge (category 1); or each of the three meets the other two in a vertex (category 2).

Categories 2 and 4 are symmetric in the three vertices, and hence give rise to only one representative. In the other four categories, the image of one of the vertices is in a distinguished state; hence these give rise to three representatives each, depending on which of the coordinate axes is mapped to the distinguished state. Therefore there are 14 representatives in all. Since each of the 24 octahedral rotations can be applied to each of the 14 representatives, there are a total of 336 different sign matrices, as stated above.

We can now use the following method to represent the sign matrices. We identify one representative $Q_1 \ldots Q_{14}$ from each of these equivalence classes, and then we can write any sign rotation matrix in the form $Q \cdot P$ for some representative Q and some octahedral rotation P. Since the octahedral rotations are easy to deal with, this greatly simplifies the analysis of the sign rotation matrices.

For any numerical matrix or sign matrix M and octahedral rotation R, the product $M \cdot R$ is simply a permutation of the columns of M, possibly with one or more change of sign, and the product $R \cdot M$ is a permutation of the rows of M, possibly with sign change. Likewise, the product $\hat{u} \cdot R$ is a permutation of the elements of $\hat{u}$, possibly with sign changes. In particular, all three products are always uniquely valued, even when $\hat{u}$ is a sign vector or M is a sign matrix.

4.2 Using Octahedral Symmetry

Having identified the octahedral rotations and a representative from each equivalence class, we can use this analysis to simplify the computation of operations on sign vectors and sign rotation matrices. In particular we define the representative sign vectors to be the vectors $\langle +, 0, 0 \rangle$, $\langle +, +, 0 \rangle$ and $\langle +, +, + \rangle$. We can then use the octahedral rotations to map any problem involving sign vectors and matrices to an equivalent problem involving the representative sign vectors and sign matrices. Thus the number of problems of the form $\hat{u} \cdot P$ is reduced from $26 \cdot 336$ to $3 \cdot 14$, a simplification by a factor of 208, and the number of problems of the form $P \cdot Q$ is reduced from $336 \cdot 336$ to $14 \cdot 14$, a simplification by a factor of 576.

Given the problem, "Compute the possible signs of $\hat{v} \cdot P$," where $\hat{v}$ is a sign vector and P is a sign matrix, we proceed as follows.

1. Let $\hat{u}$ be the representative sign vector with the same number of 0's as $\hat{v}$. Find an octahedral rotation R such that $\hat{v} = \hat{u} \cdot R$.

2. Compute the product $S = R \cdot P$. This is a rotation matrix.

3. Factor the matrix S as the product of a a rotation representative Q and a octahedral rotation R': $S = Q \cdot R'$.

4. Compute the possible values of $\hat{w} = \hat{u} \cdot Q$. Note that this is the product of a representative sign vector with a representative sign matrix.

5. Return $\hat{w} \cdot R'$. Note that $\hat{w} \cdot R' = \hat{u} \cdot Q \cdot R' = \hat{u} \cdot R \cdot P = \hat{v} \cdot P$.

Given the problem "Compute the possible signs of $P \cdot Q$" where P and Q are sign matrices, we proceed as follows:

1. Factor $P = P_1 \cdot R_1$ where P_1 is a representative matrix and R_1 is an octahedral rotation.

2. Compute $Q_1 = R_1 \cdot Q$.

3. Factor $Q_1 = Q_2 \cdot R_2$ where Q_2 is a representative sign matrix and R_2 is an octahedral rotation.

4. Compute $W = P_1 \cdot Q_2$. Note that this is the product of two representative matrices.

5. Return $W \cdot R_2$.

We have thus reduced these two problems to the following subproblems:

A. Multiplying a sign vector or a sign matrix by an octahedral rotation. As discussed above, this is simply a permutation plus changes of sign.

B. Factor a sign vector $\vec{v} = \hat{u} \cdot R$ where $\hat{u}$ is a representative sign vector and R is an octahedral rotation. This is easily precomputed; there are 26 cases.

C. Factor a sign matrix $P = P' \cdot R$ where P' is a representative matrix and R is an octahedral rotation. This can be precomputed by computing $Q' \cdot R$ for all pairs of a representative matrix P and an octahedral rotation R; there are 336 cases.

D. Multiply a representative sign vector by a representative sign matrix or multiply two representative sign matrices. These are precomputed; there are $3 \cdot 14 = 42$ and $14 \cdot 14 = 196$ cases respectively. However, the analysis here deserves its own section.

4.3 Multiplying representatives

We have thus reduced our two problems to the problem of multiplying a representative sign vector by a representative sign matrix and the problem of multiplying two representative matrices.

4.3.1 Multiplying a vector by a matrix

One obvious approach to these problems is to use the standard method for computing products, applying the sign calculus for combining signs. For example,

$$\text{let } \hat{v} = [+,+,0] \text{ and let } R4 = \begin{bmatrix} + & + & + \\ + & - & - \\ - & + & - \end{bmatrix}$$

Then multiplying in the usual way by taking the dot product of $\hat{v}$ with each column of $R4$, and using the sign arithmetic we get

$\hat{v} \cdot R4 =$
$[(+ \otimes +) \oplus (+ \otimes +) \oplus (0 \otimes -);$
$(+ \otimes +) \oplus (+ \otimes -) \oplus (0 \otimes -);$
$(+ \otimes +) \oplus (+ \otimes -) \oplus (0 \otimes -)] =$
$[+ \oplus + \oplus 0; + \oplus - \oplus 0; - \oplus + \oplus 0] = [+, I, I].$

In the above formula, we have used $\oplus$ and $\otimes$ for the addition and multiplication operators on signs and I as the symbol for "indefinite".

This calculation would suggest that the product may take on any of 9 sign values, all combinations of the three signs for each of the I values. However, this process loses information; all of these products would be attainable if $R4$ were an *arbitrary* matrix with these signs, but not if $R4$ is restricted to be a *rotation* matrix. In fact, of the 9 possible combinations, only 5 are actually possible. The calculation, however, obviously gives a *necessary* condition; no combination of signs that lies outside $[+, I, I]$ is possible.

The following theorem illustrates that the above product generates values that are in fact impossible:

Theorem 1 *If Sign*$(\hat{v}) = [+, +, 0]$ *and* M *is a rotation matrix such that Sign*$(M) = R4$*, then Sign*$(\hat{v} \cdot M)$ *is not* $[+, 0, 0], [+, 0, +], [+, -, 0]$ *or* $[+, -, +]$.

Proof: Let $\vec{v} = [v_1, v_2, 0]$ and let

$$M = \begin{bmatrix} x_1 & y_1 & z_1 \\ x_2 & -y_2 & -z_2 \\ -x_3 & y_3 & -z_3 \end{bmatrix}$$

where all the variables are positive. Thus $\hat{u} = \hat{v} \cdot M = [v_1x_1 + v_2x_2, v_1y_1 - v_2y_2, v_1z_1 - v_2z_2]$

Since M is orthogonal, the dot product of the third row with the first and second is 0 so

$$-x_2x_3 - y_2y_3 + z_2z_3 = 0 \rightarrow z_2z_3 > y_2y_3$$

$$-x_1x_3 + y_1y_3 - z_1z_3 = 0 \rightarrow y_3y_1 > z_3z_1$$

Multiplying these two inequalities we get $y_1y_3z_2z_3 > y_2y_3z_1z_3$ so $z_2/z_1 > y_2/y_1$

Now if $u_2 \leq 0$ we have $v_1y_1 \leq v_2y_2$, so $v_1/v_2 \leq y_2/y_1 < z_2/z_1$, so $v_1z_1 < v_2z_2$ so $u_3 < 0$. Therefore none of the signed vectors $[+, 0, 0], [+, 0, +], [+, -, 0]$, and $[+, -, +]$ can be the value of Sign$(\hat{u})$. ▮

Establishing that a particular sign vector $\hat{u}$ is a possible product of $\hat{v} \cdot M$ is done by finding numerical instantiations of the $\hat{u}$, $\hat{v}$ and M with the desired properties: The three instantiations must have the specified signs; $\hat{u}$ must be equal to $\hat{v} \cdot M$; and M must be an orthogonal matrix. We have implemented a search method that combines three techniques:

- If a variable has a non-zero sign, assign it a random numerical value with that sign.
- Value propagation: If there is an equation where all but one variable has been assigned a numerical value, then solve for the value of the remaining variable. We use equations of four kinds:
  - $\hat{u} \cdot M[:, j] = \hat{v}[j]$.
  - Each direction and each row and column of M has magnitude 1.
  - The dot product of any two rows/columns of M is zero.

