AN ASSERTED INEQUALITY, PROVED AND SHARPENED: THE EXACT MINIMUM OF THE THRESHOLD-ROUNDING RATIO OF ELBASSIONI AND RAY

KOREA SUPERINTELLIGENCE LABS

ABSTRACT. In their analysis of a threshold-rounding algorithm for stabbing axis-parallel segments, Elbassioni and Ray exhibit a family of fractional solutions of the standard LP relaxation whose rounding cost ratio tends to an explicit two-variable function $\gamma(\tau_x, \tau_y)$ of the two rounding thresholds, and assert that $\gamma \geq 1.89$ everywhere, with the sentence: *"We skip the technical proof since it does not yield much insight and since this can be easily checked numerically."* The assertion stands unproved in the 2021 preprint, and Appendix B of the 2024 journal version repeats it word for word, with the thresholds renamed. We prove it, and we determine the constant it approximates. The minimum of γ over $[0,1)^2$ is attained at $(t^*, \frac{1}{2})$ and at $(\frac{1}{2}, t^*)$, where t^* is the unique root in $[\frac{7}{20}, \frac{2}{5}]$ of the integer quartic $q(t) = 6000t^4 - 16040t^3 + 14600t^2 - 5080t + 577$; and that minimum lies between 1.891152196075811250 and 1.891152196075811251. So the published 1.89 is true, and short of the truth by $1.152\ldots \times 10^{-3}$. The proof of the inequality is a twenty-five-box Bernstein certificate in exact integer arithmetic, with the certified coefficient arrays computed from the source's own density rather than stored. The exact minimum needs neither a derivative nor any arithmetic in the quartic field: two divided differences turn it into polynomial identities and a handful of further rational positivity certificates, and one of those identities makes the resulting bracket for the minimum quadratically accurate, which is where the eighteen decimal places come from. Every statement below is machine-checked in Lean 4.

1. INTRODUCTION

1.1. Where the function comes from. In the *rectangle stabbing* problem one is given a set $\mathcal{R}$ of axis-aligned rectangles in $\mathbb{R}^2$, sets $\mathcal{H}$, $\mathcal{V}$ of horizontal and vertical lines and nonnegative weights, and must choose a minimum-weight set of lines meeting every rectangle. Its standard LP relaxation assigns a variable to each line and demands that the variables of the lines stabbing a given rectangle sum to at least 1. Gaur, Ibaraki and Krishnamurti [5] gave a 2-approximation algorithm for it, and Ben-David, Grant, Ma and Sharpe [6] showed that 2 cannot be improved: their Theorem 2 is an LP-relative d-approximation for hitting families of d-intervals, and their Theorem 3 produces, for every $\varepsilon > 0$, families of d-union-intervals whose gap is at least $d - \varepsilon$; rectangle stabbing is their case $d = 2$. For the special case in which every rectangle meets exactly one horizontal line — *interval stabbing* — Kovaleva and Spieksma [4] pinned the gap at $e/(e-1) = 1.582\ldots$: their algorithm STAB certifies the upper bound (their Theorem 3.4), and an explicit family of instances brings the gap arbitrarily close to it, though never to it (their Theorem 3.6).

STAB is not itself a rounding by a random threshold: it sorts the row variables of an optimal LP solution, rounds a prefix of them up, solves the residual column problem exactly, and keeps the best of the $m+1$ outcomes. It is Elbassioni and Ray [1, 2] who recast the upper bound as *threshold rounding* — round a variable up exactly when it reaches a threshold τ drawn at random, and pay for what is left with a net. The preprint offers this as "a slightly simpler (and probably more intuitive) way to see the result in" [4], and it is in that form that they extend it to two more special cases, among them SEGSTAB, where $\mathcal{R}$ consists of horizontal *and* vertical segments. Their algorithm fixes two thresholds $\tau_x, \tau_y \in (0, 1)$, rounds up the vertical variables that reach τ_x and the horizontal variables that reach τ_y, and covers the remainder with a $(1-\tau_y)$-net and a $(1-\tau_x)$-net. The analysis gives an integrality gap of at most 1.935 for SEGSTAB.

Their paper then asks how far this analysis can be pushed, and answers with an explicit witness. Under the paragraph heading *Limitation of the analysis* they fix the density

$$f(t) = \begin{cases} \frac{1}{2}, & 0 < t < \frac{1}{5}, \\ \frac{75}{4}t - \frac{13}{4}, & \frac{1}{5} \le t < \frac{2}{5}, \\ \frac{17}{4}, & \frac{2}{5} \le t < \frac{1}{2}, \\ 0, & \text{otherwise}, \end{cases} \tag{1}$$

build a fractional solution whose coordinates are distributed according to f, and observe that the ratio of the weight of the algorithm's output to the LP value can be made arbitrarily close to

$$\gamma(\tau_x, \tau_y) = \frac{1}{2\int_0^1 tf(t)\,dt}\left[\int_{\tau_x}^1 f + \int_{\tau_y}^1 f + \frac{1}{1-\tau_x}\int_0^{\tau_y} tf(t)\,dt + \frac{1}{1-\tau_y}\int_0^{\tau_x} tf(t)\,dt\right]. \tag{2}$$

The crossed pairing — $1/(1-\tau_x)$ multiplying the τ_y-integral — is theirs, and is what the displayed weight of the rounded solution produces; γ is symmetric all the same. They then write, verbatim:

> "It can be shown that $\gamma(\tau_x, \tau_y) \ge 1.89$ for any $\tau_x, \tau_y \in [0, 1]$. We skip the technical proof since it does not yield much insight and since this can be easily checked numerically."

