

# Arc Length-Based Aspect Ratio Selection

Justin Talbot, John Gerth, and Pat Hanrahan

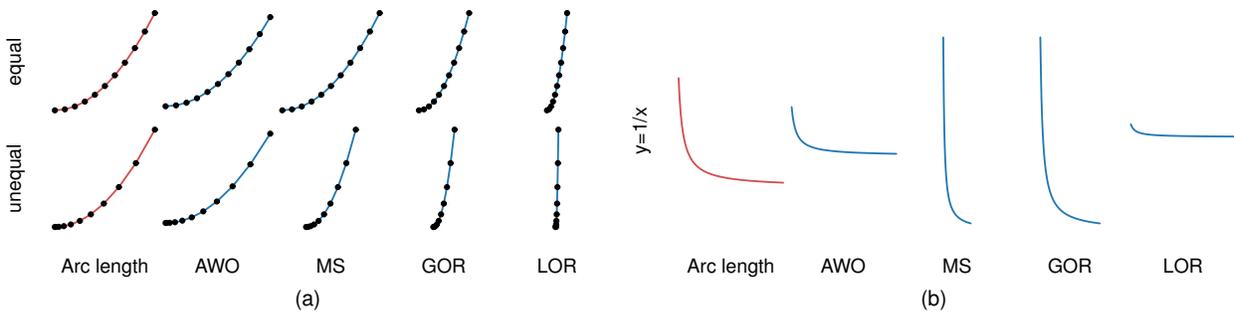


Fig. 1. Comparison of aspect ratio selection algorithms. On the left, only AWO and our new *arc length-based* method are invariant (produce the same aspect ratio) under changes to the parameterization of a curve. On the right, only our arc-length method preserves the natural symmetry of the curve  $y = 1/x$ .

**Abstract**—The aspect ratio of a plot has a dramatic impact on our ability to perceive trends and patterns in the data. Previous approaches for automatically selecting the aspect ratio have been based on adjusting the orientations or angles of the line segments in the plot. In contrast, we recommend a simple, effective method for selecting the aspect ratio: minimize the arc length of the data curve while keeping the area of the plot constant. The approach is parameterization invariant, robust to a wide range of inputs, preserves visual symmetries in the data, and is a compromise between previously proposed techniques. Further, we demonstrate that it can be effectively used to select the aspect ratio of contour plots. We believe arc length should become the default aspect ratio selection method.

**Index Terms**—Aspect ratio selection, Banking to 45 degrees, Orientation resolution.

## 1 INTRODUCTION

Our ability to perceive trends and patterns in a given visual display of data is heavily influenced by the aspect ratio. It affects densities, relative distances, and orientations within the plot—all important perceptual features that impact plot interpretation.

For data visualization tools, an important practical question is how to automatically select the aspect ratio for a given plot. Cleveland et al. proposed two methods, *median slope* (MS) [4] and *length weighted average orientation* (AWO) [3], both based on centering the orientations of a line plot’s constituent segments around 45 degrees, an approach that they call *banking to 45°*. More recently, Heer and Agrawala [7] proposed two alternative methods, *global orientation resolution* (GOR) and *local orientation resolution* (LOR), based on maximizing the orthogonality of pairs of line segments.

All four proposed methods have shortcomings. Three are not parameterization invariant, meaning that the semantically unimportant way a curve is approximated by line segments can dramatically change the selected aspect ratio (Figure 1a). None preserve semantically important symmetric shapes (Figure 1b). Most can produce extreme aspect ratios on relatively simple inputs (Figure 3). And all have been designed to work with curves that are functions of  $x$ , limiting their broader use.

To address these problems, we recommend a simple and robust op-

timization criterion for choosing the aspect ratio: minimize the arc length of the plotted curve while keeping the area of the plot constant. This approach is parameterization invariant, preserves symmetries, and produces reasonable aspect ratios on a wide range of inputs, while being fast to evaluate. Furthermore, it can easily be used with general curves in 2D, which we demonstrate by selecting good aspect ratios for contour plots. We believe that the properties of the arc length approach are a qualitative improvement over the previous approaches and that it should become the default aspect ratio selection method.

The remainder of the paper is divided into four sections. In Section 2, we review the previous work in aspect ratio selection. Next, we describe the new arc length method and discuss its connections to previous methods. In Section 4, we validate the arc length approach by demonstrating that it selects better aspect ratios on both a range of simple algebraic curves and on time series data sets; by showing that it can generalize by selecting aspect ratios for 2D contour plots; and by showing that it is computationally feasible. Finally, we conclude with discussion of the algorithm and the need for a concrete theory tying aspect ratio selection to perceptual principles.

## 2 PREVIOUS WORK

The aspect ratio question was first treated rigorously by Cleveland, McGill, and McGill [4]. Noting that changing the aspect ratio of a plot changes the perceived slopes of lines in a plot, they hypothesized that aspect ratios which maximized the *orientation resolution*, the angle between the line segments, would minimize error in the estimation of the ratio of the slopes of those line segments. A designed experiment displaying pairs of line segments supported this hypothesis.

They went on to show that in the case of two line segments in the first quadrant, the angle between them can be maximized (and thus slope ratio minimized) by selecting the aspect ratio that centers the two line segments around 45°. This led them to suggest that the aspect ratio of plots with more than two segments could be chosen by

- Justin Talbot is with Stanford University, E-mail: jtalbot@stanford.edu.
- John Gerth is with Stanford University, E-mail: gerth@cs.stanford.edu.
- Pat Hanrahan is with Stanford University, E-mail: hanrahan@cs.stanford.edu.

Manuscript received 31 March 2011; accepted 1 August 2011; posted online 23 October 2011; mailed on 14 October 2011.

For information on obtaining reprints of this article, please send email to: tvcg@computer.org.

placing the median segment slope at 1, an approach they called *median absolute slope* (MS).

