SRC Technical Note
2002 - 003

February 28, 2002

Averaging on a quotient of a sphere

Lyle Ramshaw

Systems Research Center
130 Lytton Avenue
Palo Alto, California 94301

http://www.research.compaq.com/SRC/

Averaging on a quotient of a sphere

Lyle Ramshaw

February 28, 2002

The original motivation for this note is the following problem: We would like to compute "averages" or "centers of gravity" in the quotient space S^3/H where S^3 denotes the 3-sphere (sitting in 4-space) and H is a certain subgroup of S^3 of size 48. So the points of the 3-sphere are being identified into cosets, each of size 48. This note does not directly address that original problem. Instead, it discusses the easier problem of computing averages either

- in S^n, the n-sphere itself, with no identification going on, or
- in the space $S^n/\{+1, -1\}$, the n-sphere in which each point is identified with its antipode, forming a coset of size 2.

Those easier problems have fairly satisfactory solutions.

1 The n-sphere with no identification

If there is no identification going on, things are pretty easy. Suppose that we are given a set of points on the n-sphere, each with an associated nonnegative weight. We want to compute the average of those points.

We can think of each point as a unit vector in Euclidean $(n + 1)$-space. We scale each vector by its associated weight, add the resulting vectors (as vectors), and rescale the sum to have unit norm.

The only subtle point is that the sum might turn out to be the zero vector. In that case, we have to return "undefined" as our average, since arbitrarily small changes to the input points or their weights might cause the intended average to assume any value on S^n.

Technically speaking, this need for an undefined result arises even when averaging weighted points in an affine space: When all of the weights are zero, we have no information about what point should be the average, so the only reasonable answer is "undefined". But undefinedness is a more severe problem when averaging on spheres, since it can arise even when the input includes points with positive weight.

2 The n-sphere with antipodes identified

Things get more subtle if we assume that antipodal points on the n-sphere are to be treated as the same.

How might such a situation arise? Suppose that we are working in Euclidean $(n+1)$-space. Someone gives us a set of undirected lines through the origin, where each line has an associated nonnegative weight. We want to compute the average of the given lines. Since the lines are undirected, each line corresponds to a pair of antipodal points on the unit n-sphere; so our goal becomes computing an average in the n-sphere with antipodal points identified.

2.1 The case $n = 1$: Lines through a point in the plane

Let's start with a simple case: a circle with opposite points identified. So someone has given us a set of weighted, undirected lines through the origin of a plane.

A simple way to average the lines would be to average their slopes; but this is a lousy scheme. In part, it is lousy because "slope" gives the wrong geometry. The difference in slope between two lines is not a good measure of the distance between them; it inflates the difference between close-by lines that happen to be nearly vertical. Thus, given a set of lines whose slopes are clustered around some finite value, say +1, outliers that are closer to vertical will exert too strong a pull on the average, while outliers that are closer to horizontal will pull too weakly.

That problem is easy to fix; rather than measuring a line by computing its slope, we instead compute its angle — say, the angle that the line makes with the positive x-axis, where that angle is constrained to lie in the half-open interval $[0\,..\,180)$. If we average these angles, rather than averaging slopes, we will correctly capture the local geometry.

But averaging angles is still a poor scheme, since it doesn't get the global geometry correct; it gets confused when the circle "wraps around". Consider a set of lines, all of which are nearly horizontal. Some of the lines will have angles near 0, while others will have angles near 180. When we average these angles, we might end up with any angle in the interval $[0\,..\,180)$ — bad news.

2.2 The 2-theta trick for $n = 1$

There is a special trick — the 2-theta trick — that solves the case $n = 1$ quite neatly. What makes the 2-theta trick work is a geometric coincidence: The space that results when we identify antipodal points on a circle is again a circle, just of half the length. That coincidence does not extend to higher dimensions. For example, when $n = 2$, identifying antipodal points on a

2-sphere results in a projective plane, which is quite different from a 2-sphere, even topologically.

Here is how the 2-theta trick works. We replace each of our unoriented input lines by one of the two unit vectors that lie along that line. It doesn't matter which we choose, because we next double the angle that each such vector makes with the positive x-axis. That is to say, if we temporarily view our plane as the complex plane, we replace the vector $\exp(i\theta)$ with $\exp(2i\theta)$. The equation

$$\exp(2i(\theta + \pi)) = \exp(2i\theta + 2\pi i) = \exp(2i\theta)$$

shows that it doesn't matter which vector we chose to represent each unoriented input line.