The journal version [2] — submitted in December 2021, published online in November 2023 and carried in the April 2024 issue — reproduces the same construction and the same two sentences in its Appendix B, with the thresholds renamed τ_1, τ_2. Neither the preprint nor that appendix states the true minimum of γ, and neither proves the inequality; the provenance note below says exactly how much of the journal version was read here.

1.2. **Results.** This paper supplies the proof and the constant. Throughout, $F(\tau) = \int_\tau^1 f$ and $G(\tau) = \int_0^\tau tf(t)\,dt$ for the density (1); then $\int_0^1 f = 1$, so f is a genuine probability density, and $\int_0^1 tf(t)\,dt = 57/160$, so the normaliser in (2) is 57/80 (Proposition 2.1). We work on $[0, 1)^2$: at $\tau = 1$ the expression (2) has a vanishing denominator, and Proposition 6.1 (iii) shows that this restriction is not cosmetic.

- **The asserted inequality** (Theorem 3.2): $\gamma(\tau_x, \tau_y) \ge \frac{189}{100}$ for all $\tau_x, \tau_y \in [0, 1)$.
- **The exact minimiser** (Theorems 4.2 and 4.8): the integer quartic

  $$q(t) = 6000t^4 - 16040t^3 + 14600t^2 - 5080t + 577$$

  has exactly one root t^* in $[\frac{7}{20}, \frac{2}{5}]$, and $\gamma(t^*, \frac{1}{2}) \le \gamma(\tau_x, \tau_y)$ for all $\tau_x, \tau_y \in [0, 1)$, with $\gamma(t^*, \frac{1}{2}) = \gamma(\frac{1}{2}, t^*)$. So the minimum exists, is attained, and is an algebraic number of degree at most 4 over $\mathbb{Q}$. Neither q, nor t^*, nor this value appears in [1], nor in the appendix of [2] that makes the claim; see the provenance note below for how far that reading goes.
- **The constant, to eighteen places** (Theorem 5.4):

  $$1.891152196075811250 \;\le\; \min_{[0,1)^2} \gamma \;\le\; 1.891152196075811251.$$

  So the published 1.89 is true, and short of the truth by about 1.15×10^{-3} (Corollary 5.5). The witness family of [1] obstructs this analysis at exactly $1.8911521960758112\ldots$, and at no larger constant.

1.3. **What is not claimed.** Three limits deserve to be stated at the start rather than inferred.

First, and this is the source's own caveat, $\gamma \ge 1.89$ is *not* a lower bound on the integrality gap of the LP relaxation, nor on the approximation factor of the rounding algorithm: the fractional vector built from f need not be an optimal — or even a feasible — LP solution of any instance. It bounds what this *method of analysis* can certify. The best lower bound on the gap for SegStab recorded in [1] remains $e/(e-1)$ [4]. Sharpening 1.89 to $1.8911521960758112\ldots$ therefore closes a gap in the analysis of [1, 2], not in the approximability of the problem.

Second, we say nothing about the headline constants of [1, 2] — the 1.935 for SegStab, and the unit-square constant, given as $1.98\overline{3}$ in [1] with an appendix improving it to $1.9\overline{3}$ and quoted as 29/15 for congruent-square stabbing in [3] — nor about the separate conjecture in the unit-square section of [1], that the relevant dual LP optimum is at most $\gamma^*(\infty) = (35-\sqrt{5})/20 \approx 1.6382$ for every k and every family of independent sets; that is an open conjecture rather than a declined proof, and it is not a bounded finite computation.

Third, on the provenance of the source text. The preprint [1] was read here from the live `arXiv` e-print of v1 (a single LaTeX file); as of 3 September 2026 that record still shows one submission, no journal reference and no related DOI, so the quoted sentences are the author-supplied text and not an editorial artefact. Two harmless typographical slips in it are worth recording, because a reader reconstructing f will meet them: in the display of $f(t)$ the case conditions are printed in the variable x although the function is written in t, and the third integral in the display of γ is written $\int_0^{\tau_y} t f(t)\,dx$. The journal version [2] was read on 3 September 2026 from the publisher's own article page, which serves without a subscription exactly the abstract, the endnotes, the reference list and Appendices A, B and C, and replaces the numbered main body by a purchase notice. So Appendix B was read in full and directly: it carries the density with the constants 0.5, $18.75x - 3.25$, 4.25; it carries the two sentences quoted above with τ_x, τ_y renamed τ_1, τ_2; the dx-for-dt slip survives into it, while the $f(t)$-versus-x mismatch is repaired there; and it states no minimum, exact or approximate, beyond the rounded 1.89. The numbered main body was *not* read. A proof hidden there is therefore not excluded, although an appendix that both states the claim and declines its proof would be an odd thing to leave standing beside one. The bibliographic data of [4], [5] and [6] were checked against Crossref, dblp and zbMATH and, for [6], against the official CCCG 2012 electronic proceedings; [4] was read in full from the publisher's typeset version, and the statements attributed to [4] and [6] above are theirs, not paraphrases taken from [1].

