



Interpolating Splines: Which is the fairest of them all?

Raph Levien¹ and Carlo H. Séquin²

¹Google Inc., raph.levien@gmail.com

²University of California, Berkeley, sequin@cs.berkeley.edu

ABSTRACT

Interpolating splines are a basic primitive for designing planar curves. There is a wide diversity in the literature but no consensus on a “best” spline, or even criteria for preferring one spline over another. For the case of G^2 -continuous splines, we emphasize two properties that can arguably be expected in any definition of “best” and show that any such spline is made from segments cut from a single generator curve, such as the Euler spiral.

Keywords: G^2 -continuity, interpolating splines, two-parameter splines, extensionality, locality, Euler spirals, aesthetic curves.

DOI: 10.3722/cadaps.2009.xxx-yyy

1. INTRODUCTION

A popular technique for designing curves is to fit an interpolating spline through a relatively sparse sequence of data points input by a designer. Designers want to obtain aesthetically pleasing curves even when these data points (serving as control points for the resulting spline) are quite sparse. Other popular techniques for designing planar curves include approximating splines or the manipulation of Bézier control points. However, with these methods the relationship between the input and the resulting curve is less direct.

In this paper we are concerned with the construction of smooth, at least curvature-continuous splines through a given set of sparse data points, assuming that one spline segment will be placed between every two subsequent data points. This design task domain is quite different from the problem arising in reverse engineering, where a smooth spline has to be fit as best possible to a dense but noisy data set, which may have been obtained by some scanning or sampling method.

Fairness, while clearly associated with smoothly varying curvature, is nonetheless a poorly defined concept. Even though most people think they recognize fairness (or the lack thereof) when they see it, there is no agreed-upon mathematical formalism that defines it. Fairness does not simply increase with the degree n of continuity (G^n). It also is not properly captured with the behavior of an ideal elastic strip, as in the minimum-energy spline (MEC), nor by minimizing the variation of curvature, as in the minimum-variation curve (MVC). There is no known functional that completely captures the notion of fairness, nor any agreed-upon way to measure it.

However, as we will discuss in Section 2, there are several properties that most designers agree should hold for truly good, fair, interpolating splines. By focusing on these desirable properties and the trade-offs between them, we can considerably narrow down the space in which we may look for the “ideal” spline. In particular, we will argue that we should focus our attention only on *extensional*, *two-*

parameter splines, and among those in particular on the *log-aesthetic curves* [11]. While studying these classes of curves, we have come to the remarkable discovery that any extensional two-parameter spline can be defined with its own characteristic generator curve, and that all instances of a segment of such a spline can be obtained by cutting a particular section out of this generator curve and scaling it properly to fit between two consecutive data points. Two-parameter splines also allow for a very efficient realization based on a look-up table and a Newton-type solver; thus they can easily provide real-time interactive curves even for the most demanding CAD applications.

2. SOME CLASSICAL SPLINE CURVES

Here we briefly review some important spline curves that play a key role in our discussions.

2.1 Minimum-Energy Curves (MEC)

The MEC minimizes total bending energy. It is the mathematical idealization of a flexible strip that is constrained to go through all control points, through which it can slide freely (i.e. its length is not constrained). Mathematically, the MEC is the curve that minimizes the energy functional:

$$E_{MEC}[\kappa(s)] = \int_0^l \kappa^2 ds$$

This energy functional is the integral over arc-length of the square of curvature. Thus, the MEC spline can also be defined as the curve that minimizes the L2 norm of curvature. As shown by Forsythe and Lee [6], this spline can also be defined as piecewise segments of the *rectangular elastica* between each pair of adjacent control points, with G^2 continuity across control points (that is, the curvature is continuous everywhere). The rectangular elastica (MEC) is shown, along with two other related curves, in Figure 1.




Fig. 1: Comparison of Euler's spiral, the rectangular elastica (MEC), and a simple cubic parabola.

2.2 Minimum-Variation Curves (MVC)

Starting from the premise that a circle is the perfect shape for a closed curve and thus should be represented by a functional that vanishes for that shape, Moreton and Séquin introduced the notion of a minimum variation curve [13], in which the square of the *change* of curvature is minimized rather than curvature itself. The corresponding optimization functional is thus:

$$E_{MVC}[\kappa(s)] = \int_0^l \kappa'^2 ds$$

Unlike the elastic strip representing the MEC, there is no known model for the MVC in nature.

