Dido's Problem and the Curvature-Variation Energy: Proof of the Least-Variation Conjecture and a Local Stability Theorem

Thomas Schaefer

6 Now Possible Things, Independent Researcher

8 September 2026

Section 5 repaired: fixed-coordinate curvature map

Contents:

Revision note. This file is the 8 September 2026 integration of the local-stability repair into the existing note. Section 5 from its heading through the paragraph immediately before Section 5.2 is the replacement text of that date. Two consistency edits are applied: the paragraph after Proposition 3, and the definition of Φ in Section 5.2. No change is made to Theorem 1, the numerical audits, Proposition 6's formulas, Conjecture 5, or the central least-variation claim.

Abstract

For a closed curve γ on a Riemannian surface, define the curvature-variation energy

$$J[\gamma] = \oint_\gamma \Big(\frac{\mathrm{d}\kappa_g}{\mathrm{d}s}\Big)^2 \mathrm{d}s.$$

We prove the conjecture that a solution of Dido's problem – a fixed-length curve of maximal enclosed area – minimizes J. Precisely, on every smooth Riemannian surface and in every specified admissible component of oriented embedded disk-type regions, any smoothly attained fixed-length maximizer has pointwise constant geodesic curvature and therefore attains the unconditional global minimum $J = 0$. This assertion is conditional on attainment and sufficient regularity; it neither proves that an equality-constrained maximizer exists nor classifies or uniquely selects one. Those restrictions are hypotheses of the theorem rather than exceptions to it. The proof is the classical first-variation argument, whose immediate consequence for J appears to have been obscured by a mistaken reading of Gauss–Bonnet. Thus the first result is the classical constant-geodesic-curvature characterization applied to J, not a new proof of that characterization; the new content lies in resolving the proposed least-variation implication and in the quantitative local theorem below.

We also prove a local quantitative stability theorem on an arbitrary smooth Riemannian surface. If the chosen maximizer is regular and the projected linearized curvature operator is

an isomorphism after fixing reparametrization and ambient symmetries, then in a gauge-fixed $W^{3,2}$ neighborhood

$$A_{\max}(L_0) - A(\gamma) \le C\, J[\gamma].$$

The topology, locality, and nondegeneracy are substantive conditions: $W^{3,2}$ is the space in which curvature variation and the normal-graph inverse map are simultaneously controlled; locality avoids the possible coexistence of other stable or unstable constant-curvature loops; and nondegeneracy excludes Jacobi kernel directions in which higher-order scaling replaces the linear estimate. No Morse–Bott classification of the full zero set is needed for this local theorem.

We do not claim a global localization theorem. Such a result would require, in addition, a closed or otherwise compact ambient setting, a compact smoothly stratified family of constant-curvature loops, uniform Morse–Bott estimates, and compactness of every sequence with $J \to 0$ without self-touching, multiple covering, escape to infinity, or change of topological type. Degenerate kernel directions may force metric-dependent exponents, and a Łojasiewicz–Simon formulation would add analyticity, Fredholm, gauge, and norm-comparison requirements. We state this global problem separately, document the false boundary relation $\kappa_g = \lambda - \mu K$ that originally concealed the elementary proof, and give an explicit Fourier model on surfaces of revolution illustrating the local theorem and its loss of coercivity at resonance.

1 Introduction

Dido's problem – enclose the most area with a boundary of given length – is solved in the Euclidean plane by the circle, on the appropriate isoperimetric branches of the sphere and hyperbolic plane by geodesic circles, and on general surfaces by regions whose boundaries the modern theory characterizes completely at the level of the first variation. The present note concerns a natural variational companion. For a closed curve γ of length L on a smooth Riemannian surface (M, g), parametrized by arclength s, with geodesic curvature $\kappa_g(s)$, define

$$J[\gamma] = \oint_\gamma \Big(\frac{\mathrm{d}\kappa_g}{\mathrm{d}s}\Big)^2 \mathrm{d}s.$$

J is the Dirichlet energy of the curvature function. It measures one component of, but is not equal to, physical jerk. At unit speed the Frenet equations give

$$\nabla_T^2 T = \kappa_g' N - \kappa_g^2 T, \qquad \big|\nabla_T^2 T\big|^2 = (\kappa_g')^2 + \kappa_g^4,$$

with additional normal-curvature terms possible for an ambient embedding. Thus J measures the squared curvature-variation component of intrinsic jerk. Even before extrinsic normal-curvature effects are included, the full squared covariant jerk contains the additional κ_g^4 term, so minimizing J is not the same variational problem as minimizing total physical jerk. The functional and its relatives appear in geometric design as minimum variation curves [Mor92, MS92]. J is not the elastic energy $\oint \kappa_g^2\, \mathrm{d}s$ of Euler and of Langer–Singer [LS84]; it penalizes the variation of bending, not bending itself.

The conjecture. On a Riemannian surface, the fixed-length curve of maximal enclosed area minimizes J among closed curves of that length.

On the Euclidean plane this is immediate: $J \ge 0$ with equality precisely for constant

κ_g, i.e. circles, which are also the isoperimetric optimizers. The same argument runs on any constant-curvature surface. The interesting case appears to be variable Gaussian curvature K, where one is tempted to reason as follows: Gauss–Bonnet forces $\oint \kappa_g \, \mathrm{d}s = 2\pi - \iint_D K \, \mathrm{d}A$, the right side depends on the enclosed region, hence (the temptation concludes) constant κ_g is generically impossible, $J > 0$ at the optimum, and the conjecture becomes a genuine comparison between two competing functionals.

