

Fitting quadratic curves to data points

Abstract

Fitting quadratic curves (a.k.a. conic sections, or conics) to data points (digitized images) is a fundamental task in image processing and computer vision. This problem reduces to minimization of a certain function over the parameter space of conics. Here we undertake a thorough investigation of that space and the properties of the objective function on it. Our analysis leads us to many important conclusions relevant to the performance of minimization algorithms.

Keywords: Least squares fitting, Ellipses, Conic sections, Minimization

AMS classification: 62H35, 62J99

1 Introduction

In many areas of human practice one needs to approximate a set of planar points representing experimental data or observations by an ellipse [9, 11, 12, 23] or hyperbola [10, 24] or by any conic including a parabola [3, 8, 27, 28, 29]. This task is popular in image processing and modern computer vision.

The classical least squares fit minimizes geometric distances from the observed points to the fitting curve:

$$(1) \quad \mathcal{F}(S) = \sum_{i=1}^n [\text{dist}(P_i, S)]^2 \rightarrow \min$$

where $P_1, \dots, P_n$ denote the observed points and $S \subset \mathbb{R}^2$ the fitting curve.

The geometric fit (1) has many attractive features. It is invariant under translations, rotations, and scaling, i.e., the best fitting curve does not depend on the choice of the coordinate system. It provides the maximum likelihood estimate under standard statistical assumptions [14, 18, 19]. The minimization of geometric distances is often regarded as the most desirable solution¹ of the fitting problem, albeit hard to compute in many cases.

Figure 1 shows a sample of eight points (their coordinates are given in Table 1; they are borrowed from [23]) and the best fitting ellipse obtained by (1). We will explore this example in Section 15.




Figure 1: A sample of eight points and the best fitting ellipse.

When one fits straight lines, the problem (1) has a closed form solution, and its properties have been studied deeply [15, 19]. But when one fits quadratic curves (ellipses, hyperbolas, etc.), no closed form solution exists, and the problem happens to be extremely complicated. All the practical algorithms are iterative, they require a carefully chosen initial guess, they are computationally intense, and they suffer from frequent failures (divergences and traps in local minima).

The first thorough investigation of the ellipse fitting methods was done in the middle 1990's by Gander, Golub, and Strebel [23], and they basically arrived at a conclusion that the minimization of geometric distances for ellipses was a prohibitively difficult task. Recently Sturm and Gargallo

¹In particular, it has been prescribed by a recently ratified standard for testing the data processing software for coordinate metrology [8].

[31] modified their methods in several ways, in particular they incorporated other quadratic curves (hyperbolas and parabolas) into the minimization procedure, so it could freely switch between different types of conics during iterations. But their new methods are still computationally very intense and hardly practical.

In the early 2000s, more efficient approaches to conic fitting were found by Ahn and others [3, 4, 5, 8], further modified by Aigner and Jüttler [1, 2] and by us [20]. But still all the existing methods are computationally intense and suffer from frequent failures, and they have not become universally accepted in the computer vision community yet.

Instead, the so called algebraic ellipse fitting schemes became widely popular, in which algebraic distances from data points to the ellipse are minimized. They produce ellipses that fit less accurately than those minimizing geometric distances (1). Some authors say that “the performance gap between algebraic fitting and geometric fitting is wide...” (see [8, p. 12]).

Back to the geometric fit (1), nearly all popular algorithms use standard numerical schemes for least squares minimization: Steepest Descent, Gauss-Newton, Levenberg-Marquardt, Trust Region, etc. (see, e.g., a survey in [19, Chapter 4]). They are based on the first derivatives of the objective function $\mathcal{F}$ with respect to the parameters of the fitting curve; they are all heuristic and heavily depend on a good choice of initial guess.

No body seems to have investigated general properties of the objective function $\mathcal{F}$ on the parameter space of conics: its continuity (or lack thereof), its differentiability (or lack thereof), the types of its singularities, etc. We undertake this task here. We thoroughly describe the parameter space of all conics and study the property of $\mathcal{F}$ on it. Our focus is, of course, on the features of that space and the function $\mathcal{F}$ that are most relevant to the performance of minimization algorithms. We discover many unknown and sometimes unexpected facts. In the end, we make practical conclusions and illustrate them by a numerical experiment.

2 Algebraic parameters

A *quadratic curve* (or *conic*) is defined by quadratic equation

$$(2) \quad Ax^2 + 2Bxy + Cy^2 + 2Dx + 2Ey + F = 0,$$

where A, B, C, D, E, F are parameters of the conic. Since they are the coefficients of a quadratic polynomial (i.e., an algebraic expression), they are called *algebraic parameters* of the conic. They can be regarded as components of a vector, $\mathbf{A} = (A, B, C, D, E, F)^T$.

One exceptional case ($A = B = C = D = E = F = 0$) has to be excluded, because then the respective “conic” would be the entire plane $\mathbb{R}^2$. Furthermore, the conic does not change under scalar multiplication $\mathbf{A} \mapsto \alpha \mathbf{A}$, $\alpha \neq 0$, thus we can assume that every parameter vector has length one: $A^2 + B^2 + C^2 + D^2 + E^2 + F^2 = 1$. Now the parameter space is the unit sphere $\mathbb{S}^5 \subset \mathbb{R}^6$. We may further reduce the parameter space by identifying antipodal points of $\mathbb{S}^5$, i.e., $\mathbf{A}$ and $-\mathbf{A}$, and get a projective sphere. But this has little advantage, so we will use the sphere $\mathbb{S}^5$ and simply keep in mind that antipodal points always represent the same conic.

There are two special points on $\mathbb{S}^5$: $(0, 0, 0, 0, 0, 1)$ and its antipode $(0, 0, 0, 0, 0, -1)$. They correspond to “quadratic equations” $1 = 0$ and $-1 = 0$ that have no solutions, either real or complex. We call them *poles* (like North Pole and South Pole).

For every vector $\mathbf{A} \in \mathbb{S}^5$ other than the above two poles, the quadratic equation (2) has solutions, either real or complex. If it has real solutions, those make conics. If it has only complex (non-real) solutions, then the corresponding figure in $\mathbb{R}^2$ does not exist (is an empty set), but it has a certain meaning (and name) in complex coordinates, and we will call it accordingly: imaginary ellipse, imaginary lines, etc.

3 Classification of conics

Below we list the main types of conics (real and imaginary), with examples of equations representing them:

Real non-degenerate	Real degenerate	Empty in $\mathbb{R}^2$
Ellipse $x^2 + y^2 - 1 = 0$	Intersecting lines $xy = 0$	Imaginary ellipse $x^2 + y^2 + 1 = 0$
Hyperbola $x^2 - y^2 - 1 = 0$	Parallel lines $x^2 - 1 = 0$	Im. parallel lines $x^2 + 1 = 0$
Parabola $x^2 + y = 0$	Single line $x^2 = 0$ or $x = 0$	Poles $1 = 0, \quad -1 = 0$
	Single point $x^2 + y^2 = 0$	

We note that there are two types of equations representing single lines: those with a non-zero quadratic part, such as $x^2 = 0$ (we will call them *pairs of coincident lines*) and linear equations, such as $x = 0$ (we will call them just *single lines*).

Quadratic equation (2) can be written in matrix form as

$$\begin{bmatrix} x & y & 1 \end{bmatrix} \begin{bmatrix} A & B & D \\ B & C & E \\ D & E & F \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = 0,$$

and its quadratic part $Ax^2 + 2Bxy + Cy^2$ can be written in matrix form as $\begin{bmatrix} x & y \end{bmatrix} \begin{bmatrix} A & B \\ B & C \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}$. We denote

$$\Delta = \det \begin{bmatrix} A & B & D \\ B & C & E \\ D & E & F \end{bmatrix}, \quad J = \det \begin{bmatrix} A & B \\ B & C \end{bmatrix}.$$

In textbooks, Δ is called the *determinant* and J the *discriminant* of the corresponding quadratic form. A few other useful quantities are

$$I = A + C, \quad Q = A^2 + B^2 + C^2$$

$$K = \det \begin{bmatrix} A & D \\ D & F \end{bmatrix} + \det \begin{bmatrix} C & E \\ E & F \end{bmatrix}$$

The types of conics are classified in terms of the above quantities in the following table; see [13, pp. 200–201].

Q	Δ	J	K	$\Delta \cdot I$	Type of conic	GD	AD
> 0	$\neq 0$	> 0		< 0	Ellipse	5	5
> 0	$\neq 0$	< 0			Hyperbola	5	5
> 0	$\neq 0$	0			Parabola	4	4
> 0	0	< 0			Intersecting lines	4	4
> 0	0	0	< 0		Parallel lines	3	3
> 0	0	0	0		Coincident lines	2	2
0			< 0		Single line	2	2
> 0	0	> 0			Single point	2	4
> 0	$\neq 0$	> 0		> 0	Imaginary ellipse	None	5
> 0	0	0	> 0		Im. parallel lines	None	3
0			0		Poles	None	0

The last two columns show the geometric dimension (GD) and algebraic dimension (AD) for each type of conics. The *geometric dimension* is the number of parameters needed to specify a geometric figure of the given type. For example, a single point requires two parameters - its x and y coordinates, so the geometric dimension for the “single point” type is two. A line requires two parameters (say, slope and intercept), so its geometric dimension is also two. A pair of parallel lines needs three parameters - one (common) slope and two intercepts. A parabola is completely specified by its directrix and focus, which require four parameters. An ellipse can be specified by five geometric parameters - the coordinates of its center, the lengths of its axes, and the slope of its major axis. The same applies to hyperbolas. The algebraic dimension of each type of conics characterizes the corresponding set of parameter vectors in the unit sphere $\mathbb{S}^5$.

Typically, the geometric dimension agrees with the algebraic dimension. But there is one exception: the single point type has geometric dimension 2 and algebraic dimension 4. Indeed, a single point can be specified by a quadratic equation

$$a^2(x - p)^2 + b^2(y - q)^2 + 2c(x - p)(y - q) = 0$$

with additional constraint $c^2 < a^2b^2$. It can be written as

$$\left[a(x - p) + \frac{c}{a}(y - q) \right]^2 + \left[b^2 - \frac{c^2}{a^2} \right] (y - q)^2 = 0$$

and $c^2 < a^2b^2$ guarantees that both coefficients are positive. Thus each term in the above equation must be equal to zero, which gives us $y = q$ and then $x = p$ as the only solution (a single point).

