How We Found the Test Function for Branched Stable Minimal Immersions

An AI-Assisted Computational Research Note

Gaoming Wang & **Xuwen Zhang**

August 2026

ABSTRACT

This note describes how we found a test function used in our study of branched stable minimal immersions. Starting from a capillary prototype, we combined hand calculations with AI-assisted programming, principally through GitHub Copilot and Cursor's Composer 2, to generate and test numerical experiments rapidly. These tools were used mainly for coding rather than mathematical reasoning. We expect increasingly capable AI tools to accelerate mathematical exploration and deepen, rather than replace, human understanding.

THE DISCOVERY PATH

Capillary prototype $\longrightarrow$ balanced $2/m$ tail $\longrightarrow$ small-k limit $\longrightarrow$ linear zero model $\longrightarrow$ cutoff ODE

Contents

1 Purpose and perspective 2

2 The capillary prototype 2

2.1 The boundary condition as a design equation 3

2.2 What was tried before the final family 3

2.3 The additive route to the final family 4

2.4 The design blueprint from the capillary problem 4

3 From the geometric inequality to a numerical criterion 5

4 The first computational phase: explicit formulas 6

4.1 The first numerical candidates 6

4.2 Searching over the detector P 7

5 The decisive small-parameter calculation 8

5.1 The small-k first-order expansion 9

5.2 The first-order criteria in ranks two and three 9

5.3 Rank-three scans and analytic confirmation 10

6 The local repair and the final profile 10

6.1 The linear model near a zero 10

6.2 The logarithmic transition and the rank threshold 11

6.3 The final profile and its verification 12

7 Summary and outlook 12

1 Purpose and perspective

The final test function in the paper is compact enough that its origin can be easy to miss. In its finished form it is

$$G(y) = \big(s_L^2(y) + kP(y)\big)^{1/2}, \qquad P(y) = \prod_{j=1}^{m} \phi\big(1 - \langle y, p_j\rangle\big), \tag{1.1}$$

where $L = \operatorname{span}\{p_1, \dots, p_m\}$, $s_L(y) = |\pi_{L^\perp} y|$. Throughout, we call $p_1, \dots, p_m$ the prescribed directions; they are precisely the zeros of the test function. The term *pole* is reserved for the multipole construction, its factors, and its numerical configurations. The cutoff ϕ is linear near zero and becomes a $2/m$-power away from zero. Read backwards, every part of (1.1) has a reason. The background s_L^2, the product structure, the exponent $2/m$, the linear core, the order in which T and k are selected, and the normal-rank restriction were each forced by a different calculation.

The purpose of this note is not only to explain the mathematical origin of this formula, but also to describe how intelligent assistance contributed to the research process. Our principal AI-assisted environments during the original search were a GitHub Copilot subscription and Cursor's Composer 2. Starting from hand calculations, we used these tools to generate and revise numerical tests quickly, compare alternative implementations of long formulas, scan large parameter families, and concentrate computation in the singular regimes where a proposed test function was most likely to fail.

This assistance was particularly important for numerical validation. It made the cycle from a geometric idea to an executable test sufficiently short that we could examine many more candidates than would otherwise have been practical. A negative numerical value often exposed the exact scale or region in which an ansatz broke down; a positive scan identified a candidate worthy of further analysis. Neither outcome was confused with a proof. The choice of meaningful asymptotic regimes, the interpretation of numerical failures, and the conversion of observed patterns into rigorous estimates remained the mathematical work of the researchers.

This is also a historical account of a rapidly changing technology. In the few months since the original experiments, the intelligence and reasoning ability of leading AI models have improved qualitatively. A comparable search undertaken now may be completed more quickly and with more capable assistance in symbolic reasoning, numerical design, adversarial testing, and proof organization. The workflow described here should therefore not be read as a benchmark for current AI systems. We nevertheless describe it in detail because it gives a concrete example of how AI assistance contributed to this research and of how such tools can help push mathematical understanding further.

2 The capillary prototype

The point of departure was not a ready-made formula, but a constrained search problem in our stable-capillary work. Let M^n be a capillary minimal hypersurface, let ν be its unit normal, and write

$$p_\pm = \cos\theta\, e_1 \pm \sin\theta\, e_2.$$

Here e_2 is the unit vector orthogonal to e_1 in the two-plane containing the prescribed normal directions, and we write

$$\nu_1 = \langle \nu, e_1\rangle, \qquad \nu_2 = \langle \nu, e_2\rangle.$$

We wanted one nonnegative scalar function $g = g(\nu)$ satisfying all three of the following requirements:

- a strictly positive Schoen-type inequality
  $$g\Delta_M g + |A|^2 g^2 \geq c|A|^2 \qquad \text{on } \{g > 0\}, \tag{2.1}$$
  for some $c > 0$;
- the exact zero set
  $$g(\nu) = 0 \quad \Longleftrightarrow \quad \nu \in \{p_+, p_-\};$$

- the Robin boundary condition

$$B_\theta(U) := \partial_\eta U + 2\cot\theta\, A(\eta,\eta)U, \qquad B_\theta(g^2) = 0 \qquad \text{on } \partial M \cap \{g > 0\}. \tag{2.2}$$

Here and below, η denotes the outer unit conormal of ∂M in M.