The purpose of this note is threefold. First, to prove the stated least-variation conjecture on an arbitrary smooth Riemannian surface under the natural hypothesis of smooth attainment: the optimizer has constant κ_g by a three-line first-variation argument (Section 2), and hence globally minimizes J. The Euler–Lagrange characterization is classical; the point of the proof is the consequence for the conjectured functional. Second, to dissect a false boundary relation, $\kappa_g = \lambda - \mu K$, which is seductive enough that the investigation initially accepted it, derived it, and numerically confirmed it before locating the errors that jointly manufactured it (Section 4). We record the dissection because the errors are natural enough to be committed again, by humans or by machines – indeed, as recorded in Section 4(a) and the acknowledgments, the first draft of the dissection itself misidentified one mechanism, and the error was caught only under the adversarial review protocol this note advocates. Third, to prove the local quantitative statement underneath the original conjecture: a Bonnesen-type stability inequality bounding the isoperimetric deficit by J near a regular nondegenerate maximizer (Section 5). We then separate from that theorem the genuinely global localization problem, for which we give a quadratic Fourier model on surfaces of revolution and describe a possible Łojasiewicz–Simon-type program.

The constant-geodesic-curvature characterization itself is classical and is not claimed here as new. The contributions of this note are to turn that characterization into a proof of the least-variation conjecture, prove the nonlinear local stability estimate under an explicit isomorphism hypothesis, document and correct the phantom boundary relation, and isolate the additional compactness and stratification problems required by any global localization theorem. The distinction matters: the exact least-variation conjecture and the local nondegenerate inequality are proved here; classification of all isoperimetric candidates and global localization near the entire zero set remain separate problems.

In mountaineering terms, offered without apology: the summit originally sighted – that every attained Dido optimizer minimizes curvature variation – is reached by the classical first-variation path. The shelf below it is also real: near a regular nondegenerate optimizer, J controls the area deficit. Beyond that shelf lies a different expedition, the global classification and localization of all constant-curvature loops. This note proves the first two statements and marks the additional terrain required by the third.

2 The theorem, in three lines

Let γ be a smooth embedded closed curve on (M, g) bounding a region D, and let γ_t be a normal variation with speed φ (positive outward). The classical first-variation formulas are

$$\frac{\mathrm{d}}{\mathrm{d}t} A(D_t)\Big|_{t=0} = \oint_\gamma \varphi \, \mathrm{d}s, \qquad \frac{\mathrm{d}}{\mathrm{d}t} L(\gamma_t)\Big|_{t=0} = \oint_\gamma \kappa_g \, \varphi \, \mathrm{d}s.$$

The Gaussian curvature K appears in neither; it enters at second order. If γ maximizes A subject to $L = L_0$, then for some multiplier λ and all φ,

$$\oint_\gamma (1 - \lambda\kappa_g)\,\varphi\,\mathrm{d}s = 0 \qquad \Longrightarrow \qquad \kappa_g \equiv 1/\lambda.$$

Theorem 1. *On any smooth Riemannian surface, every smooth critical point of the fixed-length maximal-area problem – in particular every isoperimetric maximizer, where one exists – has pointwise constant geodesic curvature. Consequently $J[\gamma^*] = 0$, and γ^* attains the unconditional global minimum of J among closed curves of length L_0.*

This is not new. It is the two-dimensional case of the standard fact that isoperimetric boundaries are constant-mean-curvature hypersurfaces, and on surfaces it is explicit in the work of Ritoré [Rit01a, Rit01b] and Morgan–Hutchings–Howards [MHH00]; the regularity theory guaranteeing that minimizers are smooth embedded curves is likewise standard. What deserves emphasis is the corollary that disarms the Gauss–Bonnet objection:

Corollary 2. *Gauss–Bonnet does not obstruct constancy of κ_g; it determines the value of the constant:*

$$\kappa_g \equiv \frac{2\pi - \iint_D K\,\mathrm{d}A}{L_0}.$$

Gauss–Bonnet is one scalar equation; a constant has one degree of freedom. The equation is satisfied by choice of the constant, not violated by variability of K. The "open case" of the conjecture on variable-curvature surfaces was therefore never open: the conjecture holds – trivially, and in the strongest possible sense ($J = 0$ exactly, not merely minimal) – wherever the maximizer exists.

This is a candidate-reduction theorem, not by itself a solution of the global isoperimetric problem. Let Z_{L_0} denote the embedded constant-geodesic-curvature loops of length L_0 in the chosen admissible class. Conditional on existence,

$$A_{\max}(L_0) = \max_{\gamma \in Z_{L_0}} A(\gamma).$$

Thus Theorem 1 closes Dido's problem when Z_{L_0} is classified and its members can be compared – in particular when it contains a unique member modulo isometries. This occurs in the space forms and in several rotationally symmetric settings treated by [Rit01a, MHH00]. On a general surface, the theorem reduces the candidate set without performing the global comparison.

Remark 1 (degenerate multiplier). The derivation above assumes the length constraint is regular. In the singular (Fritz John) case the constraint gradient vanishes, $\oint \kappa_g \varphi\,\mathrm{d}s = 0$ for all φ, which forces $\kappa_g \equiv 0$: the boundary is a closed geodesic. The conclusion of Theorem 1 – constancy of κ_g, hence $J = 0$ – covers this exceptional case as well. However, in the singular case the tangent space to the constraint set is no longer the mean-zero variations, so all second-order statements in this note (Proposition 3, Theorem 4, and Conjecture 5) are formulated for the regular case $\kappa_g^* \neq 0$ only; see the discussion preceding Theorem 4.

The regular-multiplier condition is especially important for any global version. At a nonzero constant curvature, the derivative of length detects the mean normal mode, so the exact-length set is locally a codimension-one Banach submanifold and that mode can be eliminated. At a closed geodesic, the derivative vanishes in every normal direction; the same constraint set then has a different local model governed first at second order. If the global zero set contains both regular loops and geodesics, no single constraint chart or uniform inverse estimate covers it. Conjecture 5 must therefore exclude the singular stratum, treat it separately, or be reformulated on a stratified constraint space.

The following conditions give the theorem its precise scope. They do not weaken the proved implication, but each becomes a genuine obstruction when one asks for the global localization statement formulated later as Conjecture 5.

Existence. The maximizer can fail to exist by several mechanisms. Where K is unbounded below, the fixed-length area supremum may be infinite: a curve of fixed length can enclose ever more area by drifting toward increasingly negative curvature. In our numerical experiments on a surface with $K \approx 2x$, optimizers fled down the curvature gradient without converging; this is failure of compactness, not a counterexample. Bounded curvature does not suffice either: a complete noncompact surface may carry arbitrarily large "balloons" attached through necks of boundary length L_0, producing unbounded enclosed area.

