Smoothly Enrobing Circles

An Implicit-Field Construction for Morphological Bridges

Abstract

We describe a method for producing a single smooth outline that "enrobes" a set of circles of arbitrary radii: a coat that follows each circle's boundary at a fixed offset, joined by smooth bridges wherever two circles are marked as connected. The construction uses signed distance fields (SDFs), a hard union via min, and a smoothed union via a quadratic smooth-maximum to round the inward cusps that a naive Boolean union would leave behind. We derive the bridge geometry in closed form, state the pinch-off condition analytically, and document three implementation defects discovered empirically during development, together with their fixes.

Contents

1 Problem statement

Given n circles $C_i = (c_i, r_i)$ with centre $c_i \in \mathbb{R}^2$ and radius $r_i > 0$, and a set of *bridges* $B \subseteq \{1, \ldots, n\}^2$ specifying which pairs should be visually connected, we want a single closed curve Γ such that:

(i) Γ lies at a constant offset $m > 0$ outside every circle not currently bridged to anything (Γ is the coat boundary);

(ii) for every bridged pair $(i, j) \in B$, Γ smoothly connects the two coats with no holes between them;

(iii) Γ has no sharp corners (it is at least C^1 everywhere).

The approach: define a scalar field $f : \mathbb{R}^2 \to \mathbb{R}$ that is negative inside the desired enrobed region and positive outside, then take Γ to be the zero level set $f^{-1}(0)$, extracted numerically by marching squares.

2 Signed distance fields and the coat

For a single circle $C = (c, r)$, the signed distance function is

$$\phi_C(p) = \|p - c\| - r, \tag{1}$$

negative inside the circle, zero on the boundary, positive outside. The *coat* of C at margin m is the circle of radius $r + m$ sharing the same centre; its SDF is simply $\phi_C(p) - m$.

For n circles with no bridges, define

$$f_0(p) = \min_{1 \le i \le n} \big(\|p - c_i\| - (r_i + m)\big). \tag{2}$$

Since the min of several SDFs is negative wherever *any* one of them is negative, $f_0^{-1}(0)$ is the union of n disjoint coat circles, as required when nothing is bridged.

3 Smooth minimum and maximum

A hard min or max between two SDFs produces a C^0 but not C^1 join: where the two surfaces cross, the resulting field has a kink, which marching squares will render as a visible corner. We replace both with quadratic, C^1-continuous variants parameterised by a blend width k:

$$\operatorname{smin}(a, b, k) = \min(a, b) - \tfrac{1}{4} k\, h(a, b, k)^2, \tag{3}$$

$$\operatorname{smax}(a, b, k) = \max(a, b) + \tfrac{1}{4} k\, h(a, b, k)^2, \tag{4}$$

$$h(a, b, k) = \max\Big(1 - \tfrac{|a-b|}{k},\, 0\Big). \tag{5}$$

The correction term is supported only where $|a - b| < k$ (i.e. near the crossing), and vanishes smoothly at $|a - b| = k$, giving a parabolic fillet of width $\approx k$ centred on the crossing. As $k \to 0$, $\operatorname{smin}, \operatorname{smax} \to \min, \max$, recovering the hard Boolean case. This is the standard quadratic smooth-min used in implicit modelling (Inigo Quilez's formulation); we use it identically for smax by symmetry ($\operatorname{smax}(a, b, k) = -\operatorname{smin}(-a, -b, k)$).

4 Bridging two circles: the phantom-circle construction

Connecting two coats with a smooth bridge — without simply taking their convex hull, which would produce straight tangent lines and sharp corners at the tangent points — is done by introducing an auxiliary *phantom circle* of radius ρ that is externally tangent to both coats, and smooth-maximum–blending its (negated, and hence carving) SDF into the field. The phantom is invisible in the final render; it exists purely to define where the bridge's inward curve comes from.

4.1 Locating the phantom centre

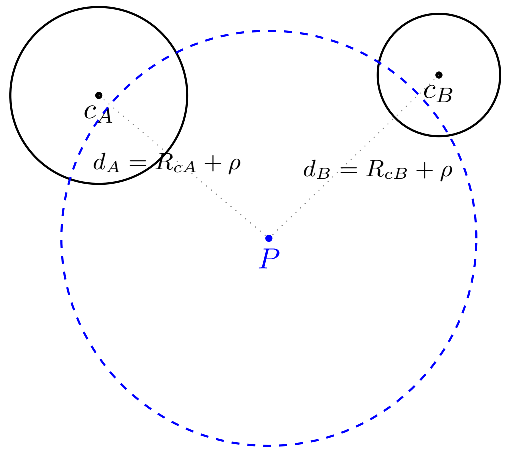

Figure 1: The phantom circle P (dashed, blue) is externally tangent to both coat circles. Its negated SDF, blended in with smax, carves the inward bridge curve where it overlaps the hard union of the two coats.