2.1 The boundary condition as a design equation

The Robin condition was not an afterthought. It was one of the main filters on every candidate. Along ∂M,

$$\nu_1 = \cos\theta, \qquad \partial_\eta \nu_1 = \sin\theta\, A(\eta,\eta), \qquad \partial_\eta \nu_2 = -\cot\theta\, A(\eta,\eta)\nu_2. \tag{2.3}$$

Consequently, if a one-variable squared profile $U = p^2(\nu_1)$ is positive on the boundary, then direct substitution into (2.2) shows that its square-root profile satisfies

$$(\log p)'(\cos\theta) = -\frac{\cos\theta}{\sin^2\theta}. \tag{2.4}$$

This logarithmic-derivative constraint, evaluated at $\nu_1 = \cos\theta$, was used to generate the one-variable candidates.

There was also an elementary but decisive closure property. The operator B_θ is linear, and hence

$$B_\theta(U_1) = B_\theta(U_2) = 0 \quad \Longrightarrow \quad B_\theta(\alpha_1 U_1 + \alpha_2 U_2) = 0 \qquad \text{for all } \alpha_1, \alpha_2 \in \mathbb{R}.$$

Thus squared profiles could be designed additively without introducing a second boundary operator.

2.2 What was tried before the final family

Before reaching the final capillary family, we tried many alternatives. We began with the natural two-pole detector

$$g_{\mathrm{pair}} = \left[(1 - \langle \nu, p_+\rangle)(1 - \langle \nu, p_-\rangle)\right]^{1/2}.$$

It has exactly the desired two zeros and satisfies $B_\theta(g_{\mathrm{pair}}^2) = 0$. On the other hand, the one-pole background

$$h_\theta = 1 - \cos\theta\, \nu_1$$

satisfies the exceptionally clean identities

$$(\Delta_M + |A|^2)h_\theta = |A|^2, \qquad B_\theta(h_\theta^2) = 0,$$

but it does not detect the two prescribed normals. Thus one candidate had the right zero set, while the other had the cleaner differential inequality.

We next tried to interpolate between these advantages through the geometric mixture

$$g = g_{\mathrm{pair}}^{\alpha} h_\theta^{1-\alpha}, \qquad 0 < \alpha < 1,$$

as well as truncated functions, power profiles, and profiles $p(\nu_1) = \exp(q(\nu_1))$ chosen to satisfy (2.4). Working initially in the acute-angle range $0 < \theta < \pi/2$, we tested, among others:

$$q(t) = -\frac{\cos\theta}{\sin^2\theta}t, \qquad q(t) = -\frac{t^a}{a\cos^{a-1}\theta\sin^2\theta}, \qquad q(t) = \log\left(1 - \frac{\cos^{2-b}\theta}{b\sin^2\theta + \cos^2\theta}t^b\right).$$

We also tested the direct power $1 - \cos^{2-a}\theta\, \nu_1^a$. The calculations explained why these formulas were not the final choice: the multiplicative mixtures created an unfavorable cross-gradient term, while the one-variable modifiers could not simultaneously provide the exact two-pole zero set and a uniform interior margin.

At this stage, hand calculations were repeatedly converted into symbolic and numerical tests. For a chosen profile and parameter, we varied the normal direction and the trace-free second fundamental form, then computed the smallest value of the normalized expression in (2.1). The experiments were especially useful for rejecting a profile when its minimum became negative and for identifying a small-parameter regime that was worth proving analytically.

2.3 The additive route to the final family

The addition rule for B_θ led directly to the final family. In addition to g_{pair}, define $D(v) = 1 - v_1^2$. The three building blocks are $D(v)$, $g_{\text{pair}}^2(v)$, and v_2^2. Using (2.3), a direct computation gives $B_\theta(D) = B_\theta(g_{\text{pair}}^2) = B_\theta(v_2^2) = 0$. Consequently, every function of the form

$$U_\alpha(v) = \alpha_1 D(v) + \alpha_2 g_{\text{pair}}^2(v) + \alpha_3 v_2^2$$

satisfies $B_\theta(U_\alpha) = 0$.