Nor does the exact constraint $L = L_0$ acquire compactness automatically on a closed surface. Perimeter is lower semicontinuous under the usual weak convergence, so a limit may lose length and leave the equality-constrained class. On S^2, smooth embedded curves of fixed length can lie in shrinking balls while enclosing area tending to zero; choosing the complementary disk then gives areas tending to $\mathrm{Area}(S^2)$ without attainment. The standard compactness theory applies to minimizing perimeter at fixed area (or to an inequality constraint), not automatically to maximizing area at exact perimeter. We therefore retain the original fixed-length question but make all such statements conditional on attainment in a specified admissible component, or on a monotone branch of the isoperimetric profile where the fixed-area and fixed-length formulations are equivalent. See [Rit01b, Ros05] for standard existence results in the classical fixed-area formulation.

This attainment condition is indispensable to the global problem. A local theorem begins with one curve and studies a neighborhood of it; a global theorem must first guarantee that sequences of nearly zero curvature variation do not drift out of every compact set or converge only after losing the exact perimeter. On a noncompact surface such sequences may escape geometrically, while even on a closed surface the exact-perimeter class can lose length under weak convergence. Consequently Conjecture 5 cannot follow from pointwise analysis of the Jacobi operator alone: it needs a separate global compactness hypothesis, or geometric assumptions from which that compactness is proved.

Regularity, and what "enclosed" means. The functional J requires $\kappa_g' \in L^2$. We therefore formulate Theorem 1 for smooth embedded critical curves. When a candidate is obtained from the classical fixed-area isoperimetric problem, the two-dimensional regularity theory surveyed in [Ros05] supplies this smoothness; an equality-constrained weak maximizer without such regularity would not yet define J.

On a closed surface, moreover, a separating curve has two complementary sides. We fix an oriented disk-type region D, not merely its unoriented boundary, and restrict variations to the corresponding component of the space of regions. On the sphere the small-cap and large-cap branches must not be conflated: the same round boundary determines both, but their local area-stability signs at fixed length are opposite. Every assertion of maximality or stability below is relative to the chosen side and admissible component.

Regularity and the choice of admissible component are likewise local data that become global constraints. In a sufficiently small C^1 tubular neighborhood, an embedded curve remains embedded, retains its orientation and topological type, and has an unambiguous normal-graph representation. Across the full curve space none of these properties is closed: embedded curves may converge to a self-touching curve, simple loops may approach a multiple cover, and the two complementary sides of a separating curve may represent different stability branches. A global localization theorem must therefore exclude these boundary strata or include them in a more complicated compactification. Conjecture 5 records this as preservation of embeddedness and topological type rather than hiding it inside the word "curve."

Uniqueness fails, and that is where the geometry lives. $J = 0$ is attained by every closed constant-κ_g loop of length L_0, and explicit rotationally symmetric examples contain both stable and unstable members. For small lengths one also expects loops concentrating near nondegenerate critical points of K, analogous one dimension down to Ye's construction of constant-mean-curvature spheres near critical points of scalar curvature [Ye91]. We use this only as motivation: [Ye91] does not prove the loop statement on surfaces. Ritoré [Rit01a] supplies the relevant classification in rotationally symmetric settings. Stability is a necessary second-order filter, not a global selector: stable nonmaximizing loops may coexist, and the final choice is made by global area comparison.

This failure of uniqueness is not a limitation of the least-variation proof: the maximizer has $J = 0$ whether or not other zero-energy loops exist. It is, however, the central obstruction to a global converse. Small J can place a curve near an unstable constant-curvature loop or a stable loop with smaller enclosed area just as readily as near the global maximizer. Thus no global estimate can generally localize small-J curves to the optimizer or bound their global area deficit. The only viable global target is localization to the full relevant zero set, whose components must then be classified and compared by information not contained in J.

Proposition 3 (candidate reduction and stability)**.** *In the regular case $\kappa_g^* \neq 0$, the isoperimetric maximizer is distinguished within the zero set $\{J = 0\}$ of length-L_0 constant-curvature loops as an area-maximal element. Its Jacobi form is*

$$Q(\varphi, \varphi) = \oint_\gamma \Big((\varphi')^2 - (\kappa_g^2 + K)\varphi^2 \Big) \, \mathrm{d}s,$$

and fixed-length area maximality implies

$$\frac{1}{\kappa_g^*} Q(\varphi, \varphi) \geq 0 \quad \text{on} \quad \Big\{ \oint \varphi \, \mathrm{d}s = 0 \Big\}.$$

Indeed, $Q = \delta^2(L - \kappa_g^* A)$, so along a variation with $L(t) \equiv L_0$ one has $A''(0) = -Q/\kappa_g^*$. For a chosen branch with $\kappa_g^* > 0$, this condition reduces to $Q \geq 0$; for $\kappa_g^* < 0$, it becomes $Q \leq 0$. The sign is determined by the selected region and outward-normal convention, so switching the enclosed side is not a harmless change of notation. The local estimate of Theorem 4 requires only $\kappa_g^* \neq 0$; the Fourier model in Section 5.3 is explicitly restricted to the positive-curvature branch. Gaussian curvature enters this second-variation test, although it does not enter the first variations of area and length.

3 Numerical verification

Because the false relation of Section 4 was originally "confirmed" numerically, we record a verification protocol with the property the original experiments lacked: a per-curve, closed-form audit.

The experiments are not part of the proof of Theorem 1. Their narrower purpose is to validate the intrinsic formulas, demonstrate that substantial variation of K is compatible with constant κ_g, and reproduce the mechanism by which the phantom relation appeared plausible. A failed optimizer, a curve sampling negligible variation of K, or a resolution-unstable result is therefore reported as inconclusive rather than as mathematical evidence.