Cleveland later suggested an alternative method, *length-weighted average orientation* (AWO), that sets the length-weighted average of the absolute segment orientations (the angle made with the horizontal) at 45° [3]. The switch from slope to orientation was motivated by the fact that our perceptual processes are sensitive to orientation, not slope. The length-weighting ensures that subdividing line segments will not change the resulting aspect ratio, providing some measure of robustness to changes in the input representation. Across a wide range of inputs, AWO results in flatter aspect ratios than MS; Cleveland asserts that these are more visually pleasing aspect ratios [2].

Later, Heer and Agrawala [7] proposed selecting the aspect ratio by maximizing the sum of squares of the angles between all pairs of segments in the plot. They called this the *global orientation resolution* (GOR) method and, unlike the previous methods, it has the nice property of being derived from Cleveland’s hypothesized perceptual justification. In practice, this approach gives very similar results to an unweighted average orientation approach at the cost of a very expensive optimization function. A computationally cheaper approach is to only consider the orientation resolution between adjacent pairs of segments. Heer and Agrawala called this approach the *local orientation resolution* (LOR). Both approaches have trouble dealing with perfectly horizontal and vertical segments, so they must be removed first. “For completeness,” Heer and Agrawala also suggest a third method, setting the *average slope* (AS) to be 1, an approach closely related to our arc length method.

Guha and Cleveland [5] recently suggested the *resultant vector* method which has a simple, tractable algebraic form that they employ to study theoretical properties of aspect ratio selection. Geometrically, the plot’s line segments are treated as vectors and positioned in the first quadrant by reflection around the  $x$  and  $y$  axes. The sum of these is the plot’s “resultant vector”. The transformation which banks it to 45 degrees determines the plot’s aspect ratio. Mathematically, this simplifies to computing the ratio of the segments’ total variation in the  $y$  and  $x$  directions. With this formulation, Guha and Cleveland have been able to derive geometric properties for the aspect ratios of select families of curves and demonstrate the statistical convergence of the aspect ratio as random noise is added to the data.

Our arc length method arises from an independent analysis and geometric interpretation of the aspect ratio selection problem. But we can show that the arc length method can be interpreted as a generalization of the resultant vector method. Because of this, many of the arc length method’s empirical advantages demonstrated in this paper are shared by the resultant vector method as well.

Visualization systems have long supported interactive aspect ratio changes (at least since Buja’s Data Viewer [1]), a powerful and underutilized analytic tool. Automatic aspect ratio selection methods are complementary to such interaction, providing reasonable defaults, enabling automated plot generation, and helping to develop insight into the relationship between aspect ratio and human perception.

Finally, Heer and Agrawala [7], in the second part of their paper, suggest using frequency space analysis to decompose a time series into multiple components each of which can be banked independently. Since the decomposition problem is independent of the aspect ratio selection problem, in this paper we assume that, if necessary, the input has already been processed using Heer and Agrawala’s method or another statistical decomposition technique.

### 3 ARC LENGTH-BASED ASPECT RATIO SELECTION

In this section we propose a new aspect ratio selection method: minimize the arc length of the plotted curve while keeping the area of the plot constant. We first briefly describe desirable properties we would like the algorithm to have. We then describe the arc length method in more detail. Finally, we compare it to previous algorithms and show how it can be derived from the optimization of *orientation resolution* by following reasoning similar to Heer and Agrawala’s GOR and LOR methods.

### 3.1 Criteria

In developing a new aspect ratio selection algorithm, we were guided by the following design criteria suggested by our experience with the methods in the previous work.

- **Scale invariant.** Changes to the input scale of one or both axes should not change the resulting aspect ratio.
- **Parameterization invariant.** The displayed curve is naturally approximated by a set of line segments. Changes to this approximation that do not drastically alter the *visual* form of the curve should not impact the aspect ratio. In particular, the method should not assume that the values are equally spaced on the  $x$  axis.
- **Robust.** Corner cases, such as horizontal, vertical, or collinear segments, should not require special handling. The addition of a small amount of noise to the curve should not change the aspect ratio.
- **Symmetric.** The most reasonable aspect ratio for a curve symmetric around  $y = x$  is one that preserves that symmetry (e.g. Figure 1b). (This is the weakest of our criteria. However, in the absence of any perceptual results to the contrary, a symmetric banking is the most parsimonious result.)
- **Fast to compute.** To be useful in interactive visualization tools, the method must be relatively fast. We would like the computation cost to be at most  $O(n)$  in the number of line segments in the plot.

### 3.2 Algorithm

The input to our algorithm is the set of line segments for which we want to compute an aspect ratio for display. The arc length method consists of simply minimizing the total length of the line segments under the constraint that the area of the plot is preserved.<sup>1</sup>

We enforce the area-preserving constraint by not looking for an aspect ratio directly. Instead, we search for a transformation that adjusts the plot’s aspect ratio while preserving area. Such transformations are called *hyperbolic rotations* or *squeeze mappings* and can be written in matrix form as:

$$\begin{pmatrix} 1/\sqrt{a} & 0 \\ 0 & \sqrt{a} \end{pmatrix}$$

(This definition follows Cleveland, who defined aspect ratio as *height/width*. Heer and Agrawala use the, perhaps more common, *width/height* definition.) We can confirm that these transformations preserve area since the determinant of the matrix is always

$$\frac{\sqrt{a}}{\sqrt{a}} - 0 = 1$$

Given  $a$ , we can multiply the coordinates of the input line segments by the corresponding squeeze mapping to produce line segments appropriate for display.

