Robust circle reconstruction with the Riemann fit

R Frühwirth[1] **and A Strandlie**[2]

[1] Institute of High Energy Physics of the Austrian Academy of Sciences,
Nikolsdorfer Gasse 18, A-1050 Wien, Austria
[2] NTNU Norwegian University of Science and Technology,
P. O. Box 191, N-2802 Gjøvik, Norway

E-mail: rudolf.fruehwirth@oeaw.ac.at, are.strandlie@ntnu.no

Abstract. Finding and fitting circles from a set of points is a frequent problem in the data analysis of high-energy physics experiments. In a tracker immersed in a homogeneous magnetic field, tracks are close to perfect circles if projected to the bending plane. In a ring-imaging Cherenkov (RICH) detector, circles of photons around the crossing point of charged particles have to be found and their radii estimated. In both cases, non-negligible background may be present that tends to complicate the pattern recognition and to bias the circle fit. In this contribution we present a robust circle fit based on a modified Riemann fit that removes or significantly reduces the effect of background points. As in the standard Riemann fit, the measured points are projected to the Riemann sphere or paraboloid, and a plane is fitted to the projected points. The fit is made robust by replacing the usual least-squares regression by a least median of squares (LMS) regression. Because of the high breakdown point of the LMS estimator, the fit is insensitive to background points. The LMS plane is used to initialize the weights of an M-estimator that refits the plane in order to suppress eventual remaining outliers and to obtain the final circle parameters. The method is demonstrated on three sets of artificial data: points on a circle plus a comparable number of background points; points on two overlapping circles with additional background; and points obtained by the simulation of tracks in a drift chamber with mirror points and additional background. The results show high circle finding efficiency and small contamination of the final fitted circles.

1. Introduction

The fitting of circular arcs or full circles to a discrete set of measurements is an important problem in high-energy physics experiments. In ring-imaging Cherenkov (RICH) detectors rings of photons have to be reconstructed; from their estimated radii the Cherenkov emission angle of the crossing particles can be computed. Significant background can be expected, and there may be overlapping rings from other particles in the same event. Tracking detectors often are embedded in homogeneous magnetic fields, resulting in a nearly circular track in the bending plane of the particle trajectories. There is usually background from other tracks in the event, and in a drift chamber there is additional irreducible background because of the left-right ambiguity of the position measurement.

Methods for finding rings in RICH detectors include, amongst others, the Hough transform [1, 2, 3, 4], track-based maximum likelihood [5], global maximum likelihood [6], trackless ring finding via Markov chain Monte Carlo [7], and deconvolution [8]. Methods in the pattern recognition literature include learning automata [9], genetic algorithms [10] and center-based clustering [11]. It should be noted, however, that some of these are not relevant to the problem studied here, as

they are designed to the recognition of entire circular objects, for instance coins or balls, in an image.

In this contribution we propose a novel method for finding and fitting rings: a robust Riemann circle fit for tagging the noise points, followed by an adaptive least-squares estimator for obtaining the best possible precision. The Riemann fit [12, 13, 14] has recently been extended to a scale and location invariant version [15], following a recommendation in [16], including the full error propagation from the measured points to the final circle parameters. In this contribution it is shown how to make the circle fit robust against background. The least-squares regression, which is the core of the Riemann fit, is replaced by a robust regression algorithm. The robust fit is used to identify the points on the circle, and the found circle serves as the starting point for the final adaptive least-squares regression. If track information is available, it can be used to define a window for preselecting the points included in the robust fit.

The paper is organized as follows. After a brief recapitulation of the basic concept of the Riemann track fit the robust least median of squares (LMS) regression is briefly described (Section 2). The full chain of the LMS fit followed by the adaptive fit is evaluated first on full circles with background (Section 3), then on overlapping circles with background (Section 4), and finally on a simulation of points in a drift chamber with mirror points and additional background (Section 5). The paper is concluded by a summary and an outlook to further work.