Now that the angles have been doubled, we can treat our vectors as points on a vanilla circle — a circle with no identifications. That problem, we know how to solve: We scale each vector by its weight, add the results as vectors, and rescale their sum to have unit norm.

It remains only to undo the effect of the angle-doubling. We note that the normalized sum will have the form $\exp(2i\phi)$, for some angle ϕ in the interval $[0\,..\,180)$, and that ϕ is the "average angle" that we output.

2.3 The moments-of-inertia trick for $n = 2$

Here is a different trick that handles the case $n = 2$, using ideas from physics.

Given an undirected line through the origin of 3-space with an associated weight w, we fasten two point masses of weight w along that line, both at unit distance from the origin, one on each side of the origin. We do this for each input line, and we view the resulting assemblage of masses as a rigid body. The center of gravity of that body will be at the origin, obviously. But it will have various moments of inertia around various axes of rotation.

A basic result of physics tells us that there exists a unique ellipsoid of unit density, centered at the origin, that has the same moments of inertia around all axes as our assemblage of point masses. That ellipsoid has 3 orthogonal axes, of various lengths. If the longest of those axes is unique, we take that line to be our average. If two or more of the axes tie for being longest, we must return "undefined", since any line in the linear space spanned by the tying axes has an equally valid claim to be the average.

2.4 The general case: Diagonalizing a quadratic form

We have seen tricks that handle the cases $n = 1$ and $n = 2$. Here is a general technique, based on diagonalizing a quadratic form, that gets the same answers as those tricks in the cases $n = 1$ and $n = 2$, but also extends to arbitrary n.

Working in $(n+1)$-space, let v_1 through v_m be unit vectors in the directions of the m input lines; as we shall see, it doesn't matter whether we choose v_i or $-v_i$ as the direction vector for the i^{th} line. Let the associated weights be w_1 through w_m. The average line that we want, I claim, is the line in the direction of that unit vector u that maximizes the sum

$$F(u) := \sum_i w_i \cos^2(\text{angle between } u \text{ and } v_i).$$

Note that the angle between u and $-v_i$ is supplementary to the angle between u and v_i, so the cosines of those angles will be negatives of each other, but the squares of those cosines will be the same. Hence, it doesn't matter whether we chose v_i or $-v_i$ as the direction of the i^{th} line.

Since u and v_i are both unit vectors, we can rewrite the sum as

$$F(u) = \sum_i w_i (u \cdot v_i)^2,$$

where the dot here indicates the dot product of vectors. This formula makes sense for any vector u; in fact, it defines F to be a quadratic form — a linear combination of squares of linear forms. Since the w_i are nonnegative, the quadratic form F is *positive semidefinite*, that is, satisfies $F(u) \geq 0$ for all vectors u.

Every quadratic form can be diagonalized. So there is some orthogonal coordinate system for our $(n+1)$-space in which $F(u)$ takes the form

$$F(u) = a_0 u_0^2 + \cdots + a_n u_n^2,$$

where $(u_0, \ldots, u_n)$ are the coordinates of the vector u. Since F is positive semidefinite, the coefficients a_0 through a_n are nonnegative, and we might as well assume that they are sorted:

$$a_0 \geq a_1 \geq \cdots \geq a_n \geq 0.$$

We are trying to maximize $F(u)$, subject to u being of unit norm. The vector $(1, 0, 0, \ldots, 0)$ pointing along the 0^{th} coordinate axis clearly achieves that maximum, and it is the unique maximal vector precisely when $a_0 > a_1$. So we compute a coordinate system that diagonalizes F; if the maximum element on the diagonal is unique, we output the corresponding axis line as our "average", else we output "undefined".

It remains to verify that this general technique gets the same answers as our tricks above in the cases $n = 1$ and $n = 2$. Here are sketches of arguments to that effect.

When $n = 2$, finding the longest axis of the unit-density ellipsoid corresponds to finding the axis of rotation around which the moment of inertia

of our rigid body is minimal. Minimizing the moment of inertia around the axis through a vector u corresponds to minimizing the sum

$$\sum_i w_i \sin^2(\text{angle between } u \text{ and } v_i).$$

But $\sin^2 + \cos^2 = 1$, so minimizing this sum is the same as maximizing $F(u)$.