In the above equation all the five parameters a, b, c, p, q are free independent variables, except they must be constrained by the requirement that the resulting parameter vector (A, B, C, D, E, F) belongs to the unit sphere $\mathbb{S}^5$. This leaves us with four degrees of freedom, thus $AD=4$.

4 Structure of the sphere $\mathbb{S}^5$

Our parameter space $\mathbb{S}^5$ is now divided into 10 domains corresponding to the conic types, plus two extra points (the poles):

$$\mathbb{S}^5 = \mathbb{D}_E \cup \mathbb{D}_H \cup \mathbb{D}_{IE} \cup \mathbb{D}_P \cup \mathbb{D}_{SP} \cup \mathbb{D}_{IL} \cup \mathbb{D}_{PL} \cup \mathbb{D}_{IPL} \cup \mathbb{D}_{CL} \cup \mathbb{D}_{SL} \cup \{\mathbf{P}_1\} \cup \{\mathbf{P}_{-1}\},$$

where the domains are coded by the names of the conic types: E for ellipses, IE for imaginary ellipses, etc., and $\mathbf{P}_{\pm 1} = (0, 0, 0, 0, 0, \pm 1)$.

Out of these, only three domains have non-empty interior, hence positive volume: these are open domains $\mathbb{D}_E$, $\mathbb{D}_H$, and $\mathbb{D}_{IE}$. The volume of these open domains can be easily estimated by a Monte Carlo experiment:

$$\text{Vol}(\mathbb{D}_E) = 20.7\%, \quad \text{Vol}(\mathbb{D}_H) = 76.8\%, \quad \text{Vol}(\mathbb{D}_{IE}) = 2.5\%$$

(as a percentage of the total volume of $\mathbb{S}^5$); see Figure 2.




Figure 2: Relative volumes of the three open domains in $\mathbb{S}^5$.

We see that hyperbolas occupy more than $3/4$ of the parameter space. The domination of hyperbolas over ellipses is also noted in [21].

We also note that imaginary ellipses occupy just a small fraction of the sphere $\mathbb{S}^5$, only 2.5% of it. This is good for practical fitting methods. Those usually wander in the parameter space searching for the minimum of the objective function. Every time they accidentally run into a forbidden or an unwanted domain, such as $\mathbb{D}_{IE}$, they have to retreat and readjust their step. Fortunately, due to the small size of $\mathbb{D}_{IE}$, this does not happen often.

Our open domains $\mathbb{D}_E$, $\mathbb{D}_H$, and $\mathbb{D}_{IE}$ are not connected, each of them consists of exactly two connected components. Indeed, each parameter vector $\mathbf{A} = (A, B, C, D, E, F)^T \in \mathbb{S}^5$ not only specifies a conic, but defines a quadratic function

$$Q(x, y) = Ax^2 + 2Bxy + Cy^2 + 2Dx + 2Ey + F$$

on the xy plane. Note that $Q(x, y) = 0$ on the conic, but $Q(x, y) > 0$ or $Q(x, y) < 0$ elsewhere. If the conic is an ellipse, we may have $Q > 0$ inside the ellipse and $Q < 0$ outside of it, or vice versa. Thus we have two types of parameter vectors $\mathbf{A} \in \mathbb{D}_E$: for one Q is positive inside the ellipse and for the other Q is negative inside the ellipse. This dichotomy causes $\mathbb{D}_E$ to consist of two pieces, we denote them by $\mathbb{D}_E^+$ and $\mathbb{D}_E^-$, depending on whether $Q > 0$ or $Q < 0$ inside the ellipse (i.e., at its center).




Figure 3: Conics with positive and negative centers.

Similarly, if $\mathbf{A} \in \mathbb{D}_H$, then the corresponding quadratic function $Q(x, y)$ may be positive or negative between the two branches of the hyperbola (i.e., at its center). So we have a partition $\mathbb{D}_H = \mathbb{D}_H^+ \cup \mathbb{D}_H^-$ into two pieces.

Lastly, if $\mathbf{A} \in \mathbb{D}_{IE}$, then the corresponding quadratic function $Q(x, y)$ cannot take zero values, so it is either entirely positive or entirely negative. Again this causes a natural partition of $\mathbb{D}_{IE}$ into two pieces, $\mathbb{D}_{IE}^+$ and $\mathbb{D}_{IE}^-$.

Note that if $\mathbf{A} \in \mathbb{D}_E^+$, then $-\mathbf{A} \in \mathbb{D}_E^-$ and vice versa. Thus the subdomains $\mathbb{D}_E^+$ and $\mathbb{D}_E^-$ are antipodes of each other on the sphere $\mathbb{S}^5$. In a sense, they are “mirror images” of each other; they have identical shapes and equal volumes. The same is true for the two parts of $\mathbb{D}_H$ and the two parts of $\mathbb{D}_{IE}$.

Next we examine the four-dimensional subdomains (hypersurfaces): $\mathbb{D}_P$ (parabolas), $\mathbb{D}_{SP}$ (single points), and $\mathbb{D}_{IL}$ (intersecting lines). They make the boundaries of the above open domains and separate their components from one another.

The hypersurface $\mathbb{D}_{\text{SP}}$ separates the open domain $\mathbb{D}_{\text{E}}$ of ellipses from the open domain $\mathbb{D}_{\text{IE}}$ of imaginary ellipses. To illustrate this fact, consider the parameter vector $\mathbf{A}_c = (-1, 0, -1, 0, 0, c)$, where c will play the role of a small variable (we will not normalize $\mathbf{A}_c$ here to keep our formulas simple). This parameter vector corresponds to the quadratic function

$$Q(x, y) = -x^2 - y^2 + c.$$

For $c > 0$, the equation $Q(x, y) = 0$ defines a small ellipse (more precisely, a small circle of radius $\sqrt{c}$), i.e., $\mathbf{A}_c \in \mathbb{D}_{\text{E}}$ for $c > 0$. For $c = 0$, it is a single point, $(0, 0)$, i.e., $\mathbf{A}_0 \in \mathbb{D}_{\text{SP}}$. For $c < 0$ it is an imaginary ellipse, i.e., $\mathbf{A}_c \in \mathbb{D}_{\text{IE}}$ for $c < 0$. As c changes from small positive values to zero and then on to small negative values, the ellipse shrinks and collapses to a single point, and then disappears (transforms into an imaginary ellipse). In the parameter space $\mathbb{S}^5$, this process corresponds to a continuous motion from the domain $\mathbb{D}_{\text{E}}$ to the domain $\mathbb{D}_{\text{IE}}$, across the hypersurface $\mathbb{D}_{\text{SP}}$.

The hypersurface $\mathbb{D}_{\text{IL}}$ separates the two components $\mathbb{D}_{\text{H}}^+$ and $\mathbb{D}_{\text{H}}^-$ of the open domain $\mathbb{D}_{\text{H}}$ of hyperbolas from each other. To illustrate this fact, consider the parameter vector $\mathbf{A}_c = (-1, 0, 1, 0, 0, c)$, where c is again a small variable. This parameter vector corresponds to the quadratic function

$$Q(x, y) = -x^2 + y^2 + c.$$

The equation $Q(x, y) = 0$ defines a hyperbola with center $(0, 0)$, unless $c = 0$, in which case it is a pair of intersecting lines, $y = \pm x$. More precisely, for $c > 0$ it is a hyperbola with positive center, because $Q(0, 0) > 0$, i.e., $\mathbf{A}_c \in \mathbb{D}_{\text{H}}^+$ for $c > 0$. For $c < 0$, it is a hyperbola with negative center, because $Q(0, 0) < 0$, i.e., $\mathbf{A}_c \in \mathbb{D}_{\text{H}}^-$ for $c < 0$. For $c = 0$, it is a pair of intersecting lines, i.e., $\mathbf{A}_0 \in \mathbb{D}_{\text{IL}}$. As c changes from small positive values to zero and then on to small negative values, the hyperbola with positive center transforms into a pair of intersecting lines and then into a hyperbola with negative center; see Figure 4. In the parameter space, this process corresponds to a continuous motion from the subdomain $\mathbb{D}_{\text{H}}^+$ to the subdomain $\mathbb{D}_{\text{H}}^-$, across the hypersurface $\mathbb{D}_{\text{IL}}$.

The hypersurface $\mathbb{D}_{\text{P}}$ separates the domain $\mathbb{D}_{\text{H}}$ of hyperbolas from the domain $\mathbb{D}_{\text{E}}$ of ellipses. To illustrate this fact, consider the parameter vector $\mathbf{A}_c = (-1, 0, c, 0, 1, 0)$, where c is again a small variable. This parameter vector corresponds to the quadratic function

$$Q(x, y) = -x^2 + cy^2 + y = -x^2 + c\left(y + \frac{1}{2c}\right)^2 - \frac{1}{4c}.$$




Figure 4: Transformation of a hyperbola with positive center into one with negative center.

For $c > 0$, the equation $Q(x, y) = 0$ defines a hyperbola, i.e., $\mathbf{A}_c \in \mathbb{D}_H$ for $c > 0$. For $c = 0$, it is a parabola $y = x^2$, i.e., $\mathbf{A}_0 \in \mathbb{D}_P$. For $c < 0$, it is an ellipse, i.e., $\mathbf{A}_c \in \mathbb{D}_E$ for $c < 0$. As c changes from small positive values to zero and then on to small negative values, the hyperbola transforms into a parabola, and then into an ellipse. In the parameter space, this process corresponds to a continuous motion from the domain $\mathbb{D}_H$ to the domain $\mathbb{D}_E$, across the hypersurface $\mathbb{D}_P$.




Figure 5: Transformation of a hyperbola into an ellipse.

A closer look at the above examples reveals that the “positive” subdomain $\mathbb{D}_H^+$ borders on the “negative” subdomain $\mathbb{D}_E^-$, and vice versa. Similarly, the “positive” subdomain $\mathbb{D}_E^+$ borders on the “negative” subdomain $\mathbb{D}_{IE}^-$, and vice versa. Thus the “sign” always changes when a parameter vector moves continuously from one subdomain to another.