2.3 Euler Spiral Splines

The Euler spiral is most easily defined as the curve where curvature increases linearly with arc length. In an Euler spiral spline, the curvature of each segment between adjacent control points is characterized by: $\kappa = \kappa_0 + s\kappa_1$, and with G^2 -continuity throughout the curve, i.e., subsequent segments are put together so that their curvature values match at the shared control point. While the curve is most correctly attributed to Euler, it is also commonly known as the Cornu spiral or clothoid [7].

2.4 Log-Aesthetic Curves

A very promising family of candidates for generator curves is the *aesthetic curve*, also known as the *log-aesthetic curve* [11]. These are defined as having a straight line when κ'/κ is plotted against κ on a log-log scale. In addition to a number of special cases (including the equiangular spiral and Nielsen's spiral), the general formula for aesthetic curves is $\kappa = s^\alpha$, where the exponent α is the parameter that distinguishes the members of this family. An intuitive motivation for the aesthetic curve is that the eye is not as sensitive to absolute variation of curvature as it is to relative ones. In areas of high curvature, high variation of curvature is more easily tolerated; but in low-curvature areas, even small variations in curvature are visible.

The literature on aesthetic curves doesn't tend to include curves with inflection points—on a log-log plot, these are off the scale. However, since it is well known that the aesthetic curve with $\alpha=1$ is the Euler spiral, curves with inflection points are known to be part of the family. In general, odd integer values for α have odd symmetry and admit inflection points, but it's possible to generalize the formula to $\kappa = \text{signum}(s)|s|^\alpha$ to admit a continuous range of exponents. A family of such curves (used to interpolate a zigzag sequence of control points) is shown in Figure 8 and will be discussed later.

3. DESIRABLE SPLINE PROPERTIES

There is widespread agreement that some key properties should hold for any truly good, interpolating spline. They are discussed here in decreasing order of importance.

3.1 Extensionality

Probably the most important characteristic of an ideal interpolating spline is the property that when additional data points are placed on a given spline, the shape of the curve will not change. If the addition of such points were to make the curve better in any way, perhaps by lowering the value of some defining cost-functional, then, clearly, the original curve could not have been optimal in the first place. Knuth called this property *extensionality* and presented it as one of a set of desirable properties for splines in a 1989 lecture on Mathematical Typography [5, pp. 39-42]. Moreton in his Ph.D. thesis [12, p. 28], and other workers, have called this property *consistency*.




Fig. 2: Extensionality failure of the circle spline

All variationally defined splines are automatically extensional, including the MEC and the MVC. But there also exist extensional splines that are not defined variationally, for instance the Euler spiral spline, first described by Birkhoff and de Boor [1] in 1964, and implemented by Mehlum in the Autokon system in the early 1970's [10].

But many splines in the literature fail to exhibit extensionality. One dramatic example is the circle spline introduced by Séquin, Lee, and Yen [15], in which each spline segment is a blend between two circular arcs, each defined by 3 subsequent data points. As shown in Figure 2, adding a new point at the center of the shown S-shape dramatically changes the curve by adding two new inflection points, because the tangent at this point must aim in the direction of the two adjacent neighbors. The Metafont [3] and Ikarus [4] splines (discussed in Section 4) are two other splines that are also not perfectly extensional. The Metafont spline can be thought of as a cubic approximation to the Euler spiral spline, and it therefore approximates the extensionality property of that spline for small angles.

3.2 Roundness

Another rarely contested aspect related to fairness is *roundness*, the property that points placed on a circle yield an exact circle as the resulting spline. It seems intuitively obvious that a circle is the fairest possible curve in such situations, but many splines do not have this property. In particular, the MEC spline deviates rather significantly from roundness, as shown in Figure 3. By contrast, the Euler spiral spline, the MVC, and the circle spline readily form a circle in this situation.




