
Preface

This book is devoted to an active research topic in modern statistics—fitting geometric contours (lines and circles) to observed data, in particular, to digitized images. In such applications both coordinates of the observed points are measured imprecisely, i.e., both variables (x and y) are subject to random errors. Statisticians call this topic the Errors-In-Variables (EIV) model. It is radically different, and much more complex, than the classical regression where only one variable (usually, y) is random.

Fitting straight lines to observed data with errors in both variables is an old problem dating back to the 1870s [1, 2, 117], with applications in general statistics, sciences, econometrics, and image processing. Its studies have a colorful history (which we overview in Chapter 1) through the twentieth century, and its most active period perhaps lasted from 1975 to 1995. By the late 1990s all the major issues in the linear EIV problem appeared to be resolved, and now this topic is no longer an active research area.

For a detailed and complete account of the linear EIV regression studies, see surveys [8, 73, 126, 127, 132, 187] and books [40], [66], as well as Chapter 10 in [128] and Chapter 29 in [111]. We note that the linear EIV problem, despite its illusive simplicity, is deep and vast; entire books, such as [66] and [40], are devoted to this subject. We only overview it as much as it is related to our main theme — fitting circles and other curves.

Fitting *nonlinear* models to data with errors in both variables has been studied by statisticians since the 1930s [50, 55]. This topic can be divided into two parts. In one, the main goal is to describe observed data by a nonlinear function $y = g(x)$, such as a polynomial, or an exponential function, etc. In those applications the x and y variables usually have different natures, measured in different units, and errors in x and y may have different magnitude. Such applications are common in statistics and econometrics. A detailed presentation of this type of nonlinear model can be found [27]; see also the latest edition [28], updated and expanded.

The second type of nonlinear EIV problems is common in image processing applications. In those, data points come from a picture, photograph, map, etc. Both x and y variables measure length and are given in the same units; the choice of the coordinate system is often quite arbitrary, hence errors in x and y have the same magnitude, on average. Fitting explicit functions $y = g(x)$ to

images is not the best idea: it inevitably forces a different treatment of the x and y variables, conflicting with the very nature of the problem. Instead, one fits geometric shapes that are to be found (or expected) on the given image. Those shapes are usually described geometrically: lines, rectangles and other polygons, circles, ovals (ellipses), etc. Analytically, the basic curved shapes—circles and ovals—are defined by implicit quadratic functions. More complicated curves may be approximated by cubic or quartic implicit polynomials [150, 176]; however, the latter are only used on special occasions and are rare in practice. Our book is devoted to fitting most basic geometric curves—circles and circular arcs—to observed data in image processing applications. This topic is different from the nonlinear EIV regression in other statistical applications mentioned above and covered in [27, 28]. Fitting ellipses to observed data is another important topic that deserves a separate book (the author plans to publish one in the future).




Figure 0.1 *The Brogar Ring on Orkney Islands [15, 178]. The stones are marked by pluses; the fitted circle is the dotted line.*

The problem of fitting circles and circular arcs to observed points in 2D images dates back to the 1950s. Its first instance was rather peculiar: English engineers and archaeologists examined megalithic sites (stone rings) in the British Isles trying to determine if ancient people who had built those mysterious structures used a common unit of length. This work started in the 1950s and continued for several decades [15, 65, 177, 178, 179]; see an example in Fig. 0.1, where the data are borrowed from [178].

In the 1960s the necessity of fitting circles emerged in geography [155]. In the 1970s circles were fitted to experimental observations in microwave engineering [54, 108]; see an example in Fig. 0.2, where the data are taken from

[15]. Since about 1980, fitting circles became an agenda in many areas of human practice. We just list some prominent cases below.

In medicine, one estimates the diameter of a human iris on a photograph [141], or designs a dental arch from an X-ray [21], or measures the size of a fetus on a picture produced by ultrasound. Archaeologists examine the circular shape of ancient Greek stadia [157], or determine the size of ancient pottery by analyzing potsherds found in field expeditions [44, 80, 81, 190]. In industry, quality control requires estimation of the radius and the center of manufactured mechanical parts [119]. In mobile robotics, one detects round objects (pillars, tree trunks) by analyzing range readings from a 2D laser range finder used by a robot [197].




Figure 0.2 *Reflection coefficients in microwave engineering [15]. The observed values are marked by stars; the fitted circle is the dashed line.*

But perhaps the single largest field of applications where circles are fitted to data is nuclear physics. There one deals with elementary particles born in accelerators and colliders. The newborn particles move along circular arcs in a constant magnetic field; physicists determine the energy of the particle by measuring the radius of its trajectory; to this end they fit an arc to a string of mechanical or electrical signals the particle leaves in the detector [45, 53, 82, 106, 107, 136, 173, 174, 175]. Particles with high energy move along arcs with large radii (low curvature), thus fitting arcs to nearly straight-looking trajectories is quite common; this task requires very elaborate techniques to ensure accurate results.