1.4. **Method.** Both halves of the argument are exact rational computations of a shape that a proof assistant can carry, and neither uses floating-point arithmetic anywhere.

For the inequality, clearing the two positive denominators $1-\tau_x$, $1-\tau_y$ and the constant normaliser turns $\gamma \geq c$ into a polynomial inequality. The breakpoints of F and G — refined once, at 7/20, to isolate the minimiser — cut $[0,1]^2$ into twenty-five boxes, on each of which F and G are single polynomials, and on each of which the cleared inequality carries a Bernstein positivity certificate. Nothing is stored: the certified coefficient array of a box is *computed* from the piece polynomials inside the proof by integer arithmetic, so there is no transcription step between the density (1) and the checked data. The device that makes this cheap is the classical one — a polynomial of bidegree (m,n) has a bihomogenisation whose monomial coefficients are its Bernstein coefficients times binomials, so "all coefficients are nonnegative" needs no basis-conversion lemma and no de Casteljau step. In a proof assistant this idea goes back to Muñoz and Narkawicz [7].

For the exact minimum the obvious route — differentiate, solve, evaluate — would need sign decisions for elements of the quartic field $\mathbb{Q}(t^*)$. Two divided differences remove the field entirely. Writing $g(\tau) = \gamma(\tau, \frac{1}{2})$, the identity

$$114(1-\tau)(1-\sigma)\big(g(\tau) - g(\sigma)\big) = (\tau - \sigma)\,q(\sigma) + 160(\tau-\sigma)^2(1-\sigma)\,B(\tau,\sigma)$$

holds identically on the middle piece for an explicit quadratic B, and it is proved by expanding both sides. At a root of q the first term vanishes and the right-hand side becomes a perfect square times B, so $g(\tau) \geq g(t^*)$ follows from a single certificate for $B \geq 0$; and because the difference is *quadratic* in $\tau - t^*$, a rational bracket of width 10^{-12} for t^* pins $g(t^*)$ to within about 8×10^{-24}, which is where the eighteen printed places come from. The first divided difference of q itself, handled the same way, gives the root uniquely. No derivative is taken anywhere in the formal development.

2. The objects

Let f be the density (1), extended by $f(t) = 0$ for $t \leq 0$ and $t \geq \frac{1}{2}$, and put

$$F(\tau) = \int_\tau^1 f(t)\,dt, \qquad G(\tau) = \int_0^\tau t f(t)\,dt.$$

Proposition 2.1 (F, G and the normaliser)**.** *For every $\tau \in [0,1]$,*

$$F(\tau)=\begin{cases}1-\frac{\tau}{2}, & 0\le\tau\le\frac15,\\ \frac58+\frac{13}{4}\tau-\frac{75}{8}\tau^2, & \frac15\le\tau\le\frac25,\\ \frac{17}{8}-\frac{17}{4}\tau, & \frac25\le\tau\le\frac12,\\ 0, & \frac12\le\tau\le1,\end{cases}\qquad G(\tau)=\begin{cases}\frac{\tau^2}{4}, & 0\le\tau\le\frac15,\\ \frac{1}{40}+\frac{25}{4}\tau^3-\frac{13}{8}\tau^2, & \frac15\le\tau\le\frac25,\\ \frac{17}{8}\tau^2-\frac{7}{40}, & \frac25\le\tau\le\frac12,\\ \frac{57}{160}, & \frac12\le\tau\le1.\end{cases}$$

In particular $\int_0^1 f=1$, so f is a probability density as the source intends, and $\int_0^1 tf(t)\,dt=\frac{57}{160}$, so the normaliser $2\int_0^1 tf$ of (2) equals $\frac{57}{80}$.

The two lists overlap at the breakpoints because F and G are continuous there; the shared values are $\frac{9}{10},\frac{17}{40},0$ for F and $\frac{1}{100},\frac{33}{200},\frac{57}{160}$ for G. That continuity is what allows every argument below to work with *closed* boxes that share their boundaries, instead of half-open boxes and a separate boundary analysis.

Definition 2.2. For $\tau_x,\tau_y\in[0,1)$ put

$$\gamma(\tau_x,\tau_y)=\frac{80}{57}\left[F(\tau_x)+F(\tau_y)+\frac{G(\tau_y)}{1-\tau_x}+\frac{G(\tau_x)}{1-\tau_y}\right].$$

By Proposition 2.1 this is exactly (2). Clearing denominators in Definition 2.2 gives $\gamma(\tau_x,\tau_y)=\gamma(\tau_y,\tau_x)$ at once, and $F,G\ge0$ on $[0,1]$.