Figure 6 summarizes the above analysis in a schematic diagram showing the structure of the parameter space, with all principal subdomains and the respective separating hypersurfaces.




Figure 6: Schematic structure of the parameter space.

5 Relations between subdomains of $\mathbb{S}^5$

Generally, domains of higher dimension terminate on domains of smaller dimension. In other words, domains of smaller dimension make boundaries of domains of higher dimension. More precisely, we say that a domain $\mathbb{D}_1$ terminates on a domain $\mathbb{D}_2$ if there is a sequence of points $P_n \in \mathbb{D}_1$ that converges to a point $P \in \mathbb{D}_2$, i.e., $P_n \rightarrow P$ as $n \rightarrow \infty$.

Figure 7 shows how our domains terminate on each other. An arrow from $\mathbb{D}_1$ to $\mathbb{D}_2$ means that $\mathbb{D}_1$ terminates on $\mathbb{D}_2$, i.e., there is a sequence of points of $\mathbb{D}_1$ that converges to a point of $\mathbb{D}_2$. The domains are named by the types of conics. A detailed description of all the subdomains can be found in our web page [32].

Figure 7 contains all parts of the parameter space: from the largest, five-dimensional (5D) open domains to the smallest, two-dimensional (2D) regions. We note that all our domains terminate on each pole, so there should be an arrow from every domain down to the bottom line (Two Poles). For simplicity, we just put one large arrow pointing to the poles.

Figure 7 shows how parameter vectors $\mathbf{A} \in \mathbb{S}^5$ may converge, from one domain to another. Since this convergence involves algebraic parameters $\mathbf{A} = (A, B, C, D, E, F)$, we will call it *algebraic convergence*. It will refer to the convergence of a sequence of parameter vectors $\mathbf{A}_n$ to a parameter vector $\mathbf{A}$ on the sphere $\mathbb{S}^5$.

On the other hand, a sequence of conics S_n may converge to a conic S , in a natural geometric sense. A precise definition of geometric conver-




Figure 7: All the subdomains in the parameter space $\mathbb{S}^5$.

gence of conics (and more general planar objects) is given in [21], it is called W-convergence there. In particular, a sequence of conics of one type may geometrically converge to a conic of another type. In [21] we presented examples of such a convergence: circles converging to a line, ellipses converging to a parabola, etc.

A natural question is: Do algebraic and geometric types of convergence agree? The answer is *yes*, but there are some notable exceptions; see below.

Theorem 1 (Convergence of conics: general case). *Suppose a sequence of parameter vectors $\mathbf{A}_n \in \mathbb{S}^5$ corresponding to real (not imaginary) conics, S_n , converges to a parameter vector $\mathbf{A}$ corresponding to a real (not imaginary) conic, S , which is not a pair of coincident lines, i.e., $\mathbf{A} \notin \mathbb{D}_{\text{CL}}$. Then $S_n \rightarrow S$ in the geometric sense.*

Theorem 2 (Divergence of conics: general case). *Suppose a sequence $\mathbf{A}_n$ of parameter vectors corresponding to real (not imaginary) conics, S_n , converges to a parameter vector $\mathbf{A}$ corresponding to an imaginary conic or to a pole,*

i.e., $\mathbf{A} \in \mathbb{D}_{\text{IPL}}$ or $\mathbf{A} = \mathbf{P}_{\pm 1}$. Then S_n moves off toward infinity, i.e., for any point $P = (x, y) \in \mathbb{R}^2$ we have $\text{dist}(P, S_n) \rightarrow \infty$ as $n \rightarrow \infty$.

We note that the limit vector $\mathbf{A}$ cannot be in the domain of imaginary ellipses $\mathbb{D}_{\text{IE}}$, because the latter is open.

Theorem 3 (The exceptional case of coincident lines). *Suppose a sequence $\mathbf{A}_n$ of parameter vectors corresponding to real (not imaginary) conics, S_n , converges to a parameter vector $\mathbf{A} \in \mathbb{D}_{\text{CL}}$ corresponding to a pair of coincident lines; the latter make a line in $\mathbb{R}^2$ which we denote by L . Then S_n gets closer and closer to L , as n grows. More precisely, for any rectangle*

$$R = \{-A \leq x \leq A, \quad -B \leq y \leq B\}$$

we have

$$\max_{P \in S_n \cap R} \text{dist}(P, L) \rightarrow 0 \quad \text{as } n \rightarrow \infty.$$

In other words, if we look “through the window” R , we will see that all the points of S_n get closer and closer to L .

The proofs of the above theorems are straightforward but quite technical, we omit them. See our web page [32] for detailed proofs. The same applies to the subsequent theorems.

On the other hand, the conics S_n in Theorem 3 may *not* converge to the line L in the geometric sense as defined in [21]. For example, S_n may be parabolas that converge to a half-line that is only a part of L . Or S_n may be hyperbolas that converge to two opposite half-lines that are only parts of L . Or S_n may be ellipses that converge to a segment of L . Or S_n may be single points that converge to a point in L . Or S_n may be any of the above but instead of converging to any part of L they may wander along L back and forth, or go off toward infinity.

For example, let S_n be defined by

$$x^2 + \frac{\alpha_n(y + C_n)^2}{1 + C_n^2} = \beta_n,$$

where $\alpha_n \rightarrow 0$ and $\beta_n \rightarrow 0$ as $n \rightarrow \infty$. Then algebraically this sequence converges to $x^2 = 0$, which is a pair of coincident lines. But geometrically S_n may be an ellipse or a hyperbola or a single point, depending on the values (and the signs) of C_n, α_n, β_n , and it may converge to various parts of L or move back and forth along L or move off toward infinity altogether.

6 Continuity of the objective function

The objective function (1) depends on the conic S , as the data points $P_1, \dots, P_n$ are fixed. If $\mathbf{A}$ is a vector of parameters describing the conic S , then $\mathcal{F}$ naturally becomes a function of $\mathbf{A}$, i.e., $\mathcal{F}$ is a function on the parameter space. More precisely, $\mathcal{F}$ is defined on all parameter vectors $\mathbf{A} \in \mathbb{S}^5$ corresponding to real conics. It is not defined for parameters corresponding to imaginary conics or poles. We denote the domain of the objective function by

$$\mathfrak{D}_{\mathcal{F}} = \mathbb{D}_{\text{E}} \cup \mathbb{D}_{\text{H}} \cup \mathbb{D}_{\text{P}} \cup \mathbb{D}_{\text{SP}} \cup \mathbb{D}_{\text{IL}} \cup \mathbb{D}_{\text{PL}} \cup \mathbb{D}_{\text{CL}} \cup \mathbb{D}_{\text{SL}}.$$

We note that $\mathfrak{D}_{\mathcal{F}}$ does not include regions $\mathbb{D}_{\text{IE}}$ and $\mathbb{D}_{\text{IPL}}$ corresponding to imaginary conics, and it does not include the poles $\mathbf{P}_{\pm 1}$.

Our first goal is to examine the continuity of $\mathcal{F}$ on $\mathfrak{D}_{\mathcal{F}}$.

Theorem 4 (Continuity of the objective function). *The objective function $\mathcal{F}$ is continuous everywhere on its domain $\mathfrak{D}_{\mathcal{F}}$ except on the region $\mathbb{D}_{\text{CL}}$ corresponding to coincident lines.*

This theorem is an immediate consequence of Theorem 1 and the results of [21] where we proved the continuity of $\mathcal{F}$ with respect to the geometric convergence. Indeed, we only need to apply a general principle: the composition of two continuous functions is a continuous function.

On the region $\mathbb{D}_{\text{CL}}$ corresponding to coincident lines the objective function $\mathcal{F}$ is badly discontinuous, according to Theorem 3 and the discussion around it. Recall that if $\mathbf{A}_n \rightarrow \mathbf{A}$ and $\mathbf{A} \in \mathbb{D}_{\text{CL}}$ corresponds to a line L , then the conics S_n corresponding to $\mathbf{A}_n$ may move back and forth along the line L or move off toward infinity. Accordingly, the values of the objective function $\mathcal{F}(\mathbf{A}_n)$ may oscillate within a wide range or diverge to infinity.

But the objects S_n must get closer and closer to L , as n grows, they just may not stretch all the way along L . This implies that the objects S_n , in the limit $n \rightarrow \infty$, cannot provide a better fit to the given points than the line L does. Thus any limit value obtained from $\mathcal{F}(\mathbf{A}_n)$ cannot be smaller than the value $\mathcal{F}(\mathbf{A})$, i.e.,

$$\liminf_{n \rightarrow \infty} \mathcal{F}(\mathbf{A}_n) \geq \mathcal{F}(\mathbf{A}).$$

Functions with the above property are said to be *lower semi-continuous*. Hence we obtain one more important fact:

Theorem 5 (Lower semi-continuity of the objective function). *The objective function $\mathcal{F}$ is lower semi-continuous on the region $\mathbb{D}_{\text{CL}}$ corresponding to coincident lines. The objective function $\mathcal{F}$ grows to infinity near the region $\mathbb{D}_{\text{IPL}}$ and near the poles $\mathbf{P}_{\pm 1}$. More precisely, if $\mathbf{A}_n \rightarrow \mathbf{A}$ and either $\mathbf{A} \in \mathbb{D}_{\text{IPL}}$ or $\mathbf{A} = \mathbf{P}_{\pm 1}$, then $\mathcal{F}(\mathbf{A}_n) \rightarrow \infty$.*

The above theorems easily imply the existence of a global minimum of $\mathcal{F}$ on $\mathfrak{D}_{\mathcal{F}}$ as follows. The domain $\mathfrak{D}_{\mathcal{F}}$ is not compact, but due to Theorem 5 we can cut out and ignore a small vicinity of the region $\mathbb{D}_{\text{IPL}}$ and the poles $\mathbf{P}_{\pm 1}$ where the function is too big. Then the remaining part of the domain $\mathfrak{D}_{\mathcal{F}}$ will be compact. And now the lower semi-continuity of $\mathcal{F}$ guarantees the existence of its global minimum. Indeed, any lower semi-continuous function on a compact domain attains its minimum.

The existence of a global minimum of $\mathcal{F}$ is proved, by a totally different approach, in [21].