2. Robust Riemann circle fit

We assume that the measured points $(u_i, v_i),\ i = 1, \ldots, N$, in the (x, y)-plane are given in Cartesian coordinates and have already been centered and scaled (see [15]). They are mapped to the "Riemann paraboloid", a second-degree surface with the equation $z = x^2 + y^2$, by the following projection:

$$x_i = u_i, \quad y_i = v_i, \quad z_i = u_i^2 + v_i^2.$$

In this way the circle in the plane with the equation

$$(u - u_0)^2 + (v - v_0)^2 = \rho^2$$

is mapped to the plane in 3D space with the equation

$$z - 2xu_0 - 2yv_0 = \rho^2 - u_0^2 - v_0^2.$$

The observed points $\boldsymbol{r}_i = (x_i, y_i, z_i)^\mathsf{T},\ i = 1, \ldots, N$, are fitted to the plane $\boldsymbol{n}^\mathsf{T}\boldsymbol{r} + c = 0$, where $\boldsymbol{r} = (x, y, z)^\mathsf{T}$, $\boldsymbol{n}$ is the unit normal vector and c is the signed distance of the plane from the origin. Let $d_i(\boldsymbol{n}, c) = \boldsymbol{n}^\mathsf{T}\boldsymbol{r}_i + c$ denote the distance of the point $\boldsymbol{r}_i$ from the plane. The standard Riemann fit is orthogonal LS regression, in which the (possibly weighted) sum of the squared distances d_i of the points from the plane is minimized:

$$(\boldsymbol{n}_{\mathrm{LS}}; c_{\mathrm{LS}}) = \arg\min_{(\boldsymbol{n};c)} \textstyle\sum_{i=1}^{N} w_i\, d_i^2.$$

The weights reflect the precision of the measured points. For the sake of simplicity, we assume equal weights $1/N$ from now on. The generalization to arbitrary weights is straightforward.

It is well known that LS regression is sensitive to background or outlying measurements. A highly robust alternative to LS regression is LMS regression [17], in which not the sum but the median of the squared distances is minimized:

$$(\boldsymbol{n}_{\mathrm{LMS}}; c_{\mathrm{LMS}}) = \arg\min_{(\boldsymbol{n};c)} \mathrm{med}_{i=1}^{N}\, d_i^2.$$

The breakdown point of the LMS regression is 0.5. This means that up to 50% of the measurements can be moved to arbitrary positions without introducing a bias into the fit.

The price to pay for the extraordinary robustness of the LMS regression is threefold: it is more difficult to compute than LS regression; it is less efficient in terms of its asymptotic variance; and there is no simple prescription for computing the covariance matrix of the estimated parameters. We have therefore used the LMS regression only for identifying the background measurements. Computing the LMS regression parameters is essentially a combinatorial problem: one has to find the subset of size $\lfloor N/2 \rfloor$ that minimizes the sum of the squared orthogonal distances to the plane. As it is hardly possible to check all such subsets, triples of points are selected randomly, defining a plane. The squared distances of all points to this plane are computed, and the plane with the smallest median of the squared distances is taken as an approximation of the LMS plane. This plane is used to compute the initial weights of the final fit, which is a redescending M-estimator with deterministic annealing [18, 19, 20], implemented as an iterated re-weighted LS regression. The final temperature of the annealing process is set to $T = 0.25$, resulting in a hard discrimination between inliers and outliers. Points having a final weight above 0.5 are tagged as inliers and are included in the final circle.

In the following sections the robust Riemann circle fit will be compared to randomized Hough transform (RHT). As in the robust Riemann fit, triples of points are selected randomly, from which the center and the radius of the circle defined by the triplet are computed and stored. After finding the most populated bin in the 2D histogram of the center coordinates, the found circle is defined as the medoid of all centers and radii in this bin. Points are tagged as inliers or outliers by a cut on their weighted radial distance to the found circle.

3. Performance with full circles plus background

We have simulated 10000 circles with superimposed background. The signal consists of points on a circle with radius $\rho = 1$, which is contained in a window of size 1.5×1.5 around the origin. The points are distributed uniformly on the circumference of the circle. In analogy to a RICH detector, where the number of photons follows a Poisson distribution, the number of points is drawn from a Poisson distribution with mean 20. Gaussian position errors with $\sigma_x = \sigma_y = 0.02$ are added to their (x, y)-positions. The background points are uniformly distributed in the window, and their number is drawn from a Poisson distribution with mean 15. This means that in about 22.3% of the cases there are at least as many background points as signal points in the window.