Everything below uses one further piece of bookkeeping. All the division points that occur — 0, $\frac15$, $\frac{7}{20}$, $\frac25$, $\frac12$, 1 — are multiples of $\frac{1}{20}$, so we write $X=20\tau$ and record the five intervals

$$J_0=[0,\tfrac15],\quad J_1=[\tfrac15,\tfrac{7}{20}],\quad J_2=[\tfrac{7}{20},\tfrac25],\quad J_3=[\tfrac25,\tfrac12],\quad J_4=[\tfrac12,1],$$

i.e. $X\in[0,4],[4,7],[7,8],[8,10],[10,20]$. On each of them

$$\Phi(X)=3200\,F(X/20),\qquad \Gamma(X)=1280000\,G(X/20) \tag{3}$$

are polynomials with *integer* coefficients: $\Phi=3200-80X$ and $\Gamma=800X^2$ on J_0; $\Phi=2000+520X-75X^2$ and $\Gamma=32000-5200X^2+1000X^3$ on $J_1\cup J_2$; $\Phi=6800-680X$ and $\Gamma=-224000+6800X^2$ on J_3; $\Phi=0$ and $\Gamma=456000$ on J_4. The extra breakpoint $\frac{7}{20}$ splits the second piece of F and G into J_1 and J_2 without changing the formulas; it exists to isolate the minimiser, and Proposition 6.1 (v) records that omitting it makes the certificate below fail.

3. The inequality that was asserted

3.1. **Clearing and boxing.** Fix integers N,D with $D>0$. For $\tau_x,\tau_y\in[0,1)$ both $1-\tau_x$ and $1-\tau_y$ are positive, so $\gamma(\tau_x,\tau_y)\ge N/D$ is equivalent to the polynomial inequality $T_{N,D}\ge0$, where

$$T_{N,D}(\tau_x,\tau_y)=18240000\,D\,(1-\tau_x)(1-\tau_y)\Big(\gamma(\tau_x,\tau_y)-\tfrac{N}{D}\Big). \tag{4}$$

In the integer variables $X=20\tau_x$, $Y=20\tau_y$, and with Φ,Γ as in (3), the quantity (4) equals the integer polynomial

(5)

$$D\Big[20\big(\Phi_i(X)+\Phi_j(Y)\big)(20-X)(20-Y)+\Gamma_j(Y)(20-Y)+\Gamma_i(X)(20-X)\Big]-45600\,N\,(20-X)(20-Y),$$

where the subscripts i,j name the pieces containing X and Y. Its bidegree in (X,Y) is at most $(4,4)$.

On the box $J_i\times J_j$, with J_i corresponding to $X\in[a_i,b_i]$ and J_j to $Y\in[a_j,b_j]$, substitute the affine box coordinates

$$X=a_i+(b_i-a_i)s,\qquad Y=a_j+(b_j-a_j)u,\qquad (s,u)\in[0,1]^2,$$

and call the resulting integer polynomial P_{ij}.

3.2. **The certificate.** For a polynomial P of bidegree at most (m,n) let

$$\widehat{P}(s_0,s_1,u_0,u_1) = (s_0+s_1)^m(u_0+u_1)^n\, P\left(\frac{s_0}{s_0+s_1}, \frac{u_0}{u_0+u_1}\right)$$

be its bihomogenisation at bidegree (m,n): a form of bidegree (m,n) in the two pairs of variables, with at most $(m+1)(n+1)$ monomials, whose coefficient at $s_0^k s_1^{m-k} u_0^l u_1^{n-l}$ is $\binom{m}{k}\binom{n}{l}$ times the (k,l)-th Bernstein coefficient of P at bidegree (m,n). Two facts are all that is used: $\widehat{P}(s,1-s,u,1-u) = P(s,u)$, and every monomial of $\widehat{P}$ is nonnegative when $s,u \in [0,1]$. Hence

Lemma 3.1 (the certificate rule)**.** *If every coefficient of $\widehat{P}$ at some bidegree (m,n) is nonnegative, then $P \geq 0$ on $[0,1]^2$.*

Raising (m,n) is degree elevation, and it only improves the test. All certificates below are taken at bidegree $(8,8)$ — for the box polynomials this is elevation by $(4,4)$ over the exact bidegree of (5), and for the auxiliary certificates of Section 4 it is more. A certificate is thus a list of at most 81 integers.

Theorem 3.2 (the inequality of [1, 2])**.** *For all $\tau_x, \tau_y \in [0,1)$,*

$$\gamma(\tau_x,\tau_y) \;\geq\; \frac{189}{100} \;=\; 1.89.$$

Proof sketch. Every $\tau \in [0,1]$ lies in one of $J_0,\dots,J_4$, and on J_i the formulas (3) compute F and G; this is where the continuity of F and G at the breakpoints is used, since the intervals share endpoints. So it suffices to prove $T_{189,100} \geq 0$ on each of the $5 \times 5 = 25$ boxes $J_i \times J_j$, and by Lemma 3.1 it suffices that each $\widehat{P_{ij}}$ at bidegree $(8,8)$ have no negative coefficient.

This last step, and only this step, rests on exhaustive computation. The finite search is: for each of the 25 ordered pairs (i,j), form the integer polynomial (5) with $N = 189$, $D = 100$ from the piece data, substitute the affine box coordinates, bihomogenise to bidegree $(8,8)$, and test the resulting list of at most 81 integers for a negative entry. All 25 lists pass, and that — the absence of a negative entry, nothing finer — is the whole of what the formal development checks. Looking at the arrays themselves, which is done here outside that development, the smallest coefficient occurring anywhere in the 25 is exactly 0, at the corner $(\tau_x,\tau_y) = (0,1)$ of the box $J_0 \times J_4$, where the factor $(1-\tau_y)$ in (4) vanishes and the cleared inequality is trivially tight; every other coefficient is strictly positive. No coefficient list is supplied as input: each is computed from (3) inside the proof. □

Remark 3.3 (why the partition is refined). The five-interval partition is not the natural one — the natural one has four intervals, cut at $\frac{1}{5}, \frac{2}{5}, \frac{1}{2}$. At the same bidegree $(8,8)$, the four-interval version of the same test *fails*, on four of its sixteen boxes; the breakpoint $\frac{7}{20}$ is what repairs it, by separating a neighbourhood of the minimiser $t^* = 0.3805\ldots$ from the rest of the middle piece. This is Proposition 6.1 (v), and it is checked, not asserted.