We select  $a$  such that the total length of the line segments is minimized. To do this we solve the optimization problem,

$$\min_{a \in (0, \infty)} \sum_{i=1}^N \left\| \frac{\Delta x_i}{\sqrt{a}}, \sqrt{a} \Delta y_i \right\| \quad (1)$$

where  $\Delta x_i$  and  $\Delta y_i$  are the lengths of the  $i$ th line segments in the  $x$  and  $y$  directions and  $\|x\|$  is the Euclidean length. Note that the lengths of the line segments will be scaled by the area-preserving aspect ratio transformation. Due to the square root in the Euclidean length measure, this optimization problem does not have a closed form solution; however,

<sup>1</sup>Constraining the area simplifies the algorithm, but isn’t strictly necessary. Arc length scales linearly with the square root of the plot area, so we could also minimize the arc length divided by the square root of the area.

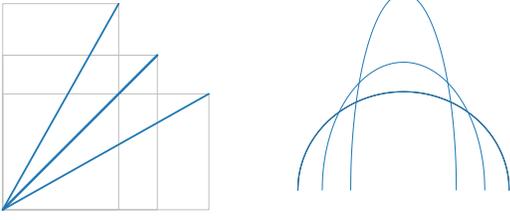


Fig. 2. For a straight line, the minimum length is achieved at 45 degrees. For an ellipse, the minimum arc-length is achieved when banked to a circle.

it has a unique minimum and can be easily and quickly solved using a variety of methods. In practice, we parameterize the optimization search with  $\log(a)$ , rather than with  $a$  directly. This makes the cost function smooth and symmetrical around the minimum and removes the need to deal with the lower bound on  $a$ .

We can gain some insight into the behavior of our proposed method by considering simple cases (Figure 2). If we stretch or squish a rectangle bounding a diagonal line while keeping the contained area the same, the line will reach a minimal length when the rectangle is square. Thus, a single diagonal line will always be banked to 45 degrees. Similarly, consider a collection of line segments approximating an ellipse. Our method will bank this to a circle since the ellipse with minimal circumference enclosing a fixed area is a circle.

### 3.3 Connection to Orientation Resolution

The arc length metric is closely connected to Cleveland et al.’s orientation resolution perception hypothesis which assumes that, to make the ratios of segment slopes easily visible, the angle between the segments should be made as large as possible.

Heer and Agrawala’s GOR method [7] attempts to directly maximize the orientation resolution by finding the aspect ratio,  $a$ , that maximizes the sum of the squares of the larger angles between all pairs of line segments:

$$\max_a \sum_i \sum_j \theta_{ij}(a)^2$$

Following similar reasoning, we propose an alternate criterion for maximizing orientation resolution:

$$\max_a \frac{\sum_i \sum_j |\sin(\theta_{ij}(a))| l_i(a) l_j(a)}{\sum_i \sum_j l_i(a) l_j(a)} \quad (2)$$

where  $\theta_{ij}(a)$  is the angle made between the two line segments  $i$  and  $j$ , and  $l_i$  and  $l_j$  are the lengths of the  $i$ th and  $j$ th line segments, all of which vary as functions of  $a$ .

In comparison to Heer and Agrawala’s approach, our formulation has two primary differences. First, our formulation includes weighting by the lengths of the line segments. This makes the measure invariant to parameterization changes, addressing a shortcoming of GOR. Second, we use  $|\sin(\theta_{ij})|$  instead of  $\theta_{ij}^2$ . Both functions reach a maximum when lines are at right angles to each other; but our choice is the same for both the larger and smaller angles made between line segments  $i$  and  $j$ ; so, unlike GOR, we do not need to specify which is selected.

The behavior of this criterion is easy to describe. It will be 0 if all the lines are parallel and reaches a maximum at 1 if all the lines are at right angles to each other. The latter case is, in practice, impossible, since we include the terms for each line segment paired with itself and this pairing obviously cannot be at right angles. However, the intuition holds that the closer this number is to 1, the more line segments are orthogonal.

Now we show that Equation (2) can be simplified to the arc length criteria. First, the numerator can be restated as a cross product since

$$|\sin(\theta_{ij}(a))| l_i(a) l_j(a) = |v_i(a) \times v_j(a)|$$

where  $v_i$  is the  $i$ th line segment represented as a vector. Next, within the family of area-preserving aspect ratio changes, the cross product is not a function of  $a$  and can be ignored in the optimization:

$$v_i(a) \times v_j(a) = \frac{\sqrt{a}}{\sqrt{a}} x_i y_j - \frac{\sqrt{a}}{\sqrt{a}} x_j y_i = x_i y_j - x_j y_i$$

Finally, the denominator can be easily turned into the square of the total arc length.

$$\sum_i \sum_j l_i(a) l_j(a) = \sum_i l_i(a) \sum_j l_j(a) = (\sum_i l_i(a))^2$$

Thus, orientation resolution (as expressed by Equation (2)) can be maximized by minimizing the total arc length.

### 3.4 Connection to Resultant Vector and Average Slope

We have presented the arc length method using the Euclidean metric to measure length. A potentially useful property of the Euclidean metric is that it is rotationally invariant. However, there may be reasons to prefer other metrics. For example, if we replace the Euclidean metric with the Manhattan metric, we can get a closed form solution for  $a$ :

$$\min_{a \in (0, \infty)} \sum_{i=1}^N \left| \frac{\Delta x_i}{\sqrt{a}} \right| + |\sqrt{a} \Delta y_i| \Rightarrow a = \frac{\sum_{i=1}^N |\Delta y_i|}{\sum_{i=1}^N |\Delta x_i|} \quad (3)$$

This is exactly equivalent to the *resultant vector* method proposed by Guha and Cleveland [5]. Thus, another geometric interpretation of the resultant vector method is that it minimizes the Manhattan length of the curve. Further, in the special case of a time series with equally spaced data points, this Manhattan metric approach is also equivalent to Heer and Agrawala’s AS method.

## 4 RESULTS

To demonstrate the effectiveness of the arc length method, we first compare the methods on a set of simple algebraic curves and then on a variety of time series data sets. The former serves to emphasize the robustness and symmetry of the arc length method. The latter shows that arc length is a consistent compromise method, producing aspect ratios that fall within the range of previous methods. Next, we show that the arc length method is also effective at selecting aspect ratios for general 2D curves. Finally, we discuss the performance of the method.

