Least Squares Fitting of Spatial Circles

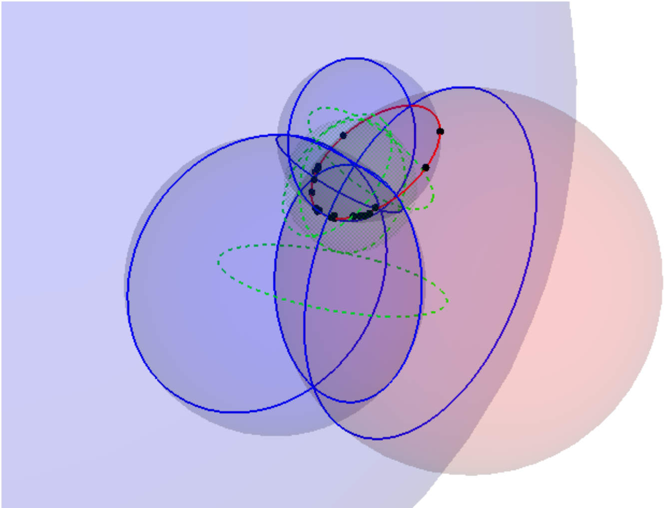

Leo Dorst (`L.Dorst@uva.nl`)
Intelligent Systems Laboratory, Informatics Institute,
University of Amsterdam, The Netherlands

FUGRO, February 1, 2013
IAS, April 16, 2013 (modified)
Santander, 2016 (modified)

1 Motivation: Accurate Fitting of Spatial Circles

FUGRO: Large international company specialized in measurement of geodata.
Focus: Accurate measurement of undersea pipes for construction and maintenance.
Have enormous 3D point clouds to be modelled (as do TNO, NLR, NFI).
Money no objection: 1 M€/day for repairs on the sea floor.

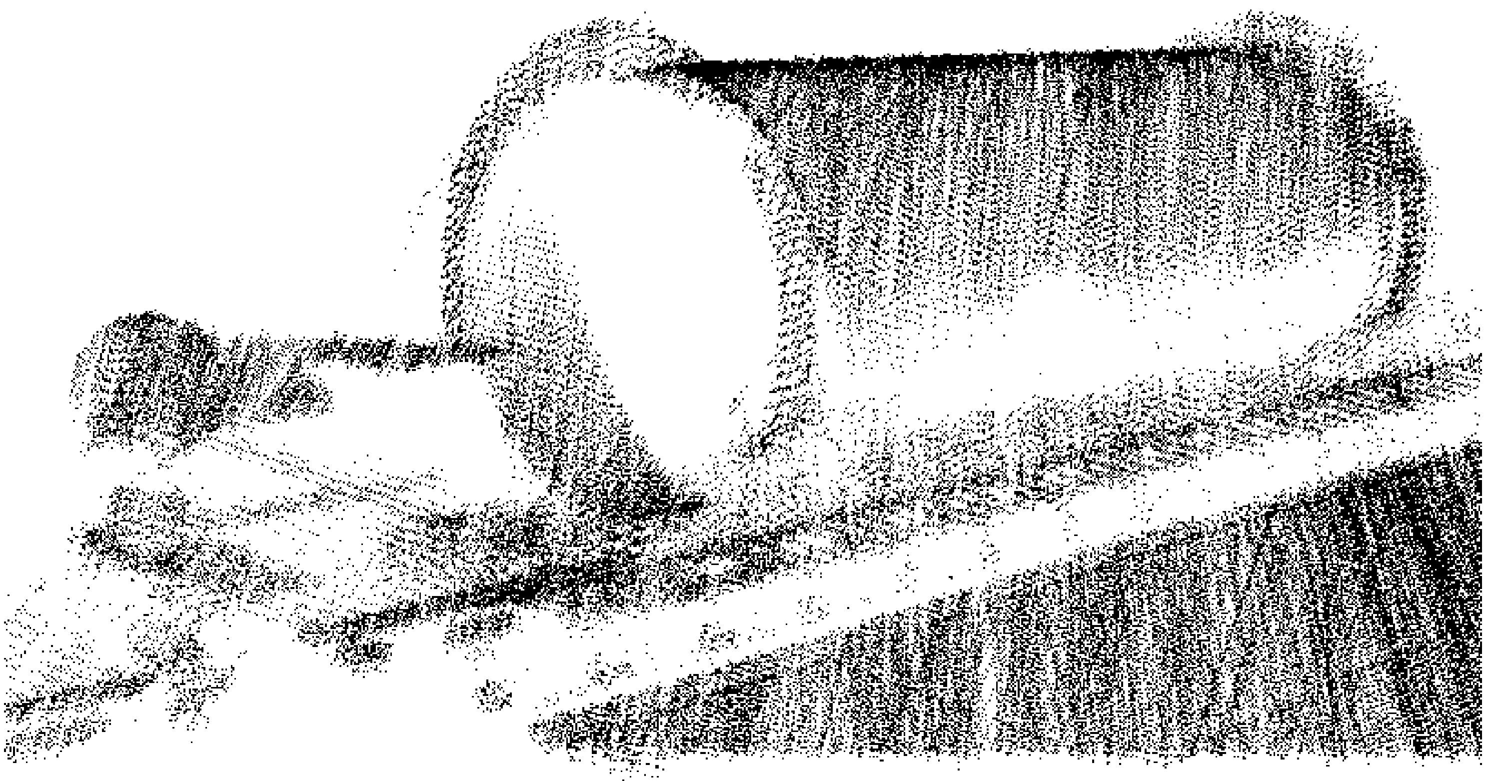

2 Fitting Circles - Overview

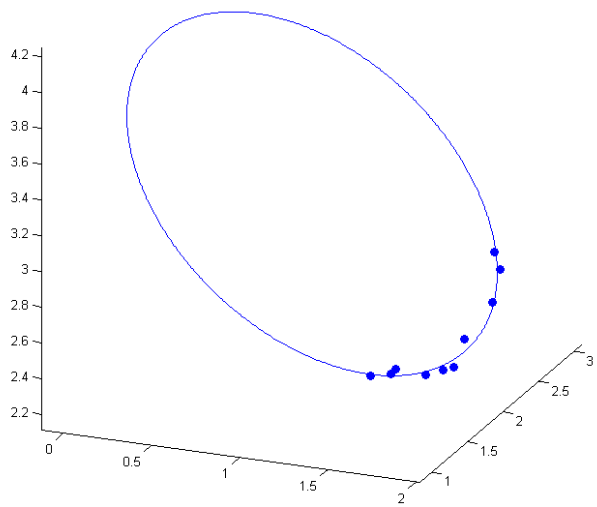

- The puzzle: how to fit a circle to 3D point data?
- Picking the right representation (CGA).
- Optimal sphere fitting solved as eigenproblem.
- Circle fitting as eigenproblem.
- Circle fitting by sphere fitting.
- Evaluation of accuracy.

- Bonus ad: What more could CGA do for you?

3 Circle definition, in geometry and algebra

We want to fit circles to point data. A circle is the intersection of a sphere and a plane.

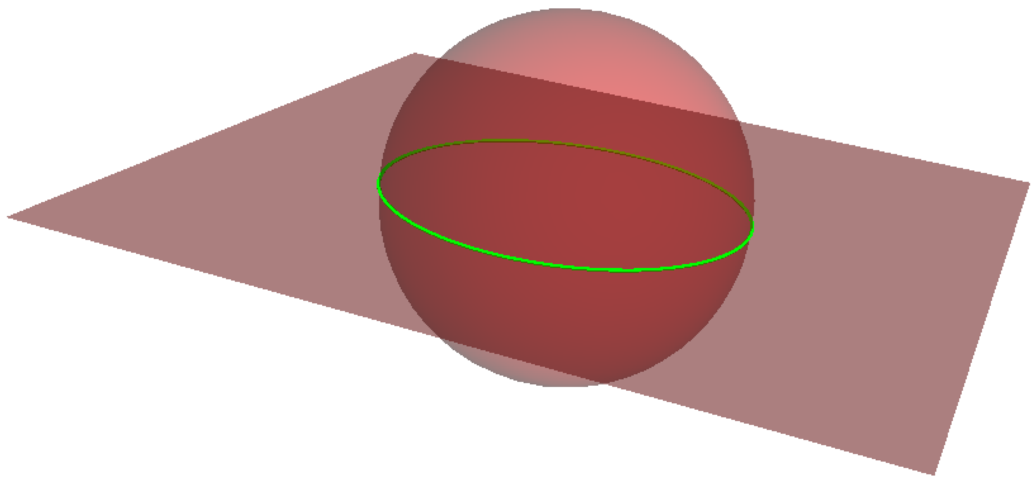

There is an algebra that directly implements this definition:

$$\kappa = \sigma \wedge \pi.$$

For the fit, use the geometric algebra of a vector space in which all elements in the fit are basic: its vectors represent spheres, including planes (spheres of infinite radius) and points (spheres of zero radius).