Fig. 3: Roundness failure of the minimum energy curve (MEC)

The classical cubic spline interpolant also fails the roundness criterion, as it is impossible to form an exact circular arc in terms of polynomial $x(t)$ and $y(t)$ parameterized functions. In addition, raw cubic splines are very sensitive to spacing of the control points, techniques are often employed to compensate. One common technique (used in Ikarus [4]) is to normalize the parameterization by the chord length. However, this adjustment destroys the extensionality of the spline.

3.3 Monotone Curvature

The human visual system is particularly sensitive to minima and maxima of curvature. This property implies that all curvature extrema should coincide with the given control points, and that the spline segments in between should exhibit monotone curvature. A fair amount of literature is concerned with constraining existing curves (such as cubic polynomials) to have monotone curvature. The Euler spiral, in which curvature varies linearly with arc length, obviously has this property. All instances in the aesthetic curve family also have monotone curvature. However, the MEC does not have this property; for symmetrical end conditions it has a tendency to increase arc length and to add a curvature maximum in the middle of the segment. The MVC also in general does not place curvature extrema to coincide with the control points.

3.4 High-Order Continuity Versus Locality

One might assume that the higher the degree of G^n continuity a curve possesses, the “fairer” that curve would look. Discontinuities of curvature are quite visible, and thus we demand unconditionally that a fair spline be at least G^2 -continuous. G^3 -discontinuities can be seen as G^2 -discontinuities in reflections on shiny surfaces, but they are not really visible in drawn plane curves. Thus from the point of view of visual aesthetics, G^2 -continuity is good enough. This is good to know, because higher-order G^n -continuity comes at a price: It typically decreases *locality*!

Locality is the property that defines how far away along the spline the effect of a small local change in the position of one data point can still be seen. In general, with interpolating splines there is a tradeoff between locality and higher-order G^n -continuity. For example, the MVC spline has G^4 -continuity, but has a slower fall-off of local perturbations than the MEC or the Euler spiral spline. This is illustrated in Figure 4, which shows the Euler spiral spline on the left (same locality properties as MEC) and an analogous G^4 -spline (same locality properties as MVC) on the right. The ripples in the G^4 curve (Fig.4 bottom right) die out more slowly. As an extreme example in the direction of maximizing locality, the degenerate “rubber-band spline” connecting subsequent control points with straight line segments, has ideal locality but worst-case fairness.



Fig. 4: Curvature behavior and resulting shapes for a G^2 -spline (left) and a G^4 -spline (right).

In general, fair, extensional, interpolating splines require perturbations to ripple out indefinitely, but with exponential die-off. Other splines (such as the circle spline) can bound the effect of perturbations to a finite number of control points, but give up fairness or extensionality or both. Note that several other approaches to interactive curve design, such as some approximating splines, have excellent locality and continuity properties, but without the direct geometrical control offered by interpolating splines. Thus, in order to get the best possible locality from interpolating splines, we should limit ourselves to G^2 -continuity.

4. TWO-PARAMETER SPLINES

A two-parameter spline is defined as one in which every interpolating curve segment between two adjacent control points is drawn from a family with a two-dimensional parameter space, modulo the transformation needed to make the endpoints of the curve coincide with the two control points. This transformation consists of scaling, rotation, translation, and mirroring; yet its parameters are not counted, as they are uniquely determined by the context of the curve segment within the spline.

One example of a two-parameter spline is the MEC spline. That means that the shape of the curve above a given chord is completely determined by the two tangent directions of the curve at either end of the chord. It is not immediately obvious from the MEC definition that the spline has the two-parameter property, but careful analysis (such as the study by Forsythe and Lee [6]) shows that it is so. Another example of a two-parameter spline is the Euler spiral spline; in this case the two parameters that control the shape of a curve segment can again be the two tangents, or more conveniently in our context, the curvature values at the two ends: κ_0 and κ_1 .

Besides the MEC and Euler spiral, a number of other interpolating splines fit into the two parameter model. The spline developed for the Ikarus font design system [4], the spline used in Metafont [3], and

the G^1 circle spline of Séquin, Lee, and Yen [15] also use curve segments drawn from a two-parameter space. Finally, while technically not an interpolating spline, the Potrace automated curve-tracing system [14] also uses a two-parameter family of primitive curves.