  – Each row/column of M is equal to plus or minus the cross product of the other two rows/columns.

- Perturbation. Let $p \cdot A = q$ be a sign equation with $q[3] = 0$, and suppose that you have found a solution $\hat{u}, M$ where $\text{Sign}(\hat{u}) = p$, $\text{Sign}(M) = A$, $\text{Sign}(\hat{u} \cdot M) = q$. Let q' be the same as q except that $q'[3] \neq 0$. Then unless $\hat{u}[i] \cdot M[i, 3] = 0$ for all $i = 1, 2, 3$, it is possible to perturb $\hat{u}$ and M to a solution of the equation $p \cdot A = q'$. Similarly, one can perturb away from 0 in other positions in q and in p and A.

As a final step, $\hat{u}$ and M can be normalized so that $\hat{u}$ and each row and column of M have magnitude 1.

We have succeeded in finding, for every "equation" of the form $\hat{u} = \hat{v} \cdot R$, either an instantiation proving that the equation is satisfiable or an algebraic proof analogous to the one above proving that the equation is unsatisfiable. Over the space of 3 representative vectors $\hat{v}$ and 14 representative matrices M, there are a total of 154 such valid equations. See [1] for details.

4.4 Multiplying two matrices

The problem of finding all possible sign matrices that can be the values of a product $P \cdot Q$ of two particular representative matrices is much harder.

Three types of constraints are easily found:

1. The result $P \cdot Q$ must be one of the 336 sign rotation matrices.
2. If we decompose the matrix P into rows then the product $P \cdot Q$ decomposes into the individual products of each row of P with Q.

$$\text{If } P = \begin{bmatrix} \hat{v}_1 \\ \hat{v}_2 \\ \hat{v}_3 \end{bmatrix} \text{ then } P \cdot Q = \begin{bmatrix} \hat{v}_1 \cdot Q \\ \hat{v}_2 \cdot Q \\ \hat{v}_3 \cdot Q \end{bmatrix}$$

   We can then use the methods described in the previous section to limit the possible values of $\hat{v}_i \cdot Q$. (The rows $\hat{v}_i$ are not in general representative vectors, but we can use the octahedral rotations to transform the problem into one with representative vectors, as described earlier.)
3. We can write $P \cdot Q = (Q^T \cdot P^T)^T$, and then divide Q^T into rows and proceed as in part 2 above. (Again P^T may not be a representative matrix, but again that does not matter.)

All three of these give necessary conditions on the possible signs of the product, and, in general, they give different constraints, so applying them all limits the possible values of the signs. However, they do not together give sufficient conditions; a result may satisfy all three conditions and yet not be a possible product.

The algebraic analysis of this problem becomes quite formidable, and we have not found algebraic arguments analogous to those in the previous section that allow us to rule out additional impossible values. Rather, we have implemented a randomized procedure for searching for instantiations; details are given in [1]. If the procedure does not find an instantiation for a sign equation after a specified number of attempts, we presume that no such instantiation exists. The randomized method we used is similar to those discussed in [12], but tailored to our particular calculus.

These cases, where a possible value cannot be excluded by the above constraints and has not been instantiated by our search procedure, thus represent a gap in our analysis; in each individual case, it seems likely that there is actually no solution, though we would certainly hesitate to claim that there are none such in the whole collection. Out of 2782 equations that satisfy the above constraints, we have found instantiations for 2604; there are thus 178 unresolved cases. If these are indeed unsatisfiable, that may be difficult to verify: each of these 178 cases consists of a system combining 18 quadratic equations with 27 linear inequalities over 27 variables; the constraints, moreover, are not convex. For the time being, we have left these as open problems.

All the solutions we have found, both of vector times matrix and of matrix times matrix, have been saved in a database, together with one sample instantiation for each sign solution.

It may be noted that the standard vertex propagation and arc propagation algorithms used in reasoning about constraint networks in any case in general calculate supersets of the true class of possible labels. Therefore if such a constraint propagator were implemented using all the compositional triples that have not been excluded, the guarantees that could be stated would be the same, and the behavior would be almost the same, as it would be had we achieved a complete characterization of the composition table.

4.5 Code

The code used to find the solutions, a database containing the solutions of the representative problems, and a program with a user interface to find the solution to a specific problem, are available online at http://www.cs.nyu.edu/QualitativeRotations/. The code is 2000 lines of Java in total. The code is currently uncommented; however, we will be happy to discuss it and provide documentation if any readers wish to work with it.

5 Complexity and expressivity results

We have shown a number of complexity and expressivity results for existential languages over this theory. Proofs are given in the appendix. We begin by defining constraint networks generally, and qualitative rotation constraint networks (QRCNs) in particular.

Definition 1 *Let $\mathcal{D}$ be a domain. A* constraint network *N over $\mathcal{D}$ is a directed graph where each vertex corresponds to a variable that ranges over $\mathcal{D}$ and is labelled with a unary relation over $\mathcal{D}$ each arc is labelled with a binary relation over $\mathcal{D}$. An* instantiation *of N is an assignment of values in $\mathcal{D}$ to the nodes of N such that all the labels are satisfied.*

Definition 2 *A k-*dimensional qualitative rotation constraint network (QRCN) *is a constraint network over the domain of k-dimensional unit vectors in which the nodes are either unlabelled or labelled with a k-dimensional sign vector and the arcs are labelled with a $k \times k$ sign rotation matrix.*

We next define three restricted classes of QRCN's:

Definition 3 *A k-dimensional rotation Φ* lies in a coordinate plane *if all but 2 of the coordinates axes are invariant under Φ.*

Thus, a three-dimensional rotation lies in a coordinate plane if the axis of rotation is one of the coordinate axes. All two-dimensional rotations by definition lie in a coordinate plane.

Let Φ is a rotation that lies in a coordinate plane, let $q \in 1...m$ be the index of a coordinate axis that is invariant under Φ and let M be the sign matrix corresponding to Φ. Then $M[q,q] = +$ and for all $j \neq q$, $M[q,j] = M[j,q] = 0$. For instance, the sign matrix for a rotation around the $\hat{y}$ axis by an angle between 0 and $\pi/2$ is

$$\begin{bmatrix} + & 0 & - \\ 0 & + & 0 \\ + & 0 & + \end{bmatrix}$$

Definition 4 *A k-dimensional QRCN is* pure *if every arc is labelled with a sign matrix corresponding to a rotation lying in a coordinate plane.*

Definition 5 *A QRCN is* non-negative *if all the coordinates of all the node variables are constrained to be non-negative.*

For example, in three-dimensions, each vertex variable $\hat{u}$ has one of the seven signs $\langle +,+,+\rangle$, $\langle +,+,0\rangle$, $\langle +,0,+\rangle$, $\langle 0,+,+\rangle$, $\langle +,0,0\rangle$, $\langle 0,+,0\rangle$ or $\langle 0,0,+\rangle$.

We can now categorize the complexity of the satisfiability problem for a number of classes of QRCN:

Theorem 2 *The problem of determining whether a QRCN is satisfiable is*

A. Polynomial time for pure non-negative k-dimensional QRCNs, $k \geq 2$.

B. NP-complete for two-dimensional QRCNs.

C. NP-complete for pure k-dimensional QRCNs, $k \geq 3$.

D. NP-hard for general k-dimensional QRCNs, $k \geq 3$.

The proofs are in appendix A.1. It is an open problem whether the satisfiability problem for general k-dimensional QRCN's is in NP. We conjecture that it is.

We next turn to two expressivity results. We begin by defining the concepts of two values being distinguishable by a constraint network, and of a value being uniquely identifiable.

Definition 6 *Let $\mathcal{D}$ be a domain and let $\mathcal{N}$ be a set of constraint networks over $\mathcal{D}$. A value $a \in \mathcal{D}$ is* distinguishable *from $b \in \mathcal{D}$ by constraint networks in $\mathcal{N}$ iff there exists a network $N \in \mathcal{N}$ and a node V in N such that*

- *there exists an instantiation of N in which the value of V is a; and*
- *there does not exist an instantiation of N in which the value of V is b.*

Definition 7 *Let $\mathcal{D}$ and $\mathcal{N}$ be as in definition 6. A value $a \in \mathcal{D}$ is* uniquely identifiable *by constraint networks in $\mathcal{N}$ if there exists a network $N \in \mathcal{N}$ and node V in N such that every instantiation of N assigns a to V.*

This is a special case of the more general notion of a *definable relation* presented in [13], definition 3.