### 4.1 Method Comparison

**Algebraic Curves.** In Figure 3, we compare arc length to the four methods recommended in the previous work on a variety of simple algebraic curves. The curves were sampled evenly along the x-axis to make the comparison to the non-length-weighted methods as fair as possible. For the GOR and LOR methods, we discarded horizontal and vertical segments as recommended by Heer and Agrawala.

AWO is the most similar to arc length, though arc length’s aspect ratios are slightly taller with higher orientation resolution. This is promising since Cleveland found AWO to generally produce superior aspect ratios. The quarter circle example, in the fifth row, demonstrates that arc length can preserve symmetry, whereas AWO produces a segment of an ellipse. As the scale parameter of the gamma ( $\Gamma$ ) distribution increases, arc length maintains roughly the same aspect ratio. In contrast, AWO gets flatter, biased by the large nearly flat region. The last two rows show where the results of arc length and AWO diverge most noticeably. AWO, GOR, and LOR all scale down the narrow triangle case to maintain 45° lines. Heer and Agrawala argue that this behavior is preferred, since it maintains ideal slopes. However, in the limit, as the triangle’s width approaches 0, this behavior will select a completely flat aspect ratio, obscuring the presence of the spike entirely. Arc length’s behavior is different; it selects nearly the same aspect ratio regardless of the triangle’s width. While this deviates from Cleveland’s 45° recommendation, it ensures that the spike will remain visible regardless of how narrow it is.

MS and GOR are susceptible to producing abnormally tall aspect ratios when there are nearly flat regions of the curve. In the gamma

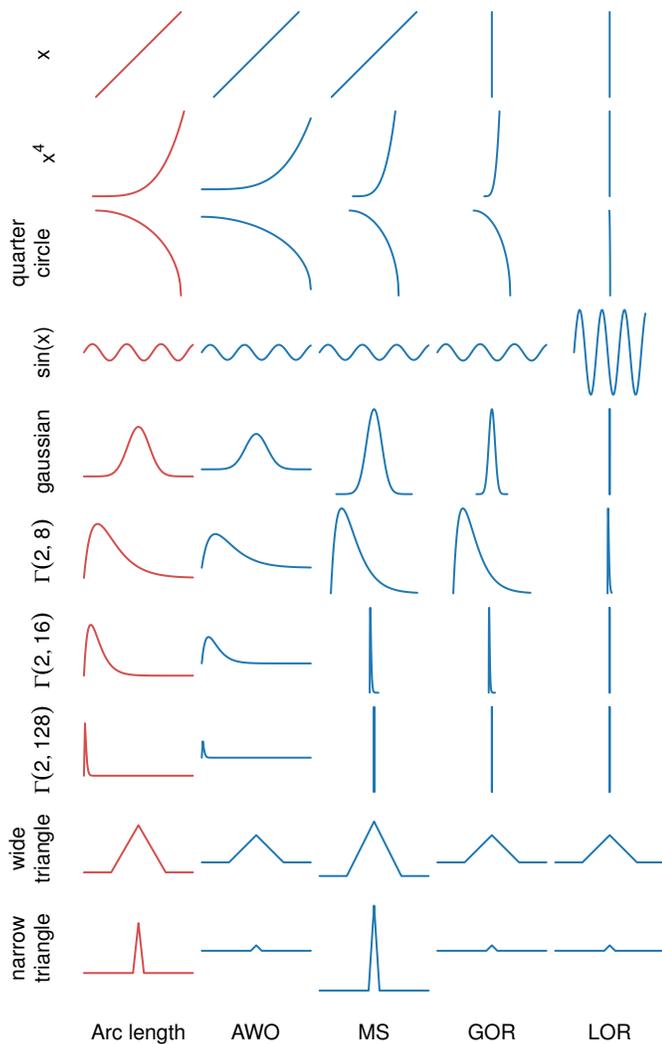


Fig. 3. Comparison of aspect ratios produced by recommended banking methods for a selection of simple algebraic curves evenly sampled along the  $x$ -axis. Arc length and AWO are similar, but arc length selects a more pleasing symmetric aspect ratio for the quarter circle. MS and GOR work well in many cases, but have degeneracies when dealing with nearly horizontal or vertical segments. LOR performs poorly across the board.

distribution cases, both methods degenerate to a 0-width aspect ratio as the shape parameter increases. Additionally, note the degenerate result of GOR on the  $y = x$  plot due to the colinearity of all the segments. Finally, the behavior of LOR is quite idiosyncratic. We cannot recommend its use for robust aspect ratio selection.

It may be possible to improve the behavior of GOR or LOR by culling not just horizontal and vertical segments, but also line segments that are within some threshold of vertical and horizontal. However, it's not obvious if this threshold can be made scale invariant. With a naïve thresholding approach, changes to the input scale of an axis will lead to changes in the input slope and, thus, to changes in the set of lines falling within the culling threshold, making the resulting algorithm not scale invariant.

**Time Series.** Figure 4 shows the log of the aspect ratio chosen by the banking heuristics for a large set of artificial and natural time series [11]. Like Heer and Agrawala, we find that the length-weighted AWO method (circle) consistently produces flatter aspect ratios than the median slope method (triangle). GOR (square) is often quite similar to median slope, while LOR (diamond) is more variable.