We work on conformal metrics $g = e^{2u}(\mathrm{d}x^2 + \mathrm{d}y^2)$ with u a sum of two asymmetric Gaussian bumps (heights 0.60 and -0.35, distinct widths and centers), with no imposed continuous symmetry and with $K = -e^{-2u}\Delta u$ varying substantially along the optimized boundary. Curves are truncated Fourier series with derivatives taken exactly from the coefficients; metric length, metric enclosed area (by a Stokes fan quadrature), and geodesic curvature (by the conformal transformation law of Section 4(b)) are computed intrinsically. Certification requires optimizer convergence, the length constraint, turning number one, full-resolution embeddedness, a per-curve Gauss–Bonnet audit, an independent ray-casting Monte Carlo area check, sufficient variation of K, and stability under resolution refinement. Gauss–Bonnet residuals are at machine precision, including for arbitrary nonoptimized test curves.

Maximizing area at fixed length $L_0 = 6$ (SLSQP, multistart, refinement to 24 Fourier modes and 2048 points), the optimizer exhibits:

quantity	value
spread of K along the curve	0.678 (from -0.174 to $+0.504$)
arclength-weighted $\mathrm{std}(\kappa_g)/\lvert\kappa_g\rvert$	1.6×10^{-5}
weighted slope of κ_g against K	-0.0000 ($R^2 = 0.0003$)
$J[\gamma^*]$	9.2×10^{-7}

The curvature of the ambient surface swings by two-thirds of a unit along the optimizer; its geodesic curvature does not move at the resolved scale. Additional randomized Gaussian and multimodal Bessel surveys are supplied as ancillary falsification tests. Their outcome categories distinguish powered confirmations from locally underpowered, optimizer-inconclusive, and resolution-inconclusive cases; no powered certified case produced a resolved departure from $J = 0$. These computations do not certify that a multistart champion is the global Dido optimizer. The hardened verifier (`dido_verifier_v3.py`), survey harness (`dido_survey.py`), forcing rules, and complete machine-readable reports accompany this note.

4 Anatomy of a phantom relation

During the investigation the following boundary condition was proposed, derived, and numerically confirmed before being identified as false:

$$\kappa_g = \lambda - \mu K \qquad \text{(phantom).}$$

It is worth recording how a false relation acquired three independent pillars of support, because each mechanism is a general-purpose trap. It is also worth recording – in the spirit of full disclosure – that the first draft of this very section misidentified mechanism (a), and the error survived until adversarial review; the corrected account below is sharper than the original.

(a) A multiplier on an identity – done inconsistently. The Gauss–Bonnet relation

$$G(\gamma) = \oint_\gamma \kappa_g \, \mathrm{d}s + \iint_D K \, \mathrm{d}A - 2\pi = 0$$

holds identically for every embedded closed curve bounding a disk-type region; it is not a constraint that can be active or inactive. Consequently its first variation vanishes identically: under a normal variation with speed φ,

$$\delta\Big(\iint_D K \, \mathrm{d}A\Big) = \oint_\gamma K\varphi \, \mathrm{d}s, \qquad \delta\Big(\oint_\gamma \kappa_g \, \mathrm{d}s\Big) = -\oint_\gamma K\varphi \, \mathrm{d}s,$$

and the two terms cancel exactly – as they must, since G is constant on the space of curves. (The second formula requires no separate computation: it follows from the identity and the first formula.) Adjoining $\mu \, G(\gamma)$ to the Lagrangian therefore contributes nothing to the Euler–Lagrange equation. The phantom relation arises in one of two ways, both erroneous. Either one constrains the bulk term $\iint_D K \, \mathrm{d}A$ alone – a genuinely different variational problem, whose Euler–Lagrange condition for $A - \lambda L - \mu \iint_D K \, \mathrm{d}A$ is indeed $1 - \lambda\kappa_g - \mu K = 0$, i.e. exactly the phantom – or one adjoins the full identity but drops the variation of the boundary term $\oint \kappa_g \, \mathrm{d}s$, an algebra slip camouflaged as a modeling choice. The portable moral: an identity has identically vanishing first variation, so the moment a multiplier on an "identity constraint" produces a new Euler–Lagrange term, a variation has been dropped somewhere.

(b) An extrinsic formula on intrinsic geometry. For a conformal metric $g = e^{2u} g_0$, the geodesic curvature transforms as

$$\kappa_g = e^{-u}(\kappa_0 + \partial_n u),$$

where n is the g_0-unit normal chosen compatibly with the orientation convention defining κ_0 (reversing the normal reverses both terms). Early numerical experiments computed κ_g with the planar formula, omitting the correction $e^{-u}\partial_n u$. The omitted term is a first derivative of the conformal factor, while $K = -e^{-2u}\Delta u$ (equivalently $\Delta u = -e^{2u}K$) is a second-derivative quantity; there is no general implication that the two are linearly related along a curve. In the bump metrics of the original experiments, however, they were strongly correlated along the optimizer, and regressing the contaminated κ_g against K recovered a spurious linear relationship with $R^2 = 0.94$ – an artifact confirming the artifact. The audit of Section 3 exists to make this failure mode impossible: a curve whose computed κ_g, $\mathrm{d}s_g$, and $K \, \mathrm{d}A_g$ do not close the Gauss–Bonnet identity has been measured with the wrong geometry, whatever its R^2.

(c) A citation that was never checked. The phantom relation entered the investigation attributed to the isoperimetric literature. The actual literature says the opposite – constancy – and says it prominently [Rit01a, Rit01b, MHH00]. We flag this with some humility: portions of the investigation were conducted with and by large language models, one of which fabricated the attribution, and the relation's subsequent "derivation" (mechanism (a)) and "confirmation" (mechanism (b)) by other models and by the investigator gave the fabrication an evidential halo that survived several rounds of otherwise genuine scrutiny. A load-bearing named theorem must be verified against a primary source before anything is built on it. This is old advice; the new observation is that systems of multiple independent-seeming reasoners can manufacture *correlated* false confirmation, because the same seductive error is natural to each of them.

The blind-then-adversarial review protocol described in the acknowledgments is our proposed countermeasure, and Section 4(a) records it catching an error in this note itself.