Write $R_{cA} = r_A + m$, $R_{cB} = r_B + m$ for the two coat radii, and let $D = \|c_B - c_A\|$. External tangency to both coats means the phantom centre P must satisfy

$$\|P - c_A\| = R_{cA} + \rho =: d_A, \qquad \|P - c_B\| = R_{cB} + \rho =: d_B. \tag{6}$$

This is the classical circle-circle intersection problem: P lies on both a circle of radius d_A centred at c_A and a circle of radius d_B centred at c_B. Writing $\hat{u} = (c_B - c_A)/D$ for the unit vector from A to B and $\hat{n} = \hat{u}^{\perp}$ for its perpendicular, the standard solution is

$$a = \frac{d_A^2 - d_B^2 + D^2}{2D}, \tag{7}$$

$$h = \sqrt{d_A^2 - a^2}, \tag{8}$$

$$P = c_A + a\hat{u} \pm h\hat{n}. \tag{9}$$

Equation (7) is obtained by subtracting the two equations $\|P - c_A\|^2 = d_A^2$ and $\|P - c_B\|^2 = d_B^2$ (expanded along the $\hat{u}$ axis) to eliminate the quadratic term; (8) then follows from Pythagoras on the right triangle with hypotenuse d_A and base a. The $\pm$ gives the two phantom positions symmetric about the line $c_A c_B$ — see §6 for how the implementation chooses between them.

4.2 Carving the bridge

The phantom's contribution to the field is

$$f \;\leftarrow\; \operatorname{smax}\big(f,\; \rho - \|p - P\|,\; s\big), \tag{10}$$

where s is the cusp-smoothing width (a separate parameter from the tube/coat construction below). The term $\rho - \|p - P\|$ is positive inside the phantom disk and negative outside; taking smax with the existing field f means: wherever the phantom disk overlaps the previously-negative (enrobed) region, the field is pulled toward zero from inside, i.e. that overlap is excluded from the shape. The result is a concave bite taken out of the union exactly along the phantom's boundary, blended smoothly into the existing coat boundaries at the two tangent points by construction (since the phantom is exactly tangent there, $\rho - \|p - P\| = 0 =$ the coat SDF at the tangent point, so the two surfaces meet with equal value and the smooth-max blend has nothing to discontinuously jump across).

4.3 Pinch-off

Equation (8) requires $h^2 = d_A^2 - a^2 \geq 0$. Two failure modes:

Circles too far apart. If $D \geq d_A + d_B$, the two tangency circles (radius d_A about c_A, radius d_B about c_B) no longer intersect — there is no phantom of radius ρ tangent to both, geometrically because ρ is too small to reach across the gap.

One coat swallows the other. If $D \leq |d_A - d_B|$, one tangency circle is entirely inside the other.

We define the *pinch-off margin*

$$\mu(\rho) = \min\big((d_A + d_B) - D,\ D - |d_A - d_B|\big), \tag{11}$$

positive exactly when a valid phantom exists. $\mu(\rho) \to 0^+$ is the threshold at which the bridge visually pinches to a point before severing entirely as ρ decreases further (equivalently, as D grows for fixed ρ). This quantity is checked numerically before the expensive field evaluation, so a bad parameter choice fails fast with a diagnostic rather than silently producing a degenerate or empty contour.

5 Bridging the gap: the tube field

The phantom-circle carve in (10) only *rounds off* an existing union; it does not by itself make the field negative in the gap between two circles that don't already overlap. A second term is needed: a "tube" along the segment $c_A c_B$ that is negative wherever the gap should read as inside the shape.

5.1 Naive tube (defective)

A first attempt used a single tube radius for the whole bridge,

$$f_{\text{tube}}(p) = \text{dist}(p, \overline{c_A c_B}) - R, \qquad R \in \{\max(R_{cA}, R_{cB}),\ \min(R_{cA}, R_{cB})\}, \tag{12}$$

where $\text{dist}(p, \overline{c_A c_B})$ is the point-to-segment distance. Both choices of R fail: taking $R = \max(R_{cA}, R_{cB})$ pushes the tube's own zero-contour *outside* the smaller circle's true coat radius near that circle's end of the bridge (the tube dominates the union there, since $\min(R_{cA}, R_{cB}) < R$), causing the rendered outline to balloon several pixels past where the smaller circle's coat should be. Taking $R = \min(R_{cA}, R_{cB})$ instead under-shoots at the *larger* circle's end for the same reason in reverse.

5.2 Tapered tube (fix)

The correct construction lets the tube radius vary linearly along the segment, matching each circle's own coat radius exactly at its own end. Parameterise the segment by $t \in [0, 1]$, $t = 0$ at c_A and

$t = 1$ at c_B:

$$t(p) = \mathrm{clip}\left(\frac{(p - c_A)\cdot(c_B - c_A)}{\|c_B - c_A\|^2},\ 0,\ 1\right), \tag{13}$$

$$R(p) = R_{cA}\,(1 - t(p)) + R_{cB}\,t(p), \tag{14}$$

$$f_{\mathrm{tube}}(p) = \mathrm{dist}(p, \overline{c_A c_B}) - R(p). \tag{15}$$

At $t = 0$, $R(p) = R_{cA}$ exactly, so the tube's zero-contour coincides with circle A's coat boundary there; symmetrically at $t = 1$. In between, $R(p)$ interpolates linearly, giving a frustum-shaped tube rather than a cylinder. This is folded into the base union as before, $f \leftarrow \min(f, f_{\mathrm{tube}})$, once per bridge before the phantom-carve pass.