We illustrate our discussion by a real-life example from archaeology. To estimate the diameter of a potsherd from a field expedition, the archaeologist traces the profile of a broken pot — such as the outer rim or base — with a pencil on a sheet of graph paper. Then he scans his drawing and transforms it into an array of pixels (data points). Lastly, he fits a circle to the digitized image by using a computer.

A typical digitized arc tracing a circular wheelmade antefix is shown in




Figure 0.3 A typical arc drawn by pencil with a profile gauge from a circular wheel-made antefix.

Fig. 0.3 (this image contains 7452 pixels). The best fitting circle found by a standard least squares procedure has parameters

$$\text{center} = (7.4487, 22.7436), \quad \text{radius} = 13.8251. \quad (1)$$

This does not seem challenging, as the arc in Fig. 0.3 is clearly visible to the naked eye, so one can even reconstruct a circle manually.




Figure 0.4: A fragment of the arc shown in Fig. 0.3.

Now suppose we can only see a small fragment of the above arc, with very few points on it. Fig. 0.4 shows a sample of merely 22 randomly chosen points from a tiny part of the original arc. Suppose we are to fit a circle to these points, without seeing the rest of the image. Is it possible?

Visually, the 22 points in Fig. 0.4 do not even form a clear circular arc, they rather look like a shapeless string. Reconstructing a circle manually from these 22 points appears an impossibility. However, the best computer algorithm returns the following parameters:

$$\text{center} = (7.3889, 22.6645), \quad \text{radius} = 13.8111. \quad (2)$$

Compare this to (1). The estimates are strikingly accurate!

The algorithm that produced the estimates (2) is the Levenberg-Marquard geometric fit (minimizing the geometric distances from the given points to the circle); it is described in Section 4.5. One may naturally want to estimate errors of the returned values of the circle parameters, but this is a difficult task for the EIV regression problems. In particular, under the standard statistical models described in Chapter 6, the estimates of the center and radius of the circle have infinite variances and infinite mean values! Thus, the conventional error estimates (based on the standard deviations) would be absurdly infinite. An approximate error analysis developed in Chapter 7 can be used to assess errors in a more realistic way; then the errors of the estimates (2) happen to be ≈ 0.1 .

We see that the problem of fitting circles and circular arcs to images has a variety of applications. It has attracted the attention of scientists, engineers, statisticians, and computer programmers. Many good (and not-so-good) algorithms were proposed; some to be forgotten and later rediscovered. For example, the Kåsa algorithm, see our Chapter 5, was published independently at least 13 times, the first time in 1972 and the last (so far) in 2006, see references in Section 5.1. But, despite the popularity of circle fitting applications, until the 1990s publications were sporadic and lacked a unified approach.

An explosion of interest in the problem of fitting circles and other geometric shapes to observed points occurred in the 1990s when it became an agenda issue for the rapidly growing computer science community, because fitting simple contours (lines, circles, ellipses) to digitized images was one of the basic tasks in pattern recognition and computer vision. More general curves are often approximated by a sequence of segments of lines or circular arcs that are stitched together (“circular splines”); see [12, 145, 158, 164, 165].

Since the early 1990s, many new algorithms (some of them truly brilliant) for fitting circles and ellipses have been invented; among those are circle fits using the Riemann sphere [123, 175] and conformal maps of the complex plane [159], “direct ellipse fit” by Fitzgibbon et al. [61, 63, 147] and Taubin’s eigenfit [176], a sophisticated renormalization procedure due to Kanatani [47, 94, 95] and a no less superb HEIV method due to Leedan and Meer [48, 49, 120], as well as the Fundamental Numerical Scheme by Chojnacki et al. [47, 48]. Chojnacki and his collaborators developed a unified approach to several popular algorithms [49] and did a remarkable job of explaining the underlying ideas.

Theoretical investigation also led to prominent accomplishments. These include consistent curve and surface fitting algorithms due to Kukush, Markovsky, and van Huffel [114, 130, 167], “hyperaccurate” ellipse fitting methods by Kanatani [102, 104], and a rather unconventional adaptation of the classical Cramer-Rao lower bound to general curve fitting problems [42, 96].

The progress made in the last 15 years is indeed spectacular, and the total output of all these studies is more than enough for a full size book on the subject. To the author’s best knowledge, no such book exists yet. The last book

on fitting geometric shapes to data was published by Kanatani [95] in 1996. A good (but rather limited) tutorial on fitting parametric curves, due to Zhang, appeared on the Internet in about the same year (and in print in 1997, see [198]). These two publications covered ellipse fitting methods existing in 1996, but not specifically circle fitting methods. The progress made after 1996 remains unaccounted for.

The goal of this book is to present the topic of fitting circles and circular arcs to observed points in full, especially accounting for all the recent achievements since the mid-1990s. I have tried to cover all aspects of this problem: geometrical, statistical, and computational. In particular, my purpose is to present numerical algorithms in relation to one another, with underlying ideas, to emphasize strong and weak points of each algorithm, and to indicate how to combine them to achieve the best performance. The book thoroughly addresses theoretical aspects of the fitting problem which are essential for understanding advantages and limitations of practical schemes. Lastly, an attempt was made to identify issues that remain obscure and may be subjects of future investigation.