7 Differentiability of the objective function

Our next goal is to examine the differentiability of $\mathcal{F}$ on $\mathfrak{D}_{\mathcal{F}}$. This is important because most popular minimization algorithms (such as the steepest descent, Newton-Raphson, Gauss-Newton, or Levenberg-Marquardt) use derivatives of $\mathcal{F}$. Some use the first order derivative of $\mathcal{F}$, others use the second order derivative, or approximations to the second order derivative.

Thus it is essential that our objective function $\mathcal{F}$ be differentiable, at least once. As $\mathcal{F}$ is the sum of squares of the distances, see (1), it will be enough to check that $[\text{dist}(P_i, S)]^2$, i.e., the square of the distance from the given point $P_i = (x_i, y_i) \in \mathbb{R}^2$ to the conic S , is differentiable with respect to the conic's parameters.

We consider a more general problem. Given a point $P = (x_0, y_0)$ and a conic S , we will investigate the differentiability (with respect to the parameters of S) of the function

$$[\text{dist}(P, S)]^2 = [\text{dist}(P, Q)]^2 = (x - x_0)^2 + (y - y_0)^2,$$

where $Q = (x, y)$ denotes the projection of P onto the conic S . To this end it will be enough to check that the coordinates x, y of the footpoint Q of the projection are differentiable with respect to the conic's parameters.

One may guess, intuitively, that whenever the point $P = (x_0, y_0)$ is kept fixed and the conic S changes continuously, the projection Q of P onto S

would change continuously and smoothly. We will prove that generally this is true. However, there are exceptional cases where the continuity breaks down.

The reason for the breakdown is that the point Q on the conic S closest to the given point P may be not unique. For example, if S is a circle and P is its center, then all the points of S are equally distant from P , hence the point Q can be chosen anywhere on the circle. Another example: S is an ellipse and P lies on the major axis near the center. Then there are exactly two points on S closest to P (they are symmetric about the major axis of S). Similar situations occur when S is a hyperbola or a parabola and P lies on its axis. See Figure 8.




Figure 8: Examples of non-unique projection points.

In these exceptional cases, if one changes such a conic S continuously, then the point Q may instantaneously “jump” from one side (or branch) of S to another.

A more subtle exceptional case occurs when P lies at the center of curvature of S at the point Q . This means that P coincides with the center of the osculating circle of the conic S at the projection point Q . Then the projection Q may be technically unique, but “barely unique”, as to the second order all the points on S close enough to Q will be equally distant from P . This is a subtle situation, we will explore it separately.

Theorem 6 (Differentiability of projection coordinates). *Let S be a conic and P a given point. Suppose (i) the point Q on the conic S closest to the given point P is unique and (ii) P is not the center of curvature of the conic S at the point Q . Then the coordinates x and y of the point Q are differentiable with respect to the conic’s parameters.*

Theorem 6 is proved by implicit differentiation, we omit details. Next we turn to the exceptional case (ii) in the above theorem, i.e., suppose P is at the center of the osculating circle of S at the point Q . Then the coordinates

x and y are not differentiable with respect to the conic's parameters, but surprisingly the distance $\text{dist}(P, S)$ is differentiable with respect to the conic's parameters (see details in [32]). Thus we get

Theorem 7 (Differentiability of distances). *Let S be a conic and P a given point. Suppose (i) the point Q on the conic S closest to the given point P is unique and (ii) P coincides with the center of curvature of the conic S at the point Q . Then the distance $\text{dist}(P, S)$ is differentiable with respect to the conic's parameters.*

8 Singularities of the objective function

Thus the objective function $\mathcal{F}$ is differentiable, unless the point P has more than one projection onto the conic S . In rare cases where the condition (i) of Theorems 6 and 7 does not hold, the objective function may not be differentiable. This happens, for instance, if S is an ellipse and one of the data points P_i happens to lie on its major axis somewhere in the middle of S (then P_i is equally distant from the two halves of the ellipse). Or if S is a circle and one of the data points P_i is its center.

To illustrate the above effect let us consider a simplified family of conics defined by

$$(3) \quad x^2 + y^2 + 2Dx + 2Ey - 3 = 0$$

where only two algebraic parameters, D and E , are variable and all the others are fixed ($A = C = 1$, $B = 0$, and $F = -3$). This is actually a family of circles with center $(-D, -E)$ and radius $R = \sqrt{D^2 + E^2 + 3}$. The distance from $P = (x_0, y_0)$ to this circle is given by

$$\text{dist}(P, S) = \left| \sqrt{(D + x_0)^2 + (E + y_0)^2} - \sqrt{D^2 + E^2 + 3} \right|.$$

Figure 9 shows the graph of this distance, as a function of D and E , plotted by MATLAB. We have set $x_0 = y_0 = 0$ and let D and E vary from -1 to 1 .

We clearly see a sharp peak on the graph at the point $D = E = 0$, exactly where the point $P = (0, 0)$ coincides with the center of the circle. The graph is a cone-shaped surface near the peak with no derivatives at the summit.

However, the cases of non-differentiability of the objective function $\mathcal{F}$ are rare. They occur when one of the data points happens to be in an unusual




Figure 9: Graph of the objective function in Example (3).

place where the distance to the conic may be computed in more than one way. Such points “confuse” the objective function and cause the failure of its differentiability.

It is important to explore what happens whenever the objective function $\mathcal{F}$ fails to be differentiable. It turns out, fortunately, that in all such cases the shape of $\mathcal{F}$ resembles a “peak” (pointing upward), as in Figure 9. It cannot have a shape of a “pothole” (pointing downward).

Indeed, suppose a data point P_i can be orthogonally projected onto the conic S in more than one way (meaning its projections on different parts or on different branches of S). Denote the footpoints of those projections by Q'_i, Q''_i , etc. Then

$$\text{dist}(P_i, S) = \min\{\text{dist}(P_i, Q'_i), \text{dist}(P_i, Q''_i), \text{ etc.}\}$$

Thus the distance is obtained as the *minimum* of several smooth functions. And here is a general fact: the minimum of several smooth functions can only have “peak-type” singularities, not “pothole-type” singularities; see a simple illustration in Figure 10.

In other words, $\mathcal{F}$ has “peaks”, or local maxima, at singular points.

Theorem 8 (Smoothness at local minima). *The objective function $\mathcal{F}$ is smooth at all its local minima. More precisely, the first order derivatives of $\mathcal{F}$, as well as those of the distances $\text{dist}(P_i, S)$, exist and are continuous at all local minima.*

Since our main goal is minimization of $\mathcal{F}$, i.e., finding its (local) minima, the singularities of $\mathcal{F}$ will not really concern us, they will not be harmful.




Figure 10: Minimum of two smooth functions.

Standard minimization algorithms, such as Levenberg-Marquardt or Trust Region, are prohibited from moving in the “wrong direction” where the function $\mathcal{F}$ increases. They will only move if they find a smaller value of $\mathcal{F}$. This restriction forces them to move away from local maxima of $\mathcal{F}$, in particular away from singular points of $\mathcal{F}$.

If an algorithm converges to a limit, then $\mathcal{F}$ has a local minimum there, and by our Theorem 8 the function $\mathcal{F}$ and the distances $\text{dist}(P_i, S)$ have continuous first order derivatives. Since the above mentioned algorithms only use the first order derivatives of the distances $\text{dist}(P_i, S)$, they should be able to find the local minimum of $\mathcal{F}$ and converge quickly.

It is also important that the change of the conic type does not affect the differentiability of $\mathcal{F}$. For example, if the given parameter vector $\mathbf{A}$ corresponds to a parabola, so that nearby parameter vectors correspond to either ellipses or hyperbolas, then the objective function is still differentiable at $\mathbf{A}$ (unless again, one of the data points happens to lie on the axis of the parabola and be equally distant from its two halves). Our proofs of Theorems 6 and 7 do not rely upon any specific type of the conic, so that they work just fine when that type changes.

The same is true if the given parameter vector $\mathbf{A}$ corresponds to intersecting lines, so that nearby parameter vectors correspond to hyperbolas with “opposite signs”, i.e., with positive center and negative center. In that case again the objective function will be differentiable at $\mathbf{A}$.

As an illustration, let us consider a simplified family of conics defined by

$$(4) \quad x^2 + Cy^2 + 2Dx + 1 = 0$$

where only two algebraic parameters, C and D , are variables and all the

others are fixed ($A = 1$, $B = E = 0$, and $F = 1$). We easily see that

$$\Delta = \begin{vmatrix} 1 & 0 & D \\ 0 & C & 0 \\ D & 0 & 1 \end{vmatrix} = C(1 - D^2), \quad J = \begin{vmatrix} 1 & 0 \\ 0 & C \end{vmatrix} = C,$$

$$I = 1 + C, \quad Q = 1 + C^2 \quad K = \begin{vmatrix} 1 & D \\ D & 1 \end{vmatrix} + \begin{vmatrix} C & 0 \\ 0 & 1 \end{vmatrix} = 1 - D^2 + C.$$

Accordingly, the conic (4) may be of the following types: *Hyperbola* whenever $C < 0$ and $|D| \neq 1$, *Ellipse* whenever $C > 0$ and $|D| > 1$, *Imaginary ellipse* whenever $C > 0$ and $|D| < 1$, *Single point* whenever $C > 0$ and $D = \pm 1$, *Intersecting lines* whenever $C < 0$ and $D = \pm 1$, *Parallel lines* whenever $C = 0$ and $|D| > 1$, *Imaginary parallel lines* whenever $C = 0$ and $|D| < 1$, *Coincident lines* whenever $C = 0$ and $D = \pm 1$. Figure 11 shows the types of the conic (4) on the CD plane.




Figure 11: The CD plane for Example (4) and respective conic types.

Next we chose five data points $P_1, \dots, P_5$ lying on the ellipse $x^2 + 6y^2 - 12x + 1 = 0$ and computed the objective function $\mathcal{F}$ for all $-10 \leq C, D \leq 10$. Below is the graph of $\mathcal{F}$, as a function of C and D , plotted by MATLAB.

The dark red part of the graph lies above the domain of imaginary ellipses $\mathbb{D}_{IE}$ where the objective function cannot be defined. Other than that, the entire graph appears to be one smooth and “glassy” surface. In particular, we clearly see that the objective function is not broken or even wrinkled at places where the conic type changes.