Theorem 3 *In dimension 3 or higher, any direction is distinguishable from any other direction by pure QCRN's.*

Theorem 4 *In dimension 3 or higher, the set of directions uniquely identifiable by pure QRCN's is dense in the unit sphere.*

Essentially, rotations in $\mathcal{Q}$ allow one to construct circles on the unit sphere centered at the coordinate axes, and the proof consists in showing that there are uniquely identifiable patterns among these circles. The proof is given in appendix A.2.

An interesting class of QRCN's in dimensions ≥ 3 for which we have practically no such results are QRCN's in general position:

Definition 8 *A sign rotation matrix is in* general position *if none of its elements are 0. A QRCN is in general position if every arc is labelled by a matrix in general position.*

For QRCN's in general position the analogue of theorem 4 fails in the strongest possible way; any constraint network defines an open neighborhood in the space of valuations and thus there are *no* uniquely identifiable directions. However we do not know anything about the analogue of theorem 3. We also do not know anything about the complexity of the satisfiability problem for QRCN's in general positions, in dimensions ≥ 3 (except that it is easily shown to be in the existential theory of real arithmetic, and thus decidable). We conjecture, with no great confidence, that no two directions with the same sign vector are distinguishable by general position QRCN's and that the decision problem is NP-complete.

6 Future Work

The immediate problems to be addressed would be to close the gaps in our analysis of signed matrix composition, and to extend the complexity/expressivity results discussed in the previous section.

It would also be helpful to be able to generate *random* solutions to a specified sign equation. It is not very important to achieve any specific distribution (e.g. uniform), but it would be desirable to have a program that at least achieves reasonable coverage; that is, for any given sign equation Q and for any instantiation I satisfying Q, in some reasonable number of trials the program outputs a solution to Q fairly close to I. The use of perturbation in the current technique for finding instantiations has the consequence that, in some cases, instantiations far from 0 may not be found.

More interestingly, the theory could be extended to other tilings of the unit sphere beyond the octahedral, though the sign calculus would no longer be applicable. Indeed a very similar theory can be developed for any of the regular (Platonic) polyhedra. Consider, for example, a fixed icosahedron centered at the origin, with individual names assigned to each vertex, edge, and face. A direction can be qualitatively characterized by identifying the vertex, edge, or face it passes through. A rotation can be characterized as follows. Pick a particular face of the icosahedron, and let $\hat{p}$, $\hat{q}$, $\hat{r}$ be the vertices of that face. Then characterize a rotation Γ in terms of the qualitative characterization of $\Gamma(\hat{p}), \Gamma(\hat{q})$, and $\Gamma(\hat{r})$.

The resultant theory is then very similar to the one we have developed here. The symmetry group of the icosahedron has 60 elements. There are still three representative qualitative directions, corresponding to vertex, edge, and face. By our count, there are 20 representative rotations in 8 categories (figure 5).

The geometric calculations involved are, of course, much harder than in our analysis.

Qualitative representations based on the Platonic solids can only achieve a fixed level of granularity, as there are only five Platonic solids. One can achieve arbitrary levels of granularity by using a finer

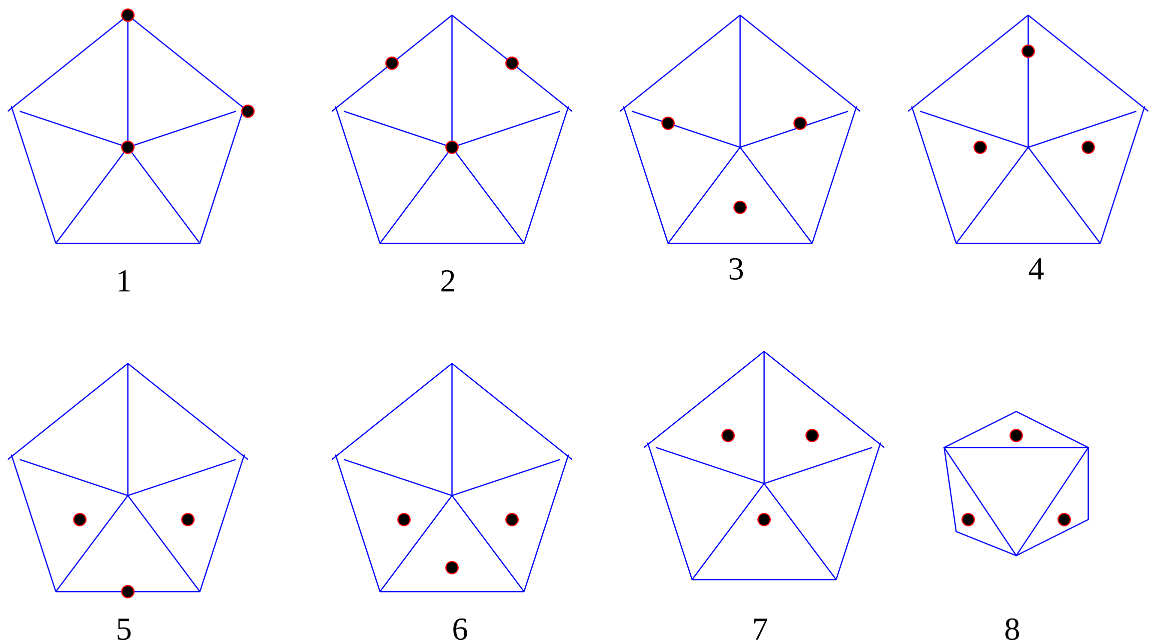

Figure 5: Categories of icosahedral rotations. The lines are edges of the standard icosahedron; the red dots are the images of $\hat{x}, \hat{y}, \hat{z}$

tiling of the unit sphere, at the cost of losing some of the symmetries. Choose a reference tiling, and three reference directions $\hat{p}, \hat{q}, \hat{r}$; characterize a direction in terms of the element of the tiling that it lies in; and characterize a rotation in terms of the characterizations of the images of $\hat{p}, \hat{q}, \hat{r}$.

A final problem is to integrate a theory of translation so that, for example, in the case of the robotic arm discussed in section 3 one can reason qualitatively about the positions of the segments as well as their orientation.

References

[1] A. Asl, A Qualitative Calculus for Three-Dimensional Rotations, MS Thesis, Computer Engineering, NYU Polytechnic, 2011.

[2] R.A. Brooks, Symbolic reasoning among 3-D models and 2-D images, *Artificial Intelligence,* Vol. 17:1-3, 1981, pp. 285-348.

[3] A.G. Cohn and J. Renz, Qualitative spatial reasoning in F. van Harmelen, V. Lifschitz, and B. Porter (eds.) *Handbook of Knowledge Representation,* Elsevier, 2007, 551-596.

[4] E. Davis, *Linear Algebra and Probability for Computer Science Applications*, CRC Press 2012.

[5] J. de Kleer and J.S. Brown, A qualitative physics based on confluences, in D. Bobrow (ed.) *Qualitative Reasoning about Physical Systems,* Cambridge, Mass.: MIT Press, 1985, 7-84.

[6] A. Frank, Qualitative spatial reasoning: cardinal directions as an example *International Journal of Geographical Information Systems,* **10**:3, 1996.

[7] Antony Galton, *Qualitative Spatial Change,* Oxford University Press, 2000.

[8] R.K. Goyal and M.J. Egenhofer, The direction-relation matrix: A representation of direction relations for extended spatial objects. *UCGIS Annual Assembly and Summer Retreat, Bar Harbor, ME,* 1997.

[9] K. Hsiao, L. Kaelbling, and T. Lozano-Pérez, Robust grasping under object pose uncertainty *Autonomous Robots,* **31**(2-3), 2011, 253-268.

[10] N. Karmarkar, A new polynomial-time algorithm for linear programming, *16th Symposium on Theory of Computing (STOC)*, 1984.

[11] H. Kim, Qualitative Kinematics of Linkages, In P. Struss and B. Faltings (eds.) *Recent Advances in Qualitative Physics,* Cambridge, Mass.: MIT Press, 1992.

[12] W. Liu and S. Li, On a Semi-Automatic Method for Generating Composition Tables, arXiv:1105.4224v1, May 2011.

[13] W. Liu and S. Li, Reasoning about cardinal directions between extended objects: The NP-hardness result, *Artificial Intelligence* **175**:18 2129-2222

[14] A. Mackworth and E. Freuder, The complexity of some polynomial network consistency algorithms for constraint satisfaction problems *Artificial Intelligence,* **25:**1, 1985, 65-74.

[15] D. McDermott and E. Davis, Planning and executing routes through uncertain territory, *Artificial Intelligence,* **22**, 1984, 107-156.