This algebra is called CGA (conformal geometric algebra).

4 Fitting Planes and Spheres

Fitting a plane to point data $\{\mathbf{p}_i\}$ to minimize $\Sigma_i((\mathbf{p}_i - \mathbf{d}) \cdot \mathbf{n})^2$ is solved:

1. Compute the centroid $\bar{\mathbf{p}} = \frac{1}{N}\Sigma_i\ \mathbf{p}_i$.
2. Compute the covariance matrix $C = \frac{1}{N}\Sigma_i\ (\mathbf{p}_i - \bar{\mathbf{p}})\,(\mathbf{p}_i - \bar{\mathbf{p}})^T$ of the relative vectors.
3. The eigenvector $\mathbf{n}$ of C with smallest eigenvalue is the normal vector of the optimal plane.
4. Construct the best fit plane with normal vector $\mathbf{n}$ passing through $\bar{\mathbf{p}}$.

Algebraically a bit clumsy but it works.

But how to fit a sphere to point data?

- Awkward nonlinear criterion 'minimize $\Sigma_i(\sqrt{(\mathbf{p}_i - \mathbf{c}) \cdot (\mathbf{p}_i - \mathbf{c})} - \rho)^2$'.
- The plane method does not generalize, in linear algebra.
- Methods known (2D circle fits [Al-S.&Ch. 2009] , generalizable to n-D) but look *ad hoc*.
- Let's invent our own using CGA. (It will actually reinvent [Pratt 1987].)

Spoiler: The best fitting circle is NOT the best sphere cut by the best plane!

5 Optimal Fitting of Spheres

Given N data point vectors $\mathbf{p}_i$ in n-D, what is the best fitting hypersphere?

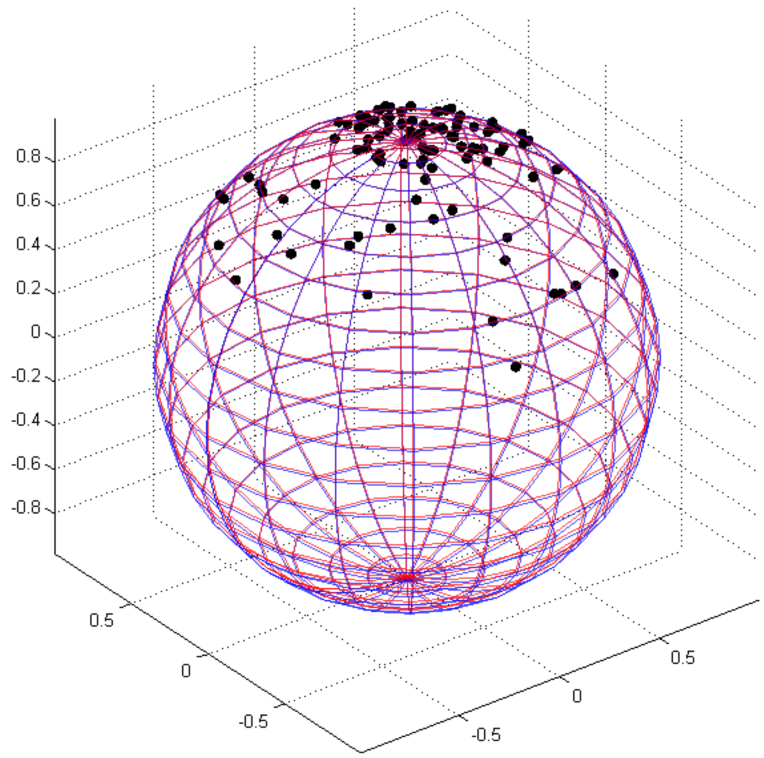

6 Optimal Sphere Fitting Solution: the Recipe

1. Put the N point data $\mathbf{p}_i$ in a data matrix $[D]$ with column i equal to $\begin{bmatrix} \mathbf{p}_i \\ 1 \\ \frac{1}{2}\|\mathbf{p}_i\|^2 \end{bmatrix}$.
2. Make a matrix $[P]$ as $[P] = [D]\,[D]^T[M]/N$, where $[M] = \begin{bmatrix} \mathbf{I}_{n\times n} & \mathbf{0} & \mathbf{0} \\ \mathbf{0}^T & 0 & -1 \\ \mathbf{0}^T & -1 & 0 \end{bmatrix}$.
3. Solve the eigenproblem for $[P]$, giving minimum eigenvalue λ_* and its eigenvector x_*.
4. Interpret the solution: normalize the eigenvector $[x_*]$ to have $[x_*]_{n-1}$ equal to 1.

   It then relates to the best-fit hypersphere parameters as $[x_*] = \begin{bmatrix} \mathbf{c} \\ 1 \\ \frac{1}{2}(\|\mathbf{c}\|^2 - \rho^2) \end{bmatrix}$.

   The first n components of $[x_*]$ give the center $\mathbf{c}$; the radius $\rho = \sqrt{\|\mathbf{c}\|^2 - 2[x_*]_\infty}$.

This sphere fitting recipe can be implemented in Matlab straightforwardly.
The circle recipe will have a similar style.

But to understand them both, we will have to dive deeper!

7 Spheres in Linear Algebra

Sphere with center $\mathbf{c}$ and radius ρ:

$$\|\mathbf{x} - \mathbf{c}\|^2 = \rho^2.$$

Introducing 3D coordinates (yuck!) $\mathbf{x} = (x_1, x_2, x_3)^T$ and $\mathbf{c} = (c_1, c_2, c_3)^T$ this is:

$$x_1^2 + x_2^2 + x_3^2 - 2(x_1c_1 + x_2c_2 + x_3c_3) + c_1^2 + c_2^2 + c_3^2 - \rho^2 = 0.$$

Setting $x_0 = x_1^2 + x_2^2 + x_3^2$ and introducing constants, it seems tidier:

$$Ax_0 + B_1x_1 + B_2x_2 + B_3x_3 + C = 0.$$

This is the basis of the classical algebraic treatment of the sphere fits.
The squares hidden in x_0 make this a bit awkward, and different from plane fitting.
A least square fit requires constrained optimization [Al-Sharadqah & Chernov 2009].

By contrast, we will set up a 5-D space that gives as sphere equation:

$$x \cdot s = 0$$

with x and s vectors. This is much more tractable.

8 One Space to Rule Them All: the Geometry of Spheres, Planes and Points

Recipe for CGA (Conformal Geometric Algebra [Anglès 1980, Hestenes 1984]):

- Embed your space $\mathbb{R}^n$ in $\mathbb{R}^{n+1,1}$ (so Minkowski space of two more dimensions)
- Choose basis with $\mathbb{R}^{n+1,1}$ with Euclidean part, plus n_o and n_∞ for the extra dimensions. Pick the metric such that $n_o \cdot n_o = n_\infty \cdot n_\infty = 0$, and $n_o \cdot n_\infty = -1$.
  $\oplus$ Recognize this from the recipe? $[M] = \begin{bmatrix} \mathbf{I}_{n\times n} & \mathbf{0} & \mathbf{0} \\ \mathbf{0}^T & 0 & -1 \\ \mathbf{0}^T & -1 & 0 \end{bmatrix}$, on basis $\{\mathbf{e}_1, \cdots, \mathbf{e}_n, n_o, n_\infty\}$.
- A point at location $\mathbf{p}$ is represented as the vector

$$p = n_o + \mathbf{p} + \tfrac{1}{2}\|\mathbf{p}\|^2 n_\infty.$$

  You may think of n_o as point at origin, n_∞ as point at infinity.
  $\oplus$ Recognize this from the recipe? $[p] = [\mathbf{p}^T, 1, \frac{1}{2}\|\mathbf{p}\|^2]^T$, on basis $\{\mathbf{e}_1, \cdots, \mathbf{e}_n, n_o, n_\infty\}$.
- This gives an isometric model with squared Euclidean distances as dot products:

$$p \cdot q = -\tfrac{1}{2}\|\mathbf{p} - \mathbf{q}\|^2$$

  For a point, $p \cdot p = 0$, so points are represented as *null vectors*.
  $\oplus$ In the implementation, $[p \cdot q] = [p]^T\,[M]\,[q] = [-\frac{1}{2}\|\mathbf{p} - \mathbf{q}\|^2]$.

9 CGA: the Geometry of Spheres, Planes and Points (continued)

- To make a sphere with center $\mathbf{c}$ and radius squared ρ^2, use the vector:

$$s = c - \tfrac{1}{2}\rho^2 n_\infty.$$

Now $0 = x \cdot s \;\Leftrightarrow\; \|\mathbf{x} - \mathbf{c}\|^2 = \rho^2$.

$\oplus$ In the implementation, $[s] = [\mathbf{c}^T, 1, \tfrac{1}{2}(\mathbf{c}^2 - \rho^2)]^T$.