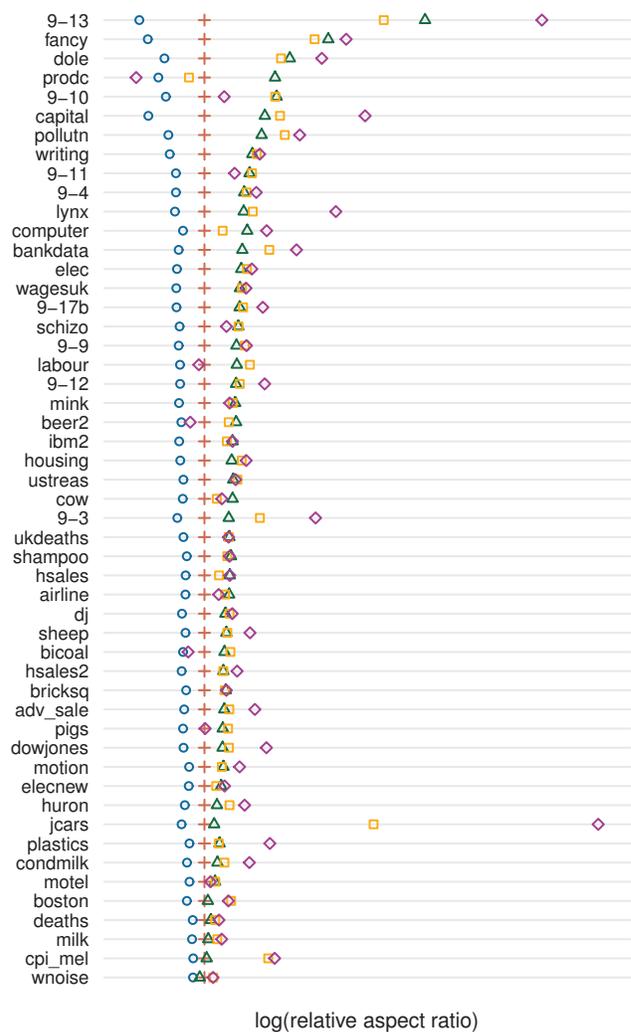


Fig. 4. Comparison of aspect ratios produced by previous methods relative to arc length for select time series data sets: AWO (circle), arc length (cross), GOR (square), MS (triangle), and LOR (diamond). The arc length method selects compromise aspect ratios between AWO and the other methods.

The arc length heuristic (cross) consistently picks an aspect ratio that falls within the range of the previous methods. This is encouraging since the previous techniques are known to produce bankings that effectively highlight trends in time series data. Further, the arc length aspect ratio often falls near the center of the range. On a log scale, the midpoint between two aspect ratios corresponds to their geometric mean, a natural compromise aspect ratio.

Figure 5 shows a selection of the time series summarized in Figure 4. As with the algebraic curves, arc length produces somewhat taller aspect ratios than AWO. The other methods generally produce much taller aspect ratios. In some cases, as in *9-13* and *jcars* the result is obviously too tall. In others, such as *dole* and *fancy*, the result of the taller aspect ratio is to highlight a lower frequency trend in the data, obscuring the higher frequency cycles of interest.

## 4.2 Curves in 2D

In previous work, banking has been only applied to plots where the curve is a function of  $x$ . However, banking heuristics can also be applied to general curves in 2D. The arc length approach has two advantages in this task. First, the arc length method's parametric invariance means that arc length does not require even sampling of the curve