The performance can be summarized by the efficiency of the circle finding and by the purity and the completeness of the reconstructed circles. The efficiency of the circle finder can be measured by counting the outliers in the residual distributions of the circle parameters, which are the coordinates of the center and the radius. In this context, an outlier is defined as a residual which is outside the interval $[m-4\,s, m+4\,s]$, where m and s are robust estimates of the location and the spread of the residual distribution, respectively. We have chosen m as the median and s as the median of the absolute deviations from the median (MAD) times 1.485, which is very close to one for large standard normal samples. According to this criterion, the efficiency is about 98%. The average bias μ, the standard error σ and the root-mean-square error RMSE of the estimated circle parameters (x_c, y_c, ρ) are listed in Table 1, along with the corresponding results of the RHT. The robust Riemann fit gives smaller RMSE for all three parameters.

The purity of a found circle is defined as the ratio of the signal inliers over all inliers. Its average over all 4899 found circles is 94.6%. The completeness is the ratio of the signal inliers over all signal points. Its average is 98.8%. The corresponding figures for the RHT are again 98% for the efficiency, 92.9% for the average purity, and 99.7% for the average completeness. The robust fit is almost twice as fast as the RHT.

Figure 1 shows two examples of circle fits. The true circle is drawn in cyan, the LMS circle

in magenta, and the final circle in black. In Example 1 on the left, which is representative for the overwhelming majority of cases, the three circles are virtually identical, although two background points that are close to the true circle are included. In Example 2 on the right, the circle is biased by the inclusion of two background points by the LMS regression (magenta dotted circle), from which the M-estimator (black circle) cannot recover.

Table 1. Average bias μ, standard error σ and root-mean-square error RMSE of the estimated circle parameters. Left: robust Riemann fit, right: randomized Hough transform.

	x_c	y_c	ρ
$\mu \cdot 10^3$	0.027	0.037	0.815
$\sigma \cdot 10^2$	0.784	0.781	0.546
RMSE $\cdot 10^2$	0.784	0.781	0.552

	x_c	y_c	ρ
$\mu \cdot 10^3$	−0.028	0.151	0.173
$\sigma \cdot 10^2$	1.140	1.147	0.811
RMSE $\cdot 10^2$	1.140	1.147	0.811

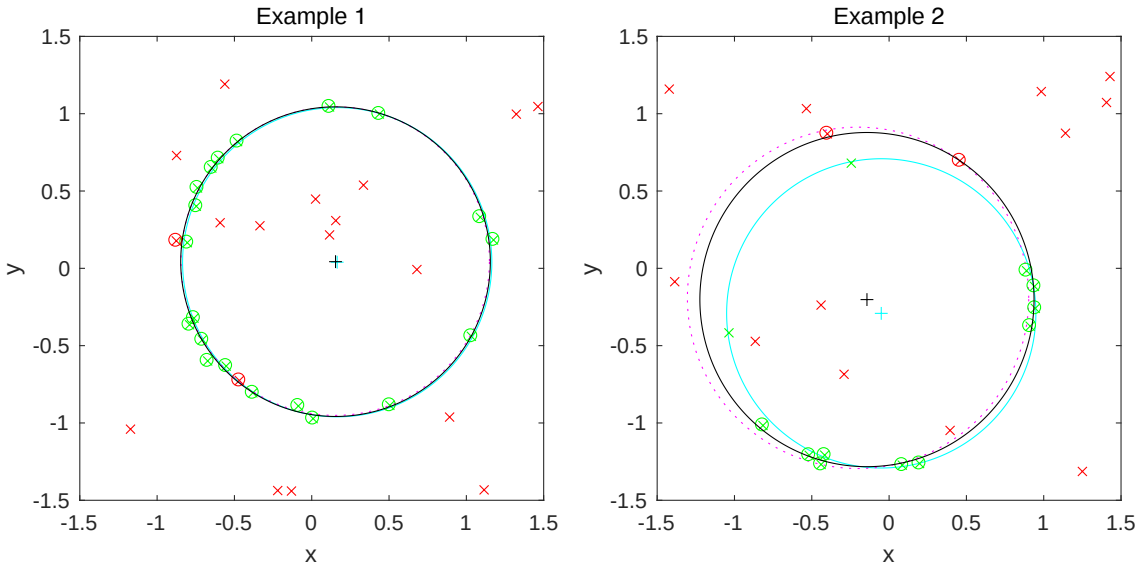

Figure 1. Two examples of fitted circles. Cyan: true circle with true center, green: signal points, red: background points, circled: points tagged as inliers by the M-estimator, magenta dotted: LMS circle, black: final circle with final center