- A plane with normal $\mathbf{n}$ through $\mathbf{p}$ is (dually) represented as the vector:

$$\pi = \mathbf{n} + (\mathbf{n} \cdot \mathbf{p})\, n_\infty.$$

$\oplus$ In an implementation, $[\pi] = [\mathbf{n}^T, 0, \mathbf{n} \cdot \mathbf{p}]^T$.

- A circle is the intersection of a sphere and a plane, or of two spheres.
It is (dually) represented as a 2-D subspace using the outer product of geometric algebra:

$$\kappa = s \wedge \pi = s_1 \wedge s_2.$$

$\oplus$ In the implementation, $\wedge$ is a certain matrix (see later).

- Perpendicularity of geometrical elements represented by x and y is algebraically: $x \cdot y = 0$.

$\oplus$ A point p on a sphere s is a small sphere perpendicular to it, so $p \cdot s = 0$.

- As a true geometric algebra, CGA has a geometric product.
This permits division by vectors and other subspaces. For vectors, $x^{-1} = x/(x \cdot x)$.

10 Distance of Point and Hypersphere

For a dual sphere $\sigma = c - \frac{1}{2}\rho^2 n_\infty$ and a point p, the CGA dot product $\sigma \cdot p$ gives a somewhat strange squared distance measure between point and sphere [Perwass & Förstner 2006], [Rockwood & Hildenbrand 2010]:

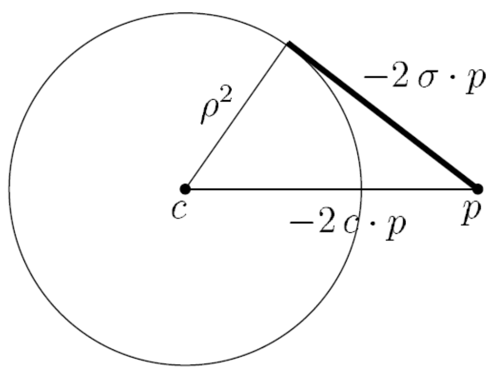

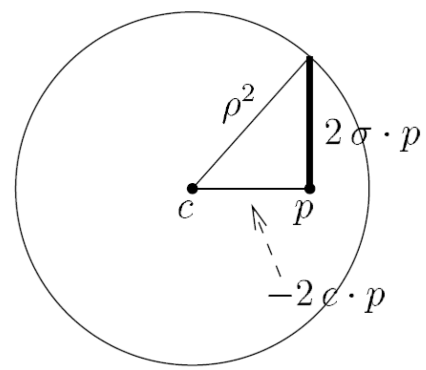

However, for point p a small signed distance δ outside the sphere:

$$\mp 2\,\sigma \cdot p = \pm\Big(d_E^2(c,p) - \rho^2\Big) = \pm\Big((\rho + |\delta|)^2 - \rho^2\Big) \approx 2\,\rho\,\delta$$

Therefore, using $\rho^2 = \sigma^2$:

$$(\sigma \cdot p)^2/\sigma^2 \approx \delta^2$$

11 An Algebraically Natural Approximate Criterion

Good approximation to sum of squares of distances δ_i of points p_i to (dual) hypersphere σ with radius $\rho = \sqrt{\sigma^2}$:

$$\Sigma_i \; (p_i \cdot \sigma)^2 / \sigma^2 = \Sigma_i \; \delta_i^2 \Big(1 + \frac{\delta_i}{\rho} + \big(\frac{\delta_i}{2\rho}\big)^2 \Big) \approx \Sigma_i \; \delta_i^2$$

(The sum is less than its parts: sum over p_i tends to cancel contribution ${\delta_i}^3/\rho$ by points inside and outside the sphere. We will get back to just how good this is.)

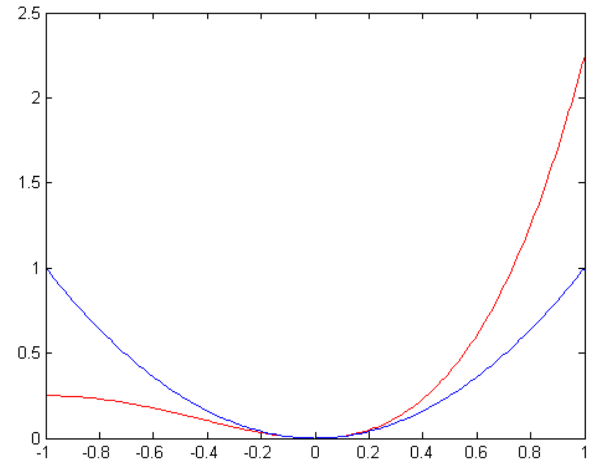

So we try to solve in $\mathbb{R}^{n+1,1}$, given conformal points p_i:

$$\boxed{\textit{Find an } x \textit{ that minimizes: } \quad \mathcal{L}(x) = \tfrac{1}{N} \Sigma_i \; (p_i \cdot x)^2 / x^2}$$

To unclutter our work we set $P[x] = \frac{1}{N} \Sigma_i \; p_i \, (p_i \cdot x)$ (a symmetric linear function):

$$\boxed{\textit{Find an } x \textit{ that minimizes: } \quad \mathcal{L}(x) = x^{-1} \cdot P[x].}$$

This is like the classical optimization with constrained norm using a Lagrangian multiplier:

$$\textit{Find an } x \textit{ and } \lambda \textit{ that minimize: } \; \mathcal{L}(x, \lambda) = x \cdot P[x] + \lambda (x \cdot x - 1).$$

12 Straightforward Solution by Coordinate-free Differentiation ∂_x

$$
\begin{aligned}
0 &= \partial_x \mathcal{L}(x) \\
&= \partial_x \left(x^{-1} \cdot P[x]\right) \\
&= -x^{-1} P[x]\, x^{-1} + \bar{P}[x^{-1}] && \text{standard GA differentiation} \\
&= (-x\, P[x] + \bar{P}[x]\, x)\, x^{-3} && \text{rearranging by linearity} \\
&= (-x\, P[x] + P[x]\, x)\, x^{-3} && \text{by symmetry of } P[\,] \\
&= 2\, (P[x] \wedge x)\, x^{-3} && \text{by definition of outer product}
\end{aligned}
$$

Multiply by the invertible vector $\frac{1}{2} x^3$ and rewrite:

$$P[x] \wedge x = 0.$$

This is an *eigenproblem*! Any solution x_* is an eigenvector of the operator

$$P[\,] : \mathbb{R}^{n+1,1} \to \mathbb{R}^{n+1,1} : \quad x \ \mapsto \ \tfrac{1}{N} \Sigma_i \, p_i \, (p_i \cdot x). \tag{1}$$

The eigenvalue is the cost of the solution, i.e. the realized mean of squared distances:

$$\mathcal{L}(x_*) = x_*^{-1} \cdot P[x_*] = \lambda_* \, x_*^{-1} \cdot x_* = \lambda_*.$$

Minimize the $\mathcal{L}$ with a real sphere x_*: pick λ_* as minimal non-negative eigenvalue of $P[\,]$.

13 Problem solved: Optimal Sphere Found

The sphere $x_* = c - \frac{1}{2}\rho^2 n_\infty$ minimizing the sum of approximate squared distances of a set of conformal points $\{p_i\}$ is the (normalized) eigenvector of minimum nonnegative eigenvalue of the linear operator $P[\,] = \Sigma_i \; p_i\,(p_i \cdot [\,])$.