6 Three or more circles: choosing a side

For exactly two circles, (9)'s $\pm$ is resolved by using *both* signs — both phantom circles are blended in, giving the familiar peanut/dumbbell shape with a bridge arc on each side of the line $c_A c_B$.

For $n \geq 3$ circles this is no longer safe: one of the two phantom choices may lie on the side of $\overline{c_A c_B}$ facing the *interior* of the cluster, and carving there would bite into circles that are not part of this particular bridge. The fix is a half-plane test: for bridge (i, j), compute the signed perpendicular offset of every *other* circle's centre from the line $\overline{c_i c_j}$,

$$\sigma_k = (c_k - \tfrac{c_i + c_j}{2}) \cdot \hat{n}, \qquad k \neq i, j, \tag{16}$$

where $\hat{n}$ is the unit normal to $\overline{c_i c_j}$, average them, $\bar{\sigma} = \frac{1}{n-2}\sum_k \sigma_k$, and choose the phantom candidate on the *opposite* side from $\bar{\sigma}$. This reliably selects the outward-facing phantom for convex-ish clusters (in particular it is exact for the triangular case verified empirically in §8), since the average position of the remaining circles is a reasonable proxy for "the interior of the cluster."

7 Contour extraction and artefact loops

The zero level set is extracted by marching squares on a sampled grid ($n_x \times n_y$ samples over the bounding box of all circles, padded by a margin). For two circles this reliably returns a single connected loop. For three or more circles with multiple bridges, the tube fields from different bridge pairs can cross near the cluster's centroid, producing a second, small, disconnected zero-crossing distinct from the true outline — visually a tiny spurious loop in the middle of the cluster (Figure 2).

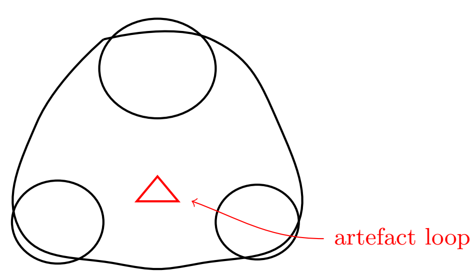

Figure 2: Three mutually-bridged circles: the real outline (black) and a small spurious artefact loop (red) near the centroid, where the three bridge tubes overlap.

Since the true outline is always the longer of the two contours in every case observed, the implementation discards every loop except the one with the most sample points,

$$\Gamma = \arg \max_{\gamma \in \{\text{contour loops}\}} |\gamma|. \tag{17}$$

This is a pragmatic rather than provably-general fix; it is correct whenever the artefact is small relative to the main outline, which held in every configuration tested.

8 Worked example: three mutually bridged circles

Take $C_0 = ((0, 150), 70)$, $C_1 = ((-160, -60), 55)$, $C_2 = ((160, -60), 50)$, all pairs bridged at $\rho = 200$, $m = 14$. The half-plane test for bridge $(0, 1)$ checks circle 2's position relative to the line through c_0, c_1: numerically, c_2 lies on the $\sigma > 0$ side, so the implementation selects the phantom candidate on the $\sigma < 0$ side — placing it up and to the *left*, outside the triangle, rather than the alternative candidate which would sit inside the triangle near c_2. The same logic applied to the other two bridges places their phantoms below-right and below-left respectively, so all three carves bow outward and the resulting outline is the rounded-triangle shape of Figure 2, not a self-intersecting or inward-collapsing curve.

9 Summary of parameters

Symbol	Name	Effect
m	coat margin	Distance the outline sits outside each circle's true radius.
ρ	phantom radius	Per-bridge. Large $\rho \Rightarrow$ flat, taut connector; small $\rho \Rightarrow$ deep, curved waist; too small $\Rightarrow$ pinch-off (§4.3).
s	cusp smoothing	Width of the quadratic fillet at the phantom/coat junction (§3).
n_x, n_y	grid resolution	Sampling density for marching squares; trades contour smoothness against compute cost.
`step`	SVG subsampling	Points are subsampled before spline fitting on export, to bound output path length without reintroducing visible faceting.

10 Implementation notes

The reference implementation evaluates (2) and the per-bridge terms as vectorised NumPy operations over the full sample grid, then calls a standard marching-squares routine (via Matplotlib's contouring backend) on the resulting array. The exported SVG path is built by fitting a closed quadratic-Bézier spline through the midpoints of consecutive (subsampled) contour points — this keeps the curve C^1 in the rendered output even though the marching-squares points are themselves piecewise-linear.