The lowest (darkest) point of the graph is at $C = 6$, $D = -6$, where the objective function achieves its global minimum $\mathcal{F} = 0$ (corresponding to the ellipse $x^2 + 6y^2 - 12x + 1 = 0$ that passes through all our five data points).




Figure 12: Graph of the objective function for Example (4).

Thus the objective function $\mathcal{F}$ is differentiable on the entire green and yellow area of the diagram shown in Figure 11, including separating hypersurfaces $\mathbb{D}_{\text{IL}}$ (intersecting lines, shown as IL) and $\mathbb{D}_{\text{P}}$ (parabolas, shown as P).

9 Effects of the singularities of $\mathcal{F}$

The objective function is not defined for imaginary conics, in particular it is not defined on the red disks $\mathbb{D}_{\text{IE}}^{\pm}$ in Figure 6 (shown as $\text{IE}^{\pm}$). The objective function is defined on the subdomains $\mathbb{D}_{\text{E}}^{\pm}$ and the bordering hypersurface $\mathbb{D}_{\text{SP}}$, but then it stops; it does not extend into $\mathbb{D}_{\text{IE}}^{\pm}$. The graph of $\mathcal{F}$ is smooth over $\mathbb{D}_{\text{E}}^{\pm}$ but then it abruptly terminates. Could this cause trouble for our main purpose – the minimization of $\mathcal{F}$?

There is a good reason why the above termination does *not* cause trouble: the objective function $\mathcal{F}$ actually grows near the hypersurface $\mathbb{D}_{\text{SP}}$, as we explain below. Its growth discourages minimization procedures from moving in the direction of $\mathbb{D}_{\text{SP}}$. All decent minimization algorithms (Levenberg-Marquardt, Trust Region, etc.) keep moving only as long as the value of the objective function decreases at each iteration. If the objective function does not decrease, the algorithm retreats, its step is recalculated, and recalculation is repeated until the algorithm finds a place where the objective function decreases. If the algorithm comes close to the hypersurface $\mathbb{D}_{\text{SP}}$, then it will

have to turn around and move away from $\mathbb{D}_{\text{SP}}$ just in order to find smaller values of the objective function.

Now why does the objective function $\mathcal{F}$ increase near the hypersurface $\mathbb{D}_{\text{SP}}$? Recall that the parameter vectors $\mathbf{A} \in \mathbb{D}_{\text{E}}$ that are near $\mathbb{D}_{\text{SP}}$ correspond to very small ellipses. As $\mathbf{A}$ gets closer to $\mathbb{D}_{\text{SP}}$, the corresponding ellipse shrinks and converges to a single point. Obviously, this does not help to fit the given data points any better. Making the ellipse smaller and smaller only increases the distances from that ellipse to all the data points located outside the ellipse. Admittedly, it reduces the distances to the data points inside the ellipse, but this is just a second order effect. If the ellipse has small size ε , then the contribution of the “interior” data points to the objective function is of order ε^2 , so its further reduction does not help much. On the other hand, distances to the “exterior” data points would grow by increments of order ε , causing the overall increase of the objective function.

Less formally, one can just visually compare a tiny ellipse E_0 in Figure 13 with a large ellipse E_1 that cuts through E_0 . It is clear that the data points inside E_0 are nearly equally close to both ellipses, while the data points outside E_0 are much closer to E_1 than to E_0 .




Figure 13: A tiny ellipse E_0 and a larger ellipse E_1 .

Another potentially dangerous area in the parameter space $\mathbb{S}^5$ is the vicinity of the domain $\mathbb{D}_{\text{CL}}$ corresponding to coincident lines. In that area the objective function $\mathcal{F}$ becomes highly irregular and badly discontinuous.

Again an argument similar to the above shows that the objective function tends to grow near the domain $\mathbb{D}_{\text{CL}}$, which discourages minimization algorithms from approaching this domain. Indeed, recall that if a parameter vector $\mathbf{A}$ is close to the domain $\mathbb{D}_{\text{CL}}$, i.e., $\mathbf{A} \approx \mathbf{A}_0 \in \mathbb{D}_{\text{CL}}$, then the conic S corresponding to $\mathbf{A}$ is close to the line L_0 corresponding to $\mathbf{A}_0$. More precisely, S wholly lies in a narrow strip around L_0 . As $\mathbf{A} \rightarrow \mathbf{A}_0$, the entire

conic S gets closer and closer to L_0 (though it may not stretch all the way along L_0).

Now again it is easy to see that by squeezing the conic S toward the line L_0 one cannot improve the overall fit unless *all* the data points lie on L_0 . The situation where all the data points are collinear is very exceptional and can be detected and handled differently, without even applying general fitting procedures. If not all the data points are collinear, then squeezing the conic is detrimental to the overall fit.

Figure 14 shows a very narrow ellipse E_0 (blue) that is almost squeezed to a line. It also shows two arcs of a much longer ellipse E_1 (green). We do not see it fully – the figure only shows one (nearly straight) arc passing right through E_0 and the other (nearly straight) arc lying above E_0 . Obviously, E_1 is closer to some data points than E_0 is.




Figure 14: A narrow ellipse E_0 and a wider ellipse E_1 .

The above argument basically shows that a single line (or any object stretching along a single line) cannot provide a good fit for typical (non-collinear) sets of data points. Many other conics achieve a better fit, including parallel lines, intersecting lines, long ellipses (that are close to two parallel lines), or similar hyperbolas, etc. Respectively, the objective function $\mathcal{F}$ tends to decrease if the parameter vector moves away from the domain $\mathbb{D}_{SL}$ corresponding to single lines. Thus the minimization algorithms are not likely to move toward this domain, i.e., it should not bother us.

On the contrary, a pair of parallel lines may provide quite a good fit for some sets of data points. It is not apparent at all that ellipses or hyperbolas would provide a fit better than a pair of parallel lines. In fact, it is possible that the best fit is achieved only by a pair of parallel lines. Respectively, we believe the minimization algorithms are quite likely to approach the domain $\mathbb{D}_{PL}$ or wander around in its vicinity. Hence this this domain is essential.

To summarize, we list all the domains where the minimization algorithms are likely to “maneuver” searching for the best fitting conic and where the

best fit can be found: Ellipses $\mathbb{D}_E$, hyperbolas $\mathbb{D}_H$, parabolas $\mathbb{D}_P$, intersecting lines $\mathbb{D}_{IL}$ and parallel lines $\mathbb{D}_{PL}$. We formalize this in the following statement:

Theorem 9 (Essential domain). *For any set of data points $P_1, \dots, P_n$ the global minimum of the objective function $\mathcal{F}$ belongs to the union*

$$(5) \quad \mathfrak{D}_{\mathcal{F}, \text{ESS}} = \mathbb{D}_E \cup \mathbb{D}_H \cup \mathbb{D}_P \cup \mathbb{D}_{IL} \cup \mathbb{D}_{PL}.$$

If the objective function $\mathcal{F}$ has multiple global minima, then at least one of them belongs to the above union. This union cannot be shortened, i.e., for any conic S in this union of domains there exists a data set for which S provides the (unique) best fit.

We call $\mathfrak{D}_{\mathcal{F}, \text{ESS}}$ the *essential domain*, or the essential part of the domain $\mathfrak{D}_{\mathcal{F}}$, of the function $\mathcal{F}$. The above theorem basically says that all the other parts of the parameter space $\mathbb{S}^5$ can be ignored for the purpose of minimization of the objective function. On those parts $\mathcal{F}$ is either not defined or tends to grow.

10 Local minima

In the previous section we described the objective function $\mathcal{F}$ on its natural domain $\mathfrak{D}_{\mathcal{F}} \subset \mathbb{S}^5$ focusing on its continuity and differentiability.

We showed that $\mathcal{F}$ is continuous and differentiable everywhere except certain bad places where $\mathcal{F}$ either has “peaks” (local maxima) or somehow tends to grow. Since our main goal is minimization of $\mathcal{F}$, those bad places should not cause trouble. Standard minimization algorithms, such as Levenberg-Marquardt, can only go where the objective function decreases, so they are bound to move away from bad places. Since the parameter space is compact, they cannot move off (i.e., diverge) toward infinity. Therefore they are bound to converge to a minimum of $\mathcal{F}$ (local or global).

Ideally, the minimization procedures should converge to the global minimum of $\mathcal{F}$ and not be distracted by its local minima. Here we investigate the local minima of $\mathcal{F}$ and assess their potentially distractive role.

In the simplest case of $n = 5$ data points there is always an interpolating conic, i.e., a conic passing through all the five points. Thus if $n = 5$, the objective function takes its global minimum $\mathcal{F} = 0$. And, quite surprisingly, it has *no local minima*! This is a mathematical fact that we proved in [21].

For $n > 5$ data points, local minima are possible. They occur even if the points are observed without noise, i.e., if all the n points lie on a conic. We have investigated local minima for data points placed on an ellipse by a numerical experiment (its detailed description is given in our web page [32]). We used the ellipse with semi-axes $a = 2$ and $b = 1$ and placed n points (equally spaced) along the entire ellipse or along a certain arc of the ellipse. This is done by using an internal angular parameter $\varphi \in [-\pi, \pi]$. We choose φ_{ini} and φ_{end} (representing the endpoints of the arc) and set

$$x_i = a \cos \varphi_i, \quad y_i = b \sin \varphi_i, \quad \varphi_i = \varphi_{\text{ini}} + \frac{1}{n}(i - 0.5)(\varphi_{\text{end}} - \varphi_{\text{ini}})$$

for $1 \leq i \leq n$. We observed that when the data points are placed along the entire ellipse, i.e., $\varphi_{\text{ini}} = -\pi$ and $\varphi_{\text{end}} = \pi$, the objective function had no local minima. The same is true when the data points are placed along the right half of the ellipse, i.e., $\varphi_{\text{ini}} = -\pi/2$ and $\varphi_{\text{end}} = \pi/2$.

But when the data points are placed along the upper half of the ellipse, i.e., $\varphi_{\text{ini}} = 0$ and $\varphi_{\text{end}} = \pi$, the objective function does have local minima, see below.