5 The shelf: a proved local quantitative inequality

The exact least-variation conjecture identifies the right pair of quantities but, because its conclusion is $J = 0$, leaves a sharper question: does small positive curvature variation control the area lost from a specified optimum? Under the following local regularity and nondegeneracy hypotheses, the answer is yes on an arbitrary smooth Riemannian surface.

Fix an admissible component of oriented embedded disk-type regions and a length L_0 for which a smooth maximizer γ^* exists. Write its constant geodesic curvature as κ_*, and assume the length constraint is regular, so $\kappa_* \neq 0$. The region and its outward normal remain fixed by this choice; the local argument does not require changing the enclosed side or assuming $\kappa_* > 0$.

Parametrize γ^* by its arclength $s \in \mathbb{R}/L_0\mathbb{Z}$. In a tubular neighborhood, nearby unparametrized curves are represented by normal graphs

$$\gamma_f(s) = \exp_{\gamma^*(s)}\big(f(s)\,N(s)\big), \qquad f \in H^3(\mathbb{R}/L_0\mathbb{Z}),$$

where $H^j = W^{j,2}$ and N is the outward unit normal. All Sobolev spaces below use the fixed reference coordinate s and measure $\mathrm{d}s$. Define

$$\rho_f(s) = \big|\partial_s\gamma_f(s)\big|_g, \qquad \mathrm{d}\ell_f = \rho_f\,\mathrm{d}s, \qquad k(f)(s) = \kappa_g(\gamma_f)(s).$$

Here $k(f)$ is the curvature pulled back by the normal-graph parametrization, not by a curve-dependent arclength reparametrization. After shrinking the graph neighborhood, there are constants $0 < m \le M < \infty$ such that

$$m \le \rho_f(s) \le M$$

uniformly. In these coordinates the intrinsic energy is exactly

$$J[\gamma_f] = \int_0^{L_0} \left|\frac{\partial_s k(f)}{\rho_f}\right|^2 \rho_f\,\mathrm{d}s = \int_0^{L_0} \frac{|\partial_s k(f)|^2}{\rho_f}\,\mathrm{d}s. \tag{5.1}$$

The length derivative at zero is

$$\mathrm{D}L_0[h] = \kappa_* \int_0^{L_0} h\,\mathrm{d}s.$$

Since $\kappa_* \neq 0$, the Banach implicit-function theorem solves the exact-length constraint $L(\gamma_f) = L_0$ for the constant mode. Its tangent space at zero consists of mean-zero normal variations. If ambient isometries generate nontrivial normal variations, fix a local transverse slice to their orbit. We assume this slice is a smooth submanifold in the normal-graph H^3 chart and that nearby curves under consideration have representatives in it. Reparametrization is already removed by the normal-graph representation. Denote the resulting length-and-symmetry slice by $\mathcal{S}$.

Let

$$P_0 u = u - \frac{1}{L_0}\int_0^{L_0} u\,\mathrm{d}s, \qquad H^1_0 = \Big\{u \in H^1 : \int_0^{L_0} u\,\mathrm{d}s = 0\Big\},$$

and define the fixed-coordinate curvature defect

$$F(f) = P_0 k(f) \in H^1_0. \tag{5.2}$$

The mean in (5.2) is the mean with respect to the reference measure $\mathrm{d}s$. It is generally different from the intrinsic arclength mean

$$\bar{\kappa}_f = \frac{1}{L_0}\int_0^{L_0} k(f)\,\rho_f\,\mathrm{d}s = \frac{2\pi - \displaystyle\int_{D_f} K\,\mathrm{d}A}{L_0}.$$

This distinction causes no change in the zero set or the differentiated curvature: $F(f) = 0$ precisely when the geodesic curvature is constant, and $\partial_s F(f) = \partial_s k(f)$. At the constant-curvature reference curve, subtracting either mean also gives the same projected linearization on exact-length variations.

In the absence of ambient symmetry directions, set $Y = H^1_0$ and $\Pi_Y = I$. When such directions have been removed from the domain, choose a fixed closed target complement $Y \subset H^1_0$ to the corresponding cokernel and a bounded projection $\Pi_Y : H^1_0 \to Y$. Define

$$\widehat{F} = \Pi_Y F : \mathcal{S} \longrightarrow Y.$$

With the outward-normal convention used for the first variation of length, its derivative at zero is

$$\widehat{\mathcal{L}}h = \mathrm{D}\widehat{F}_0[h] = \Pi_Y P_0\big[-h'' - \big(K|_{\gamma^*} + \kappa_*^2\big)h\big]. \tag{5.3}$$

We call the maximizer nondegenerate on the chosen slice when

$$\widehat{\mathcal{L}} : T_0\mathcal{S} \longrightarrow Y$$

is a bounded isomorphism. This is an explicit hypothesis, not a consequence of maximality. Without symmetries it is invertibility of the projected elliptic operator on mean-zero variations; with symmetries it is the corresponding invertibility after symmetry reduction.

Write $A_{\max}(L_0) = A(\gamma^*)$ and

$$\delta(\gamma) = A(\gamma^*) - A(\gamma).$$

Theorem 4 (local quantitative curvature-variation inequality)**.** *Let γ^* be a smoothly attained, regular fixed-length area maximizer in the specified admissible component. Assume the local slice described above exists and the projected linearized curvature operator* (5.3) *is a bounded isomorphism. Then there are a neighborhood U of zero in $\mathcal{S}$ and a constant $C > 0$ such that every graph γ_f, $f \in U$, satisfies*

$$0 \le \delta(\gamma_f) \le C\,J[\gamma_f].$$

The same conclusion holds for nearby curves having representatives in this slice under ambient isometries, because length, area, and J are invariant under those isometries. The neighborhood is an $H^3 = W^{3,2}$ neighborhood in the normal-graph chart.