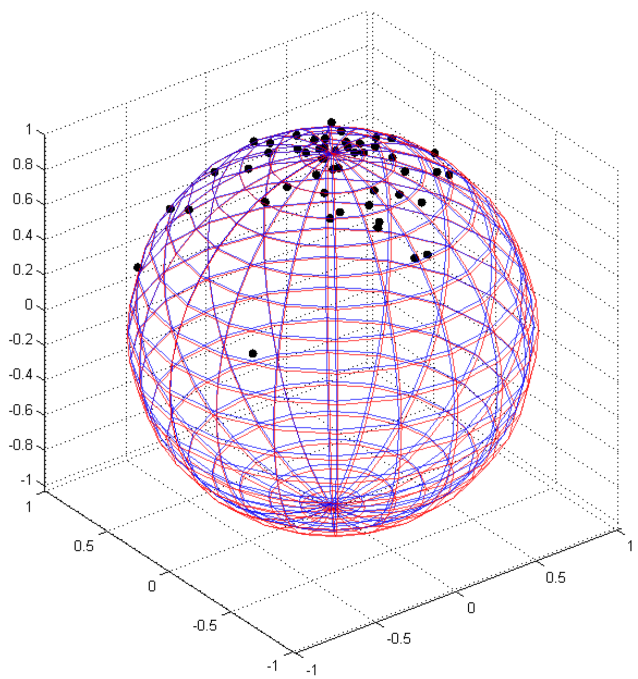

`sphere_fit(50,0.01,0.5)`

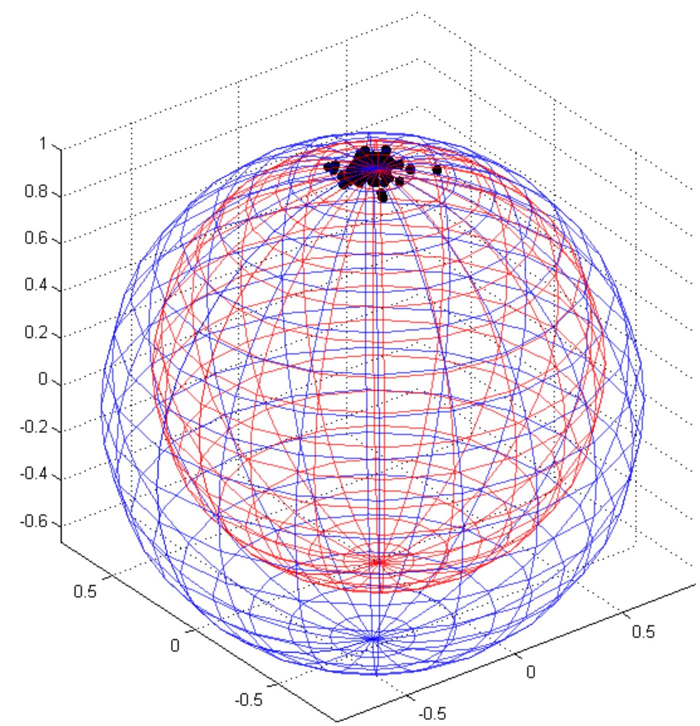

`sphere_fit(100,0.01,0.1)`

14 Optimal Sphere Fitting Solution: the Recipe (As Before)

1. Put the N point data $\mathbf{p}_i$ in a data matrix $[D]$ with column i equal to $\begin{bmatrix} \mathbf{p}_i \\ 1 \\ \frac{1}{2}\|\mathbf{p}_i\|^2] \end{bmatrix}$.

2. Make a matrix $[P]$ as $[P] = [D]\,[D]^T[M]/N$, where $[M] = \begin{bmatrix} \mathbf{I}_{n\times n} & \mathbf{0} & \mathbf{0} \\ \mathbf{0}^T & 0 & -1 \\ \mathbf{0}^T & -1 & 0 \end{bmatrix}$.

3. Solve the eigenproblem for $[P]$, giving minimum eigenvalue λ_* and its eigenvector x_*.

4. Interpret the solution: normalize the eigenvector $[x_*]$ to have $[x_*]_{n-1}$ equal to 1.

   It then relates to the best-fit hypersphere parameters as $[x_*] = \begin{bmatrix} \mathbf{c} \\ 1 \\ \frac{1}{2}(\|\mathbf{c}\|^2 - \rho^2) \end{bmatrix}$.

   The first n components of $[x_*]$ give the center $\mathbf{c}$; the radius $\rho = \sqrt{\|\mathbf{c}\|^2 - 2[x_*]_\infty}$.

This sphere fitting recipe can be implemented in Matlab without any knowledge of CGA.

15 Fitting Hypercircles - an Approximate Least Squares Measure

Let X be a factorizable bivector representing a (dual) hypercircle. Then X could be factorized as product of its carrier hyperplane π and a hypersphere σ perpendicular to it: $X = \pi \wedge \sigma$ with $\pi \cdot \sigma = 0$. Then, trying the same formula:

$$-(p \cdot X)^2/X^2 = -\big(p \cdot (\pi \wedge \sigma)\big)^2/(-\pi^2\sigma^2) = (p \cdot \pi)^2/\pi^2 + (p \cdot \sigma)^2/\sigma^2.$$

This is the sum of the exact squared distance to the carrier plane, plus the approximate squared distance to the carrier hypersphere. Very reasonable measure to minimize.

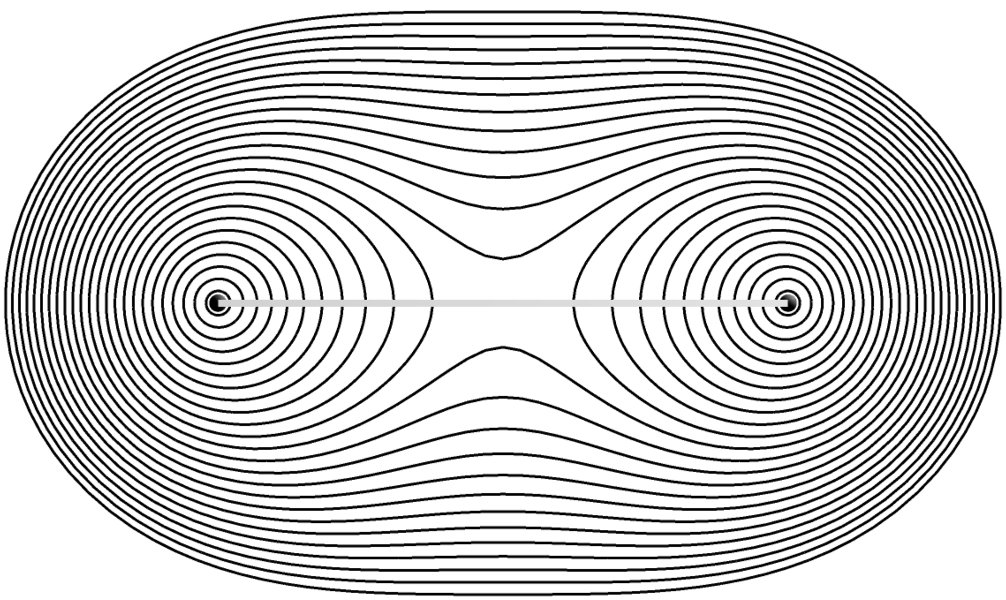

Cross section of equidistance lines of $-(p \cdot X)^2/X^2$ for circle X.

Figure [Perwass 2009].

16 Fitting Hypercircles - Lagrange Multiplier

We need to demand that bivector X is a hypercircle. It must be factorizable; in n-D, that is the constraint:

$$X \wedge X = 0.$$

Enforce scalarity by 4-vector Lagrange multiplier μ and scalar product $*$, giving:

$$\boxed{\textit{Find } X \textit{ and } \mu \textit{ that minimize} \quad \mathcal{L}(X, \mu) = \frac{X * P'[X]}{X * X} + \mu * (X \wedge X)}$$

with P' the 'conformal inertia' operator:

$$P' : \bigwedge^2 \mathbb{R}^{n+1,1} \to \bigwedge^2 \mathbb{R}^{n+1,1} : \quad X \quad \mapsto \quad \tfrac{1}{N} \Sigma_i \, p_i \wedge (p_i \cdot X). \tag{2}$$

Solution to this problem, after a bit of work:

The optimal dual hypercircle X_* is the eigenbivector for smallest non-negative eigenvalue $\mathcal{L}_*$ of the operator P', also satisfying $X \wedge X = 0$.

17 Matlab Implementation of Circle Fitting in 3D (But We'll Do Better Soon)

- For $n = 3$, we use the 10-D bivector basis $\{e_{23}, e_{31}, e_{12} \,|\, e_{o1}, e_{o2}, e_{o3} \,|\, e_{1\infty}, e_{2\infty}, e_{3\infty} \,|\, e_{o\infty}\}$.
- Compute the 10×10 matrix of $[P']$:

$$[P'] = \Sigma_i \left[\begin{array}{c|c|c|c} -[\mathbf{p}_i^\times]^2 & -\frac{1}{2}\mathbf{p}_i^2\,[\mathbf{p}_i^\times] & [\mathbf{p}_i^\times] & 0 \\ \hline -[\mathbf{p}_i^\times] & [\mathbf{p}_i^\times]^2 + \frac{1}{2}\mathbf{p}_i^2\,[\mathbf{I}_{3\times 3}] & [\mathbf{I}_{3\times 3}] & -[\mathbf{p}_i] \\ \hline [\mathbf{p}_i^\times] & \frac{1}{4}\mathbf{p}_i^4\,[\mathbf{I}_{3\times 3}] & [\mathbf{p}_i^\times]^2 + \frac{1}{2}\mathbf{p}_i^2\,[\mathbf{I}_{3\times 3}] & -\frac{1}{2}[\mathbf{p}_i^3] \\ \hline 0^T & \frac{1}{2}[\mathbf{p}_i^3]^T & [\mathbf{p}_i]^T & -\mathbf{p}_i^2 \end{array}\right],$$

where $[\mathbf{p}_i^\times]$ is the cross product matrix, and $[\mathbf{p}_i^3] = \mathbf{p}_i^2\,[\mathbf{p}_i]$.

- Compute the eigenvector of smallest nonnegative eigenvalue of this matrix using a linear algebra package.
- Check if it is a 2-blade. *To our (initial) surprise, all eigenvectors are eigen-2-blades!*
- Unpack the resulting 10D eigen(bi)vector into meaningful circle parameters (tedious but straightforward, details later).