[16] R. Moratz, F. Dylla and L. Frommberger, A relative orientation algebra with adjustable granularity, *Proc. Workshop on Agents in Real-Time and Dynamic Environments, IJCAI,* 2005.

[17] R. Moratz and M. Ragni, Qualitative spatial reasoning about relative point position, *Journal of Visual Languages and Computing*, **19**:1, 2008, 75-98.

[18] T. Mossakowski and R. Moratz, Qualitative reasoning about relative direction of oriented points, *Artificial Intelligence* 180-181, 2012, 34-45.

[19] D. Randell, Z. Cui, and A.G. Cohn, A spatial logic based on regions and connection. *KR*, 1992, 165-176.

[20] J. Renz and D. Mitra, Qualitative direction calculi with arbitrary granularity, *Proc. PRICAI*, *LNAI* 3157, 2004, 65-74.

[21] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach* (3rd edn), Prentice Hall, 2010.

[22] A. Scivos and B. Nebel, The finest of its class: The practical natural ternary calculus $\mathcal{L}R$ for qualitative spatial reasoning in C. Freksa et al. (eds.) *Spatial Cognition IV: Reasoning, Action, Interaction,* LNAI vol..3343, Springer, 2005, 283-303.

[23] M. Westphal et al. Guiding the Generation of Manipulation Plans by Qualitative Spatial Reasoning. *Spatial Cognition and Computation.* **11:**(1) 2011, 75-102.

[24] D. Wolter and J.H. Lee, Qualitative reasoning with directional relations, *Artificial Intelligence,* **174**, 2010, 1498-1507.

[25] X. Zhang et al. Reasoning with Cardinal Directions: An Efficient Algorithm, *AAAI*, 2008, 387-392.

Appendix A: Proofs of theorems

This appendix contains the proofs of the theorems discussed in section 5.

Appendix A.1. Complexity

We begin by establish the NP-completeness of the satisfiability problems for two-dimensional QRCNs in which every arc is labelled with the sign matrix

$$\begin{bmatrix} + & - \\ + & + \end{bmatrix}$$

It is convenient to associate unit vectors with their angle from the $\hat{x}$ axis, and thus with a number in the interval $[0, 2\pi)$. We can then restate the problem as follows:

> Given:
>
> - A set of variables $U = \{u_1 \dots u_m\}$ ranging over the half open interval $[0, 2\pi)$.
> - A collection of constraints of the form $0 < (u_j - u_i) \mod 2\pi < \pi/2$. That is, the positive angle from u_i to u_j is positive and less than 90°.
>
> Is there a valuation of the variables satisfying all the constraints?

For the remainder of section A.1, a "constraint" means a constraint of this form, unless otherwise specified. The "qualitative two-dimensional rotation problem" is the satisfiability problem for a system of such constraints. We prove that this problem is NP-complete.

A.1.1: Preliminaries

We abbreviate the constraint $0 < (u_j - u_i) \mod 2\pi < \pi/2$ in the form $u_i \rightarrow u_j$, and view a system of such constraints as defining a directed graph, where the vertices are the variables and the arcs are the constraints.

Any value $u_i \in [0, 2\pi)$ has a unique expression in the form $u_i = k_i \cdot \pi/2 + f_i$ where $k_i \in \{0, 1, 2, 3\}$ and $0 \leq f_i < \pi/2$.

Lemma 1 *Let $u_i = k_i \cdot \pi/2 + f_i$ and $u_j = k_j \cdot \pi/2 + f_j$. Then*

$$0 < (u_j - u_i) \text{ mod } 2\pi < \pi/2 \text{ iff either } [k_j = k_i \text{ and } f_i < f_j] \text{ or } [k_j = k_i + 1 \text{ mod } 4 \text{ and } f_j < f_i]$$

Proof: Immediate.

The *cardinal directions* are the values $0, \pi/2, \pi$, and $3\pi/2$.

Throughout this section, let $0 < \epsilon < \pi/10$ be a fixed small value. The proof actually goes through for any value of ϵ between 0 and $\pi/4$, but in visualizing the construction, it is helpful to think of ϵ as quite small.

Definition 9 *A valuation $\phi(u)$ over the set of variables U is an* ϵ-valuation *if, for every $u \in U$, $\phi(u)$ is within ϵ following one of the cardinal directions. That is, $0 < \phi(u) - k\pi/2 < \epsilon$ for some $k \in \{0, 1, 2, 3\}$.*

Lemma 2 *If a system of constraints has a satisfying valuation then for any $\epsilon > 0$ it has a satisfying ϵ-valuation.*

Proof: Let $\psi(u)$ be a satisfying valuation. For $i = 1 \ldots m$ let $\phi(u_i) = k_i \cdot \pi/2 + f_i$. Sort the collection of numbers f_i and let σ be the corresponding permutation of indices. Thus $\sigma(i) < \sigma(j)$ if and only if $f_i < f_j$.

Define $\phi(u_i) = k_i\pi/2 + \epsilon\sigma(i)/(n+1)$. It is easily verified that ϕ is an ϵ valuation and that $\phi(u_i)$ and $\phi(u_j)$ satisfy the condition of lemma 1 if and only if $\psi(u_i)$ and $\psi(u_j)$ satisfy the same condition. ∎

Corollary 1 *The qualitative two-dimensional rotation problem is in NP.*

Proof: Take $\epsilon = \pi/8$. If there exists a satisfying valuation then the construction in lemma 2 shows that there exists a valuation where every variable has a value of the form $k_i\pi/2 + (m\pi/10(n+1))$ for $m = 1 \ldots n$. This valuation is a witness; checking that the conditions of lemma 1 are satisfied is just a matter of checking the inequalities.

Corollary 2 *The satisfiability problem for two-dimensional QRCN's is in NP.*

Proof: We can modify the proof of lemma 2 to include rotations that are an integer multiple of $\pi/2$ or that lie in one of the intervals $(\pi/2, \pi)$, $(\pi, 3\pi/2)$ or $(3\pi/2, 2\pi)$. The argument that a valuation satisfying a two-dimensional QCRN exists if and only if an ϵ valuation exists is essentially the same as the proof of lemma 2. By the same argument as in corollary 1, such an ϵ valuation is a witness; it is straightforward to check that it satisfies all the constraints. ∎

In view of lemma 2 we can restrict attention to ϵ-valuations, and we do so for the remainder of this section. We also eliminate the explicit reference to the valuation, and use the symbol u_i to refer either to the variable or its value, in the usual way.

Given an ϵ-valuation, we can write each value u_i in the form $u_i = q_i \cdot \pi/2 + e_i$, where $q_i \in \{0, 1, 2, 3\}$ and $0 < e_i < \epsilon$. The number q_i is called the *quadrant* of u_i.

Note that in an ϵ-valuation, if $u_i \rightarrow u_j$ then one of two things are true:

- $q_j = q_i$ and $0 < e_j - e_i < \epsilon$. In this case, we say there is a *short arc* from e_i to e_j.
- $q_j = q_i + 1 \mod 4$ and $-\epsilon < e_j - e_i < 0$. In this case, we say there is a *long arc* from e_i to e_j.

A.1.2: Qualitative two-dimensional rotation is NP-hard

We now prove that qualitative two-dimensional rotation is NP-hard. The proof is by reduction from graph 4-coloring. We build up the construction in pieces.

Lemma 3 *Consider a system of constraints that forms a five cycle $u \rightarrow v \rightarrow w \rightarrow x \rightarrow y \rightarrow u$. Any ϵ-valuation satisfying this system:*

a. goes once around the unit circle.

b. has one short arc and four long arcs.

Proof: Since it starts and ends with u, it must go around the unit circle at least once. Since each of the arcs is less than $\pi/2$, their total must be less than $5\pi/2$, so it cannot go around the unit circle twice. Since it goes once around the unit circle, it must transition from one quadrant to the next exactly four times; these are the four long edges.

In fact the above generalizes to *any* cycle of edges that goes once around the circle

Lemma 4 *Let* $u_1 \rightarrow \ldots \rightarrow u_n \rightarrow u_1$ *be a cycle of n-edges and consider a* ϵ*-valuation in which this cycle goes once around the unit circle. Then there are 4 long edges in the cycle (those that transition from one quadrant to the next) and* $n-4$ *short edges (those that stay in the same quadrant).*

Proof: Since the cycle goes once around the unit circle, it transits from one quadrant to the next exactly four times. By definition of an ϵ-valuation, these are the long edges. The rest are short edges. ∎

We now give a construction with many variables that goes once around the unit circle.