4. Performance with overlapping circles plus background

We have simulated 10000 pairs of overlapping circles that are fully contained in the square $[-2.5, 2.5] \times [-2.5, 2.5]$. The primary circle is centered at the origin and has the radius $\rho_1 = 1$. The number of points on this circle is drawn from a Poisson distribution with mean $N_1 = 20$. The center of the secondary circle is uniformly distributed in the square $[-1.5, 1.5] \times [-1.5, 1.5]$, its radius ρ_2 is drawn from a uniform distribution in $[0.5\,\rho_1, \rho_1]$, and the number of points is drawn from a Poisson distribution with mean $N_2 = 15$. The points on the circles are smeared with Gaussian errors as before. The background points are uniformly distributed in the window, and their number is drawn from a Poisson distribution with mean $N_b = 10$.

After the fit, for each of the circles the sum of the weights of its point are computed, and the circle with the larger sum is tagged as found. The efficiency of the circle finding as well as

the purity and completeness of the found circles are defined as in Section 3. The total circle finding efficiency, measured as above, is 94.7%. Of the 4735 found circles inside the efficiency cut, 2498 are primary circles and 2237 are secondary circles. The purity averaged over all circles within the efficiency cut is 94.2%, the completeness is 98.7%. There is no significant difference between primary and secondary circles. The corresponding figures for the RHT are 97.1% for the efficiency, 92.8% for the average purity, and 99.8% for the average completeness. The average bias μ, the standard error σ and the root-mean-square error RMSE of the estimated circle parameters (x_c, y_c, ρ) are listed in Table 1, along with the corresponding results of the RHT. Again, the robust Riemann fit gives smaller RMSE for all three parameters, but the RHT is slightly more efficient. As the robust fit needs more iterations with larger background, the RHT is slightly faster in this case.

Figure 2 shows an example of a successful fit. The left hand panel shows all points including the background, the right hand panel shows which circle has been found and which points are tagged as inliers. After removing the points on the found circle the procedure can be repeated in order to find the other circle.

Table 2. Average bias μ, standard error σ and root-mean-square error RMSE of the estimated circle parameters. Left: robust Riemann fit, right: randomized Hough transform.

	x_c	y_c	ρ
$\mu \cdot 10^3$	0.168	0.081	0.736
$\sigma \cdot 10^2$	0.810	0.818	0.579
RMSE $\cdot 10^2$	0.810	0.818	0.583

	x_c	y_c	ρ
$\mu \cdot 10^3$	0.253	0.031	0.291
$\sigma \cdot 10^2$	1.112	1.131	0.813
RMSE $\cdot 10^2$	1.113	1.131	0.814