18 This Works

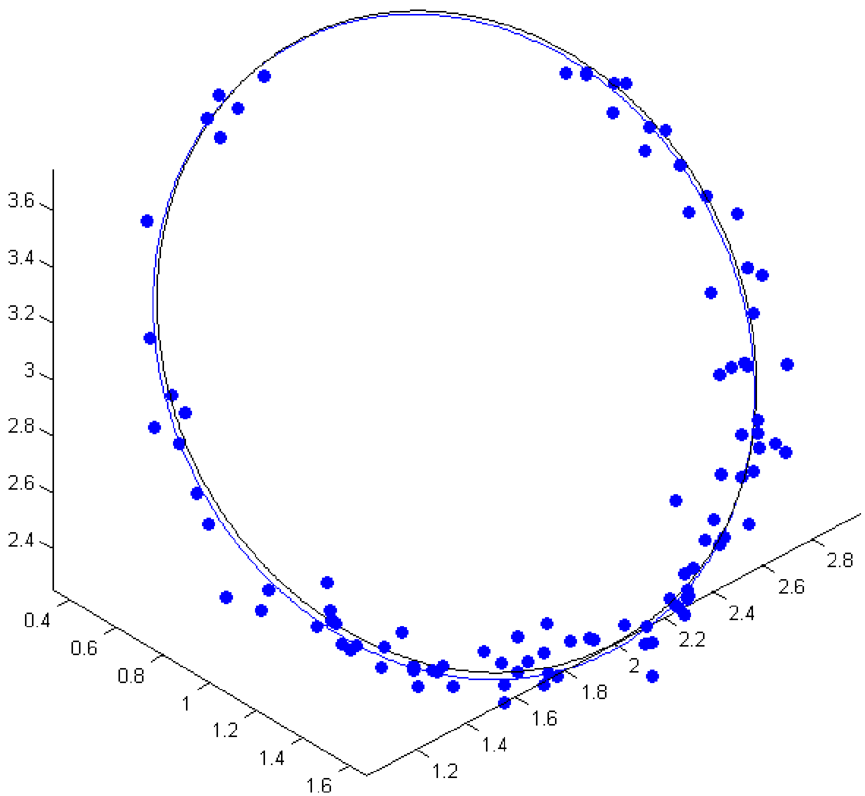

`circle_fit(100,0.1,0.2)`

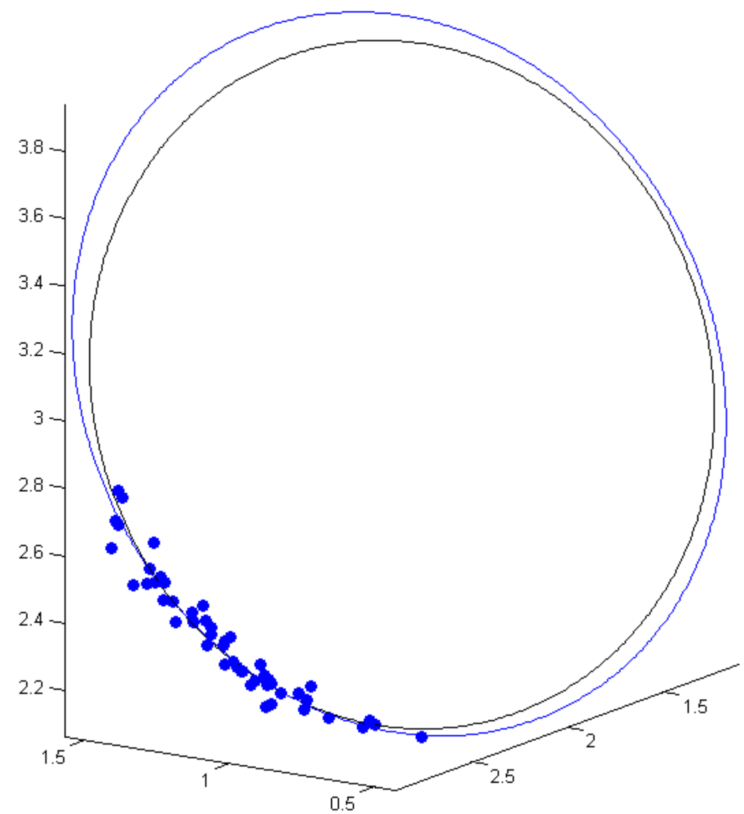

`circle_fit(50,0.05,0.05)`

Black is ground truth circle for noisy point generation; blue is best fit circle.

19 Why Is Every Eigenbivector of P' a 2-Blade?

Lemma: *Let x_j and x_k be eigenvectors of P, with eigenvalues λ_j and λ_k, respectively. Then $x_j \wedge x_k$ is an eigen-2-blade of P' with eigenvalue $(\lambda_j + \lambda_k)$.*

Proof:

$$
\begin{aligned}
P'[x_j \wedge x_k] &= \Sigma_i\, p_i \wedge \big(p_i \cdot (x_j \wedge x_k)\big) && \text{definition of } P' \\
&= \Sigma_i\, p_i \wedge \big(\, (p_i \cdot x_j)\, x_k - x_j\, (p_i \cdot x_k)\big) && \text{dot product on 2-blade} \\
&= P[x_j] \wedge x_k - P[x_k] \wedge x_j && \text{definition of } P' \\
&= (\lambda_j + \lambda_k)\, x_j \wedge x_k. && \text{eigenvectors!} \qquad \square
\end{aligned}
$$

Lemma: *All eigenbivectors of P' are 2-blades.*
Proof:

The operator P has an eigenbasis of $(n+2)$ vectors. From these we can construct $\binom{n+2}{2}$ 2-blade elements. Each of those is an eigenbivector (even an eigen-2-blade) of P'. But P' should have an eigenbasis of $\binom{n+2}{2}$ 'eigenvectors' in its bivector space. It follows that we can construct each eigen(bi)vector of P' from eigenvectors of P, so all are 2-blades. $\square$

Small print: this holds strictly only if eigenvalues of P are all different; that happens in practice.

20 The Eigenvectors of $[P]$ Represent Orthogonal Spheres

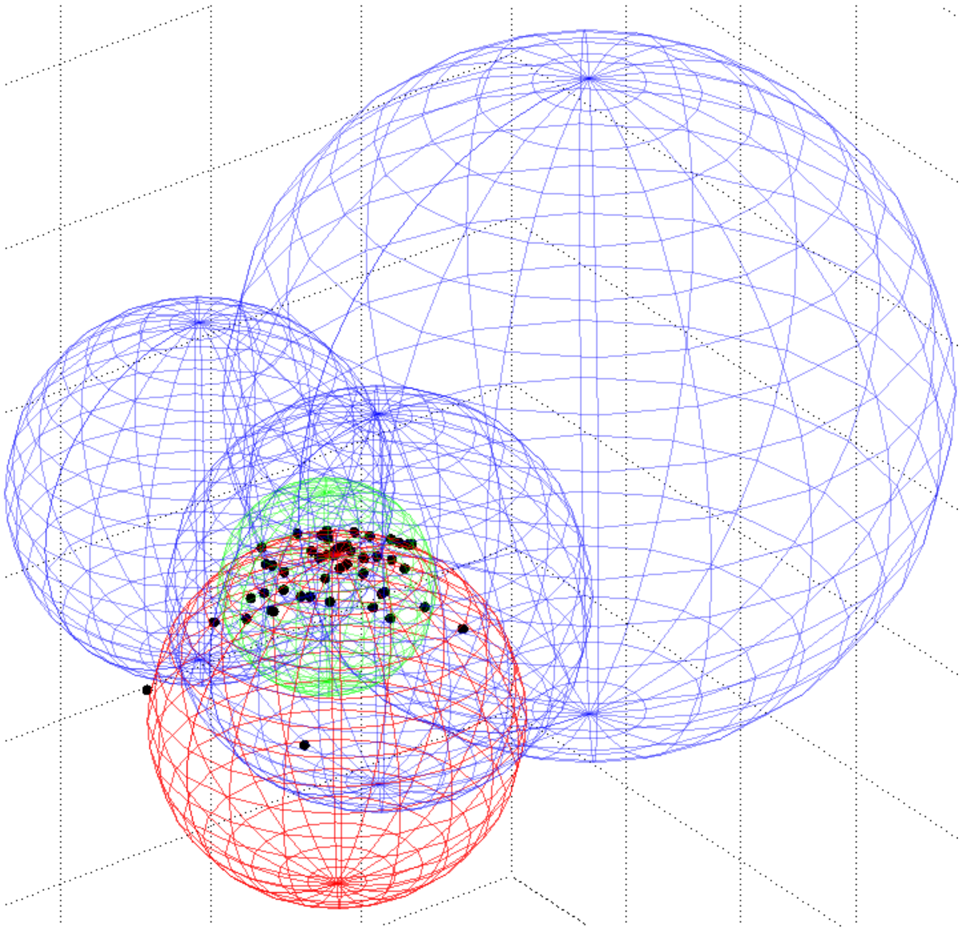

$P[\,]$ is a symmetric operator, so its eigenvectors form an orthonormal basis for $\mathbb{R}^{n+1,1}$.