Figure 15 shows $n = 6$ data points placed along the upper half of the ellipse $x^2/4 + y^2 = 1$ (whose semi-axes are $a = 2$ and $b = 1$). The global minimum of $\mathcal{F}$ is achieved by the interpolating ellipse (red), and two local minima of $\mathcal{F}$ correspond to two hyperbolas (blue and green). The first column of the numerical output in the top right corner gives the values of $\mathcal{F}$ at all these three minima. The two hyperbolas are obviously symmetric to each other with respect to the y axis, this is why $\mathcal{F}$ takes the same value at each. The second column (PER) gives the percentages of random initial guesses from which the minimization routine converged to each minimum.

Note that from an overwhelming majority (96%) of randomly generated initial conics the minimization procedure converged to the global minimum. We explain this phenomenon below.

In another example, with $n = 6$ data points along a quarter of the same ellipse (i.e., we chose $\varphi_{\text{ini}} = 0$ and $\varphi_{\text{end}} = \pi/2$), the function $\mathcal{F}$ happened to have one global minimum (the interpolating ellipse, red) and one local minimum (the blue hyperbola); see Figure 16. Again, from an overwhelming majority (99%) of randomly generated initial conics the minimization procedure converged to the global minimum.

In yet another example, with $n = 8$ points placed along the upper half of the same ellipse, the function $\mathcal{F}$ happened to have one global minimum (the




Figure 15: One global minimum (red ellipse) and two local minima (blue and green hyperbolas).

interpolating ellipse) and seven (!) local minima, which include two ellipses and five hyperbolas. The corresponding figures are included in our web page [32]. The important fact is, again, that from an overwhelming majority (94%) of randomly generated initial conics the minimization procedure converged to the global minimum. Only 6% of random initial conics fell into the vicinities of local minima.

We have also run similar tests for data points with added Gaussian noise (at levels $\sigma = 0.05$ and $\sigma = 0.1$). We observed quite a similar picture: one global minimum and several (up to 8) local minima. The largest number of local minima in a single data set we found was 8.

The global minimum is usually an ellipse (close to the “true” ellipse on which the original, unperturbed points are placed) when the noise is small. For larger noise the global minimum is often a branch of a hyperbola that is close to the elliptic arc containing the original points.

On the contrary, the local minima are mostly hyperbolas whose both branches run through the corridor containing the perturbed points. Less frequently local minima are long narrow ellipses whose both halves run through the corridor containing the perturbed points. Figure 17 shows the picture schematically.

A more detailed report on the number of observed local minima and their types (hyperbolas and ellipses) can be found in our web page [32].

In most of our tests, local minima exist, but hard to find by minimization procedures. Only from a very small percentage of randomly generated initial




Figure 16: One global minimum (red ellipse) and one local minimum (blue hyperbola).




Figure 17: Grey corridor contains noisy points scattered around the red elliptic arc. Two branches of the blue hyperbola and two branches of the long narrow green ellipse run through the corridor.

guesses those converges to a local minimum. This indicates that local minima are really tiny narrow pits (cavities) in the graph of the objective function, like in Figure 18.

Such functions are characterized by very large second derivatives at local minima and very large first derivatives near them. Large derivatives cause the function change rapidly, so that it climbs out of the local minima quickly. To this extent we can prove the following:

Theorem 10 (Narrowness of local minima). *Suppose our data points lie in a rectangle (corridor) R and a conic S corresponding to a parameter vector $\mathbf{A} \in \mathbb{S}^5$ has two branches crossing R at a small distance ε from each other. Then the first and second derivatives of the objective function $\mathcal{F}$ at $\mathbf{A}$ are of order $1/\varepsilon$. If only one branch of the conic crosses R , then both derivatives of $\mathcal{F}$ at S are of order 1.*




Figure 18: A function with one wide global minimum (red arrow) and four narrow local minima (blue arrows).

In the examples described above we have computed the second derivatives of $\mathcal{F}$ (more precisely, the eigenvalues of its Hessian matrix) numerically. We found that at all the local minima those derivatives were about 10-100 times larger than the corresponding derivatives at the (unique) global minimum.

To summarize, we have two largely different cases. First, if data points are sampled along the entire ellipse, or an elliptic arc with high curvature, the objective function tends to have one global minimum and no local minima.

Second, if data points are sampled along an elliptic arc with low curvature, they appear in a corridor around the arc like the one shown in Figure 17. In that case the best fit (the global minimum of $\mathcal{F}$) is given by a single elliptic arc (or, occasionally, by a single hyperbolic arc) stretching through the corridor. At the same time there might be distractive fits (local minima of $\mathcal{F}$) made by hyperbolas or long narrow ellipses whose both branches run through the corridor; see again Figure 17. However, those minima tend to be small and narrow, so that the chance of falling into one of them is low.

11 Geometric parameters of conics

So far we used algebraic parameters of conics. In practical applications, however, most authors prefer more natural geometric parameters. Here we describe those and compare them to the algebraic parameters.

Ellipses can be described by five geometric parameters. The most popular choice is: coordinates of the center (x_c, y_c) , semi-axes a and b (where usually a denotes the major semi-axis and b the minor semi-axis, i.e., $a \geq b$), and the angle of tilt, α , of the major axis. There are natural restrictions on the semi-axes, $a \geq b > 0$. The angle α has period π and one often requires $\alpha \in [0, \pi)$.

The full ellipse equation in these parameters is rather complicated:

$$(6) \quad \frac{[(x - x_c) \cos \alpha + (y - y_c) \sin \alpha]^2}{a^2} + \frac{[-(x - x_c) \sin \alpha + (y - y_c) \cos \alpha]^2}{b^2} = 1.$$

Alternatively one can define the ellipse in parametric form:

$$x(t) = x_c + a \cos t \cos \alpha - b \sin t \sin \alpha$$

and

$$y(t) = y_c + a \cos t \sin \alpha + b \sin t \cos \alpha,$$

where $0 \leq t \leq 2\pi$ is an inner parameter on the ellipse. The inner parameter t plays an important role in some fitting algorithms; see [23, 28, 31].

The five parameters (x_c, y_c, a, b, α) were used by many authors. In earlier publications they were used in conjunction with the inner parameter t on the ellipse; see [23, 28]. In later publications, they were used without the inner parameter; see [3, 6, 7]. They generally work well, but they involve a peculiar problem described in the next section.

Ellipses can be described by another set of geometric parameters, which involve its foci: coordinates of both foci (x_1, y_1) and (x_2, y_2) , and the major semi-axis a . Now the equation of the ellipse is

$$(7) \quad \sqrt{(x - x_1)^2 + (y - y_1)^2} + \sqrt{(x - x_2)^2 + (y - y_2)^2} = 2a,$$

based on the fact that the sum of the distances from any point of the ellipse to its foci is constant. There is a natural restriction on the semi-axis:

$$2a \geq \sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2}.$$

These parameters were used in [17], where the authors called (7) Kepler's definition of ellipse, so (x_1, y_1, x_2, y_2, a) can be called Kepler's parameters. These parameters also work well, but the peculiar singularity mentioned above gets even worse now, see the next section.

Hyperbolas can be described by five geometric parameters, similar to ellipses: coordinates of the center (x_c, y_c) , semi-axes a and b (there are no constraints on a and b , except $a, b > 0$), and the angle of tilt, α , of the major axis. The hyperbola equation is

$$\frac{[(x - x_c) \cos \alpha + (y - y_c) \sin \alpha]^2}{a^2} - \frac{[-(x - x_c) \sin \alpha + (y - y_c) \cos \alpha]^2}{b^2} = 1.$$

It is possible to define the hyperbola in parametric form, too:

$$x(t) = x_c + a \cosh t \cos \alpha - b \sinh t \sin \alpha$$

and

$$y(t) = y_c + a \cosh t \sin \alpha + b \sinh t \cos \alpha,$$

where $-\infty < t < \infty$ is an inner parameter on the hyperbola. These parameters were used in [3].

Alternatively, hyperbola can be described by geometric parameters involving its foci: coordinates of both foci (x_1, y_1) and (x_2, y_2) , and the major semi-axis a . Now the equation of the hyperbola is

$$\left| \sqrt{(x - x_1)^2 + (y - y_1)^2} - \sqrt{(x - x_2)^2 + (y - y_2)^2} \right| = 2a,$$

based on the fact that the difference of the distances from any point of the hyperbola to its foci is a constant. We are not aware of any published work using these parameters. In fact very few authors studied the problem of hyperbola fitting (see [3, 8, 10, 24]), so advantages and disadvantages of the above parametrization schemes are yet to be investigated.

Parabolas can be described by four geometric parameters: coordinates of the vertex (x_c, y_c) , focus distance p to the directrix, the angle of tilt, α . The parabola equation is

$$[-(x - x_c) \sin \alpha + (y - y_c) \cos \alpha]^2 = 2p[(x - x_c) \cos \alpha + (y - y_c) \sin \alpha].$$

These parameters were used in [3].

12 Geometric versus algebraic parameters

An odd problem with ellipse geometric parameters arises when the ellipse turns into a circle. In this case the angle α becomes a redundant parameter that cannot be determined. Indeed, changing the angle α amounts to rotating the circle around its center, which does not affect it at all. This impedes the performance of fitting algorithms as we describe below.

Most popular fitting procedures (Gauss-Newton, Levenberg-Marquardt, Trust Region, etc.) use first derivatives of the objective function $\mathcal{F}$ with respect to the conic parameters. Those are computed through the Jacobian matrix $\mathbf{J}$ whose components are $\partial d_i / \partial \theta_j$, where θ_j , $j = 1, \dots, k$ denote the

parameters of the conic ($k = 5$ for ellipses and hyperbolas and $k = 4$ for parabolas) and d_i , $i = 1, \dots, n$, denote the distances from the data points $P_1, \dots, P_n$ to the conic. The Jacobian matrix has size $n \times k$, and most of the algorithms involve the inversion of the $k \times k$ square matrix $\mathbf{J}^T \mathbf{J}$ (or its modification, depending on the method used). Thus it is important that the matrix $\mathbf{J}$ has rank k , i.e., full rank.