Remark 3.4 (an auxiliary sweep). The same machinery proves a strictly better bound away from the minimiser, and Section 4 needs it. Let $\mathcal{B}$ be the fourteen boxes $J_i \times J_j$ with $0 \leq i,j \leq 3$ other than $J_2 \times J_3$ and $J_3 \times J_2$. Then $\gamma \geq \frac{757}{400} = 1.8925$ on every box of $\mathcal{B}$, by fourteen further certificates of the same shape at the same bidegree. The threshold $\frac{757}{400}$ sits in a narrow window: above the true minimum $1.89115\ldots$, and below $\gamma(\frac{2}{5}, \frac{1}{2}) = \frac{1079}{570} = 1.89298\ldots$, a value that γ does take on one of those fourteen boxes. The margin on that side is about 4.8×10^{-4}. Of this remark, what the formal development proves is the bound $\gamma \geq \frac{757}{400}$ on $\mathcal{B}$; the two comparisons that locate the threshold in its window are exact rational arithmetic carried out outside it.

4. The exact minimum

Throughout this section

$$q(t) = 6000t^4 - 16040t^3 + 14600t^2 - 5080t + 577,$$

and $g(\tau) = \gamma(\tau, \frac{1}{2})$. Note $q(\frac{1}{5}) = \frac{657}{25} > 0$, $q(\frac{7}{20}) = -\frac{4071}{400} < 0$ and $q(\frac{2}{5}) = \frac{201}{25} > 0$.

4.1. **Root isolation without a derivative.** The first divided difference of q is the symmetric polynomial

$$V(\tau,\sigma) = 6000(\tau^3+\tau^2\sigma+\tau\sigma^2+\sigma^3) - 16040(\tau^2+\tau\sigma+\sigma^2) + 14600(\tau+\sigma) - 5080,$$

and $q(\tau) - q(\sigma) = (\tau-\sigma)V(\tau,\sigma)$ is an identity of polynomials, proved by expansion.

Lemma 4.1. $V(\tau,\sigma) \geq 125$ *for all* $\tau,\sigma \in J_2 = [\frac{7}{20}, \frac{2}{5}]$*. Consequently* q *is strictly increasing on* J_2*.*

Proof sketch. One certificate. Written in the integer variables $X = 20\tau$, $Y = 20\sigma$ and then in the box coordinates of $J_2 \times J_2$ — as in Section 3, and as for the two certificates below — the integer polynomial $8000\,V - 10^6$ bihomogenises at bidegree $(8,8)$ to a form with no negative coefficient; Lemma 3.1 gives $V \geq 125$. The slack 125 is built into the certified polynomial so that the conclusion is *strict* monotonicity, which is what the uniqueness argument needs. Strict monotonicity then follows from the divided-difference identity alone, with no derivative. □

Theorem 4.2 (the minimiser is well defined)**.** *q has exactly one root t^* in* $[\frac{7}{20}, \frac{2}{5}]$.

Proof sketch. Existence is the intermediate value theorem applied to the continuous q with $q(\frac{7}{20}) < 0 < q(\frac{2}{5})$; uniqueness is Lemma 4.1. Only Lemma 4.1 rests on computation, and the computation is a single coefficient array. □

Root isolation for one integer quartic is of course textbook, by Sturm sequences or Descartes' rule; it is recorded here because the statement of Theorem 4.8 quantifies over "the root", and this is what makes that description denote.

4.2. **The second divided difference.** On the middle piece the reduction of γ to the line $\tau_y = \frac{1}{2}$ is governed by

Lemma 4.3 (the divided-difference identity)**.** *Put*

$$B(\tau,\sigma) = -\tfrac{25}{2}\tau^2 - 25\tau\sigma - \tfrac{75}{2}\sigma^2 + \tfrac{201}{8}\tau + \tfrac{201}{4}\sigma - \tfrac{127}{8}.$$

Then for all $\tau,\sigma \in [\frac{1}{5}, \frac{2}{5}]$,

$$114(1-\tau)(1-\sigma)\big(g(\tau)-g(\sigma)\big) = (\tau-\sigma)\,q(\sigma) + 160(\tau-\sigma)^2(1-\sigma)\,B(\tau,\sigma). \tag{6}$$

Identity (6) is proved by substituting the middle-piece formulas of Proposition 2.1 for $F(\tau)$, $G(\tau)$, $F(\sigma)$, $G(\sigma)$ and for $F(\frac{1}{2}) = 0$, $G(\frac{1}{2}) = \frac{57}{160}$, clearing the denominators $1-\tau$, $1-\sigma$, and expanding. It carries no computation beyond that expansion. Its content is that q and B are the first and second divided differences of g: comparing the coefficients of $(\tau-\sigma)$ and of $(\tau-\sigma)^2$ on the two sides gives

$$q(\sigma) = 114(1-\sigma)^2 g'(\sigma), \qquad B(\sigma,\sigma) = \frac{q'(\sigma)}{320(1-\sigma)}.$$

Those two relations explain where q and B come from; they are used nowhere below, and no proof in this paper differentiates anything.