5.1 Proof of Theorem 4

We first establish the regularity needed for the inverse-function theorem entirely in the fixed reference coordinate. In tubular coordinates, the curvature has the quasilinear form

$$k(f) = a(s, f, f')f'' + b(s, f, f'),$$

with smooth coefficients and a bounded away from zero in absolute value on a sufficiently small neighborhood. In one dimension, H^1 is a multiplication algebra, and the usual Sobolev multiplication and smooth superposition estimates imply that

$$k : H^3 \longrightarrow H^1$$

is continuously differentiable. The density ρ_f depends smoothly on (s, f, f'); length and enclosed area are twice continuously differentiable, as are their restrictions to the smooth constraint slice. Since P_0 and Π_Y are fixed bounded linear maps, $\widehat{F}$ is continuously differentiable on $\mathcal{S}$. No differentiability of composition with a variable arclength parametrization is used.

At zero, $\widehat{F}(0) = 0$ and $\mathrm{D}\widehat{F}_0 = \widehat{\mathcal{L}}$ is an isomorphism. The Banach inverse-function theorem, applied in a local chart on $\mathcal{S}$, therefore gives

$$\|f\|_{H^3} \le C_1\|\widehat{F}(f)\|_{H^1} \le C_2\|F(f)\|_{H^1} \tag{5.4}$$

after shrinking the neighborhood. Here the slice-chart norm and the ambient graph norm are locally comparable. The local inverse has bounded derivative on a sufficiently small neighborhood of zero, which gives the first inequality; boundedness of Π_Y gives the second.

Because $F(f)$ has zero mean with respect to $\mathrm{d}s$, Poincaré's inequality on the fixed reference circle yields

$$\|F(f)\|_{H^1}^2 \le \left[1 + \left(\frac{L_0}{2\pi}\right)^2\right] \int_0^{L_0} |\partial_s k(f)|^2 \,\mathrm{d}s \le M\left[1 + \left(\frac{L_0}{2\pi}\right)^2\right] J[\gamma_f], \tag{5.5}$$

where the last step uses (5.1) and $\rho_f \le M$. Equations (5.4) and (5.5) give

$$\|f\|_{H^3}^2 \le C_3 J[\gamma_f]. \tag{5.6}$$

Finally, the area functional restricted to $\mathcal{S}$ is twice continuously differentiable and has zero first derivative at zero, because γ^* is a constrained critical point. Taylor's theorem in a slice chart, local boundedness of its second derivative, and maximality give

$$0 \le A(\gamma^*) - A(\gamma_f) \le C_4\|f\|_{H^3}^2. \tag{5.7}$$

Combining (5.6) and (5.7) proves the asserted inequality. Shrinking the neighborhood preserves embeddedness, the chosen side, and disk-type topology, since $H^3(S^1)$ embeds continuously into $C^1(S^1)$. The proof is local in the H^3 normal-graph topology; it does not assert the same quantifier over an entire C^1 neighborhood. □

The area-deficit inequality is one-sided: high-frequency perturbations can make the ratio of curvature-variation energy to area deficit arbitrarily large. This does not preclude a two-sided comparison with squared H^3 graph distance. Indeed, continuous differentiability of

$k : H^3 \to H^1$, together with $k(0) = \kappa_*$, gives

$$\|\partial_s k(f)\|_{L^2} \le \|k(f) - k(0)\|_{H^1} \le C_5 \|f\|_{H^3}.$$

Using $\rho_f \ge m$ in (5.1) and combining with (5.6), we obtain constants $c, C' > 0$ such that

$$c\,\|f\|_{H^3}^2 \le J[\gamma_f] \le C' \|f\|_{H^3}^2$$

on the chosen slice. Thus J is locally equivalent to squared H^3 graph distance under the nondegeneracy hypothesis, while it need not be equivalent to the area deficit. When symmetries are present, this distance is evaluated after passing to the chosen transverse representative.

5.2 The separate global problem and its conditions

Theorem 4 concerns one regular nondegenerate maximizer and one neighborhood. It neither requires nor supplies a classification of every constant-curvature loop. A global statement would have to replace the one invertible linearization by uniform control over the entire zero set and would have to prove that every small-J curve reaches a neighborhood in which such local analysis applies. We record that different problem as a conjecture rather than treating it as an unproved part of Theorem 4.

Conjecture 5 (global localization, compact Morse–Bott form)**.** *Assume (M, g) is closed, and let Z_0 be a compact C^1 submanifold, modulo reparametrization and ambient isometries, of embedded constant-geodesic-curvature loops of length L_0 in the chosen admissible class. Suppose that at every $\gamma_0 \in Z_0$ the kernel of the projected linearized operator $\mathcal{L}$ equals $T_{\gamma_0} Z_0$, with uniform Morse–Bott estimates. Suppose further that every sequence of admissible embedded curves with $J[\gamma_j] \to 0$ has, after reparametrization, a subsequence converging in $W^{3,2}$ to Z_0 without loss of embeddedness or topological type. Then there exist $\varepsilon_0, C > 0$ such that every admissible curve of length L_0 with $J[\gamma] \le \varepsilon_0$ satisfies*

$$\mathrm{dist}_{W^{3,2}}(\gamma, Z_0) \le C\sqrt{J[\gamma]}.$$

Closedness and ambient compactness. Closedness is not cosmetic. On a noncompact surface a sequence with $J \to 0$ can move out of every compact set while its local geometry remains controlled, so no finite collection of normal-graph charts captures all small-J curves. A global theorem on a noncompact surface would need a replacement such as bounded geometry plus a no-escape hypothesis. Even on a closed surface, compactness of the base manifold does not by itself make the exact-length embedded curve space compact, as the length-loss examples of Section 2 show.

Compact structure of the zero set. The local theorem needs only one isolated zero modulo symmetries. Conjecture 5 assumes instead that the relevant zero set Z_0 is a compact C^1 submanifold with a uniform tubular neighborhood. On a general metric the constant-curvature loops may form isolated points, components of different dimensions, bifurcating branches, or singular intersections. Establishing the asserted structure is therefore a classification problem, not a routine consequence of the Euler–Lagrange equation. Without it there is no single normal bundle on which to patch the local inverse estimates.

Uniform Morse–Bott nondegeneracy. At one curve, Theorem 4 assumes that the projected curvature operator is invertible after symmetry reduction. Along a family of zeros, the appropriate replacement is that its kernel consist exactly of tangent directions to that family and that the inverse on normal directions remain uniformly bounded. Pointwise kernel identities are insufficient: the smallest nonzero eigenvalue may approach zero along Z_0, causing the constants in the local estimates to diverge. Compactness of Z_0 helps only after continuity, a uniform tubular gauge, and the Morse–Bott structure have actually been established.