They represent $(n+2)$ orthogonal spheres! Such spheres have been studied before [Raynor 1934].

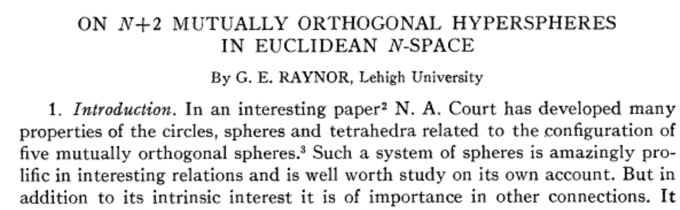

ON $N+2$ MUTUALLY ORTHOGONAL HYPERSPHERES IN EUCLIDEAN N-SPACE

By G. E. RAYNOR, Lehigh University

1. *Introduction*. In an interesting paper[2] N. A. Court has developed many properties of the circles, spheres and tetrahedra related to the configuration of five mutually orthogonal spheres.[3] Such a system of spheres is amazingly prolific in interesting relations and is well worth study on its own account. But in addition to its intrinsic interest it is of importance in other connections. It

These spheres intersect orthogonally in circles.
By the lemma on eigenvalues, the two best spheres give the best circle!

21 On Second Thoughts, The 5-Sphere Basis Is Not Surprising

The usual orthonormal basis $\{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3, e_+, e_-\}$ of $\mathbb{R}^{n+1,1}$ consists of 3 dual coordinate planes, and a real and imaginary dual sphere. By a conformal versor (with $\binom{n+2}{2}$ DoF), these can be transformed into other spheres without affecting their orthogonality.

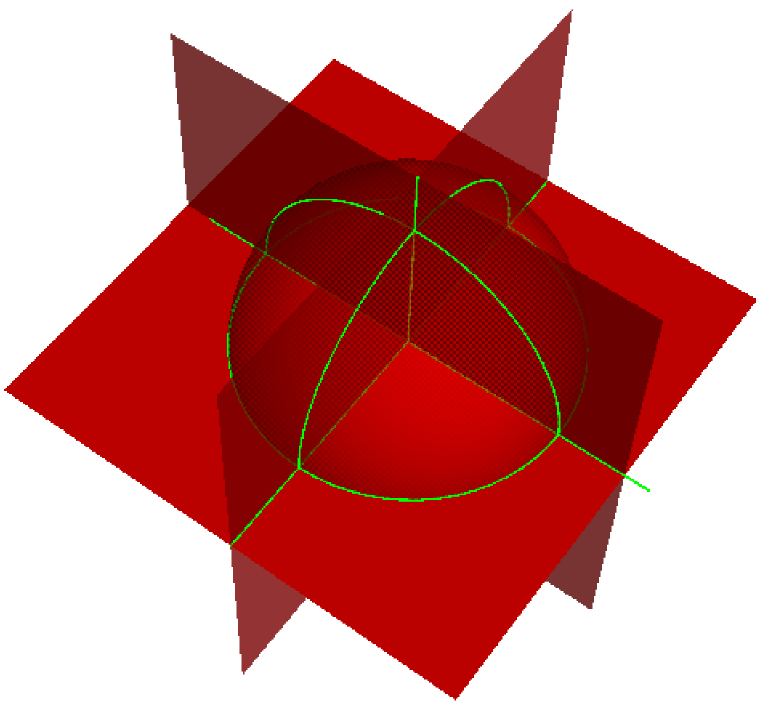

$\xrightarrow{V}$

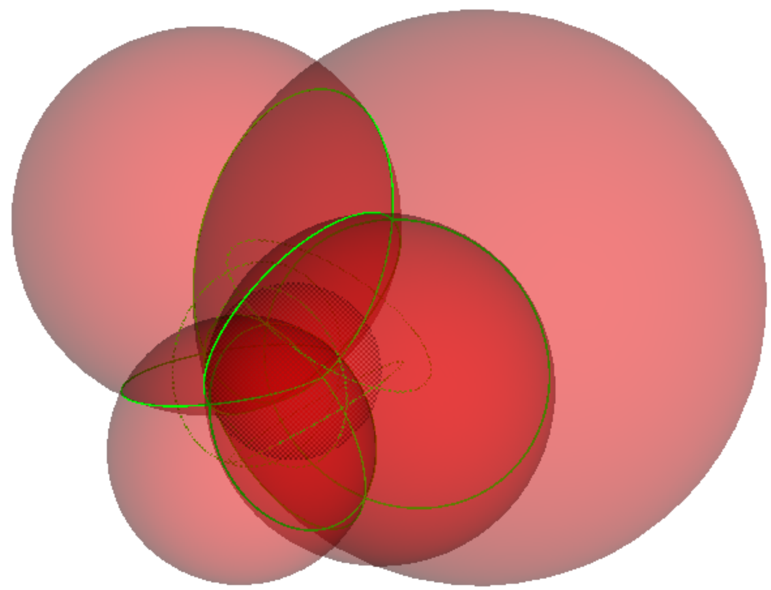

22 Algorithm for Optimal (Hyper-)Circle Fitting

No need to construct $\binom{n+2}{2} \times \binom{n+2}{2}$-dimensional matrix $[P']$ and solve its eigenproblem. Use the $(n+2) \times (n+2)$ matrix $[P]$ instead! (For 3D, 5×5 rather than 10×10.)

- Set up $[P] = [D]^T [D] [M]$ as before.
- Solve the eigenproblem for $[P]$, and save the two eigenvectors x_1 and x_2 with smallest non-negative eigenvalues.
- Compute the intersection $x_1 \wedge x_2$ of the two hyperspheres x_1 and x_2. On the vector and bivector bases given before, this employs an $\binom{n+2}{2} \times (n+2)$ matrix:

$$[y \wedge x] = \begin{bmatrix} [\mathbf{y}^\times] & \mathbf{0} & \mathbf{0} \\ y_o\,[1] & -\mathbf{y} & 0 \\ -y_\infty\,[1] & 0 & \mathbf{y} \\ \mathbf{0}^T & -y_\infty & y_o \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ x_o \\ x_\infty \end{bmatrix}.$$

- Interpret the eigenbivector components as hypercircle parameters.

23 Unpacking the Hypercircle Parameters (Gory but Straightforward)

The general expression for a (dual) circle κ in CGA, as the intersection of a hyperplane π with normal vector $\mathbf{n}$ containing c, and a sphere σ with radius ρ around c:

$$\begin{aligned}
\kappa &= \pi \wedge \sigma \\
&= \alpha \big(\mathbf{n} + (\mathbf{c} \cdot \mathbf{n})\, n_\infty\big) \wedge \big(n_o + \mathbf{c} + \tfrac{1}{2}(\mathbf{c}^2 - \rho^2)\, n_\infty\big) \\
&= \alpha \Big(\mathbf{n} \wedge n_o + \mathbf{n} \wedge \mathbf{c} - (\mathbf{c} \cdot \mathbf{n})\, n_o \wedge n_\infty + \big(\tfrac{1}{2}(\mathbf{c}^2 - \rho^2)\, \mathbf{n} - (\mathbf{c} \cdot \mathbf{n})\, \mathbf{c}\big)\, n_\infty\Big).
\end{aligned}$$

- Thus $-\alpha\mathbf{n}$ can be retrieved immediately as the components of $\{e_{oi}\}$, normalization then splits it in α and $\mathbf{n}$ if necessary.
- The Euclidean e_{ij} and $e_{o\infty}$ parts then give the outer and inner product of $\mathbf{c}$ and $\mathbf{n}$. Using the matrix implementation of the geometric product:

$$\begin{bmatrix} \mathbf{n}^T \\ \mathbf{n}^\times \end{bmatrix} [\mathbf{c}] = \begin{bmatrix} \mathbf{n} \cdot \mathbf{c} \\ (\mathbf{n} \wedge \mathbf{c})^* \end{bmatrix}$$

  we can solve for $\mathbf{c}$. (Effectively, geometric division in 3D, as implemented in linear algebra.)

- With α, $\mathbf{n}$ and $\mathbf{c}$ known, ρ^2 can be derived from the $\{e_{i\infty}\}$ component vector $\mathbf{v}_\infty$ as $\rho^2 = \|\mathbf{c}\|^2 - 2\mathbf{n} \cdot \mathbf{v}_\infty/\alpha - 2(\mathbf{c} \cdot \mathbf{n})^2$.

24 Unpacking the Circle Parameters with CGA Software

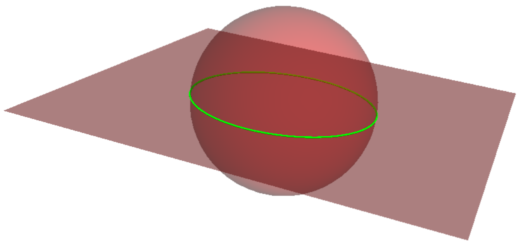

Given a (dual) circle κ, retrieve its parameters $\mathbf{n}$, $\mathbf{c}$, ρ.