Now the above indetermination of the parameter α causes the Jacobian matrix become singular (rank deficient) – more precisely, one of its columns is filled with zeros (as noted on page 2290 in [3]). A rigorous proof of this fact is given in our web page [32].

This singularity becomes more severe if the ellipse is not just a circle, but the *best fitting circle*, i.e., the circle that fits the data best among all circles. In that case the first order partial derivatives of the objective function $\mathcal{F}$ with respect to the coordinates x_c and y_c turn zero (precisely because the circle has an optimal center!). To see what else goes wrong, let us replace the semi-axes a and b with two parameters $u = a + b$ and $v = a - b$. Now when taking the partial derivative of the objective function $\mathcal{F}$ with respect to u we are supposed to keep all the other parameters fixed, in particular we keep $v = a - b = 0$ fixed, hence the ellipse remains a circle. Therefore the first order partial derivative of $\mathcal{F}$ with respect to u turns zero as well (precisely because the best fitting circle has an optimal radius!). Thus only one first order partial derivative of $\mathcal{F}$ (the one with respect to $v = a - b$) remains different from zero. This implies that the Jacobian of the objective function $\mathcal{F}$ is not only singular, its rank drops from five to one!

The severe singularity of the Jacobian was first noted in [23], and for this reason the authors had to avoid the best fitting circle as an initial guess for their iterative ellipse fitting algorithms (which otherwise seemed quite a logical choice to them). One can argue that this singularity should not really cause much trouble – the Levenberg-Marquardt algorithm works with singular matrices. Even the Gauss-Newton step can go through if one applies SVD (this was noted on page 2290 in [3]). On the other hand, when four out of five partial derivatives of $\mathcal{F}$ turn zero, the iterative procedure gets severely constrained, its freedom is extremely limited. It is likely to change $v = a - b$ but keep other parameters, $u = a + b$, x_c , y_c , and α , unchanged, or changed little (and rather randomly). In other words, the procedure just elongates the ellipse in the direction specified by α and squeezes it in the orthogonal direction. One also should note that if the best fitting circle is chosen as the initial guess, the parameter α is undetermined and has to be

assigned arbitrarily. Thus the elongation of the ellipse occurs in the direction arbitrarily set by the user during the initialization. As a result, the first step of the iterative procedure is likely to be quite arbitrary and awkward; and it may take the process some time to recover and find the best fitting ellipse.

Singularities related to circles were noticed by other authors, too. Some of them impose the restriction $a \neq b$ to avoid singularities; see [3, 29, 30]. Still others note that the statistical accuracy of the corresponding parameter estimates deteriorates (their variances grow) when the ellipse is close to a circle [4]. We also noticed in our computer experiments that whenever $a \approx b$, the angle α tends to change erratically, which may destabilize the performance of the fitting procedure.

The above singularity gets even worse with Kepler's parameters (7). Let the ellipse be a circle, i.e., let $x_1 = x_2$ and $y_1 = y_2$. Then the Jacobi matrix again becomes singular, and furthermore its rank drops from 5 to 3. Indeed, consider a family of ellipses with foci $(x_1 + s, y_1 + r)$ and $(x_2 - s, y_2 - r)$ for which the fifth parameter a is fixed. When $s = r = 0$, we have the original circle. By elementary geometry, when s and r are small, the ellipse is at distance $\mathcal{O}(s^2 + r^2)$ from the original circle (in the Hausdorff metric). Therefore the derivatives (with respect to s and r) of the distances $d_i = \text{dist}(P_i, S)$ from the data points P_i to the ellipse S turn zero at $s = r = 0$, causing a double-singularity for the Jacobi matrix (the loss of two dimensions in its rank). A more formal argument is given in our web page [32].

A particularly bad situation occurs when the ellipse coincides with the best fitting circle. Then the entire Jacobian matrix turns zero completely. In other words, the derivatives of the distances $d_i = \text{dist}(P_i, S)$ with respect to *all* the parameters of the ellipse vanish. Therefore, the best fitting circle is a stationary point for the objective function in Kepler's parameters. We provide a proof of this rather strange fact in our web page [32]. If the iterative procedure starts at the best fitting circle or arrives at it by chance, it will stall instantly and will not progress anywhere.

We should note that initializing iterative ellipse fitting algorithms with the best fitting circle is a very old tradition; see, e.g., [23, 3]. Some authors claim that the best fitting circle provides the most robust choice for the initial ellipse; see [8, Section 1.1.3]. As we have seen, in the natural geometric parameters this choice causes a severe singularity that affect the first step of the subsequent iterative process, from which the recovery may be slow. In Kepler's parameters, this choice leads to an instant failure.

It is remarkable that singularities like above do *not* occur with algebraic

parameters $\mathbf{A} \in \mathbb{S}^5$. One needs to keep in mind that there are six algebraic parameters, which are determined up to a scalar factor. Thus the Jacobian matrix $\mathbf{J}$ has size $n \times 6$, and its maximal rank is five, i.e., $\text{rank } \mathbf{J} \leq 5$. As it happens, its rank is 5 for all typical data sets, regardless of the conic:

Theorem 11 (Full rank of $\mathbf{J}$ in algebraic parameters). *Let $Q_1, \dots, Q_n$ denote the projections of the data points $P_1, \dots, P_n$ onto a conic S with parameter $\mathbf{A} \in \mathbb{S}^5$. If there are at least five distinct points among $Q_1, \dots, Q_n$, then $\text{rank } \mathbf{J} = 5$, i.e., the Jacobian matrix has maximal rank.*

Proof. Let $\mathcal{P} = Ax^2 + 2Bxy + Cy^2 + 2Dx + 2Ey + F$ denote the quadratic polynomial defining the conic (2). Let P denote a data point, d the distance from P to a conic S with parameters A, B, C, D, E, F , and Q the projection of P onto the conic; so that $d = \text{dist}(P, Q)$. Suppose we change one of the parameters A, B, C, D, E, F (we will denote that variable parameter by θ). Then the conic S moves, which causes the distance d change. Then the partial derivative of d with respect to θ is given by

$$\frac{\partial d}{\partial \theta} = \frac{\mathcal{P}_\theta}{\sqrt{\mathcal{P}_x^2 + \mathcal{P}_y^2}}$$

where $\mathcal{P}_\theta, \mathcal{P}_x, \mathcal{P}_y$ denote the first order partial derivatives of $\mathcal{P}$ with respect to θ, x, y , respectively, and all these derivatives are taken at the projection point Q (not at the data point P). The above formula was proved in [20].

Now each row of the Jacobian matrix $\mathbf{J}$ is proportional to

$$(x_i^2, 2x_i y_i, y_i^2, 2x_i, 2y_i, 1)$$

where (x_i, y_i) denote the *projection* of the i th data point P_i onto the conic. The denominator $\sqrt{\mathcal{P}_x^2 + \mathcal{P}_y^2}$ is the same for the entire row, so it is irrelevant.

Now the rank of the matrix $\mathbf{J}$ is the dimension of its row space. Clearly it equals 5 if and only if there is a unique conic interpolating all the projection points $Q_1, \dots, Q_n$ (i.e., our conic S). It is well known that for any five distinct points the interpolating conic is unique, which completes the proof. $\square$

13 Can one fit ellipses only?

In many practical applications one needs to fit ellipses only. It may seem as in such cases one can just restrict the parameter space to the domain $\mathbb{D}_E$ of

ellipses and minimize the objective function over that domain (alternatively, one may just use geometric parameters of ellipses). This, however, may cause serious complications. It was shown in [21] that the objective function $\mathcal{F}$ restricted to $\mathbb{D}_E$ may *not* attain its minimum. In other words, the best fitting ellipse may *not* exist. Trying to find something that does not exist is an exercise in futility.

This is a serious deficiency of ellipses that was thoroughly investigated in [21], both theoretically and numerically. It was shown that the best fitting ellipse fails to exist quite often, and this problem cannot be ignored. It was actually noted by other authors, too [12, 25, 26].

If the best fitting ellipse fails to exist, then for any ellipse E there will be another ellipse E' providing a better fit, in the sense $\mathcal{F}(E') < \mathcal{F}(E)$. If one constructs a sequence of ellipses that fit the given data points progressively better and on which the objective function $\mathcal{F}$ converges to its infimum, then those ellipses will grow in size and converge to something different than an ellipse (most likely, to a parabola; see [12, 21, 26]).

In practical terms, one usually runs a computer algorithm that executes an iterative procedure such as Gauss-Newton or Levenberg-Marquardt. It produces a sequence of ellipses E_m (here m denotes the iteration number) such that $\mathcal{F}(E_m) < \mathcal{F}(E_{m-1})$, i.e., the quality of approximations improves at every step, but those ellipses would keep growing in size and approach a parabola.

Then one has two options. The first is to admit that the minimization procedure diverges (which seems quite natural if one uses geometric parameters, as those would keep grow indefinitely). The second is to accept the limiting parabola as the best fit, as suggested in [12, 26]. But if one's goal is to find the best fitting ellipse, the second option is not satisfactory, and one may choose an ellipse by an alternative procedure (e.g., by Direct Ellipse Fit [22] or by some heuristic tricks [25]).

We argue that it is beneficial to include all the conics into the minimization procedure, even if one's final goal is to return an ellipse, no matter what. The inclusion of all conics has the following advantages:

- If the best fitting ellipse does *not* exist, then the minimization procedure would converge to a conic of another type (most likely, a hyperbola). This would imply that the best fitting ellipse does not exist and save frustrating attempts to find one and prevent divergence.

- If the best fitting ellipse *does* exist, then the minimization procedure should converge to it. However if one uses all the conics, rather than ellipses only, the procedure may find a shortcut through the domain of hyperbolas and arrive at the best ellipse faster. Moreover, sometimes the only way to arrive at the best fitting ellipse is by moving through hyperbolas for a few iterations!

Our example shown in Figures 11 and 12 presents exactly the situation described above. If one starts, say, with an ellipse $C = 1$, $D = 2$ (in the upper yellow rectangle in Figure 11 or on the yellow plateau in Figure 12), then the *only* way to converge to the best ellipse $C = 6$, $D = -6$ is to move around through the green area in Figure 11 corresponding to hyperbolas – there is no direct path down to the lower yellow rectangle, as it is blocked by imaginary ellipses (the red rectangle). We address this issue again in Sections 14 and 15.