At either prescribed normal $p_\pm$, we have $D(p_\pm) = v_2^2(p_\pm) = \sin^2\theta$ and $g_{\text{pair}}(p_\pm) = 0$. Thus $U_\alpha(p_\pm) = 0$ precisely when $\alpha_1 + \alpha_3 = 0$. Taking $\alpha_3 = -\alpha_1$, and using

$$g_{\text{pair}}^2 = \sin^2\theta(1 - v_1^2 - v_2^2) + (\cos\theta - v_1)^2,$$

we obtain

$$U_\alpha(v) = (\alpha_1 + \alpha_2 \sin^2\theta)(1 - v_1^2 - v_2^2) + \alpha_2(\cos\theta - v_1)^2.$$

If $\alpha_2 > 0$ and $C_\alpha = \alpha_1 + \alpha_2 \sin^2\theta > 0$, this expression is nonnegative on the unit sphere and vanishes exactly when $v_1 = \cos\theta$ and $v_2 = \pm\sin\theta$, that is, precisely at p_+ and p_-. Dividing by the harmless positive constant C_α, the family can be written using the more transparent parameter $k = \frac{\alpha_2}{C_\alpha} > 0$ as

$$g_{\theta,k}^2 = 1 - v_1^2 - v_2^2 + k(\cos\theta - v_1)^2.$$

Its first part confines the normal to the two-plane containing $p_\pm$, but cannot distinguish points inside that plane. The second part selects the two prescribed normals, at the cost of introducing the difficult mixed terms. The parameter k controls their size: the margin left by the first part absorbs them when k is in the appropriate range.

We then tested this family numerically with very small values of k. The experiments suggested that a positive lower bound should indeed exist, although the observed margin decreased more rapidly as θ became small. This numerical evidence identified a plausible regime, but did not prove uniform positivity. The subsequent rigorous calculation verified analytically that the required positive lower bound exists.

2.4 The design blueprint from the capillary problem

The capillary construction suggested more than one successful formula. It provided a general architecture for the next stage of the search. At this point we regarded

$$S + kP \tag{2.5}$$

only as a prototype, with the following division of roles.

- *The background S.* This term should force the normal into a preferred low-dimensional set, usually a low-dimensional span. It should also provide the robust part of the differential inequality. By itself, however, it need not distinguish the prescribed normal directions inside that set.
- *The detector P.* This term should separate the individual prescribed directions left indistinguishable by S, and ultimately produce the desired finite zero set. Its derivatives are typically much harder to control than those of the background.
- *The small parameter k.* The coefficient of P should be chosen small enough that the complicated terms introduced by the detector can be absorbed by the margin supplied by S. The parameter is therefore part of the design, rather than a final cosmetic adjustment.
- *Freedom in the outer form.* The general problem no longer carries the capillary Robin boundary condition. We were therefore not restricted to taking the test function to be exactly a square root of $S + kP$. Powers such as $(S + kP)^\gamma$, product forms, normalized products, and related compositions were all legitimate candidates, provided that they retained the correct zero set and differential positivity.

This became the blueprint for the subsequent experiments: first choose a simple S that enforces low-dimensionality, then vary the detector P, the parameter k, and the outer power or product structure. Symbolic and numerical tests could then identify which versions preserved a positive margin and, equally importantly, where each proposed design failed.

3 From the geometric inequality to a numerical criterion

We selected orthonormal Euclidean directions $e_1, \dots, e_\ell$, set

$$x = (v_1, \dots, v_\ell)^{\mathsf{T}}, \qquad v_i = \langle v, e_i \rangle,$$

and began with an arbitrary positive function

$$g^2 = G(v_1, \dots, v_\ell) = G(x). \tag{3.1}$$

The aim was to check directly whether

$$|A|^2 g^2 + g\Delta_M g \geq c|A|^2$$

could hold with $c > 0$.

Since M is minimal, its Gauss-map components satisfy

$$\Delta_M v_i = -|A|^2 v_i, \qquad 1 \leq i \leq \ell.$$

Write

$$G_i = \frac{\partial G}{\partial v_i}, \qquad G_{ij} = \frac{\partial^2 G}{\partial v_i \partial v_j}.$$

Applying the chain rule to (3.1) gives

$$|A|^2 g^2 + g\Delta_M g = W|A|^2 - \sum_{i,j=1}^{\ell} \alpha_{ij} \langle \nabla v_i, \nabla v_j \rangle, \tag{3.2}$$

where

$$W = G - \frac{1}{2} \sum_{i=1}^{\ell} v_i G_i, \qquad \alpha_{ij} = -\frac{1}{2} G_{ij} + \frac{G_i G_j}{4G}. \tag{3.3}$$