On the other hand, the MVC is an example of splines that do not belong to this two-parameter class. They require a four-parameter domain for their full characterization, e.g., a segment could be characterized by its tangent angles, as well as its curvature values, at both ends.

5. GENERATOR CURVES FOR TWO-PARAMETER SPLINES

The central, and rather surprising result of a forthcoming Ph.D. thesis [8] is this: For any extensional two-parameter spline there exists a generator curve so that all curve segments between adjacent control points can be cut from that generator curve, subject to scaling, rotation, translation, and mirror image transformations to make the endpoints of the cut segment align with the endpoints in the spline. Thus, any spline in this family can be defined simply in terms of the generator curve. In some ways, this definition circles back to one of the earliest formulations of nonlinear splines. Birkhoff and de Boor wrote in 1964, referring to two approximations to MEC splines, “After looking carefully into the relevant equations, one realizes that the schemes of [7] and [2] do not approximate mechanical splines more closely than other curves—nor does it seem particularly desirable to have them do so. For example, they approximate equally well to Hermite interpolation by segments of Euler’s spirals, joined together with continuous curvature.” [1, p. 171]. Here we propose that the Euler spiral is but one of a family of desirable curves for constructing extensional interpolating splines.

Theorem 1: *In an extensional, two parameter spline, there exists a curve such that for any spline segment (parameterized by angles θ_0 and θ_1 there exist two points on the curve (s_0 and s_1) such that the segment of the curve between them, when transformed by rotation, scaling and translation so that the endpoints coincide, also coincides along the length of the segment.*




Fig. 5: Incremental construction of the generator curve of a two-parameter spline

A typical section of curve is shown in Figure 5. For some endpoint angles, the generated curve segment is the section from s_0 to s_1 . Because of the extensionality, for any s_1 along this curve, the subdivided curve from s_0 to s_1 (which is the result of applying θ_0 and θ_1 endpoint constraints) must be consistent with the original segment. Modulo scaling, the curvature κ and its derivatives κ' and κ'' must of course also match between these two segments.

The proof for this theorem revolves around the quantity κ'/κ^2 , which is also the limit of $6(\theta_1 - \theta_0)/(\theta_0 + \theta_1)^2$ as θ_0 and θ_1 approach zero, i.e. as the length of the curve segment becomes infinitesimal. This quantity represents the amount of curvature variation for a segment of a set amount of curvature - it is invariant to similarity transformations including uniform scaling. As such, for a given generator curve, it uniquely identifies a point on the curve (modulo periodicity, if the curve is periodic as opposed to monotone in curvature).

We conjecture that for any two-parameter spline that is also extensional, the *second* derivative of curvature, suitably normalized for scale invariance, κ''/κ^3 , is uniquely determined given a particular value of κ'/κ^2 . If there two curve segments with different values for this second derivative, then it is

possible to cut subsegments from each with identical θ_0 and θ_1 values. However, value of κ'/κ^3 remains invariant even after such cutting, so there are two different values for this. Yet, the two parameter property states that there exists only a single curve for a given θ_0, θ_1 pair, - a contradiction!

Thus, there is a process for constructing the generator curve. Start with arbitrary κ and κ' values sampled from some point on the curve, then continually derive κ'' based on the lookup process above: find a curve segment that has the correct value for κ'/κ^2 , then sample κ''/κ^3 at that point. Continually integrate (in the sense of an ordinary differential equation), and the result is a curve.

The final component of the argument is that for any given θ_0, θ_1 pair, there exists a segment cut from this curve such that it is equal to the spline segment. Determine κ'/κ^2 for one side (s_0) of the spline segment, and find that point on the generator curve. Continue extending the curve until the other side is reached. Since at each point κ''/κ^3 must be equal between the two curves, the curves themselves must also be equal.

Both the two-parameter and extensionality conditions are necessary for the generator curve property to hold, and there are counterexamples when either is not present. The Ikarus and Metafont splines are two-parameter curves, but not extensional, and they do not have a single generator curve. The Minimum Variation Curve is extensional but does not fit into a two-parameter space, and therefore also does not have a single generator curve.