Lemma 4.4. $0 \leq B(\tau,\sigma)$ *for all* $\tau,\sigma \in J_2$*.*

Proof sketch. One certificate, as in Lemma 4.1: the integer polynomial $3200\,B$ in the box coordinates of $J_2 \times J_2$ bihomogenises at bidegree $(8,8)$ with no negative coefficient. □

Corollary 4.5. *If $t \in J_2$ and $q(t) = 0$, then $g(t) \leq g(\tau)$ for every $\tau \in J_2$.*

Proof. Put $\sigma = t$ in (6). The term $(\tau-t)q(t)$ vanishes, so $114(1-\tau)(1-t)(g(\tau)-g(t)) = 160(\tau-t)^2(1-t)B(\tau,t)$. The right-hand side is a square times a positive factor times $B \geq 0$ (Lemma 4.4), and the left-hand coefficient $114(1-\tau)(1-t)$ is positive on J_2. □

This is the whole trick of the section: at a root of q the difference $g(\tau) - g(t)$ becomes a perfect square times a certified-nonnegative factor, so no arithmetic in $\mathbb{Q}(t^*)$ is ever required.

4.3. **The two-dimensional reduction.** It is *not* true that $\tau_y = \frac{1}{2}$ is optimal for every τ_x — see Remark 4.9 — so the reduction to the line has to be regional. It splits three ways, of which only one needs a certificate.

Lemma 4.6 (above the half). *If $0 \le \tau_x < 1$ and $\frac{1}{2} \le \tau_y < 1$ then $\gamma(\tau_x, \frac{1}{2}) \le \gamma(\tau_x, \tau_y)$.*

Proof. For $\tau_y \ge \frac{1}{2}$ one has $F(\tau_y) = 0$ and $G(\tau_y) = \frac{57}{160}$ constant, so τ_y enters γ only through $G(\tau_x)/(1-\tau_y)$, which is nondecreasing in τ_y because $G(\tau_x) \ge 0$. No computation at all. ☐

Lemma 4.7 (the strip). *Let*

$$E(\tau_x, \tau_y) = \tfrac{17}{16}\big(\tau_y + \tfrac{1}{2}\big)(1-\tau_y) + G(\tau_x)(1-\tau_x) - \tfrac{17}{8}(1-\tau_x)(1-\tau_y).$$

Then for $\tau_x \in J_2$, $\tau_y \in J_3$,

$$\tfrac{57}{80}(1-\tau_x)(1-\tau_y)\big(\gamma(\tau_x,\tau_y) - \gamma(\tau_x, \tfrac{1}{2})\big) = (2\tau_y - 1)\, E(\tau_x, \tau_y), \tag{7}$$

and $E \le 0$ there; hence $\gamma(\tau_x, \frac{1}{2}) \le \gamma(\tau_x, \tau_y)$ on $J_2 \times J_3$.

Proof sketch. Identity (7) is an expansion, like (6). Since $2\tau_y - 1 \le 0$ on J_3, the conclusion is exactly $E \le 0$, and that is one certificate: $-25600000\, E$ in the box coordinates of $J_2 \times J_3$ bihomogenises at bidegree $(8, 8)$ with no negative coefficient. The factorisation in (7) is the point of the lemma: the unfactored difference vanishes identically along $\tau_y = \frac{1}{2}$, so a certificate for it would be degenerate, whereas E is strictly negative on the closed box. ☐

Theorem 4.8 (the exact minimum). *Let t^* be the root of Theorem 4.2. Then*

$$\gamma(t^*, \tfrac{1}{2}) \ \le\ \gamma(\tau_x, \tau_y) \qquad \textit{for all } \tau_x, \tau_y \in [0, 1),$$

and $\gamma(t^*, \frac{1}{2}) = \gamma(\frac{1}{2}, t^*)$. In particular the minimum of γ over $[0,1)^2$ exists, is attained at $(t^*, \frac{1}{2})$ and at $(\frac{1}{2}, t^*)$, and equals $\gamma(t^*, \frac{1}{2})$.

Proof sketch. First the line. For $\tau \in J_2$ apply Corollary 4.5. For $\tau \in [\frac{1}{2}, 1)$ one has $F(\tau) = 0$, $G(\tau) = \frac{57}{160}$ and hence $g(\tau) = 1 + \frac{1}{2(1-\tau)} \ge 2$. For τ in the three remaining intervals J_0, J_1, J_3 the point $(\tau, \frac{1}{2})$ lies in a box of the sweep of Remark 3.4, so $g(\tau) \ge \frac{757}{400}$; and $g(t^*) \le g(\frac{3}{8}) = \gamma(\frac{3}{8}, \frac{1}{2}) = \frac{11499}{6080} < \frac{757}{400}$ by Corollary 4.5 again, since $\frac{3}{8} \in J_2$. So $g(t^*) \le g(\tau)$ for all $\tau \in [0, 1)$.