At the same time the book is geared toward the end user. It is written for practitioners who want to learn the topic or need to select the right tool for their particular task. I have tried to avoid purely abstract issues detached from practice, and presented topics that were deemed most important for image processing applications.

I assume the reader has a good mathematical background (being at ease with calculus, geometry, linear algebra, probability and statistics) and some experience in numerical analysis and computer programming. I am not using any specific machine language in the book, though the MATLAB® code of all relevant algorithms may be found on our Web page [84].

The author is deeply indebted to his former supervisor at the Joint Institute for Nuclear Research (Russia), G. Ososkov, for his constant guidance in the studies on the circle fitting problem. The author is grateful to K. Kanatani for his strong support in research and especially in the design of this book. The author thanks his graduate students C. Lesort and A. Al-Sharadqah for their devotion to the subject and their help in preparing the manuscript and posting the computer code on the Web. Lastly, the author is partially supported by National Science Foundation, grant DMS-0652896.

The book is organized as follows, see diagram in Fig. 0.5. Chapter 1 is an introduction to the Errors-In-Variables regression analysis and gives its brief history (mostly in the context of the linear model). Chapter 2 summarizes the solution of the linear EIV problem and highlights its main properties (geometric and statistical). These two chapters do not deal with circles or arcs.

Chapter 3 gives the theory of fitting circles by least squares. It addresses the existence and uniqueness of the solution, describes various parametrization

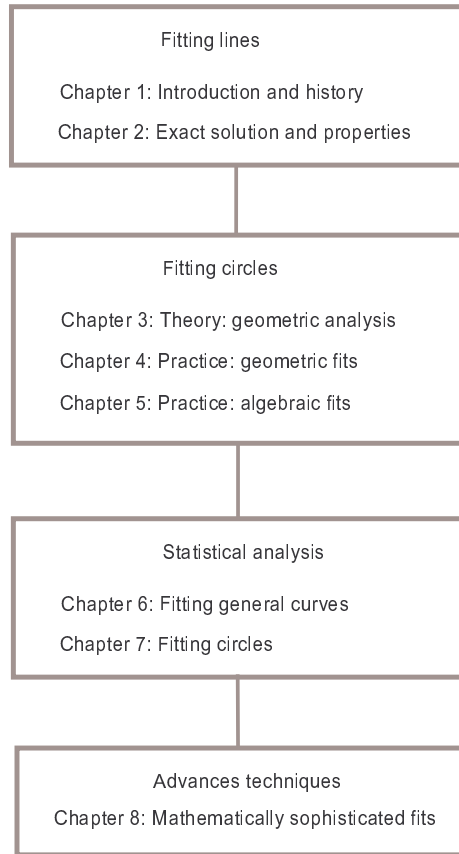


Figure 0.5: *The structure of the book.*

schemes for circles, and analyzes the shape of the objective function to be minimized (culminating in the important Two Valley Theorem).

Chapters 4 and 5 are devoted to practical circle fitting methods.

In Chapter 4 circles are fitted by minimizing geometric distances from observed points to the fitting circle, which is a classical (or geometric) fit. This is a nonlinear problem that has no closed form solution, so all algorithms are iterative, thus computationally intensive and subject to occasional divergence. We describe all popular schemes, in a historic perspective, emphasizing relations between one another, highlighting their advantages and drawbacks.

Chapter 5 deals with simplified circle fits, so called *algebraic fits*. They are fast, noniterative, and do not suffer from divergence. However, they are (in many cases) less accurate than the geometric fits of Chapter 4. Algebraic fits

are often used in mass data processing (especially in nuclear physics), where speed is of paramount importance. Algebraic fits are also used for initializing iterative geometric fitting procedures.

The reader interested in only practical algorithms can find all the relevant information in Chapters 3–5.

Chapters 6 and 7 make a sharp turn and plunge into statistical analysis of curve fitting methods. This is theoretical material, but I have tried to relate it to practice and explain all constructions and conclusions in practical terms. Chapter 6 is devoted to general nonlinear EIV regression, i.e., it covers arbitrary curves. Chapter 7 focuses on the specific task of fitting circles and circular arcs.

Chapters 6 and 7 may be of interest to professional statisticians.

Lastly, Chapter 8 presents a sample of “exotic” circle fits, including some mathematically sophisticated procedures — they make use of complex numbers and conformal mappings of the complex plane. This chapter is best for scientists with a solid mathematical background. The ideas behind methods of Chapter 8 are quite intriguing and resulting fits look very promising. This may be a starting point for future development of this subject.

Most illustrations in the book have been prepared with MATLAB. MATLAB® is a registered trademark of The MathWorks, Inc. For product information, please contact:

The MathWorks, Inc.
3 Apple Hill Drive
Natick, MA 01760-2098 USA
Tel: 508 647 7000
Fax: 508-647-7001
E-mail: info@mathworks.com
Web: www.mathworks.com