The remaining task was to find a constant s such that

$$\sum_{i,j=1}^{\ell} \alpha_{ij} \langle \nabla v_i, \nabla v_j \rangle \leq s|A|^2.$$

The remaining term in (3.2) was organized in a frame adapted to the selected directions. Put

$$e_i^{\top} = e_i - v_i \nu, \qquad 1 \leq i \leq \ell,$$

and let

$$\pi = \operatorname{span}\{e_1^{\top}, \dots, e_\ell^{\top}\} \subset T_p M.$$

At an interior point $|x| < 1$, the vectors $e_1^{\top}, \dots, e_\ell^{\top}$ are linearly independent. To keep this exploratory reduction as simple as possible, we imposed the additional ansatz that π is spanned by principal-curvature directions. We may then choose an orthonormal principal frame $\tau_1, \dots, \tau_n$ such that $\pi = \operatorname{span}\{\tau_1, \dots, \tau_\ell\}$, and write

$$A(\tau_a, \tau_b) = \lambda_a \delta_{ab}, \qquad 1 \leq a, b \leq n. \tag{3.4}$$

Equation (3.4) gives

$$\nabla v_i = -\sum_{a=1}^{\ell} \lambda_a \langle e_i, \tau_a \rangle \tau_a, \qquad \sum_{a=1}^{\ell} \langle e_i, \tau_a \rangle \langle e_j, \tau_a \rangle = \delta_{ij} - v_i v_j. \tag{3.5}$$

Consequently,

$$\sum_{i,j=1}^{\ell} \alpha_{ij}\langle \nabla v_i, \nabla v_j\rangle = \sum_{a=1}^{\ell} B_{aa}\lambda_a^2,$$

where the coefficients required by the numerical calculation are given directly by

$$B_{aa} = -\frac{1}{2}\sum_{i,j=1}^{\ell} G_{ij}\langle e_i,\tau_a\rangle\langle e_j,\tau_a\rangle + \frac{1}{4G}\left(\sum_{i=1}^{\ell} G_i\langle e_i,\tau_a\rangle\right)^2. \tag{3.6}$$

For simplicity, assume $n > \ell$ and $\lambda_{\ell+1} = \cdots = \lambda_n$. By minimality, the sharp choice of s is the smallest number such that

$$\sum_{a=1}^{\ell} B_{aa}\lambda_a^2 \le s\left(\sum_{a=1}^{\ell}\lambda_a^2 + \frac{(\lambda_1+\cdots+\lambda_\ell)^2}{n-\ell}\right).$$

Define

$$p(t) = \prod_{a=1}^{\ell}(t - B_{aa}), \qquad P(t) = (n-\ell)p(t) + tp'(t). \tag{3.7}$$

The sharp value of s is the largest real root of P; thus we take

$$s = \max\{t \in \mathbb{R} : \ P(t) = 0\}. \tag{3.8}$$

Hence

$$|A|^2 g^2 + g\Delta_M g \ge (W - s)|A|^2.$$

Numerical verification routine

A numerical verification routine takes

$$x \in \{|x| \le 1 : \ G(x) > 0\}$$

and samples the coefficients $\langle e_i, \tau_a\rangle$ subject to the second identity in (3.5). It evaluates G, G_i, G_{ij}, constructs $W, \alpha, B_{aa}, p, P, s$ by (3.3)–(3.8), and computes

$$\mathfrak{m}_{n,\ell} = W - s.$$

The candidate passes the sampled test precisely when the sampled minimum of $\mathfrak{m}_{n,\ell}$ is positive.

Warning 3.1. *The numerical experiments were used only to identify which test functions were promising. Even when a candidate passed every sampled test, the required uniform inequality still had to be verified by rigorous mathematical calculation.*

4 The first computational phase: explicit formulas

The purpose of this phase was to determine a basic form for the test function, and especially to decide which choices of the background S and detector P in (2.5) were most effective. We focused first on $\ell = 2$, asking which functions $G(v_1, v_2)$ could survive numerical screening.

4.1 The first numerical candidates

Put $S = 1 - v_1^2 - v_2^2$. The first ten tested candidates are collected below. Most belong to the $\ell = 2$ search; the final two are the immediate higher-rank product variants tested in the same phase.

Table 4.1: The first ten candidate formulas and their numerical screening outcomes.