Definition 10 *For any* $m \geq 2$ *the* simple $5m$-cycle *is defined as the following system of variables and constraints*

- *There are* $5m$ *variables arranged in a cycle:* $u_0 \rightarrow u_1 \rightarrow \ldots \rightarrow u_{5m-1} \rightarrow u_0$ *This is the* large cycle *of the construction*
- *The variables* $u_{mi}, i = 0 \ldots 4$ *are ordered in a 5-cycle:* $u_0 \rightarrow u_m \rightarrow u_{2m} \rightarrow u_{3m} \rightarrow u_{4m} \rightarrow u_0$*. These are the* keystone *variables.*
- *The variables* u_j *are constrained to be within* $\pi/2$ *of the preceding keystone. That is, for* $i = 0 \ldots 4$*, for* $1 \leq j \leq m-1$*, there is an arc* $u_{mi} \rightarrow u_{mi+j}$*.*

Figure 6 shows the simple 15-cycle.

Lemma 5 *In any valuation satisfying the simple* $5m$*-cycle, the large cycle goes once around the unit circle.*

Proof: The 5-cycle of keystone variables must go once around the unit circle, and the remaining variables are constrained to lie between them in sequence. ∎

We next give a construction that forces a pair of short edges to lie in different quadrants.

Definition 11 *Let* $a \rightarrow b$ *and* $c \rightarrow d$ *be arcs in a system of constraints.*
The cycle $a \rightarrow b \rightarrow e \rightarrow f \rightarrow c \rightarrow d \rightarrow g \rightarrow h \rightarrow a$ *is called a* separator *for the two arcs* $a \rightarrow b$ *and* $c \rightarrow d$*.*

Figure 7 illustrates a separator.

Lemma 6 *Suppose that a system of constraints contains two arcs* $a \rightarrow b$ *and* $c \rightarrow d$ *and contains a separator for these. Consider any* ϵ*-valuation in which* $a \rightarrow b$ *and* $c \rightarrow d$ *are both short arcs. Then these two arcs are in different quadrants.*

Proof: Note that the separator is an 8-cycle, and thus cannot go twice around the unit circle, since every arc is less than $\pi/2$. Therefore, it goes once around the unit circle. By lemma 4, four of the arcs are long. Therefore at least one of the arcs between b and c is long and at least one of the arcs between d and a is long. ∎

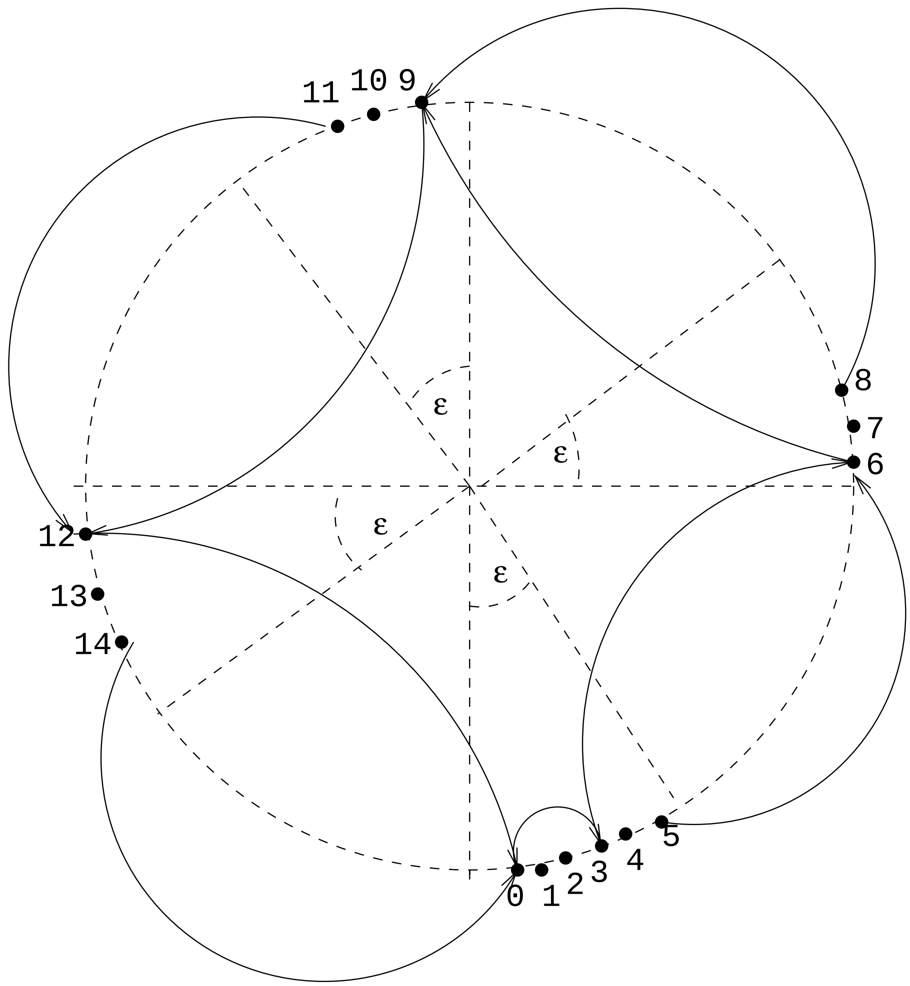

Short arcs are not shown, except 0→3

Figure 6: Simple 15-cycle

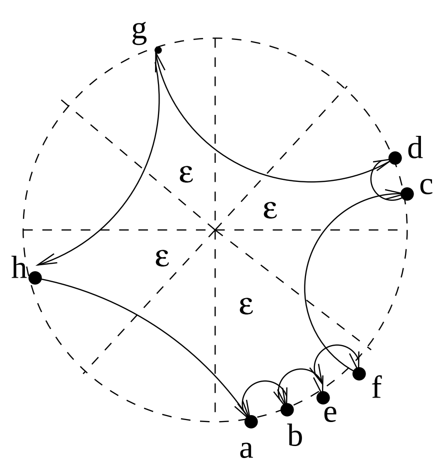

Figure 7: Separator

Lemma 7 *Let $a \rightarrow b \rightarrow e \rightarrow f \rightarrow c \rightarrow d \rightarrow g \rightarrow h \rightarrow a$ be a separator for $a \rightarrow b$ and $c \rightarrow d$. Let i be 1,2, or 3. Let ϕ be an ϵ-valuation over a, b. Then there exists an ϵ-valuation ϕ' over $\{a, \ldots, h\}$ extending ϕ in which in which $c = (a + i\pi/2)$ mod 2π and $d = (b + i\pi/2)$ mod 2π*

Proof: Without loss of generality, assume that a is in the 0 quadrant (just to save writing "mod 2π" a lot of times). For brevity, we omit the function ϕ' and just write x for $\phi'(x)$. There are six cases:

1.A $0 < a < b < \epsilon$, $i = 1$. Let $c = a + \pi/2$, $d = b + \pi/2$, $e = (2b + \epsilon)/3$, $f = (b + 2\epsilon)/3$, $g = (a + 2b)/3 + \pi$, $h = (2a + b)/3 + 3\pi/2$. (Shown in figure 7).

1.B $0 < a < b < \epsilon$, $i = 2$. Let $c = a + \pi$, $d = b + \pi$, $e = (b + 2\epsilon)/3$, $f = (2b + \epsilon)/3 + \pi/2$, $g = (b + 2\epsilon)/3 + \pi$, $h = (2b + \epsilon)/3 + 3\pi/2$.

1.C $0 < a < b < \epsilon$, $i = 3$. Let $c = a + 3\pi/2$, $d = b + 3\pi/2$, $e = (b + 2\epsilon)/3 + \pi/2$, $f = (2b + \epsilon)/3 + \pi$, $g = (2b + \epsilon)/3 + 3\pi/2$, $h = (b + 2\epsilon)/3 + 3\pi/2$.

2.A $0 < a < \epsilon$, $\pi/2 < b < a + \pi/2$, $i = 1$. Let $c = a + \pi/2$, $d = b + \pi/2$, $e = (2b + c)/3$, $f = (b + 2c)/3$, $g = (a + 2\epsilon)/3 + \pi$, $h = (2a + \epsilon)/3 + 3\pi/2$.

2.B $0 < a < \epsilon$, $\pi/2 < b < a + \pi/2$, $i = 2$. Let $c = a + \pi$, $d = b + \pi$, $e = (2a + \epsilon)/3 + \pi/2$, $f = (a + 2\epsilon)/3 + \pi/2$, $g = (2a + \epsilon)/3 + 3\pi/2$, $h = (a + 2\epsilon)/3 + 3\pi/2$.