Sequential compactness and preservation of embeddedness. The Poincaré estimate controls the nonconstant part of curvature, and for embedded disk boundaries on a closed surface Gauss–Bonnet also controls its mean in terms of L_0, $\|K\|_\infty$, and $\mathrm{Area}(M)$. This suggests $W^{3,2}$ control of arclength parametrizations, but it does not automatically place a limit in the same admissible stratum. A sequence may self-touch, converge to a multiple cover, change topological type, or lose the chosen enclosed side. The sequential hypothesis in Conjecture 5 is exactly the bridge from global smallness of J to one of the local charts; assuming only pointwise Morse–Bott nondegeneracy would leave that bridge missing.

Degenerate directions and the exponent. If Morse–Bott fails, the $\sqrt{J}$ rate need not survive. In a fixed normal-graph chart, let $\Phi(f) = P_0 k(f)$, with P_0 and k defined as in Section 5. Interpret the following normal-form expansion in H^1, along an admissible exact-length path in that chart, with the notation $\gamma_0 + t\psi$ abbreviating a path whose leading graph displacement is $t\psi$. If along a kernel direction ψ not tangent to a constant-curvature branch the first nonvanishing normal form is

$$\Phi(\gamma_0 + t\psi) = t^m q + O(t^{m+1}),$$

with q nonconstant and $m \geq 2$, then $J(\gamma_0 + t\psi) \asymp t^{2m}$ while the distance to the zero set is of order t, giving exponent $1/(2m)$. The generic quadratic case would give $1/4$, but even that rate requires verification that the quadratic coefficient is nonzero. Thus degeneracy may produce different exponents at different components, precluding a universal global power without further finite-order or analytic structure.

Analyticity and the Łojasiewicz–Simon program. A metric-dependent error bound may be approachable through a Łojasiewicz–Simon-type argument, but that theory is not installed merely by naming it. One would need an analytic metric or an appropriate analytic functional setting, a fixed Banach-manifold gauge, a Fredholm Hessian, and a precise comparison of norms. Moreover, J is not the squared norm of the constrained area gradient: the gradient-like residual is the nonconstant part of κ_g, and J controls it through Poincaré. Converting the relevant gradient inequality into a distance estimate in $W^{3,2}$ is additional analysis. We therefore present this as a program for degenerate global localization, not as an application of an established theorem.

Any partial result establishing this program for analytic metrics, for a single compact Morse–Bott component, or under a proved no-escape condition would be a natural sequel. Such results would extend the range over which the local inverse estimates can be patched; they are not prerequisites for either theorem proved here.

Localization to the zero set, not to the maximizer. Finally, Conjecture 5 cannot generally be strengthened to proximity to the maximizer. Small J near another constant-curvature loop is compatible with an $O(1)$ global area deficit, whether that loop is unstable or merely a nonglobal stable optimum. The sharp possible global conclusion is localization to the full relevant zero set Z_0. Selecting an area-maximizing component from that set remains a global comparison problem to which J, vanishing everywhere on Z_0, contributes no information.

5.3 Quadratic-order model on surfaces of revolution

On a surface of revolution $g = \mathrm{d}r^2 + a(r)^2 \mathrm{d}\theta^2$, let γ_0 be a constant-curvature parallel at $r = r_0$, with $a_0 = a(r_0)$, $L_0 = 2\pi a_0$, $\kappa_0 = a'(r_0)/a_0 \neq 0$, $K_0 = -a''(r_0)/a_0$, and $q_0 = K_0 + \kappa_0^2$. For normal graphs $r(\theta) = r_0 + f(\theta)$ with $f(s) = \sum_n f_n e^{2\pi ins/L_0}$ in arclength $s = a_0\theta$, the linearized curvature defect is $\delta\kappa = -f'' - q_0 f$, diagonal in Fourier with eigenvalues

$$\lambda_n = \left(\frac{2\pi n}{L_0}\right)^2 - q_0.$$

Choose the side for which $\kappa_0 > 0$. Since $Q = \delta^2(L - \kappa_0 A)$, elimination of the mean mode by the length constraint gives, with the stated Fourier normalization,

$$\delta_{\mathrm{quad}}[f] = \frac{L_0}{2\kappa_0} \sum_{n \neq 0} \lambda_n |f_n|^2, \qquad J_{\mathrm{lin}}[f] = L_0 \sum_{n \neq 0} \left(\frac{2\pi n}{L_0}\right)^2 \lambda_n^2 |f_n|^2,$$

where $f_0 = \Psi(f_{\neq 0}) = O(\|f_{\neq 0}\|_{H^1}^2)$. Genuine symmetry modes, when present, are quotiented before imposing strict positivity; let $\mathcal{N}$ denote the remaining nonzero Fourier modes.

Proposition 6 (quadratic coercivity)**.** *If $\lambda_n > 0$ for every $n \in \mathcal{N}$, then the quadratic models satisfy*

$$\delta_{\mathrm{quad}}[f] \leq C\, J_{\mathrm{lin}}[f], \qquad C = \frac{1}{2\kappa_0} \bigg/ \inf_{n \in \mathcal{N}} \left(\frac{2\pi n}{L_0}\right)^2 \lambda_n > 0,$$

about every strictly stable, nonresonant constant-curvature parallel of a surface of revolution, modulo symmetries. The infimum is attained at small $|n|$ because $(2\pi n/L_0)^2 \lambda_n \to \infty$. It loses positivity at a resonance $(2\pi n/L_0)^2 = q_0$, exactly where the inverse estimate required by the nonlinear argument fails and higher-order exponents may arise.

Proposition 6 is not needed to complete the proof of Theorem 4; it is the Fourier-diagonal form of that theorem's linear coercivity mechanism. It makes the nondegeneracy hypothesis computable in a symmetric setting and shows exactly how the local constant deteriorates as a resonant eigenvalue approaches zero. At resonance Theorem 4 does not apply, and the higher-order alternatives described in Section 5.2 become relevant.