Candidate $G = g^2$	Numerical screening
$G = 1 - v_1^2 - v_2^2 + k(\cos\vartheta - v_1)^2$	✓ Passed
$G = S + k v_1^2 v_2^2$	✗ Failed
$G = S^2 + k(v_1 - v_2)^2(v_1 + v_2 - 1)^2$	✗ Failed
$G = \left[S^2 + k(v_1 - v_2)^2(v_1 + v_2 - 1)^2\right]^{1/2}$	✗ Failed
$G = \left[\prod_{\pm}\left(1 - \frac{\sqrt{3}}{2}v_1 \pm \frac{1}{2}v_2\right)(1 + v_2)\right]^{2/3}$	✓ Passed
$G = \left[S^{3/2} + k\prod_{\pm}\left(1 - \frac{\sqrt{3}}{2}v_1 \pm \frac{1}{2}v_2\right)(1 + v_2)\right]^{2/3}$	✗ Failed
$G = S + k(1 - v_1^2)(1 - v_2^2)$	✗ Failed
$G = \left[S^2 + k(1 - v_1^2)(1 - v_2^2)\right]^{1/2}$	✗ Failed
$G = \left[(1 - v_1)(1 - v_2)(1 - v_3)\right]^{2/3}$	✓ Passed
$G = \left[(1 - v_1^2)(1 - v_2^2)(1 - v_3^2)(1 - v_4^2)\right]^{1/4}$	✗ Failed

These numerical screenings ruled out many natural alternatives and led us to focus on candidates of the form $G = S + kP$, where $S = 1 - v_1^2 - v_2^2$ and P vanishes at the prescribed directions. In particular, several simple choices, including $P = v_1^2 v_2^2$ and $P = (1 - v_1^2)(1 - v_2^2)$, failed the numerical screening. This made clear that P had to be designed more carefully in order to obtain a viable candidate.

4.2 Searching over the detector P

The next round of experiments focus on $\ell = 3, 4$. Thus the background was

$$S = 1 - v_1^2 - v_2^2 - \cdots - v_\ell^2, \qquad G = S + kP.$$

We varied the structure of P, testing products of affine pole factors, fractional powers, repeated factors, and tilted factors.

Table 4.2: Further choices of P and their numerical screening status.

Candidate P	Numerical screening
Rank-three tests: $S = 1 - v_1^2 - v_2^2 - v_3^2, \quad t \in [0, 2\pi]$	
$P = (1 - v_1)(1 - v_2)$	✓ Passed
$P = \left[(1 - v_1)(1 - v_2)(1 - v_1/\sqrt{2} - v_3/\sqrt{2})\right]^{2/3}$	✗ Failed*
$P = \left[(1 - v_1)(1 - v_2)(1 - v_1\cos t - v_3\sin t)\right]^{2/3}$	✗ Failed*
$P = \left[(1 - v_1)^2(1 - v_2)\right]^{2/3}$	✗ Failed*
$P = \left[(1 - v_1)(1 - v_2)(1 - v_3)(1 - v_1\cos t - v_3\sin t)\right]^{1/2}$	✗ Failed*
$P = \left[(1 - v_1)^2(1 - v_2)(1 - v_3)\right]^{1/2}$	✗ Failed*
$P = \left[(1 - v_1)^2(1 - v_2)(1 - v_2\cos t - v_3\sin t)\right]^{1/2}$	✗ Failed*
$P = \left[\prod_{i=1}^{2}(1 - v_i)(1 - v_i\cos t - v_3\sin t)\right]^{1/2}$	✗ Failed*
$P = \left[(1 - v_1^2)(1 - v_2^2)(1 - v_3^2)\right]^{1/3}$	✗ Failed*
Additional $n = 6$ tests: $S = 1 - v_1^2 - v_2^2 - v_3^2 - v_4^2$	
$P = \frac{1}{4}\left[(1 - v_1)(1 - v_2)(1 - v_3)(1 + v_3)\right]^2$	✗ Failed
$P = (1 - v_1)(1 - v_2)(1 - v_3)$	✗ Failed

Candidate P	Numerical screening
$P = (1 - v_1)(1 - v_2)$	✗ Failed

Interpreting the starred failures. Every formula marked "Failed*" failed for the same localized reason: negative values were detected only in thin neighborhoods of one or more prescribed directions. On the sampled regions bounded away from the zeros, the numerical margin remained positive. As k decreased, the bad neighborhoods became smaller and were therefore easy to miss at fixed numerical resolution. An unstarred "Failed" means that negative values were already visible in the retained general scan, rather than only in a zero neighborhood.

Thus the experiments were almost entirely negative, but the location of the failure was encouraging: away from the zeros, the starred candidates retained a positive margin. This suggested that the global product structure might still be useful and that only the local vanishing model had to be repaired. Even a sheeting theorem for three arbitrary prescribed directions would already have been a worthwhile result, but at this stage **no choice** of P passed all the relevant tests.

A second pattern identified the exponent itself. Every starred candidate shared the balanced power structure

$$P = \prod_{j=1}^{m} L_j^{2/m}, \qquad L_j = 1 - \langle v, p_j \rangle,$$

with repeated or tilted factors allowed.