5.1 Constructing New Generator Curves

Above we have demonstrated how a given dependency between κ and κ' in an extensional two-parameter spline can lead to the construction of a corresponding generator curve. Now we turn the question around: Under what circumstances can an arbitrary curve serve as a generator for an extensional two-parameter spline? It is critical to choose the right kind of curve, since the properties of the generator curve directly influence the properties of the resulting spline.

Inflection points: To start with, not all curves, including many spirals, contain an inflection point. Such curves can be used as the generator curve for a spline, but the solver will fail to find a solution if the target curve contains inflection points. Thus the generator curve should contain an inflection point with opposite signs of curvature on either side. Most single spirals, including the equiangular spiral, the Nielsen's spiral, and so on, are thus not broadly suitable as generator curves. However, with the previously mentioned tweak to the signs, so that the curve is defined as $\kappa = \text{signum}(s)|s|^\alpha$, all curves in the aesthetic curve family admit inflection points.

Roundness: To be able to generate circular spline segments, the generator curve must also contain a truly circular region, or asymptotically approach one. For instance the rectangular elastica does not have such a region, and the MEC spline therefore deviates rather significantly from roundness, as was shown in Figure 3. The Euler spiral spline, on the other hand, readily forms a circle. This is because the curvature variation within the Euler spiral tends to zero in the limit as arc-length tends towards infinity. In particular, for a spiral to approach a circle asymptotically, the curvature variation for a segment scaled to unit curvature must approach zero. Since curvature is proportional to the inverse of scale, and curvature variation to the inverse square, this property is mathematically implied by:

$$\lim_{s \rightarrow \infty} \frac{\kappa'}{\kappa^2} = 0$$

This property holds for nearly any well-behaved curve with monotone curvature, including the aesthetic curves with the implicit equation $\kappa = s^\alpha$ for any positive exponent α .

Locality: In general, the weaker the curvature variation near the inflection point of the generator curve, the better the locality of the resulting spline. The aesthetic curve family varies from very steep to very shallow curvature variation in this region across its parameter space. For this concrete family of curves, we can accurately quantify this relationship. The falloff factor, plotted against the exponent

of the aesthetic curve, is shown in Figure 6. For $\alpha=1$, which represents the Euler spiral, the amplitude of the ripple diminishes by 0.268 from one crest to the next one. This ratio can be determined analytically by using the small-angle approximation and solving for $\kappa_0\theta_1 = \kappa_1\theta_0$ for the endpoints of a segment, leading to a value of $1/(2+\sqrt{3})$ for the Euler spiral. Note also that the MEC spline has the same locality properties as the Euler spiral spline. For the preferred exponent of 1.78 (see Section 7), the falloff factor is 0.173, noticeably smaller and thus better. It is possible to achieve even better locality at the expense of higher curvature variation, if that is desired.




Fig. 6: Perturbation fall-off rate as a function of the exponent in the aesthetic curve

6. PRACTICAL IMPLEMENTATION

Having an optimal interpolating spline would be of primarily academic interest, if it were expensive or impractical to compute. One reason variationally defined splines are not popular is that numerical solvers tend to be slow. For example, Moreton's conjugate gradient solver for MEC and MVC curves took 18 and 75 seconds (600 million and 2.5 billion CPU cycles), respectively, to compute the optimal solution to a 7 point spline [12, p. 99].

Fortunately, the general extensional two-parameter spline admits an extremely fast and simple implementation. Since the curve segments are defined by two parameters, for any given generator curve, it is practical to precompute a lookup table indexed by the tangent angles at the endpoints. The lookup table stores the curvatures at the endpoints (used for enforcing curvature continuity across control points), as well as the position and tangent angle of a midpoint, used for subdivision while rendering the curve. With this precomputation done only once, drawing the construction algorithm is only slightly more complex than the standard de Casteljau subdivision for rendering Bezier curves; the only difference is the use of a lookup table rather than an algebraic formula to determine the midpoint.