2.C $0 < a < \epsilon$, $\pi/2 < b < a + \pi/2$, $i = 3$. Let $c = a + 3\pi/2$, $d = b - \pi/2$, $e = (a + 2\epsilon)/3 + \pi/2$, $f = (2a + \epsilon)/3 + \pi$, $g = (2d + a)/3$, $h = (d + 2a)/3$.

It is straightforward to check that in each case the separator constraints are satisfied.

We now have all the pieces, and we can put them together into the construction.

Lemma 8 *There is a polynomial-time reduction from graph 4-coloring to qualitative two-dimensional rotation.*

Proof: Let G be the graph to be 4-colored. Let m be the number of vertices in G. We construct the following system of constraints:

- For each vertex v of G, construct a simple $5m$-cycle with variables $v_0 \ldots v_{5m-1}$.
- For each edge $v - w$, for each pair of corresponding arcs in the large cycle $v_i \rightarrow v_{i+1}$ and $w_i \rightarrow w_{i+1}$, create a separator between these two arcs.

We need to show that:

I. If the system of constraints is satisfiable, then the graph is 4-colorable.

Proof: If the system is satisfiable, it has an ϵ-valuation.

I claim that, in that valuation, there is at least one index i (in fact, at least m indices) for which $v_i \rightarrow v_{i+1}$ is a short arc for all the variables. Proof: Each variable v can supply long arcs for only 4 indices; hence at most $4m$ indices are associated with any long arcs; hence m indices are associated with only short arcs.

Let i be an index with only short arcs. Let $u - v$ be an edge in graph G. By construction, there is a separator for $u_i \rightarrow u_{i+1}$ and $v_i \rightarrow v_{i+1}$. By lemma 6 u_i and v_i are in different quadrants. Hence, if we color each vertex u in the graph by the quadrant of u_i, we have a four-coloring of G.

II. If the graph is 4-colorable then the system of constraints is satisfiable.

Proof: Assume that the colors are $0 \ldots 3$. If vertex u is colored 0, then

- for $i = 0 \ldots 2m-1$, $u_i \leftarrow (i+1)\epsilon/(2m+1)$
- for $i = 2m \ldots 3m-1$, $u_i \leftarrow \pi/2 + (i+1-2m)\epsilon/(m+1)$
- for $i = 3m \ldots 4m-1$, $u_i \leftarrow \pi + (i+1-3m)\epsilon/(m+1)$
- for $i = 4m \ldots 5m-1$, $u_i \leftarrow 3\pi/2 + (i+1-4m)\epsilon/(m+1)$

It is easily verified that all the constraints of the $5m$ cycle are satisfied.

If vertex u is colored $c = 1$, 2, or 3, then rotate this assignment by $c\pi/2$. Clearly the $5m$ cycle constraints are still satisfied. Moreover since every pair of corresponding arcs for variables connected by arcs are place $i\pi/2$ apart for $i = 1$, 2, or 3, by lemma 7 the separator variables can be assigned so as to satisfy the separator constraints. ∎

Theorem 5 *The qualitative two-dimensional rotation problem is NP-complete.*

Proof: Corollary 1 and lemma 8.

Corollary 3 *The satisfiability problem for three-dimensional QRCN's is NP-hard.*

Proof: The qualitative two-dimensional rotation problem is equivalent to the special case of the three-dimensional QRCN where vertices are unlabelled, and every rotation is a positive rotation between 0 and $\pi/2$ around the $\hat{z}$ axis; that is, the sign matrix of the rotation is

$$\begin{bmatrix} + & - & 0 \\ + & + & 0 \\ 0 & 0 & + \end{bmatrix}$$

∎

Whether this problem is in NP, we do not know, though we conjecture that it is.

A.1.3: Tractability

Theorem 6 *The satisfiability problem over the class of pure non-negative k-dimensional QRCN's is decidable in polynomial time.*

Proof: Note that for $k = 2$ a pure non-negative QRCN is simply a collection of order constraints over the angles, which can be solved by topological sort. We therefore assume that $k \geq 3$.

We give the proof for the case $k = 3$; the proof for $k > 3$ is essentially identical.

1. A positive vector $\vec{u} = \langle u_x, u_y, u_z \rangle$ is a unit vector if and only if $u_x^2 + u_y^2 + u_z^2 = 1$.

2. Let $\hat{u} = \langle u_x, u_y, u_z \rangle$ and $\hat{v} = \langle v_x, v_y, v_z \rangle$ be positive unit vectors. Then there is a positive rotation around the $\hat{x}$ axis mapping $\hat{u}$ into $\hat{v}$ if and only if $v_x^2 = u_x^2$, $v_y^2 < u_y^2$ and $v_z^2 > u_z^2$. Likewise there is a positive rotation around the $\hat{y}$ axis mapping $\hat{u}$ into $\hat{v}$ if and only if $v_y^2 = u_y^2$, $v_z^2 < u_z^2$ and $v_x^2 > u_x^2$; and there is a positive rotation around the $\hat{z}$ axis mapping $\hat{u}$ into $\hat{v}$ if and only if $v_z^2 = u_z^2$, $v_x^2 < u_x^2$ and $v_y^2 > u_y^2$.

Therefore, the QRCN is equivalent to a set of linear inequalities over the quantities u_x^2, u_y^2, u_z^2, where u ranges over all the variables. But the satisfiability of a set of linear inequalities (linear programming) is decidable in polynomial time [10]. ∎

Theorem 7 *The satisfiability problem for a pure k-dimensional QRCN is NP-complete.*

Proof: That the problem is NP-hard follows from lemma 8, since the 2-dimensional qualitative rotation problem is just the special case of the pure QRCN problem, where the arcs are restricted to be rotations around the $\hat{z}$ axis.

To show that the problem is NP-easy, note that if the solver "guesses" the sign vector of each of the vector variables, then the rest of the content of the network can be expressed as a set of linear inequalities over the squares of the coordinates of the vectors, as in the proof of theorem 6. ∎

At this point we have all the significant upper and lower complexity bounds for theorem 2 of section 5; the remaining bounds follow from one problem being a special case of another.

Appendix A.2: Expressivity results

In this section, we prove the following two expressivity results:

Theorem 3: *In dimension $k \geq 3$, any direction is distinguishable from any other direction by pure QRCNs.*

Theorem 4: *In dimension $k \geq 3$, the set of directions uniquely identifiable by pure QRCNs is dense in the unit sphere.*

Note that at most countably many directions can be uniquely identifiable since there are only countably many different networks, each with finitely many nodes.

Clearly, a uniquely identifiable direction is distinguishable from any other direction.

Note that the situation in two dimensions is very different; it is easily proved that no two directions with the signs [+,+] are distinguishable by two-dimensional QRCNs and *a fortiori* that no such direction is uniquely identifiable. Proof: Any monotonic bijection from the interval $(0, \pi/2)$ to itself preserves the signs of all rotations.

We give the proofs of these in three-dimensions; since the three-dimensional case is a special case of the higher-dimensional case, this also establishes the result for higher dimensions.

We begin by constructing a particular network $N1$ that allows us to construct the vector $[\sqrt{2}/2, \sqrt{2}/2, 0]$.

Let $N1$ be the network with

- Three nodes $\hat{a}, \hat{b}, \hat{c}$.
- The sign labels $\hat{a} = [+, 0, +]$; $\hat{b} = [0, +, +]$; $\hat{c} = [+, +, 0]$.
- Arcs $\hat{a} \rightarrow \hat{b}$, $\hat{b} \rightarrow \hat{c}$, $\hat{c} \rightarrow \hat{a}$ with the following labels.

$$\hat{b} = \hat{a} \cdot \begin{bmatrix} 0 & 1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} \qquad \hat{c} = \hat{b} \cdot \begin{bmatrix} 0 & 0 & -1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix} \qquad \hat{a} = \hat{c} \cdot \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 1 & 0 \end{bmatrix}$$

That is: $\hat{a}$ is on the longitudinal arc from $\hat{x}$ to $\hat{z}$. $\hat{b}$ is on the longitudinal arc from $\hat{y}$ to $\hat{z}$. $\hat{c}$ is on the equatorial arc from $\hat{x}$ to $\hat{y}$. Also $\hat{b}$ is the image of $\hat{a}$ under a rotation around the $\hat{z}$ axis of $\pi/2$.