Observation 4.1 (Why the balanced exponent stood out)**.** *Beyond the formulas listed in the table, we numerically tested many products with other exponents and several non-power modifications. These additional candidates developed negative values away from the zero set. By contrast, the observed failures of the balanced $2/m$-power products were confined to shrinking neighborhoods of the zeros.*

The data now pointed both to a candidate exponent and to a rank restriction. While the balanced products behaved consistently in ranks two and three, the rank-four experiments were markedly less stable: new negative directions appeared away from the zeros, and neither decreasing k nor adding more factors removed them. We therefore restricted the next stage of the search to $\ell \leq 3$ and replaced the open-ended search for new formulas by an analysis of the necessary conditions on P. In particular, we examined the regime $k \downarrow 0$ to determine whether the exponent $2/m$ was indeed sufficient on every compact set away from the zeros, to understand why the bad neighborhoods shrank, and to isolate the additional local condition that would eventually be needed at a zero. This change of viewpoint led directly to the next section.

5 The decisive small-parameter calculation

The next phase no longer searched over unrelated formulas. We fixed

$$P(x) = \prod_{j=1}^{m} L_j^{2/m}, \qquad L_j = 1 - \langle x, p_j \rangle, \tag{5.1}$$

and asked whether this balanced power was always viable away from its zeros when k was sufficiently small.

5.1 The small-k first-order expansion

Let $S = 1 - v_1^2 - \cdots - v_\ell^2$. At a fixed point with $S > 0$, direct expansion of (3.3) and (3.6) gives

$$W = 1 + kw, \qquad w = P - \frac{1}{2}\sum_{i=1}^{\ell} v_i P_i, \qquad B_{aa} = 1 + k\mu_a + O(k^2), \tag{5.2}$$

where

$$\mu_a = -\frac{1}{2}\sum_{i,j=1}^{\ell} P_{ij}\langle e_i, \tau_a\rangle\langle e_j, \tau_a\rangle - \frac{\left(\sum_i v_i\langle e_i, \tau_a\rangle\right)\left(\sum_j P_j\langle e_j, \tau_a\rangle\right)}{S} - \frac{P\left(\sum_i v_i\langle e_i, \tau_a\rangle\right)^2}{S^2}.$$

These were precisely the expressions implemented in the first-order numerical code: the input was P, its first two derivatives, the point x, and the frame coefficients $\langle e_i, \tau_a\rangle$.

Put

$$p_0(t) = \prod_{a=1}^{\ell}(t - \mu_a).$$

Substitution of (5.2) into the threshold polynomial (3.7) shows that

$$s = 1 + k\rho_+ + O(k^2), \qquad \rho_+ \text{ is the largest real root of } p_0'. \tag{5.3}$$

The ambient dimension n disappears from this first-order root; only the active rank ℓ remains. Define the gap coefficient

$$\varrho = w - \rho_+.$$

Then

$$W - s = k\varrho + O(k^2), \qquad \frac{s}{W} = 1 - k\varrho + O(k^2). \tag{5.4}$$

Consequently, the negative first-order term occurs in the normalized ratio s/W, whereas the unnormalized threshold in (5.3) has the positive first-order expansion. The clean sufficient condition is

$$\varrho > 0.$$

5.2 The first-order criteria in ranks two and three

For $\ell = 2$, equation (5.3) immediately gives

$$\rho_+ = \frac{\mu_1 + \mu_2}{2}, \qquad \varrho = w - \frac{\mu_1 + \mu_2}{2}.$$

Equivalently, the first-order test is simply

$$2W - (B_{11} + B_{22}) = 2k\varrho + O(k^2) > 0.$$

For the balanced product (5.1), direct substitution and a Cauchy–Schwarz estimate give $\varrho > 0$ away from the zeros. Thus the balanced exponent $2/m$ passes the rank-two first-order test.

For $\ell = 3$, put

$$\sigma_1 = \mu_1 + \mu_2 + \mu_3, \qquad \sigma_2 = \mu_1\mu_2 + \mu_1\mu_3 + \mu_2\mu_3.$$

The critical-root equation becomes

$$3\rho^2 - 2\sigma_1\rho + \sigma_2 = 0, \qquad \rho_+ = \frac{\sigma_1 + \sqrt{\sigma_1^2 - 3\sigma_2}}{3}. \tag{5.5}$$

Accordingly, $\varrho > 0$ is equivalent to placing w to the right of the larger root. A convenient pair of scalar checks is

$$3w - \sigma_1 > 0, \qquad 3w^2 - 2\sigma_1 w + \sigma_2 > 0. \tag{5.6}$$

This was the form used in the rank-three tests.