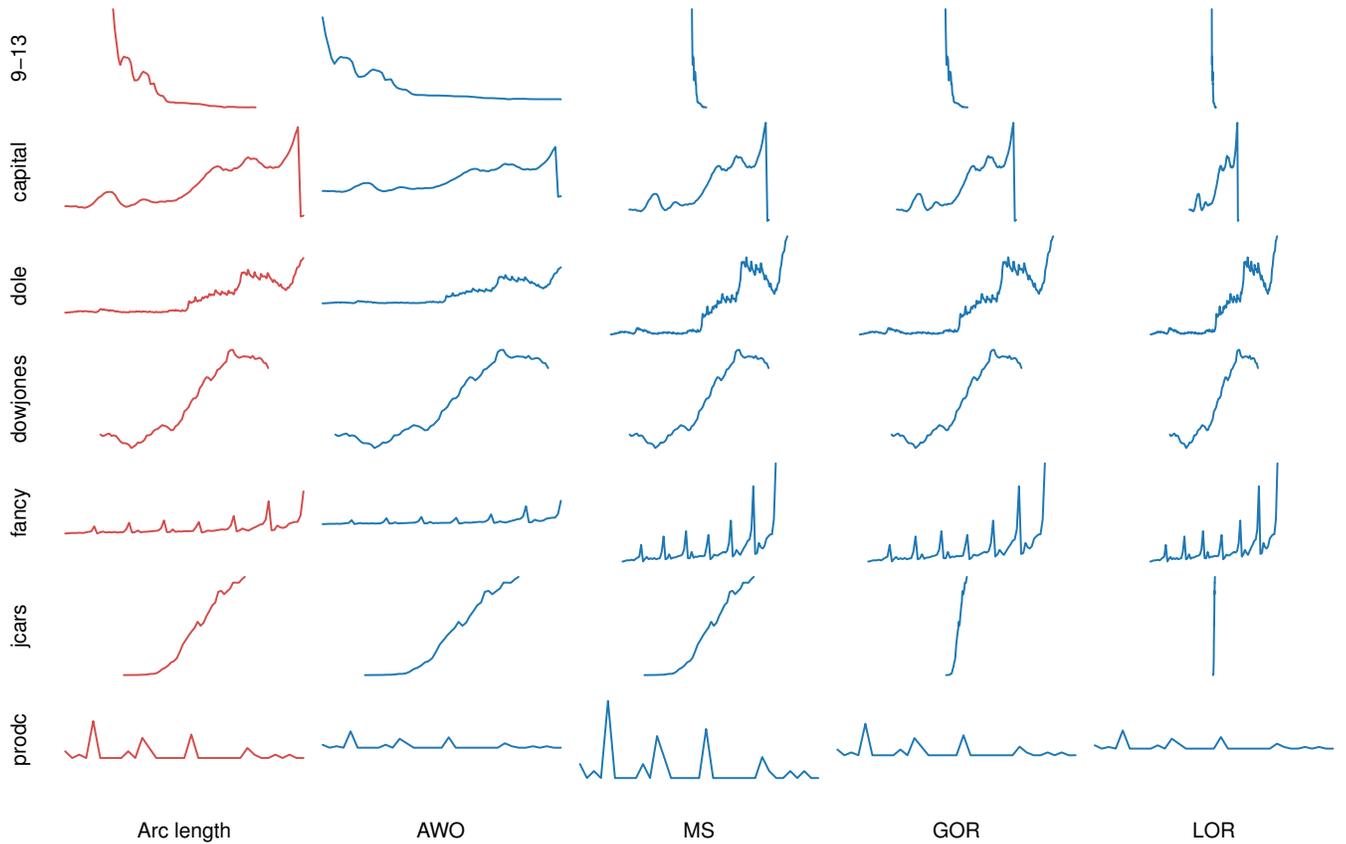


Fig. 5. Select time series from Figure 4 which show some of the largest differences in selected aspect ratios across the recommended methods. As with the algebraic curves, arc length is similar to AWO, but produces generally taller aspect ratios. The other methods continue to produce overly tall aspect ratios in some cases (*9-13* and *jcars*). In the *dole* and *fancy* time series, MS, GOR, and LOR select tall aspect ratios that emphasize the low frequency trend in the data, but obscure the details of the higher frequency patterns. Arc length and AWO do a better job of compromising between the low and high frequency patterns.

along the  $x$ -axis, an inconvenient parameterization for 2D curves. Second, since it makes curves symmetric around  $y = x$  if possible, circles will be banked to circles.

**Contour Plots.** While a number of plot types arise in data visualization featuring general curves in 2D, perhaps the most common is the contour plot. We can use the arc length method to select an aspect ratio for contour plots by simply using the line segments that make up the contours as input to the method.

Figure 6 demonstrates this approach on a number of real data sets. We use R’s `kde2d` [16] grid-based kernel density estimation routine to fit a nonparametric model of the density and then we extract contour lines. This process results in line segments of nearly equal length, again making the comparison to the non-length-weighted methods as fair as possible.

Unlike the simple curve case, where arc length was similar to AWO, in these contour plots, arc length behaves essentially the same as MS. Both methods bank circles to circles producing what we think are the most visually pleasing aspect ratios. AWO is very similar, but produces ellipses even for very simple cases such as the bivariate gaussian. GOR’s behavior is mixed, sometimes producing results similar to AWO, and sometimes producing extremely flat aspect ratios. For space reasons, we have not shown LOR which produces consistently poor aspect ratios on these data sets.

To see how banking contour plots differs from banking curves, Figure 7 shows a real data set of the education level and fertility rate of the French-speaking provinces of Switzerland in about 1888 [12]. In the top plot, we have assumed that  $y$  is a function of  $x$  and fit a LOESS curve to the data. Using the arc length method on the LOESS curve

results in an aspect ratio that emphasizes the negative relationship between fertility and education, and reduces the impact of the outliers. For comparison, in the bottom plot, we assume that the  $x$  and  $y$  variables are independent, and thus, we are interested in seeing their joint density. Using the arc length method on the contour lines of the density estimate, we select an aspect ratio that highlights the roughly bivariate normal distribution of the main cluster of points and visually emphasizes the outliers.

**Other transformations.** If the plotted curve is a function of  $x$ , then aspect ratio changes are the only transformation that generally makes sense. But for general curves in 2D other transformations such as rotations or skews might be useful. This may happen when the axes are not inherently meaningful, as when some form of dimensionality reduction has created arbitrary axes.

Again, this can easily be done using the arc length measure. For example, we can expand the set of transformations we consider in the optimization to include rotations (which are also area preserving) by modifying the optimization criterion:

$$\min_{a, \theta} \sum_{i=1}^N \left\| \frac{R(\Delta x_i, \theta)}{\sqrt{a}}, \sqrt{a} R(\Delta y_i, \theta) \right\|$$

where  $R(x_i, \theta)$  and  $R(y_i, \theta)$  are the  $x$  and  $y$  components of the  $i$ th line segment after rotation by  $\theta$ . Thus, the optimization will be over both  $a$ , the aspect ratio, and,  $\theta$ , the rotation angle, finding the best rotated squeeze mapping transformation.

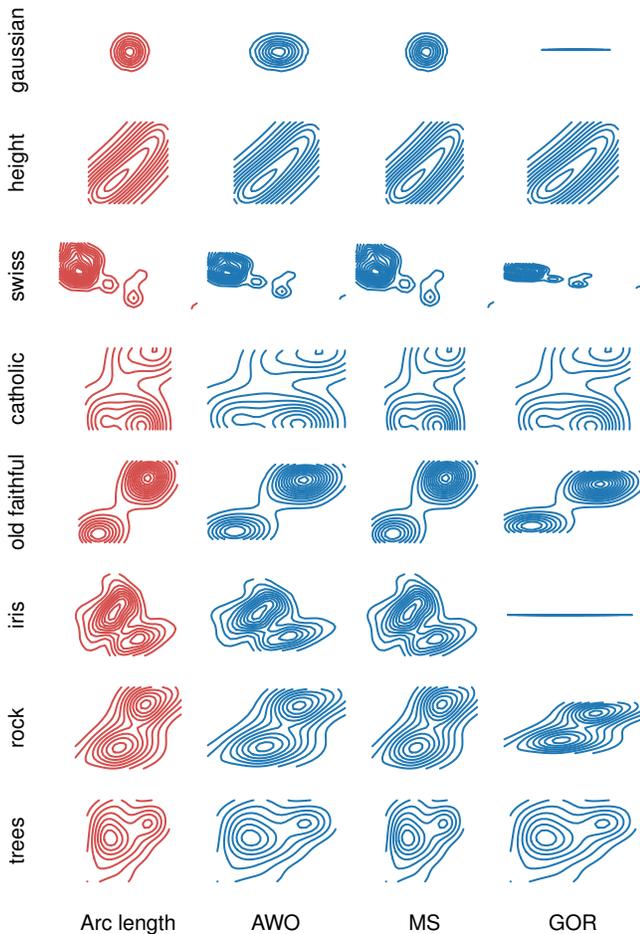


Fig. 6. Banking contour plots using arc length and three of the previously proposed methods. Arc length and MS are roughly the same, nicely banking circles to circles. AWO is similar, but its preference for flatter angles results in ellipses. GOR’s problems with near horizontal or vertical segments cause, in some cases, extremely flat aspect ratios.