Now the square. If $\tau_y \ge \frac{1}{2}$, Lemma 4.6 and the line bound give $\gamma(\tau_x, \tau_y) \ge g(\tau_x) \ge g(t^*)$; the case $\tau_x \ge \frac{1}{2}$ is the same after swapping the arguments, which is legitimate by symmetry. Otherwise both thresholds lie in $[0, \frac{1}{2}]$, hence in one of J_0, J_1, J_2, J_3 each, giving sixteen cases. Fourteen of them are covered by the sweep of Remark 3.4, which gives $\gamma \ge \frac{757}{400} > \frac{11499}{6080} \ge g(t^*)$. The remaining two are $J_2 \times J_3$ and its transpose, and there Lemma 4.7 together with Corollary 4.5 gives $\gamma(\tau_x, \tau_y) \ge \gamma(\tau_x, \frac{1}{2}) \ge g(t^*)$. The final identity $\gamma(t^*, \frac{1}{2}) = \gamma(\frac{1}{2}, t^*)$ is the symmetry of γ, which needs no hypothesis on t^* at all.

The claims resting on exhaustive computation are exactly four positivity tests, each a single coefficient array at bidegree $(8, 8)$ computed inside the proof: the fourteen-box sweep of Remark 3.4 (fourteen arrays tested together), $B \ge 0$ on $J_2 \times J_2$, $E \le 0$ on $J_2 \times J_3$, and — through Theorem 4.2 — $V \ge 125$ on $J_2 \times J_2$. Everything else is polynomial identity and case analysis. ☐

Remark 4.9 (the reduction has to be regional). The line $\tau_y = \frac{1}{2}$ does not minimise γ in τ_y for a general τ_x, so Lemma 4.6 cannot be extended and the fourteen-box sweep cannot be dispensed with. For instance $\gamma(\frac{9}{10}, \frac{1}{2}) = 6$, while $\gamma(\frac{9}{10}, \frac{1}{40}) = \frac{11269}{5928} = 1.90097\ldots$ (This computation is an exact rational evaluation carried out outside the formal development, and is quoted here only to explain the shape of the argument.)

5. The value of the minimum

Identity (6) gives more than the location of the minimum: it makes the value *quadratically* stable under a rational bracket for t^*. One more certificate supplies the upper bound on B that the estimate needs.

Lemma 5.1. *$B(\tau, \sigma) \leq 3$ for all $\tau, \sigma \in J_2$.*

Proof sketch. One certificate: $9600 - 3200\,B$ in the box coordinates of $J_2 \times J_2$ bihomogenises at bidegree $(8, 8)$ with no negative coefficient. $\square$

Lemma 5.2 (the bracket)**.** $\dfrac{380515916462}{10^{12}} < t^* < \dfrac{380515916463}{10^{12}}$.

Proof. Both endpoints lie in J_2, and q is negative at the left one and positive at the right one — two exact rational evaluations of an integer quartic. Strict monotonicity (Lemma 4.1) converts the two sign conditions into the two strict inequalities. $\square$

Lemma 5.3 (the sandwich)**.** *Let $\ell \in J_2$ with $\ell < t^* < \ell + 10^{-12}$. Then*

$$g(\ell) - \frac{312}{41} \cdot 10^{-24} \;\leq\; g(t^*) \;\leq\; g(\ell).$$

Proof. The upper bound is Corollary 4.5. For the lower bound take $\tau = \ell$, $\sigma = t^*$ in (6); the first term vanishes, so

$$114(1-\ell)(1-t^*)\big(g(\ell) - g(t^*)\big) = 160(\ell - t^*)^2(1-t^*)B(\ell, t^*).$$

On J_2 one has $1 - t^* \leq \frac{13}{20}$ and $B \leq 3$ (Lemma 5.1), so the right-hand side is at most $160 \cdot 10^{-24} \cdot \frac{13}{20} \cdot 3$; and $1 - \ell, 1 - t^* \geq \frac{3}{5}$ gives $114(1-\ell)(1-t^*) \geq 41$. Dividing yields the stated bound. $\square$

The gap in Lemma 5.3 is quadratic in the bracket width: twelve digits of t^* buy about twenty-four digits of $g(t^*)$. Eighteen are printed below, because that is where the numerals stop being readable, and not because the method stops there.

Theorem 5.4 (the constant, to eighteen places)**.** *With t^* as in Theorem 4.2,*

$$\frac{1891152196075811250}{10^{18}} \;\leq\; \gamma(t^*, \tfrac{1}{2}) \;\leq\; \frac{1891152196075811251}{10^{18}}.$$

By Theorem 4.8, these are bounds for $\min_{[0,1)^2} \gamma$.