5.3 Rank-three scans and analytic confirmation

The numerical work tested (5.6) for coordinate, symmetric, clustered, tilted, and random prescribed directions. Every entry in Table 5.1 concerns the first-order gap

$$\varrho = \partial_k (W - s)|_{k=0^+} = w - \rho_+.$$

For each listed configuration, we numerically evaluated the normalized boundary gap ϱ/P and the first-order gap ϱ. The randomized rows record separate scans for $m = 2, \dots, 5$.

Table 5.1. Rank-three first-order scans for the balanced power product.

Pole configuration	m	$2/m$	Boundary $\min(\varrho/P)$	Sampled $\min \varrho$
Coordinate triple	3	2/3	0.5006385	0.0999129 ✓
Tetrahedral four	4	1/2	0.5000182	0.3344756 ✓
Octahedral six	6	1/3	0.5000031	0.3157721 ✓
Random four	4	1/2	0.5009658	0.0426013 ✓
Random six	6	1/3	0.5000017	0.0427318 ✓
Random eight	8	1/4	0.5010803	0.1085162 ✓
Randomized two-pole (60 tries)	2	1	—	2.41×10^{-1} ✓
Randomized three-pole (60 tries)	3	2/3	—	2.06×10^{-1} ✓
Randomized four-pole (60 tries)	4	1/2	—	2.28×10^{-1} ✓
Randomized five-pole (60 tries)	5	2/5	—	2.92×10^{-1} ✓

All four randomized families passed the first-order scans. The values in the table were numerical evidence only, but together they supported the stability of the balanced $2/m$ product away from the zero set. The computation also showed why samples taken too close to a zero could be misleading at fixed k: the first-order expansion is not uniform there.

After these scans, a detailed analytic calculation reduced (5.6) to a covariance estimate. It showed that w lies strictly to the right of the larger root in (5.5) for every set of distinct prescribed directions. Consequently, ϱ has a positive minimum on every compact set separated from the zeros. Continuity and (5.4) then give

$$W - s > 0$$

there whenever $k > 0$ is chosen sufficiently small.

This completed the away-from-zero part for $\ell \le 3$, but it did not solve the original problem. Near a prescribed direction,

$$P(x) \sim C_j t^{2/m}, \qquad t = 1 - \langle x, p_j \rangle, \qquad C_j > 0,$$

and for every fixed positive k the uniform lower bound could still fail in a thin neighborhood of that zero. Decreasing k shrank the bad region but never removed it. The remaining task was therefore sharply defined: keep the successful $2/m$ tail away from the zeros and replace its local vanishing model. This led to the linear calculation in the next section.

6 The local repair and the final profile

6.1 The linear model near a zero

Once the balanced $2/m$-power had been shown to work away from its zeros, the remaining question was local. Near the prescribed direction p_1, put

$$t = 1 - \langle x, p_1 \rangle, \qquad P(x) = t^{\alpha} P_{\mathrm{reg}}(x), \qquad P_{\mathrm{reg}}(p_1) > 0.$$

The other pole factors are smooth and positive there, so the leading behavior is determined by the single exponent α. This reduced the first question to a direct one-pole calculation: which powers retain a positive first-order gap as $t \downarrow 0$?

The fixed-k one-pole calculation ruled out exponents that were too small. If $C = P_{\mathrm{reg}}(p_1) > 0$, its leading term is

$$W - s = Ct^{\alpha-1}\left(\frac{\alpha\big((n-1)\alpha - (n-2)\big)}{2n} + O(t)\right). \tag{6.1}$$

Thus a strictly positive leading margin requires

$$\alpha > \frac{n-2}{n-1}.$$

The natural dimension-independent choice is therefore $\alpha = 1$, corresponding to the linear model

$$\phi(t) \sim ct \qquad (t \downarrow 0).$$

A direct substitution of this model restores the positive near-zero margin that was lost by the pure $t^{2/m}$ power.

6.2 The logarithmic transition and the rank threshold

The next question was whether the local exponent 1 could be connected to the successful outer exponent $2/m$. This is a different $k \downarrow 0$ calculation from the fixed-k zero model above. An arbitrary smooth cutoff need not work, because its second derivative appears in the first-order threshold. For a positive profile ϕ, we therefore introduced its logarithmic slope and curvature

$$\xi = \frac{\mathrm{d}\log\phi}{\mathrm{d}\log t}, \qquad \eta = \frac{\mathrm{d}\xi}{\mathrm{d}\log t}.$$

Thus $\xi = 1$ is the linear model, $\xi = a$ is a power t^a, and $\xi = 0$ is a plateau. Applying the small-parameter calculation to the one-pole transition gives a second-order differential inequality for (ξ, η). Here the ambient dimension n disappears; only the active rank ℓ remains. The two endpoint regimes give the necessary conditions

$$\eta > -\xi^2 + \frac{3\ell-4}{2(\ell-1)}\xi - \frac{1}{2} \quad (\ell \geq 2), \qquad \eta > \frac{\ell-3}{2(\ell-2)}\xi - \frac{1}{2}\xi^2 \qquad (\ell \geq 3).$$

Setting $\eta = 0$ shows that a terminal power t^a must satisfy

$$a > \frac{\ell-3}{\ell-2} \qquad (\ell \geq 3). \tag{6.2}$$

The rank dependence became transparent in this ODE.