### 4.3 Performance

The proposed methods all have relatively simple performance characteristics. MS, AS, and Manhattan arc length (resultant vector) have closed form solutions, which can be computed in  $O(n)$  time, where  $n$  is the number of line segments in the input. AWO, LOR, and arc length can be efficiently optimized using a criterion that is  $O(n)$ . GOR uses an optimization criterion that is  $O(n^2)$ .

We tested the performance of the various banking algorithms using implementations written in R and R’s built-in `optimize` function. We found it impractical to use GOR to bank time series with more than about 250 segments (doing so took more than 1 second), well within the realm of what someone would want to plot (especially for contour data). On the other hand, the  $O(n)$  optimization methods were able to bank 50,000-100,000 segments within 1 second. Arc length was about twice as fast as AWO and LOR since it does not require the evaluation of a transcendental function. The closed form solutions were a further 10-20 times faster.

Given the low performance of an interpreted language such as R, we hesitate to draw overly broad conclusions about performance. Implementations in a compiled language would likely be orders of magnitude faster. Even so, the relative results should be reliable. Both the closed form methods and the  $O(n)$  optimization methods appear to be computationally feasible for interactive applications even in demanding conditions. A bit more care should be taken when choosing to use GOR in order to ensure that it will scale to the desired data size.

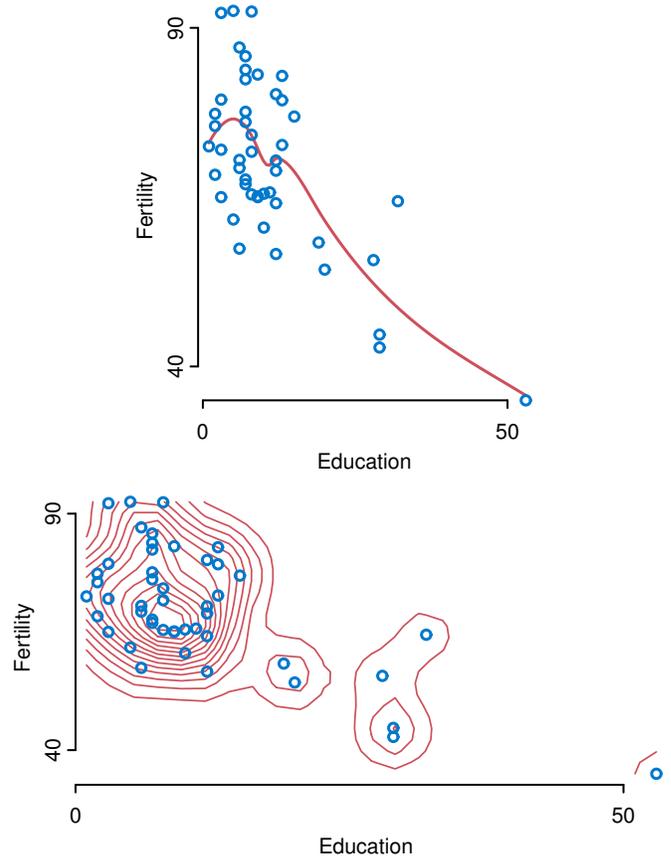


Fig. 7. Banking the same data set with two different dependency assumptions. The top plot is banked using a LOESS curve assuming Fertility is dependent on Education which emphasizes the negative relationship between the two variables. The bottom plot is banked using the contour lines of a 2D density estimation assuming Fertility and Education are independent variables. The resulting aspect ratio highlights the outliers from the main cluster of points.

## 5 DISCUSSION AND FUTURE WORK

We have presented an arc length-based approach to aspect ratio selection which has the following properties:

- It can be derived directly from an equation maximizing orientation resolution, a perceptual criteria for aspect ratio proposed by Cleveland, and it is a parsimonious formulation—a function of length alone, without reference to slope or orientation.
- It is scale invariant, does not need special handling for vertical, horizontal, or colinear segments, and in contrast to GOR, LOR, and MS, it is invariant to the parameterization of the curve.
- For time series-like curves, it produces compromise aspect ratios near the geometric mean of previous methods. In cases where GOR, LOR, and MS fail, arc length and AWO continue to behave well. For contour plots, it performs similarly to MS, banking circles to circles. In contrast, AWO produces ellipses and GOR and LOR can fail.
- It is faster to compute than AWO and substantially more scalable than GOR.

We think these empirical properties of the arc length method recommend it for use in preference to the previous work.

However, we have not presented any *perceptual* reasons to prefer our aspect ratio selection method over others. Developing an appropriate perceptual theory that can address the “correctness” of an aspect ratio selection method remains an open problem. In his work, Cleveland has suggested two possible theoretical foundations for aspect ratio

theory, neither of which appears to be completely satisfactory.