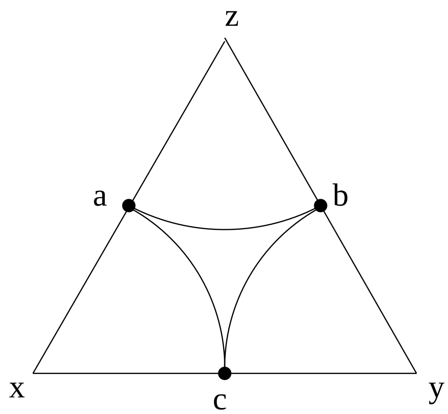

Figure 8: Lemma 9

Direction $\hat{c}$ is the image of $\hat{b}$ under a rotation around the $\hat{y}$ axis of $\pi/2$. Direction $\hat{a}$ is the image of $\hat{c}$ under a rotation around the $\hat{x}$ axis of $\pi/2$.

The construction is shown in figure 8. In reading figures 8, 9, and 10 it should be kept in mind that, though the triangle x, y, z is displayed as an equilateral triangle, it is in fact an octant on a sphere, and that the angles at x, y, and z are all right angles

Throughout this, for any directions $\hat{a}, \hat{b}$, the notation $\angle \hat{a}\hat{b}$ is the angle between $\hat{a}$ and $\hat{b}$ as measured from the center of the sphere (the origin). This of course is equal to the length of the geodesic from $\hat{a}$ to $\hat{b}$ on the sphere, and so obeys the triangle inequality.

Lemma 9 *In the network $N1$ defined above: Direction $\hat{a}$ is the midpoint of the longitudinal arc from $\hat{x}$ to $\hat{z}$; that is, $[\sqrt{2}/2, \sqrt{2}/2, 0]$. Direction $\hat{b}$ is the midpoint of the longitudinal arc from $\hat{y}$ to $\hat{z}$. Direction $\hat{c}$ is the midpoint of the equatorial arc from $\hat{x}$ to $\hat{y}$.*

Proof: Because of the rotations we have $\angle \hat{z}\hat{a} = \angle \hat{z}\hat{b}$; $\angle \hat{y}\hat{b} = \angle \hat{y}\hat{c}$; $\angle \hat{x}\hat{c} = \angle \hat{x}\hat{a}$. Also $\angle \hat{z}\hat{a} + \angle \hat{x}\hat{a} = \pi/2$; $\angle \hat{z}\hat{b} + \angle \hat{y}\hat{b} = \pi/2$; $\angle \hat{x}\hat{c} + \angle \hat{y}\hat{c} = \pi/2$.

So all these angles are equal to $\pi/4$. ∎

For any two orthogonal directions $\hat{a}, \hat{b}$ and for any $r \in [0, 1]$, let $\hat{\phi}(\hat{a}, \hat{b}, r)$ be the direction on the arc from $\hat{a}$ to $\hat{b}$, a fraction r of the way along; that is,

$$\angle \hat{a}, \hat{\phi}(\hat{a}, \hat{b}, r) = r \cdot \pi/2;$$
$$\angle \hat{b}, \hat{\phi}(\hat{a}, \hat{b}, r) = (1 - r) \cdot \pi/2.$$

Definition 12 *A real-valued function Γ is* contracting *if, for all a, b, $|\Gamma(a) - \Gamma(b)| < |a - b|$. (In other words, it obeys a Lipschitz condition with constant less than 1.)*

It is immediate that any contracting function is continuous.

Lemma 10 *There exists a function $\Gamma : (0, 1) \rightarrow (0, 1)$ and a network $N_{x,y}$ with two distinguished nodes P and Q with the following properties:*

A. Γ is continuous and strictly monotonic; that is, if $a < b$ then $\Gamma(a) < \Gamma(b)$.

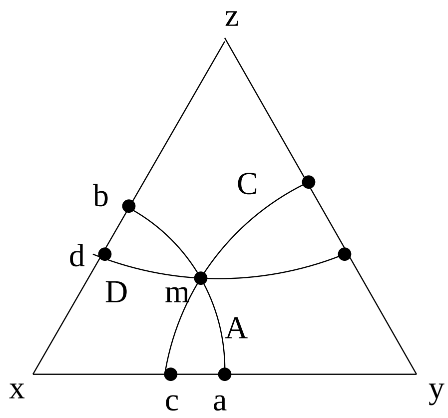

Figure 9: Lemma 10

B. $lim_{r \to 0^+} \Gamma(r) = 0$.

C. $lim_{r \to 1^-} \Gamma(r) = 1/2$.

D. Γ *is contracting.*

E. *For any* $r \in (0,1)$ *if* P *has the value* $\hat{\phi}(\hat{x}, \hat{y}, r)$ *then* Q *must have the value* $\hat{\phi}(\hat{x}, \hat{y}, \Gamma(r))$ *in any assignment satisfying* N*.*

Condition (e) above essentially states that the network N computes the function Γ.

Proof: The network implements the following system of geometric constraints (figure 9):

1. Let $\hat{a}$ be the input value $\hat{\phi}(\hat{x}, \hat{y}, r)$.
2. Draw the circular arc A through $\hat{a}$ centered at $\hat{x}$. Let $\hat{b}$ be the intersection of A with the longitudinal arc from $\hat{x}$ to $\hat{z}$.
3. Let $\hat{m}$ be the midpoint of the arc of A between $\hat{a}$ and $\hat{b}$.
4. Draw the circular arc C through $\hat{m}$ centered at $\hat{y}$, and the circular arc D through $\hat{m}$ centered at $\hat{z}$. Let $\hat{c}$ be the intersection of B with the equatorial arc from $\hat{x}$ to $\hat{y}$ and let $\hat{d}$ be the intersection of C with the longitudinal arc from $\hat{x}$ to $\hat{z}$.
5. Let $\Gamma(r) = \angle\hat{x}\hat{c}$.

It is geometrically obvious that Γ is continuous and monotonic (part A above) and that $\lim_{r \to 0+} \Gamma(r) = 0$ (part B). For part C, note that as $r \to 1^-$, the arc A converges to the longitudinal arc from $\hat{y}$ to $\hat{z}$ and $\hat{m}$ converges to the midpoint of that arc. Hence $\hat{c}$ and $\hat{d}$ converge to midpoints of C and D respectively, so $\Gamma(r)$ converges to $1/2$.

To show that Γ is contracting (part D), let $\hat{a}$ and $\hat{a}'$ be points on the $\hat{x} - \hat{y}$ arc, with $\hat{a}'$ between $\hat{a}$ and $\hat{y}$. Carry out the above construction twice, first from $\hat{a}$ to midpoint $\hat{m}$ to output point $\hat{c}$, then from $\hat{a}'$ to midpoint $\hat{m}'$ to output point $\hat{c}'$ (figure 10). Since arcs A and A' are concentric circles on the sphere centered at $\hat{x}$, we have $\angle\hat{m}\hat{m}' = \angle\hat{a}\hat{a}'$ Since $\hat{m}, \hat{m'}$, and $\hat{y}$ are not collinear (do not lie on a geodesic) we have $\angle\hat{m}\hat{m}' > \angle\hat{m}\hat{y} - \angle\hat{m}'\hat{y}$. Since $\hat{m}$ and $\hat{c}$ lie on a circle centered at $\hat{y}$ and so do $\hat{m}'$ and $\hat{c}'$ we have $\angle\hat{y}\hat{m} = \angle\hat{y}\hat{c}$ and $\angle\hat{y}\hat{m}' = \angle\hat{y}\hat{c}'$. Since $\hat{y}$, $\hat{c}$ and $\hat{c}'$ lie on a geodesic, we

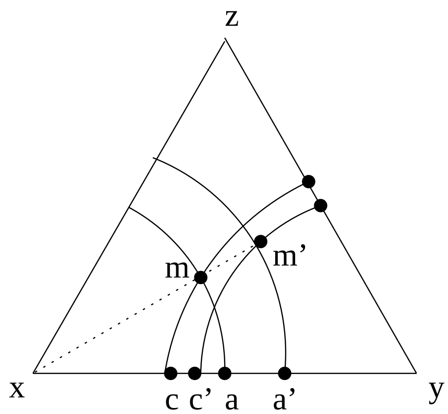

Figure 10: Γ is contracting

have $\angle \hat{c}\hat{c}' = \angle \hat{y}\hat{c} - \angle \hat{y}\hat{c}'$. Putting all this together, we have $\angle \hat{c}\hat{c}' < \angle \hat{a}\hat{a}'$, which established that Γ is contracting.

To implement this construction (part E), we define the network $N_{x,y}$ as follows.

The nodes correspond to $\hat{a}$, $\hat{b}$, $\hat{c}$, $\hat{d}$, and $\hat{m}$.

The sign constraints are $sg(\hat{a}) = sg(\hat{c}) = [+,+,0]$. $sg(\hat{b}) = sg(\hat{d}) = [+,0,+]$. $sg(\hat{m}) = [+,+,+]$.