(1) *Rank two.* The radial condition, which is the full logarithmic condition in rank two, becomes

$$\eta + \xi^2 - \xi + \frac{1}{2} > 0.$$

It allows ξ to decrease from 1 all the way to 0 in a finite logarithmic interval. After rescaling, this transition can be placed in an arbitrarily small pole cap. Outside the caps every factor is constant, so P itself is constant and the rank-two construction becomes especially simple.

(2) *Rank three.* A convenient sufficient condition for the full one-pole ODE is

$$0 < \xi \leq 1, \qquad \eta \geq -c\xi^2, \qquad 0 < c < \frac{7}{32}.$$

This condition prevents ξ from reaching 0 in finite logarithmic time, explaining the failure of the earlier linear-to-plateau attempts. It does, however, allow ξ to decrease from 1 to any fixed positive exponent, in particular to $2/m$. The logarithmic transition has a fixed finite length once m is fixed; shrinking its scale places the entire transition in an arbitrarily small physical neighborhood of the prescribed direction.

(3) *Rank four.* Condition (6.2) requires a terminal exponent strictly greater than $1/2$. Thus the balanced exponent $2/m$ has zero limiting margin when $m = 4$, and it lies below the admissible range when $m > 4$. This local obstruction already shows why the same test-function design encounters a genuine difficulty in a rank-four sheeting problem.

6.3 The final profile and its verification

For the rank-three construction, choose a small scale $T > 0$ and a smooth nondecreasing profile such that

$$\phi_T(t) = \begin{cases} T^{2/m-1}t, & 0 \le t \le T, \\ C_m t^{2/m}, & t \ge e^L T, \end{cases},$$

where $L < \infty$ is the logarithmic length of the chosen transition. It may be arranged that throughout the transition

$$\frac{2}{m} \le \xi \le 1, \qquad \eta \ge -\frac{1}{8}\xi^2.$$

The resulting multipole tilt function is

$$G_{\mathbf{p}}(v) = \left(1 - v_1^2 - v_2^2 - v_3^2 + k\prod_{j=1}^{m} \phi_T\big(1 - \langle v, p_j\rangle\big)\right)^{1/2}.$$

To guard against floating-point error in the very small sampled minima, all entries in the last column were computed using 90-digit arithmetic.

TABLE 6.1. Finite-k numerical verification of the final power-tail profile.

Pole configuration	m	T	k	**Sampled** $\min(W - s)$
Coordinate triple	3	1.35×10^{-4}	5.53×10^{-5}	4.52×10^{-5} ✓
Tetrahedral four	4	2.47×10^{-6}	1.20×10^{-8}	5.24×10^{-6} ✓
Five-point cluster	5	3.90×10^{-11}	4.10×10^{-15}	1.17×10^{-16} ✓
Five points with a close pair	5	1.77×10^{-12}	6.40×10^{-16}	6.39×10^{-11} ✓
Octahedral six	6	8.28×10^{-10}	1.20×10^{-18}	2.78×10^{-9} ✓

All sampled margins were positive. This supported the candidate without proving a uniform lower bound, so the rigorous proof used three regions.

- *Near a zero*, freeze the nonvanishing factors and use the direct linear one-pole calculation.
- *Away from the zeros*, use the strict $k \downarrow 0$ gap from the previous section.
- *In a transition annulus*, rescale and use the logarithmic ODE; this is laborious but no longer exploratory.

Fix T first to control the pole caps and transition errors, then choose $kT^{2/m-1} \ll 1$.

7 Summary and outlook

The construction developed through four successive ideas: the capillary prototype suggested a coercive background plus a small detector; numerical screening isolated the balanced exponent $2/m$; the one-pole calculation forced a linear model at each zero; and the logarithmic ODE supplied the transition between these two regimes. The computations guided the search, while the final conclusions still required rigorous mathematical proof.

AI-assisted symbolic and numerical tools greatly accelerated the testing, comparison, and correction of candidate formulas. We expect increasingly capable AI systems to contribute even more to mathematical research—not by replacing mathematical understanding, but by helping us explore ideas more quickly and deepen that understanding.