14 Conclusions

We presented a thorough investigation of the algebraic parameter space for 2D quadratic curves (conics) and the objective function (1) whose minimization is required in order to find the best fitting conic for a given set of data points.

We fully described the structure of the parameter space and its regions and subregions corresponding to different types of conics. We also described under what conditions the objective is continuous and differentiable. We carefully identified its discontinuities and other singularities and determined what effect those may have on the performance of minimization procedures. We also investigated local minima of the objective function.

Our overall conclusion is that the algebraic parameters are well suited for the use in minimization procedures which aim at finding the best fitting conic. They have a number of advantages over more popular geometric parameters:

- The equation of a conic in the algebraic parameters is simple (2), the corresponding derivatives are simple, too. On the contrary, geometric parameters involve trigonometric functions (6) or radicals (7).
- The Jacobian matrix for the algebraic parameters always has full rank for typical data sets (Theorem 11). On the contrary, the one for ge-

x	1	2	5	7	9	3	6	8
y	7	6	8	7	5	7	2	4

Table 1: A benchmark example with eight points [23].

	Failure rate	Avg. iter.	Cost per iter. (flops)
GGs	26%	60	1710
LMG	11%	20	1640
LMA	0%	16	1880
TRA	0%	16	4480

Table 2: Comparison of four ellipse fitting methods.

ometric parameters occasionally develops singularities, sometimes severe, which may distract minimization algorithms.

- The parameter space $\mathbb{S}^5$ is compact, thus there is no way for the minimization algorithms to keep increasing estimated values of the parameters indefinitely causing divergence.
- Using algebraic parameters allows the minimization procedures move freely and smoothly between conics of different types. This flexibility allows them to take shortcuts and converge faster. Sometimes in order to find the best fitting conic of the right type the algorithm has to go through conics of other types for a few iterations and then come back to those of the right type (Figure 12).

15 Numerical experiment

To illustrate our conclusions we have run a numerical experiment with eight data points shown in Figure 1, whose coordinates are given in Table 1. This is perhaps the most popular benchmark example introduced in [23] and used in [3] and other papers. The best fitting ellipse is known to have center $(2.6996, 3.8160)$, axes $a = 6.5187$ and $b = 3.0319$, and angle of tilt $\theta = 0.3596$.

We have compared four fitting algorithms. One is the original ellipse

fitting scheme proposed by Gander, Golub, and Strebel [23], we call it GGS. It uses the internal ellipse parameter t as described in Section 11. Another is the most popular (see, e.g., [16]) Levenberg-Marquardt algorithm using the geometric parameters of the ellipse, we call it LMG. One more is the Levenberg-Marquardt algorithm using our algebraic parameters, we call it LMA. It is not restricted to ellipses only, it roams over the entire parameter space $\mathbb{S}^5$ thus using all conics. The last method is Trust Region, again with the algebraic parameters, we call it TRA.

All the methods were initialized by randomly generated ellipses. To get an initial ellipse, we just picked 5 points randomly in the square $0 \leq x, y \leq 10$ and used the ellipse interpolating those 5 points (if the interpolating conic was not an ellipse, we picked another set of 5 points). After running these fitting methods from 10^5 random initial guesses, we found that the GGS method failed to converge in 26% of the runs, and the LMG failed to converge in 11% of the runs. The LMA and TRA never failed to converge, obviously due to the compactness of the algebraic parameter space $\mathbb{S}^5$ (there is nowhere to diverge!). We also noticed the LMA and TRA, even though starting with an initial ellipse, often had to go through hyperbolas for a few iterations before they arrived at the best fitting ellipse. Thus using algebraic parameters and all types of conics (rather than ellipses only) is beneficial, even if the final fit has to be an ellipse (or is known to be an ellipse).

In those cases where our algorithms converged, the GGS method took 60 iterations, on the average, the LMG took 20 iterations, and both LMA and TRA took 16 iterations. Thus working with algebraic parameters also reduces the number of iterations. We note, however, that due to high complexity of the Trust Region method, it is 2-3 times more costly, per iteration, than other fits, we would not practically recommend it. Table 2 summarizes our results.

Figures 19 and 20 show the progress of our methods as they all start from the best fitting circle. By the 4th iteration, the LMA and TRA have moved half-way toward the best fitting ellipse, while the LMG and especially GGS are far behind. By the 8th iteration, the LMA and TRA have almost converged to the best ellipse, the LMG was just slightly behind, and the GGS was far behind.

Overall, the results of our numerical experiments are in full agreement with our theoretical conclusions.




Figure 19: Progress after 4 iterations.

Foundation, grant DMS-1301537.

References

- [1] Aigner, M. & Jüttler, B. *Robust computation of foot points on implicitly defined curves*, In: Editors Daehlen, M. et al., *Mathematical Methods for Curves and Surfaces*, Tromso 2004, Nashboro Press, 2005, pp. 1–10.
- [2] Aigner, M. & Jüttler, B. *Gauss-Newton type techniques for robustly fitting implicitly defined curves and surfaces to unorganized data points*, In: *Shape Modeling International*, 2008, pp. 121–130.
- [3] Ahn, S. J., Rauh, W. & Warnecke, H. J. *Least-squares orthogonal distances fitting of circle, sphere, ellipse, hyperbola, and parabola*, *Pattern Recog.*, **34** (2001), 2283–2303.
- [4] Ahn, S. J., Rauh, W. & Recknagel, M. *Least squares orthogonal distances fitting of implicit curves and surfaces*, In: *LNCS* **2191** (2001), 398–405.
- [5] Ahn, S. J., Rauh, W. & Cho, H. S. *Orthogonal distances fitting of implicit curves and surfaces*, *IEEE trans. PAMI*, **24** (2002), 620–638.




Figure 20: Progress after 8 iterations.

- [6] Ahn, S. J., Rauh, W., Cho, H. S., and Warnecke, H. J. *Orthogonal distance fitting of implicit curves and surfaces*, IEEE Trans. PAMI, **24** (2002), 620-638.
- [7] Ahn, S. J. *Least Squares Orthogonal Distance Fitting of Curves and Surfaces in Space*, LNCS, **3151**, Springer, 2004.
- [8] Ahn, S. J. *Least Squares Orthogonal Distance Fitting of Curves and Surfaces in Space*, In: LNCS **3151**, Springer, Berlin, 2004.
- [9] Albano, A. *Representation of digitized contours in terms of conic arcs and straight-line segments*, Comp. Graph. Image Proc. **3** (1974), 23-33.
- [10] Bliss, C. and James, A. *Fitting the rectangular hyperbola*, Biometrics **22** (1966), 573-602.
- [11] Biggerstaff, R. H. *Three variations in dental arch form estimated by a quadratic equation*, J. Dental Res. **51** (1972), 1509.
- [12] Bookstein, F. L. *Fitting conic sections to scattered data*, Comp. Graph. Image Proc. **9** (1979), 56-71.
- [13] Breyer, W. H. *CRC Standard Mathematical Tables and Formulas*, CRC Press, Boca Raton, FL, 28th edition, 1987.

- [14] Chan, N. N. *On circular functional relationships*, J. R. Statist. Soc. B, **27** (1965), 45–56.
- [15] Cheng, C.-L. & Van Ness, J. W. *Statistical Regression with Measurement Error*, Arnold, London, 1999.
- [16] Faber, P. and Fisher, R. *A Buyer’s Guide to Euclidean Elliptical Cylindrical and Conical Surface Fitting*, Proc. British Machine Vision Conference BMVC01 (Manchester, Sept. 2001), 521–530.
- [17] Chernov, N., Ososkov, G. & Silin, I. *Robust fitting of ellipses to non-complete and contaminated data*, Czech. J. Phys. **50** (2000), 347–354.
- [18] Chernov, N. and Lesort, C. *Statistical efficiency of curve fitting algorithms*, Comp. Stat. Data Anal., **47** (2004), pp. 713–728.
- [19] Chernov, N. *Circular and linear regression: Fitting circles and lines by least squares*, Chapman & Hall/CRC Monographs on Statistics and Applied Probability **117**, 2010.
- [20] Chernov, N. and Ma, H. *Least squares fitting of quadratic curves and surfaces*, In: Computer Vision, Editor S. R. Yoshida, Nova Science Publishers, 2011, pp. 285–302.
- [21] Chernov, N., Huang, Q., and Ma, H. *Does the best fitting curve always exist?* ISRN Probability and Statistics (2012), Article ID 895178.
- [22] Fitzgibbon, A. W., Pilu, M., and Fisher, R. B. *Direct Least Squares Fitting of Ellipses*, IEEE Trans. PAMI, **21** (1999), 476–480.
- [23] Gander, W., Golub, G. H. & Strebel, R. *Least squares fitting of circles and ellipses*, BIT, **34** (1994), 558–578.
- [24] Hey, E. N. *The statistical estimation of a rectangular hyperbola*, Biometrics **16** (1960), 606–617.
- [25] Matei, B. and Meer, P. *Reduction of Bias in Maximum Likelihood Ellipse Fitting*, In: 15th ICCVPR, 2000, Barcelona, Spain, **3**, pp. 802–806.
- [26] Nievergelt, Y. *Fitting conics of specific types to data*, Linear Algebra Appl., **378** (2004), 1–30.

- [27] Sampson, P. D. *Fitting conic sections to very scattered data: an iterative refinement of the Bookstein algorithm*, Comp. Graphics Image Proc. **18** (1982), 97–108.
- [28] Späth, H. *Orthogonal least squares fitting by conic sections*, In: Recent Advances in Total Least Squares techniques and Errors-in-Variables Modeling, SIAM, 1997, pp. 259–264.
- [29] Späth, H. *Least-squares fitting of ellipses and hyperbolas*, Comput. Statist. **12** (1997), 329–341.
- [30] Späth, H. *Orthogonal distance fitting by circles and ellipses with given area*, Comput. Stat. **12** (1997), 343–354.
- [31] Sturm, P. and Gargallo, P. *Conic fitting using the geometric distance*, Proc. Asian Conf. Comp. Vision, Tokyo, Japan, **2** (2007), 784–795.
- [32] [http : //www.math.uab.edu/~chernov/cl](http://www.math.uab.edu/~chernov/cl)