View the circle as formed by intersection of a sphere σ with a plane π:

$$\sigma \wedge \pi = \kappa. \tag{3}$$

First, let us find the plane of κ. Just wedge the point at infinity onto it:

$$\pi^* = n_\infty \wedge \kappa^*, \qquad \text{which means that } \; \pi = n_\infty \cdot \kappa. \tag{4}$$

Its normal vector is $\mathbf{n}$, easy to read out as Euclidean part after normalization as $\pi/\sqrt{\pi^2}$.

Now find the center and radius of the encompassing sphere:

$$\sigma \cdot \pi = 0 \quad \text{(impose orthogonality!)}. \tag{5}$$

Adding those equations, $\sigma\,\pi = \kappa$. Geometric product invertible, so:

$$\sigma = \kappa/\pi = \kappa/(n_\infty \cdot \kappa).$$

This sphere is normalized, so read off Euclidean part as $\mathbf{c}$, and $\rho^2 = \sigma^2$.

25 The Geometry of Fitting Spheres and Circles (in 3D)

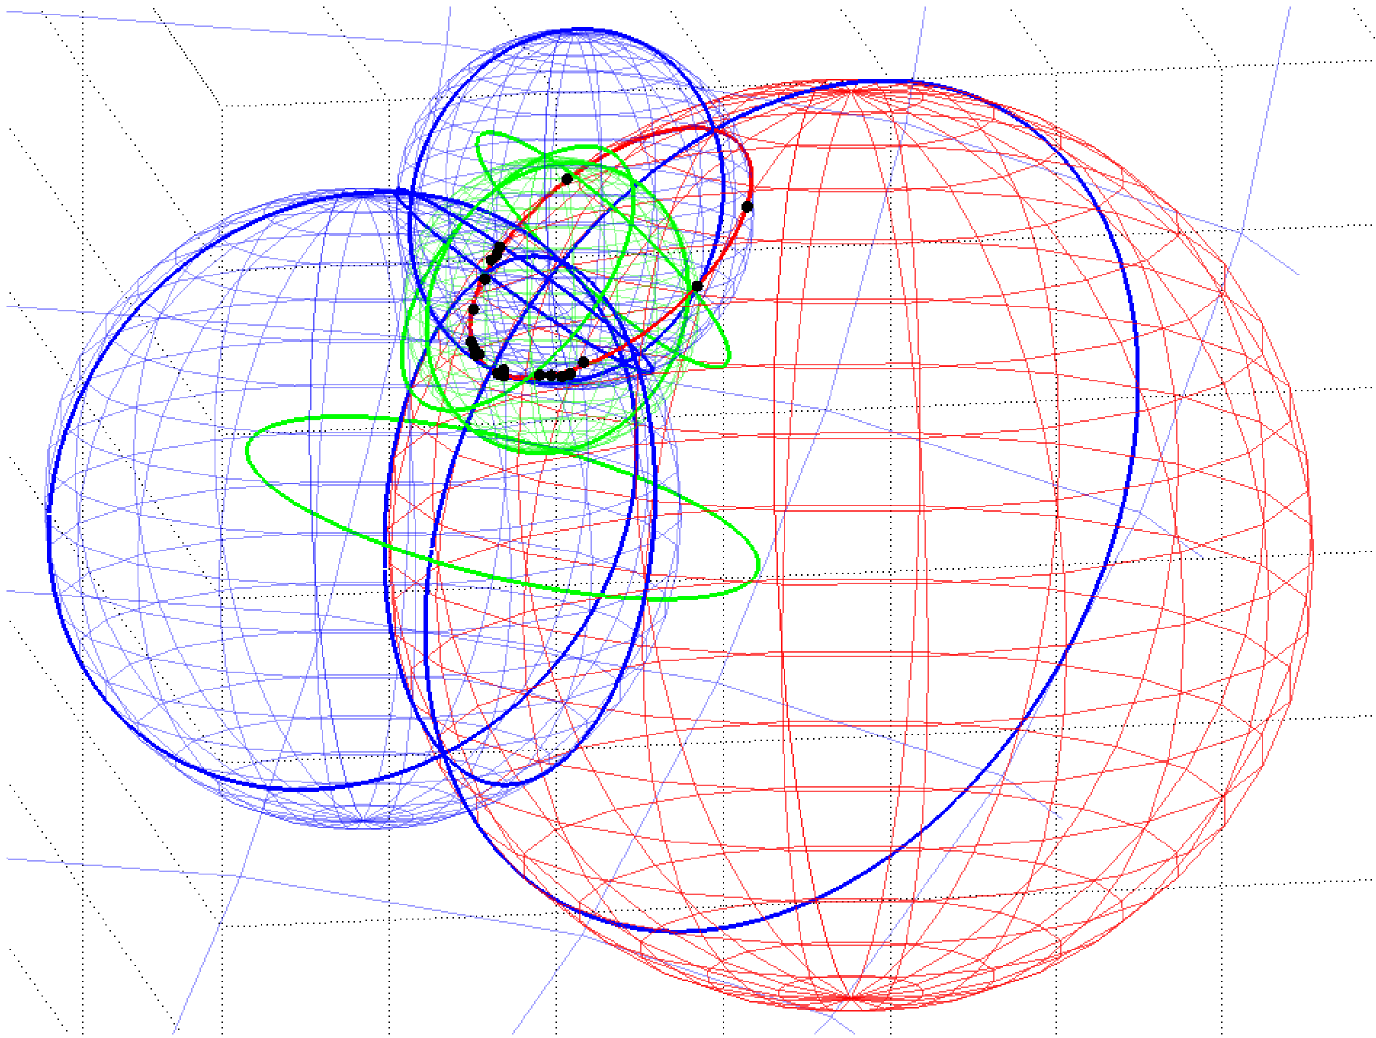

Matlab

26 The Geometry of Fitting Spheres and Circles (in 3D)

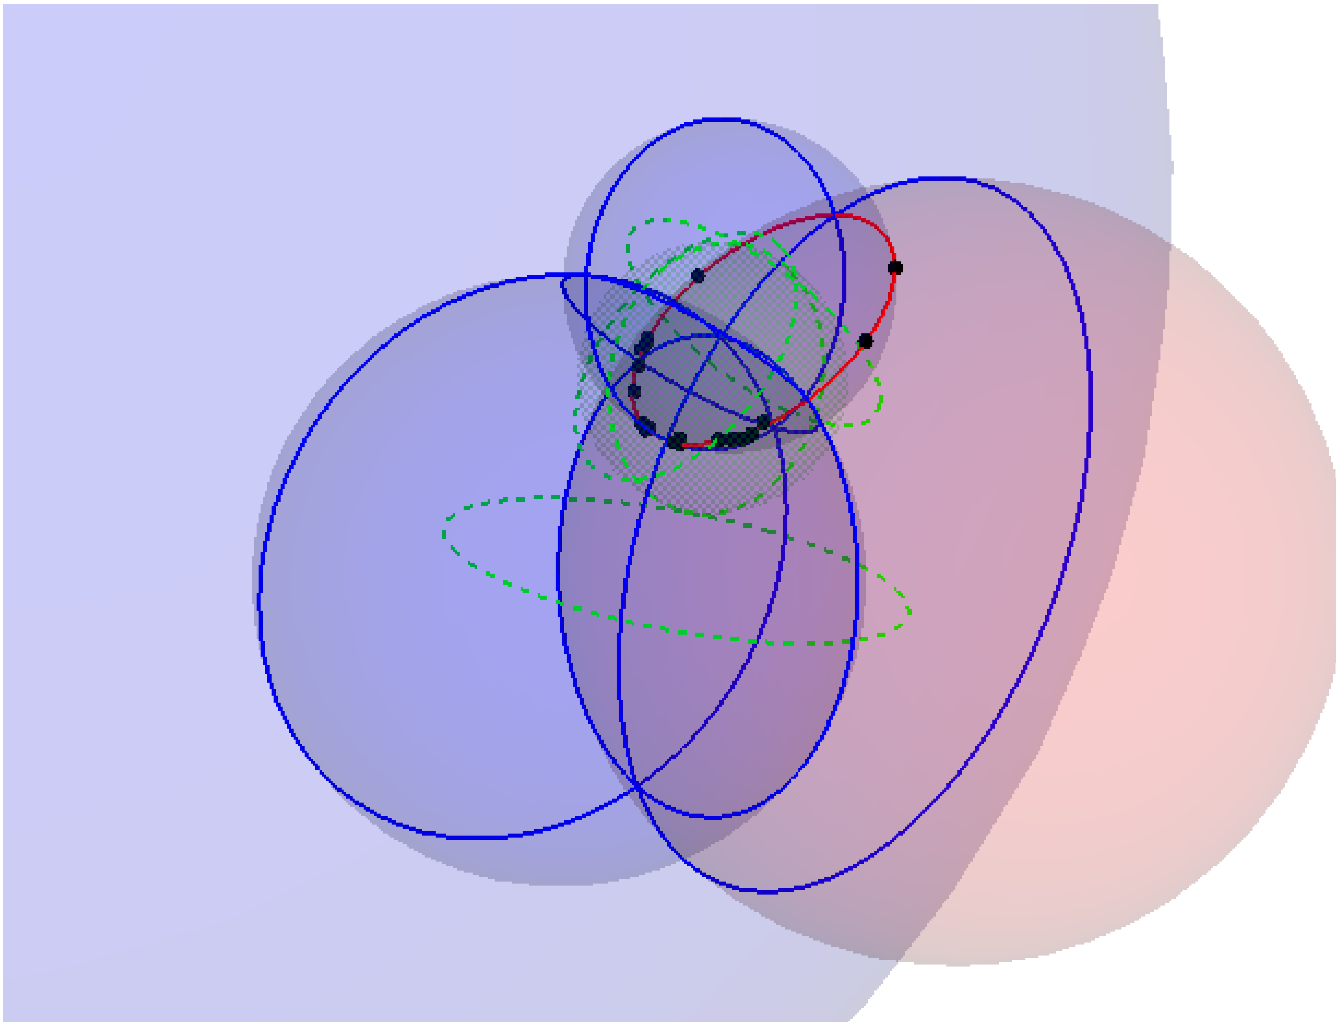

Demonstration in `GAViewer: ganew/sphere fit/sphere_eigen()`

27 The Fits Are Optimal (Though Not 'Hyperaccurate')

- Good overview of 2-D circle fitting methods in [Al-Sharadqah & Chernov 2009].
  Our hypersphere is n-D version of 2D circle 'algebraic fit' from [Pratt 1987].
- Optimal in MSE accuracy: achieves KCR lower bound of variance. Optimal in speed.
- Especially: our fit is as optimal as the fit according to geometric least squares.
  (which is 20 times slower, due to e.g Levenberg-Marquardt)
- For very large number of points $N > 1000$, there exists a hyperaccurate fit (see [Al-S&Ch]).
  Surprise: it is *not* geometric least squares, that is biased!
- Elegance: In LA, Pratt fit gives a generalized eigenproblem. In CGA, pure eigenproblem.
- Relationship of circle fit to sphere fit is new.
  Best 2D circle fit in 3D is intersection of two best orthogonal spheres.
  Best 2D circle fit in 3D is *not* the best circle in the best plane!
- We have extended our method to k-spheres in n-D, for JMIV (submitted March 2013).
  (Extra for 3D: optimally fit point pair without splitting the data.)
- Plane and line fits can be done in CGA too, and also lead to pure eigenproblems.

Journal of Mathematical Imaging and Vision, 2014, DOI 10.1007/s10851-014-0495-2

Total Least Squares Fitting of k-Spheres in n-D Euclidean Space Using an $(n+2)$-D Isometric Representation

Leo Dorst

Abstract We fit k-spheres optimally to n-D point data, in a geometrically total least squares sense. A specific practical instance is the optimal fitting of 2D-circles to a 3D point set.

Among the optimal fitting methods for 2D-circles based on 2D (!) point data compared in Al-Sharadqah and Chernov (Electron. J. Stat. 3:886–911, 2009), there is one with an algebraic form that permits its extension to optimally fitting k-spheres in n-D. We embed this 'Pratt 2D circle fit' into the framework of conformal geometric algebra (CGA), and doing so naturally enables the generalization. The procedure involves a representation of the points in n-D as vectors in an $(n+2)$-D space with attractive metric properties. The hypersphere fit then becomes an eigenproblem of a specific symmetric linear operator determined by the data. The eigenvectors of this operator form an orthonormal basis representing perpendicular hyperspheres. The intersection of these are the optimal k-spheres; in CGA the intersection is a straightforward outer product of vectors.