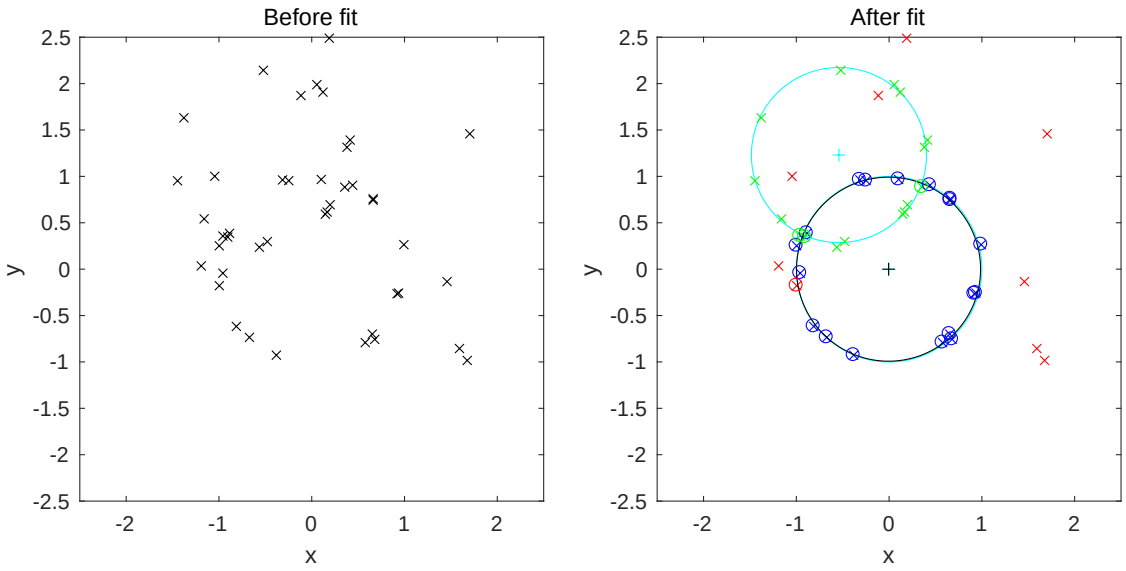

Figure 2. An example of a successful fit with overlapping circles. Left: all points before the fit. Right: after the fit; cyan: true circles with true centers, blue: signal points of the primary circle, green: signal points of the secondary circle, red: background points, circled: points tagged as inliers by the M-estimator, black: final circle with final center

5. Performance with mirror points

We have simulated tracks in a simplified and down-scaled model of the central drift chamber of the Belle II experiment [21], without the stereo layers. The model has 40 layers of sense wires, the cell size is 2 cm, and the point resolution is 150 μm. Each true point has a mirror point. In each layer, there is the probability p_{N} of an additional noise point plus its mirror point in either the signal wire itself or in the two adjacent wires on each side of the signal wire. The LMS regression is the same as in the previous sections. The M-estimator has been modified to include competition between a point and its mirror point, which is effectively an EM algorithm with deterministic annealing [22] and a special case of the Elastic Planes algorithm in [12].

Table 3 shows the average purity and the average completeness of the reconstructed track for three values of p_{N}. The completeness is close to one in all cases, the purity drops to about 94% at the highest noise level, as more noise points are included in the tracks. An example of a track at the highest noise level is shown in Figure 3. This track has 17 noise points and their mirror points. In layer 2, the signal point is missing and its mirror point is included. Three noise points are included in layers 6, 15 and 28.

Table 3. Average purity and average completeness of tracks with mirror points at three noise levels.