An example curvature lookup table, for the Euler spiral, is shown in Figure 7. The horizontal axis is the tangent angle, in radians, on the left side, the vertical axis is the tangent angle on the right, and the value is the resulting curvature on the left side; the other value follows by symmetry. Moreover, the table itself is anti-symmetric, which can further reduce storage requirements. Note that the relationship is nearly linear for small angles. The slope of that linear relationship is closely related to locality, as described in Section 5.1. Because these values are so smoothly behaved, a simple

representation as a two-dimensional lookup table with 32x32 entries, and using bicubic interpolation for lookups, gives a maximum curvature error of 10^{-4} across the range of the function. This error scales with the fourth power of the linear size of the table; thus, for example, if higher precision is desired, a 64x64 table could reduce the error to about 6×10^{-6} .




Fig. 7: Contour plot of curvature lookup table for Euler spiral

The global solver uses a Newton-style approach to enforce the curvature continuity constraints across all control points. The partial derivatives of curvature deltas with respect to tangent angles form a tri-diagonal matrix, which can easily be solved in linear time. Each iteration of the global solver does essentially the same as the matrix solver in the Metafont spline. Newton iteration converges extremely quickly; for most configurations of control points; three or four iterations are sufficient to attain high precision. The overall algorithm is efficient enough to offer excellent interactive performance even when implemented in JavaScript for execution in a standard web browser. The numerical code is approximately 230 lines of JavaScript, and a set of 32x32 lookup tables requires an additional 200k of uncompressed data. No doubt, careful tuning could reduce the data requirement further.

7. USER STUDY AND RESULTS

In order to find out where designers and graphic artists place their preferences in the trade-off between minimizing curvature variation and locality, a quick on-line user study was conducted. A survey was presented in a thread on the Typophiles web forum, at <http://typophile.com/node/52821>. The users were shown the family of such curves displayed in Figure 8, which all interpolate the same zigzag sequence of data points. Among these, the minimal bending energy is obtained with the exponent α near 0.9; thus the Euler spiral, with the exponent $\alpha=1$ also yields very low bending energy. The on-line study then asked respondents to select the “fairest” curve from the choices shown in Figure 8. The results are plotted in Figure 9. We can see that the mean preferred exponent is 1.78; this is a curve for which the bending energy is 6% higher than the Euler spiral. These results confirm the belief that the concept of fairness is rather fuzzy, and that it varies significantly among different designers and graphic artists. Nonetheless, the preference for an exponent in the range 1.5 to 2 is marked.




Fig. 8: Family of aesthetic curves with $\kappa = s^\alpha$ [α : 0.5 - 4.0]




Fig. 9: Results of survey on aesthetic curve preference

An example of a log-aesthetic spline in a practical typography application is shown in black in Figure 10. It has been overlaid with an Euler spiral spline shown in blue. The black curve is more naturally suited to the typographic application than the blue one, as can be seen primarily in areas where the control points are very sparse. The log-aesthetic spline transitions more quickly from regions of low to high curvature, which is a frequent requirement in designing letterforms. By contrast, the Euler spiral spline tends to be more “loopy.” Certainly it could be used to describe the desired shape for this letterform, but more control points would be needed.

On the other hand, we have observed that the aesthetic curve spline is somewhat less stable than the Euler spiral; small perturbations to the input can sometimes result in quite noticeable changes to the curve, as the position of an unconstrained inflection point may move significantly in response to a small movement of one of the control points. Perhaps this lesser stability is part of the price to be paid for the improved fairness and locality properties.




Fig. 10: Aesthetic spline (black) and Euler spiral spline (blue) overlaid

8. SUMMARY AND CONCLUSIONS

Interpolating splines are a useful and important primitive for designing curves. However, choosing one from the multiplicity in the literature is not easy. There is no universal consensus on what makes the best spline, or even how fairness should be measured. We have argued that any spline to be considered a “best” candidate must satisfy several properties. Foremost it must be extensional. Further, in our view of the trade-offs between higher-order continuity versus improve locality, we believe that the “best” spline should belong to the domain of G^2 -continuous two-parameter curves. Such splines can then be defined simply in terms of a generator curve, and the properties of this curve map straightforwardly to the corresponding spline.