When $n = 1$, the 2-theta trick ends up computing, as its average, the unit vector u whose angle ϕ maximises

$$\sum_i w_i \cos(2\theta - 2\phi).$$

But $\cos(2x) = 2\cos^2(x) - 1$; so maximizing this sum is the same as maximizing

$$F(u) = \sum_i w_i \cos^2(\theta - \phi).$$

3 Cosets of size greater than 2

So we have a clear definition of what it means to average on an n-sphere and an efficient algorithm for computing those averages in two situations:

- when no identifications are going on, or
- when each point is identified with its antipode.

Unfortunately, the original problem motivating this note involved identifying points on S^3 into cosets of size 48, rather than size 2. What can we say about that?

Before we can talk meaningfully about identifying points more than antipodally, we must restrict n to have one of the special values 1, 3, or 7. The key thing about those values is that the spheres S^1, S^3, and S^7 are multiplicative groups:

- S^1 is the group of unit-norm complex numbers,
- S^3 is the group of unit-norm quaternions, and
- S^7 is the group of unit-norm octonions (a.k.a. Cayley numbers).

(To complete this list, we should include $S^0 = \{+1, -1\}$, the group of unit-norm real numbers; but that case is not interesting here.)

The advantage of a group structure is that it gives us a coherent way to identify points into cosets of sizes larger than 2. Given a finite subgroup G

of S^n, say of size k, we can form the set of left (or right) cosets S^n/G. Those cosets won't themselves form a group unless G is normal. But we can still construct the cosets and try to compute averages in the space S^n/G of all cosets.

The case $n = 1$ is quite straightforward. For each positive integer k, the group S^1 of unit-norm complex numbers has a unique finite subgroup of order k, and it is both normal and cyclic — call it C_k. If we want to compute averages in the quotient group S^1/C_k, we can simply use the "k-theta trick", multiplying angles by k in the same way that the 2-theta trick multiplies them by 2. This exploits the coincidence that the group S^1/C_k is isomorphic to the group S^1; both groups are simply circles.

The case $n = 3$ is more subtle. We know how to compute averages in S^3 itself and in the quotient $S^3/\{+1, -1\}$. That quotient is a group, by the way, since the subgroup $\{+1, -1\}$ is normal; in fact, the quotient $S^3/\{+1, -1\}$ is better known as $SO(3)$. But what about modding out by some larger subgroup G — which won't be normal?

3.1 General comments on S^3/G, where $|G| > 2$

By analogy with the quadratic-form-based method above, it is tempting to define the "average" in a quotient S^3/G as the unit vector u that maximizes some sum of the form

$$F(u) := \sum_i w_i \, r_G(u, v_i),$$

where r_G is some real-valued function that measures how well correlated the two points u and v_i in S^3 are, modulo the subgroup G. Exploiting the group structure on S^3, we surely want the value of $r_G(x, y)$ to depend only upon the "difference" between x and y, which is either $x^{-1}y$ or yx^{-1}, depending upon whether we are using left or right cosets. Thus, our hoped-for sum takes the form, say,

$$F(u) := \sum_i w_i \, r_G(u^{-1}v_i),$$

where $r_G \colon S^3 \to [0\,..\,1]$ is some function that

- has a local maximum of 1 at each element of G, and
- takes the value 0 at all points that lie midway between two (or more) elements of G.

When G is a cyclic subgroup of order k, Li Zhang and I think that we see how to construct a reasonable function r_G; indeed, our function seems closely

related to $\cos(k\theta)$. On the other hand, it may be numerically a bit tricky to compute the vector that maximizes the resulting sum, since our function r_G is algebraic of degree k. But no more details about that here.

When the group G is not cyclic, things are much worse; indeed, I am afraid that this entire approach seems doomed. Consider, for example, the group

$$Q := \{1, -1, i, -i, j, -j, k, -k\}$$

of quaternions units, which has order 8. If we grow these eight points on S^3 into Voronoi cells, the resulting cells are elliptic analogs of cubes, and they fit together like the eight cubical faces of a tesseract. A cube in Euclidean 3-space has dihedral angles of 90 degrees, of course. These elliptic cubes are so large that their dihedral angles are 120 degrees — which means that three of them fit together as we go around an edge.

But that means real trouble: Each of the three planar faces incident to such an edge is a boundary between two Voronoi cells; so, in the methodology above, the function r_Q should be identically zero along that planar face. But surely we want the function r_Q to be real analytic. If it is real analytic and zero along that face, it must be zero on the entire plane through that face — which means that it must be zero at the center of the third Voronoi cell. But we want r_Q to be 1 at the center of each Voronoi cell — a dilemma! Perhaps we have to resort to defining r_Q piecewise? That seems unsatisfactory; but further analysis must be relegated to future work.