p_{N}	0.0	0.2	0.4
purity	0.994	0.966	0.942
completeness	0.994	0.998	0.998

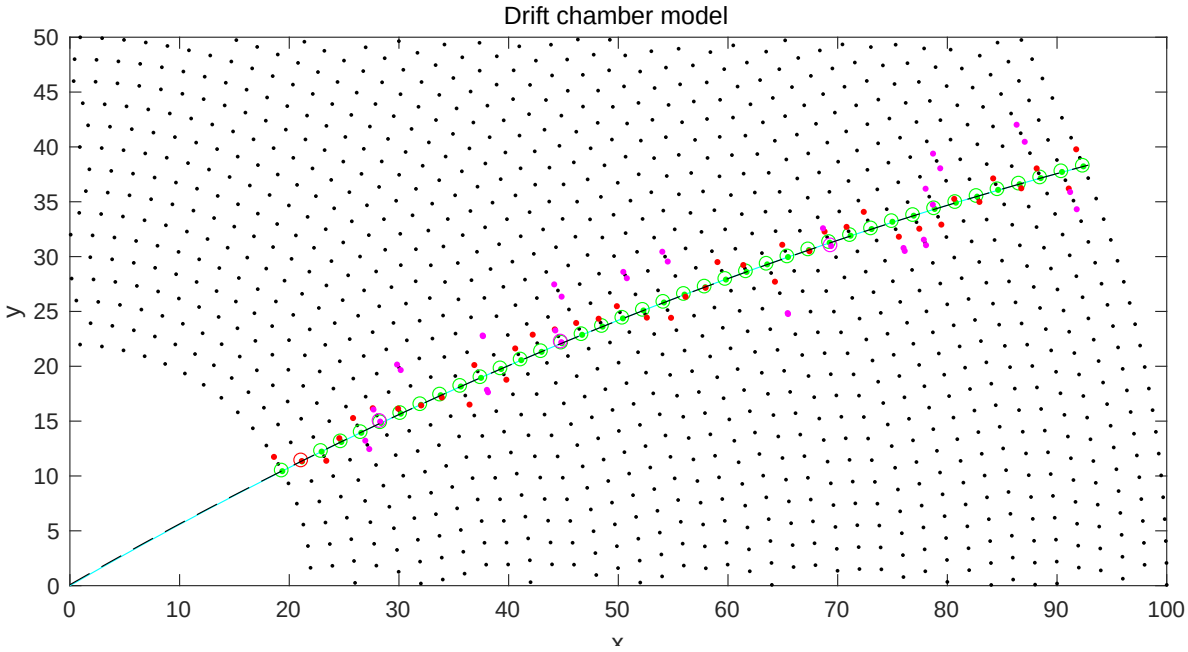

Figure 3. An example of a track at noise level $p_{\mathrm{N}} = 0.4$. Cyan: true track, green: true points, red: mirror points, magenta: noise points, circled: points included in the track, black: final track.

6. Summary and outlook

It has been shown that circle finding against a substantial background can be implemented by a robust Riemann circle fit based on LMS regression in combination with an M-estimator. Its performance is comparable to that of the randomized Hough transform. As LMS regression is computationally quite costly, further work will explore other robust regression methods such as least trimmed squares [17] or L_1 regression. Alternatively, the results of the randomized Hough transform can serve as the initial values for the M-estimator, in particular in the case of several overlapping circles. Finally, the fitting of tracks with mirror points will be extended to the case of crossing or overlapping tracks.

References

[1] Adamczewski J et al 2014 *Nucl. Instr. Meth. Phys. Res.* A **766** 250
[2] Volpe G on behalf of the ALICE Collaboration 2014 *Nucl. Instr. Meth. Phys. Res.* A **766** 259
[3] Abbon P et al 2008 *Nucl. Instr. Meth. Phys. Res.* A **595** 233
[4] Wilkinson G 2008 *Nucl. Instr. Meth. Phys. Res.* A **595** 228
[5] Starič M 2008 *Nucl. Instr. Meth. Phys. Res.* A **595** 237
[6] Buszello C P on behalf of the LHCb RICH Collaboration 2008 *Nucl. Instr. Meth. Phys. Res.* A **595** 245
[7] Lester C G 2006 *Nucl. Instr. Meth. Phys. Res.* A **560** 621
[8] Morháč M, Matoušek V 2011 *Nucl. Instr. Meth. Phys. Res.* 2011 A **639** 249
[9] Cuevas E et al 2012 *IET Image Processing* **6** 1124
[10] Ayala-Ramirez V et al 2006 *Pattern Recognition Letters* **27** 652
[11] Scitovski R and Marošević T 2015 *Pattern Recognition Letters* **52** 9
[12] Strandlie A, Wroldsen J, Frühwirth R and Lillekjendlie B 2000 *Comput. Phys. Commun.* **131** 95
[13] Strandlie A and Frühwirth R 2002 *Nucl. Instr. Meth. Phys. Res.* A **480** 734
[14] Strandlie A, Wroldsen J and Frühwirth R 2002 *Nucl. Instr. Meth. Phys. Res.* A **488** 332
[15] Strandlie A and Frühwirth R 2017 *Nucl. Instr. Meth. Phys. Res.* A **867** 72
[16] Chernov N 2010 *Circular and linear regression* (Boca Raton: CRC Press)
[17] Rousseeuw P J and Leroy A M 1987 *Robust Regression and Outlier Detection* (New York: Wiley & Sons)
[18] Huber P J 2004 *Robust Statistics: Theory and Methods* (New York: Wiley & Sons)
[19] Li S Z 1996 *Pattern Recognition* **29** 159
[20] Frühwirth R and Waltenberger W 2008 *Austrian J. Stat.* **37** 301
[21] Taniguchi N on behalf of Belle II collaboration 2017 *J. Instrum.* **12** C06014
[22] Ueda N and Nakano R 1998 *Neural Netw.* **11** 271