The first, orientation resolution [4], was used in this paper (Section 3.3) to derive the arc length metric. It was also used by Heer and Agrawala to justify their GOR and LOR methods. Cleveland's original banking study [4] showed that orientation resolution was important in slope comparisons between pairs of line segments. More recent perceptual results cited by Loffler [10] support the orientation resolution hypothesis. For example, Heeley and Buchanan-Smith [6] showed that angle judgments depend principally on the angle size (best near 0, 90, 180, and 270 degrees) and are largely rotationally invariant. However, experiments by Snippe and Koenderink [15] appear to show that angle discrimination may depend on orientation. Further, Kennedy, Orbach and Loffler showed that both angle discrimination [8] and angle estimation accuracy [9] are significantly better for angles embedded in an isosceles rather than scalene triangle. These latter results suggest that orientation resolution may be insufficient. Perception of angles is influenced by their embedding within a geometric figure, which is not accounted for by the orientation resolution hypothesis.

The second theoretical foundation proposed by Cleveland is curvature [3], or the change in orientation of the curve per unit of arc length. In the perception literature, Whitaker and McGraw [17] have tied curvature aspect ratio to better visual discrimination across multiple scales. Curvature is also promising due to its wide use in computer-aided geometric design as a criteria in defining "fair" curves [14]. Cleveland [3] showed that banking to 45° maximizes local curvature. Similarly, it is easy to show that minimizing the local arc length maximizes local curvature because the numerator of the curvature formula is invariant to changes in aspect ratio and the denominator is a monotonic transformation of arc length. However, in the process of developing the arc length method, we spent a lot of time attempting to derive an aspect ratio selection approach directly from curvature without success. Our curvature-based methods tended to select extreme aspect ratios which resulted in high curvature at a few locations along the curve and very low curvature everywhere else, similar to the LOR results presented in this paper. The problem appears to be that curvature is a local property of the curve whereas aspect ratio selection requires a global view of the entire curve.

While the arc length method does not immediately suggest a better perceptual foundation for aspect ratio selection, it does provide three interesting directions for future research. First, and perhaps most surprisingly, it demonstrates that effective aspect ratios can be selected without explicit reference to either orientation or angle, suggesting that they may not be strictly necessary elements of a sound aspect ratio theory. Second, the arc length method reveals, somewhat counterintuitively, that good aspect ratios minimize the visual space taken by the data and maximize white space. This property suggests interesting perceptual hypotheses that might lead to more insight into aspect ratio selection. For example, more white space may reduce our perceptual tendency to filter out high frequencies, thus ensuring that we can see high frequency patterns of interest. Or it may be true that minimum length lines lead to quicker, more accurate, graph reading. Finally, the arc length method emphasizes the importance of symmetry and regularity in displayed curves, properties known to trigger hyper-acuity discrimination for shapes [13]. If this acuity applies to curves it could prove an important asset when the visual analysis task depends on recognizing regularity or detecting small deviations from it.

Substantial future work remains to be done to close the gap between perceptual theory and the proposed practical methods for aspect ratio selection. Understanding how the arc length method extends to other plot types, such as 3D plots, may provide insight into this problem. However, well-designed experimental work will likely be necessary to fully understand the perceptual impact of aspect ratio selection.

## ACKNOWLEDGMENTS

The first author thanks Phil Smith; an early conversation with him directly inspired this work. The first author was supported by FODAVA grant 0937123.

## REFERENCES

- [1] A. Buja, C. Hurley, and J. A. McDonald. Elements of a viewing pipeline for data analysis. 1988.
- [2] W. S. Cleveland. *The Elements of Graphing Data*. Hobart Press, Summit, New Jersey, U.S.A., 1985, 1994.
- [3] W. S. Cleveland. A model for studying display methods of statistical graphics. *Journal of Computational and Graphical Statistics*, 2(4):323–343, 1993.
- [4] W. S. Cleveland, M. E. McGill, and R. McGill. The shape parameter of a two-variable graph. *Journal of the American Statistical Association*, 83(402):289–300, 1988.
- [5] S. Guha and W. S. Cleveland. Perceptual, mathematical, and statistical properties of judging functional dependence on visual displays. Technical report, Purdue University Department of Statistics, 2011.
- [6] D. Heeley and H. Buchanan-Smith. Mechanisms specialized for the perception of image geometry. *Vision Research*, 36(22):3607–3627, 1996.
- [7] J. Heer and M. Agrawala. Multi-scale banking to 45 degrees. *IEEE Transactions on Visualization and Computer Graphics*, 12:701–708, 2006.
- [8] G. J. Kennedy, H. S. Orbach, and G. Loffler. Effects of global shape on angle discrimination. *Vision Research*, 46(8-9):1530–1539, 2006.
- [9] G. J. Kennedy, H. S. Orbach, and G. Loffler. Global shape versus local feature: An angle illusion. *Vision Research*, 48(11):1281–1289, 2008.
- [10] G. Loffler. Perception of contours and shapes: Low and intermediate stage mechanisms. *Vision Research*, 48(20):2106–2127, 2008. Vision Research Reviews.
- [11] S. G. Makridakis, S. C. Wheelwright, and R. J. Hyndman. *Forecasting: Methods and Applications*. Wiley, New Jersey, third edition, 1998. ISBN 978-0-471-53233-0.
- [12] F. Mosteller and J. W. Tukey. *Data Analysis and Regression: A Second Course in Statistics*. Addison-Wesley, Reading Mass., 1977.
- [13] D. Regan and S. Hamstra. Shape discrimination and the judgement of perfect symmetry: Dissociation of shape from size. *Vision Research*, 32(10):1845–1864, 1992.
- [14] N. S. Sapidis, editor. *Designing fair curves and surfaces: shape quality in geometric modeling and computer-aided design*. 1994.
- [15] H. P. Snippe and J. J. Koenderink. Discrimination of geometric angle in the fronto-parallel plane. *Spatial Vision*, 8(3):309–228, 1994.
- [16] W. N. Venables and B. D. Ripley. *Modern Applied Statistics with S*. Springer, New York, fourth edition, 2002. ISBN 0-387-95457-0.
- [17] D. Whitaker and P. V. McGraw. Geometric representation of the mechanisms underlying human curvature detection. *Vision Research*, 38(24):3843–3848, 1998.