Based on evidence from existing splines in the literature, two parameter splines have in general better locality than splines defined from a larger parameter space. It is straightforward to achieve G^2 continuity with two parameters, and no higher degree of continuity is required for the perception of fairness in drawn curves. Thus, it is reasonable to expect that any “best” spline would belong to the domain of two-parameter curves. These curves can then be represented succinctly by a representative generator curve, and it is only necessary to pre-calculate a 2-dimensional look-up table that gives two curvature values for any pair of tangent angles. It is a reasonable bet that a high-efficiency implementation will make these near-optimal curves much more popular and more widely used.

We have thoroughly characterized the design space for interpolating splines, and it seems likely that our spline is in fact nearly ideal for two-dimensional curve design tasks, although not quite provably so. There are only a few places to look for promising candidates for an even better spline. First, within the two-parameter framework, it is conceivable that a better generating curve could be found, but we have no inkling where one would have to look for such a curve. Second, while the evidence strongly suggests that two-parameter splines have better locality than curves from a higher-dimensional parameter space, there is no solid proof for this relationship, and there might be some ingenious way of reconciling more parameters, and thus higher-order continuity, with very good locality.

A most promising family for the generator curve is the log-aesthetic curve, defined by the implicit equation $\kappa = s^c$; its properties match almost exactly to most of the desirable properties expected in a “fair” spline. We have quantified a tradeoff between fairness and locality, which is covered by the range of exponents of the aesthetic curve. The Euler spiral spline, the best interpolating spline known previously, is just one point along this spectrum. User testing showed that a moderately higher value for the exponent actually improves both perceived “fairness” as well as locality. The aesthetic curve with an exponent near 1.78 is thus a superior spline, and our analysis of the design space strongly suggests that it is among the best possible options. We may finally have a plausible answer to the question of which interpolating spline is “the fairest of them all.”

REFERENCES

- [1] Birkhoff, G.; de Boor, C. R.: Piecewise polynomial interpolation and approximation. Proc. General Motors Symp. 1964, pages 164–190, 1965.
- [2] Fowler, A. H.; Wilson, C. W.: Cubic spline, a curve fitting routine. Technical Report Y-1400, Oak Ridge Natl. Laboratory, September 1962.
- [3] Hobby, J. D.: Smooth, easy to compute interpolating splines. Technical Report STAN-CS-85-1047, Stanford University, 1985.
- [4] Karow, P.: Digital Formats for Typefaces. URW Verlag, April 1987.
- [5] Knuth, D. E.: Digital typography, volume 78 of CSLI lecture notes. CSLI Publications, Stanford, Calif., 1999.
- [6] Lee, E. H; Forsythe, G. E.: Variational study of nonlinear spline curves. SIAM Rev., 15(1):120-133, 1973.
- [7] Levien, R. L: The Euler spiral: a mathematical history. Tech Report EECS-2008-11, UC Berkeley, 2008.
- [8] Levien, R. L.: From spiral to spline; optimal techniques for interactive curve design. Ph. D. Thesis, UC Berkeley, 2009 (forthcoming).
- [9] MacLaren, D. H.: Formulas for fitting a spline curve through a set of points. Technical Report 2, Boeing Appl. Math. Report, 1958.
- [10] Mehlum, E.: Nonlinear splines. Computer Aided Geometric Design, pages 173–205, 1974.
- [11] Miura, K. T; Sone, J.; Yamashita, A.; Kaneko, T.: Derivation of a general formula of aesthetic curves. In In: Proceedings of the Eighth International Conference on Humans and Computers, pages 166-171, 2005.
- [12] Moreton, H. P.: Minimum Curvature Variation Curves, Networks and Surfaces for Fair Free-Form Shape Design. PhD thesis, University of California, Berkeley, Berkeley, California, 1992.
- [13] Moreton, H. P.; Séquin, C. H.: "Functional Optimization for Fair Surface Design," Proc. ACM SIGGRAPH'92, Chicago, July 1992, and Computer Graphics, Vol 26, No 2, pp 167-176, (1992).
- [14] Selinger P.: Potrace: a polygon-based tracing algorithm, 2003.
- [15] Séquin, C. H.; Lee, K.; Yen, J.: Fair, G2- and C2-continuous circle splines for the interpolation of sparse data points. Computer-Aided Design, 37(2):201-211, 2005.