The natural companions in the literature are the Bonnesen-style inequalities surveyed by Osserman [Oss79], which control the planar deficit by geometric asymmetry; the sharp quantitative isoperimetric inequality of Fusco–Maggi–Pratelli [FMP08], which controls asymmetry by the deficit; and the Łojasiewicz–Simon inequality [Sim83] as a guide for the proposed analytic framework around Conjecture 5. Theorem 4 is a statement of the same genus with the curvature-variation energy J as the controlling quantity, and to our knowledge J has not been placed in this role. Beyond intrinsic interest, Theorem 4 gives a local certificate of near-optimality computable from boundary data alone – relevant wherever curvature-fairing and largest-enclosure objectives coexist (path planning, fairing, grain-boundary and cell-packing models).

6 Where this leaves Dido

The least-variation conjecture is proved in its natural geometric scope. On every smooth Riemannian surface, wherever the fixed-length problem is attained by a smooth embedded boundary in the specified admissible component, that optimizer has constant geodesic curvature and hence globally minimizes J with value zero. The theorem is conditional on smooth attainment because without an optimizer there is no curve to which the proposed property can be attributed. It proves the property asserted of every optimizer; it does not claim a universal existence theorem or a classification of the optimizer itself.

The quantitative local statement is also proved: near a regular nondegenerate maximizer, in the natural gauge-fixed $W^{3,2}$ topology, the area deficit is bounded by CJ. This conclusion uses no classification of the global zero set and no Morse–Bott hypothesis. Proposition 6 exposes the same coercivity explicitly on surfaces of revolution.

What remains open is different in kind. A global localization theorem must control escape and loss of compactness, classify the possibly singular zero set, make the normal inverse estimates uniform, preserve embeddedness and topological type under limits, and replace the square-root law by higher-order or Łojasiewicz exponents at degeneracy. Those tasks constitute Conjecture 5, not an omitted final step in either proof given here.

Acknowledgments

The conjecture originated in a dream (T.S., January 2026) and was explored with a multi-model AI orchestration framework. Model-generated suggestions included the first-variation proof, candidate numerical protocols, the degenerate-multiplier observation, and several formulations of the quantitative question. A later blind-then-adversarial review exposed both the phantom relation and an error in the first draft of Section 4(a), which had failed to vary the complete Gauss–Bonnet identity consistently. A further adversarial pass identified the exact-perimeter compactness problem, the second-variation sign dependence, the missing hypotheses in the global localization statement, and the distinction between curvature variation and physical jerk; these corrections are incorporated in the present version. A subsequent independent collective-model review prompted the sharper priority statement in the abstract, the explicit κ_g^4 clarification, and the permanent-archive plan below.

The 8 September 2026 repair of Section 5 replaces a mixed-measure curvature defect and a variable-arclength identification with a fixed-coordinate curvature map, a density comparison for the intrinsic energy, and an explicit nondegeneracy hypothesis that does not flip the enclosed side. That repair does not change Theorem 1, the numerical audits, or Proposition 6.

All model contributions were treated as unverified suggestions rather than as mathematical sources. Cited results were checked against the underlying literature, numerical claims are accompanied by reproducible audits, and final mathematical responsibility rests with the human author. The full transcripts are preserved for a separate study of correlated multi-agent reasoning failures; they are not part of the evidentiary basis for the theorems stated here.

References

[FMP08] N. Fusco, F. Maggi, A. Pratelli, The sharp quantitative isoperimetric inequality, *Ann. of Math.* (2) **168** (2008), 941–980.

[LS84] J. Langer, D. A. Singer, The total squared curvature of closed curves, *J. Differential Geom.* **20** (1984), 1–22.

[MHH00] F. Morgan, M. Hutchings, H. Howards, The isoperimetric problem on surfaces of revolution of decreasing Gauss curvature, *Trans. Amer. Math. Soc.* **352** (2000), 4889–4909.

[Mor92] H. P. Moreton, *Minimum Variation Curves and Surfaces for Computer-Aided Geometric Design*, Ph.D. thesis, University of California, Berkeley, 1992.

[MS92] H. P. Moreton, C. H. Séquin, Functional optimization for fair surface design, *Computer Graphics (SIGGRAPH '92)* **26** (1992), 167–176.

[Oss79] R. Osserman, Bonnesen-style isoperimetric inequalities, *Amer. Math. Monthly* **86** (1979), 1–29.

[Rit01a] M. Ritoré, Constant geodesic curvature curves and isoperimetric domains in rotationally symmetric surfaces, *Comm. Anal. Geom.* **9** (2001), 1093–1138.

[Rit01b] M. Ritoré, The isoperimetric problem in complete surfaces of nonnegative curvature, *J. Geom. Anal.* **11** (2001), 509–517.

[RR02] M. Ritoré, A. Ros, Some updates on isoperimetric problems, *Math. Intelligencer* **24** (2002), 9–14.

[Ros05] A. Ros, The isoperimetric problem, in: *Global Theory of Minimal Surfaces*, Clay Math. Proc. 2, Amer. Math. Soc., 2005, 175–209.

[Sim83] L. Simon, Asymptotics for a class of non-linear evolution equations, with applications to geometric problems, *Ann. of Math.* (2) **118** (1983), 525–571.

[Ye91] R. Ye, Foliation by constant mean curvature spheres, *Pacific J. Math.* **147** (1991), 381–396.

Reproducibility. The hardened verification code (`dido_verifier_v3.py`), the ensemble survey (`dido_survey.py`), their forcing rules, and the machine-readable reports accompany this note. The supplemental intrinsic ring-torus regression (`dido_torus.py`, with its tests and machine-readable result) is included but is not used in either proof. The implementation uses Python with NumPy and SciPy and performs intrinsic per-curve Gauss–Bonnet, topology, independent-area, convergence, and resolution audits. Upon publication, a versioned snapshot of the source, configurations, tests, and reports will be deposited in Zenodo or an equivalent permanent repository; the resulting DOI will be cited in the version of record.