The resulting optimal fitting procedure can easily be implemented using a standard linear algebra package; we show this for the 3D case of fitting spheres, circles and point pairs. The fits are optimal (in the sense of achieving the KCR lower bound on the variance).

We use the framework to show how the hyperaccurate fit hypersphere of Al-Sharadqah and Chernov (Electron. J. Stat. 3:886–911, 2009) is a minor rescaling of the Pratt fit hypersphere.

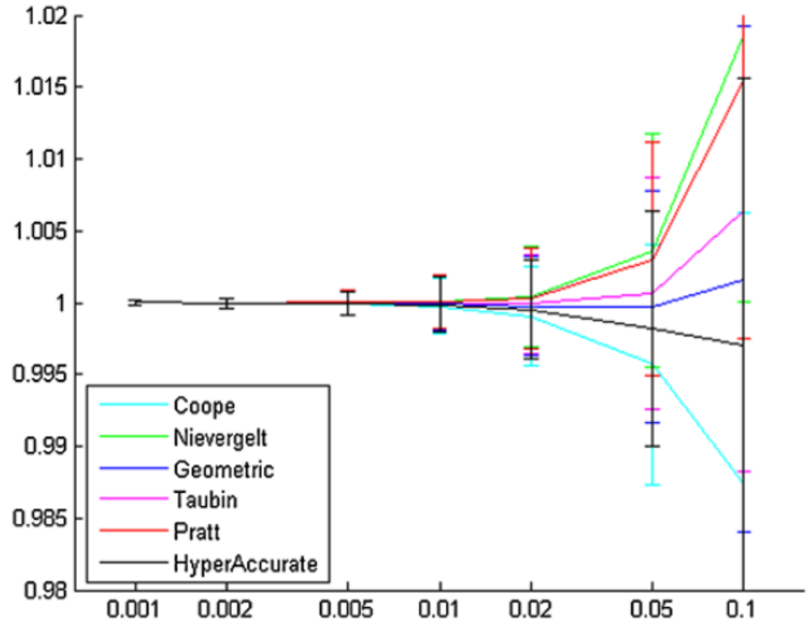

Fig. 8 Radius determination as a function of the radial noise standard deviation σ_{radial}, for spheres based on 100 data points, generated from a unit sphere, with angular standard deviation of 1 radian. With 50 trials per fit, we show average and standard deviation. Note the scale, all fits perform well (Color figure online)

References

1. A. Al-Sharadqah and N. Chernov. Error analysis for circle fitting algorithms. *Electron. J. Statist.*, 3:886–911, 2009.
2. Chris Doran and Anthony Lasenby. *Geometric Algebra for Physicists*. Cambridge University Press, 2003.
3. L. Dorst, D. Fontijne, and S. Mann. *Geometric Algebra for Computer Science: An Object-oriented Approach to Geometry*. Morgan Kaufman, 2009.
4. Michael G. Eastwood and Peter W. Michor. Some remarks on the Plücker relations. *Rendiconti del Circolo Matematico di Palermo*, II-63:85–88, 2000.
5. D. Hestenes and G. Sobczyk. *Clifford Algebra to Geometric Calculus*. Reidel, 1984.
6. K. Kanatani. Optimal estimation. In K. Ikeuchi, editor, *Encyclopedia of Computer Vision*. Springer, to be published 2012.
7. C. Perwass and W. Förstner. Uncertain geometry with circles, spheres and conics. In R. Klette, R. Kozera, L. Noakes, and J. Weickert, editors, *Geometric Properties from Incomplete Data*, volume 31 of *Computational Imaging and Vision*, pages 23–41. Springer-Verlag, 2006.
8. V. Pratt. Direct least-squares fitting of algebraic surfaces. *Computer Graphics*, 21:145–152, 1987.
9. G.E. Raynor. On $n+2$ mutually orthogonal hyperspheres in Euclidean n-space. *American Mathematical Monthly*, 41(7):424–438, 1934.
10. Alyn Rockwood and Dietmar Hildenbrand. Engineering graphics in geometric algebra. In E. Baryo-Corrochano and G. Scheuermann, editors, *Geometric Algebra Computing*, pages 53–67. Springer, 2010.

28 Note on Hyperaccuracy (Kanatani 2012, 'my best work')

Kanatani says: In geometric data processing we do not want estimators with good asymptotic behavior in the limit of infinite data, and/or large variance (the classical approach in estimation): *we have finite/minimal amount of data N, of usually rather small variance σ.*

The Mean Square Error of a *consistent* estimator (which returns the true value when $\sigma = 0$) can be shown to be:

$$\text{MSE} = \text{variance} + \text{bias}^2 \approx O(\sigma^2/N) + O(\sigma^4).$$

An *optimal estimator* minimizes the variance (achieves the 'Kanatani-Cramèr-Rao lower bound').

For large enough N, the bias term may become important, even for small σ.
A *hyperaccurate estimator* makes the $O(\sigma^4)$ term (the 'essential bias') equal to zero.

For 2D circles, a hyperaccurate estimator has been found in 2009 by Al-Sharadqah & Chernov, prompting Kanatani to develop a general theory recently.

For 2D circles, it manifests itself when $\sigma = 0.05\rho$ for $N > 1000$.
For smaller N, all optimal estimators are equivalent. Our fit is optimal.