The arcs are

$$\hat{b} = \hat{a} \cdot \begin{bmatrix} + & 0 & 0 \\ 0 & 0 & + \\ 0 & + & 0 \end{bmatrix}$$

That is, $\hat{b}$ is the image of $\hat{a}$ under a positive rotation around the x-axis.

$$\hat{d} = \hat{c} \cdot \begin{bmatrix} + & 0 & 0 \\ 0 & 0 & + \\ 0 & + & 0 \end{bmatrix}$$

That is, $\hat{d}$ is the image of $\hat{c}$ under a positive rotation around the x-axis. Note that $\angle \hat{c}\hat{x} = \angle \hat{d}\hat{x}$, and therefore $\angle \hat{c}\hat{y} = \angle \hat{d}\hat{z}$.

$$\hat{m} = \hat{a} \cdot \begin{bmatrix} + & 0 & 0 \\ 0 & + & + \\ 0 & + & + \end{bmatrix}$$

That is, $\hat{m}$ is an image of $\hat{a}$ under a rotation of between 0 and $\pi/2$ around the x-axis, and thus lies on arc A.

$$\hat{m} = \hat{c} \cdot \begin{bmatrix} + & 0 & + \\ 0 & + & 0 \\ - & 0 & + \end{bmatrix}$$

That is, $\hat{m}$ is an image of $\hat{c}$ under a rotation of between $-\pi/2$ and 0 around the y-axis, and thus lies on arc C.

$$\hat{m} = \hat{d} \cdot \begin{bmatrix} + & - & 0 \\ + & + & 0 \\ 0 & 0 & + \end{bmatrix}$$

That is, $\hat{m}$ is an image of $\hat{a}$ under a rotation of between 0 and $\pi/2$ around the z-axis, and thus lies on arc D.

Note that $\angle \hat{m}\hat{y} = \angle \hat{c}\hat{y} = \angle \hat{d}\hat{z} = \angle \hat{m}\hat{z}$, so $\hat{m}$ is the midpoint of arc A, as required.

∎.

Obviously we can define networks $N_{y,x}$, $N_{x,z}$, $N_{z,x}$, $N_{y,z}$, and $N_{z,y}$ analogously to satisfy lemmas analogous to lemma 10.

Definition 13 *Let* $\Gamma : (0,1) \mapsto (0,1)$ *be a function. Define* $\Gamma'(x) = 1 - \Gamma(1-x)$. *A number* r *is said to be* constructible by Γ *if it can be expressed in the form* $r = \Phi_n(\Phi_{n-1}(\ldots(\Phi_2(\Phi_1(1/2))\ldots))$ *where each* Φ_i *is either* Γ *or* Γ'.

Note that $\Gamma'(1-x) = 1 - \Gamma(x)$.

Lemma 11 *If* r *is constructible, then* $1-r$ *is constructible.*

Proof: Let $r = \Phi_n(\Phi_{n-1}(\ldots(\Phi_2(\Phi_1(1/2))\ldots))$ Using the above observation and induction on n, if one replaces every occurrence of Γ by Γ' and vice versa, then the result is a construction of $1-r$. ∎.

Lemma 12 *Let* $\Gamma : (0,1) \mapsto (0,1)$ *be a function satisfying the following constraints:*

- Γ *is monotonically increasing and contracting.*
- $lim_{r \longrightarrow 0^+} \Gamma(r) = 0$
- $lim_{r \longrightarrow 1^-} \Gamma(r) = 1/2$

Then the set of all constructible numbers is dense in $[0,1]$.

Proof by contradiction. Suppose not; then there exists some open interval (a,b) containing such no constructible numbers. Let us extend Γ to the interval [0,1] by defining $\Gamma(0) = 0$, $\Gamma(1) = 1/2$. Consider the set of all maximal open intervals containing no constructible numbers. Since no two such intervals can overlap, there can be only finitely many with size greater than $b-a$; hence there is a largest open interval containing no constructible numbers (not necessarily unique). Renaming, let (a,b) be this largest interval. Since $1/2$ is a constructible number, it must not be in the interval (a,b); hence, either $a \geq 1/2$ or $b \leq 1/2$. If $a \geq 1/2$, then the interval $(1-b, 1-a)$ is of equal size, and likewise must contain no constructible numbers and is less than $1/2$. Thus we can assume without loss of generality that $b \leq 1/2$. Since a and b are both less than $1/2$, then by the above conditions and the continuity of Γ, they are both in the image of Γ; that is, there exist values c and d such that $a = \Gamma(c)$, $b = \Gamma(d)$. If x is in the interval (c,d) then $\Gamma(x)$ is in (a,b) so $\Gamma(x)$ is not constructible, so x is not constructible. Thus (c,d) contains no constructible numbers. But since Γ is contracting, $d-c > b-a$, which contradicts the assumption that (a,b) was the largest interval containing no constructible numbers. ∎

Lemma 13 *Let* r *be any number constructible by* Γ. *Then there exists a network* $H_{xy}(r)$ *and a node* P *that construct the vector* $\hat{\phi}(\hat{x}, \hat{y}, r)$.

Proof The number r can be expressed in the form $r = \Phi_n(\Phi_{n-1}(\ldots\Phi_2(\Phi_1(1/2))\ldots))$, where each $\Phi_i(x)$ is either $\Gamma(x)$ or $\Gamma'(x)$.

Using lemmas 9 and 10, construct a sequence of $n+1$ networks $Q_0, Q_1 \ldots Q_n$ where

- Q_0 constructs $\hat{\phi}(\hat{x}, \hat{y}, 1/2)$
- If Φ_i is Γ then Q_i is a network of the form N_{xy}. If Φ_i is Γ' then Q_i is a network of the form N_{yx}. In either case the input node of Q_i is the output node of Q_{i-1}.

The overall network $H_{xy}(r)$ is the union of the Q_i; the node P is the output of Q_n. ∎

Lemma 14 *Let r and s be two numbers constructible by Γ, where $r+s > 1$. Let $\hat{v}$ be the intersection with sign $[+, +, +]$ of the circle of radius r centered at $\hat{x}$ and the circle of radius s centered at $\hat{y}$. Then there is a network that constructs $\hat{v}$.*

Proof: Use lemma 13 to construct the points $\hat{a} = \hat{\phi}(\hat{x}, \hat{y}, r)$ and $\hat{b} = \hat{\phi}(\hat{x}, \hat{y}, s)$. Assert that point $\hat{v}$ is both the image of $\hat{a}$ under a rotation around the x-axis towards $\hat{z}$ and also the image of $\hat{b}$ under a rotation around the y-axis towards $\hat{z}$. ∎.

We can now prove the theorems.

Theorem 4 *The set of directions that are uniquely identifiable in the class of pure positive QRCNs is dense within the non-negative octant.*

Proof: Immediate from lemmas 14 and 12. ∎

Corollary 4 *The set of directions that are uniquely identifiable in the class of pure QRCNs (not necessarily positive) is dense within the unit sphere.*

Proof: The construction in lemma 14 can be modified to apply to the remaining seven octants, simply by changing the sign associated with the vector $\vec{v}$. ∎

Theorem 3 *Any two directions are distinguishable in the class of pure QRCNs.*

Proof: Let $\hat{a} \neq \hat{b}$ be two directions. If they have different sign signatures, then that sign distinguishes them. Assume they have the same signs; without loss of generality, assume their signs are all non-negative. It must be the case that either $\angle\hat{x}\hat{a} \neq \angle\hat{x}\hat{b}$ or that $\angle\hat{y}\hat{a} \neq \angle\hat{y}\hat{b}$. (If both are equal, then they are both at the intersection of the same circles around $\hat{x}$ and $\hat{y}$; but there is only one such intersection in the positive octant.)

Assume without loss of generality that $\angle\hat{x}\hat{a} < \angle\hat{x}\hat{b}$. Using lemma 12 let r be a constructible number between $\angle\hat{x}\hat{a}$ and $\angle\hat{x}\hat{b}$. Using lemma 13, construct a network N whose output is $\hat{r} = \hat{\phi}(\hat{x}, \hat{y}, r)$. Let $\hat{p} = \hat{\phi}(\hat{x}, \hat{y}, \angle\hat{x}\hat{a})$. Supplement N with a node A for $\hat{a}$ and node P for $\hat{p}$ and with the arcs stating that $\hat{a}$ is the image of $\hat{p}$ under a rotation around the x-axis and that $\hat{r}$ is the image of $\hat{p}$ under a *positive* rotation around the z-axis. Then the network is satisfied by $\hat{a}$ at node A, but not by $\hat{b}$ at node A. ∎.