Proof sketch. Take $\ell = 380515916462 \cdot 10^{-12}$, which is legitimate by Lemma 5.2, and evaluate $g(\ell)$ exactly: it is a rational number, obtained from the middle-piece formulas of Proposition 2.1. The upper bound is $g(t^*) \leq g(\ell)$ followed by one rational comparison; the lower bound subtracts the certified gap $\frac{312}{41}10^{-24}$ of Lemma 5.3 and makes one more rational comparison. Both comparisons are exact arithmetic on explicit fractions. The only exhaustive computations behind this theorem are the three coefficient arrays already used: $B \geq 0$, $B \leq 3$ and $V \geq 125$, all on $J_2 \times J_2$. $\square$

Corollary 5.5. $0.001152196075811250 \leq \displaystyle\min_{[0,1)^2} \gamma - \frac{189}{100} \leq 0.001152196075811251$. *The constant* 1.89 *published in* [1, 2] *is therefore correct, and short of the truth by* $1.152\ldots \times 10^{-3}$*; the witness family of* [1] *obstructs that analysis at exactly* $1.8911521960758112\ldots$ *and no higher.*

Remark 5.6 (the shape of the line function). Outside the formal development, and quoted here only for orientation, exact rational arithmetic locates two critical points of g in the open middle piece: one near 0.2343280980647466, a local maximum with $g \approx 1.9220287118$, and t^* near 0.3805159164624639, the global minimiser. Both are roots of q. These numbers played the role of a target for the formal proof and are not asserted by any theorem above.

6. Controls

A one-sided bound proved by a positivity test is exactly the kind of statement a misplaced quantifier can satisfy vacuously. The following are checked with the same definitions as the theorems they guard.

Proposition 6.1. (i) *$\gamma \geq \frac{19}{10}$ is false on $[0,1)^2$, and so is $\gamma \geq \frac{757}{400}$; both are refuted at the single point $\gamma(\frac{3}{8}, \frac{1}{2}) = \frac{11499}{6080} = 1.8912828\ldots$ In particular the sweep threshold of Remark 3.4 is genuinely a regional bound and not a global one.*

(ii) *$\gamma(0,0) = \frac{160}{57}$, $\gamma(\frac{1}{2}, \frac{1}{2}) = 2$ and $\gamma(\frac{99}{100}, \frac{99}{100}) = 100$: the function is not identically anything.*

(iii) *$\gamma(x,x) = \dfrac{1}{1-x}$ for $\frac{1}{2} \leq x < 1$; hence $\gamma(1-\frac{1}{n}, 1-\frac{1}{n}) = n$ for every integer $n \geq 2$, γ is unbounded above, and the hypothesis $\tau_x, \tau_y < 1$ in Theorem 3.2 and Theorem 4.8 cannot be dropped.*

(iv) *Raising the thresholds breaks the certificates as it must: the 25-box test at $\frac{190}{100}$ fails, the fourteen-box test at $\frac{758}{400}$ fails, and the fourteen-box threshold $\frac{757}{400}$ fails on the single box $J_2 \times J_3$.*

(v) *The refinement of Remark 3.3 is load-bearing: on the unrefined four-interval partition, the test for $\gamma \geq \frac{189}{100}$ at the same bidegree $(8,8)$ fails.*

(vi) *Tampering is detected: replacing one coefficient of a passing array by -1, or adjoining a single negative coefficient to it, makes the test reject.*

Parts (iv), (v) and (vi) are six Boolean computations that all return *false*; they are the controls on the instrument itself, and they are stated because a positivity test that never rejects proves nothing.

7. Verification

Every statement of Sections 2–6 — Proposition 2.1, Lemmas 3.1, 4.1, 4.3, 4.4, 4.6, 4.7, 5.1, 5.2, 5.3, Corollary 4.5, and Theorems 3.2, 4.2, 4.8, 5.4 together with Proposition 6.1 — has been formally verified in Lean 4 (toolchain `v4.32.0`) against *Mathlib* [8, 9]; the development is seven modules and about 1600 lines, contains no `sorry` and declares no axiom by hand, and checks in 7 min 51 s of wall time and 4.84 CPU-minutes at a peak of 7.0 GB resident, of which the *Mathlib* import is about 6.6 GB. Everything except the positivity tests is free of compiled evaluation: Proposition 2.1 in all four of its parts, the identity between the built certificate polynomial and the cleared inequality (4), the divided-difference identities (6) and (7), the refutations, point evaluations and unboundedness of Proposition 6.1 (i)–(iii), and the symmetry $\gamma(t^*, \frac{1}{2}) = \gamma(\frac{1}{2}, t^*)$ of Theorem 4.8 depend on no axioms beyond `propext`, `Classical.choice` and `Quot.sound`. What is delegated to compiled evaluation, through Lean's `native_decide`, is exactly six Boolean tests of the form "this computed integer coefficient list has no negative entry": (1) the twenty-five-box test for Theorem 3.2; (2) the fourteen-box sweep of Remark 3.4; (3) $E \leq 0$ on $J_2 \times J_3$; (4) $B \geq 0$ and (5) $B \leq 3$ on $J_2 \times J_2$; and (6) $V \geq 125$ on $J_2 \times J_2$. Each theorem depends on precisely the ones its proof uses — Theorem 3.2 on (1) alone, Theorem 4.2 on (6) alone, Theorem 4.8 on (2), (3) and (4), and Theorem 5.4 on (4), (5) and (6) — and Proposition 6.1 (iv)–(vi) carries six further compiled Booleans, all returning *false*. Those twelve Booleans are the only source of extra axioms anywhere in this paper. No coefficient list is supplied as input at any point: each is computed inside the proof from the integer piece polynomials (3), so there is no transcription step between the density (1) and the checked data, and no dynamically loaded library is involved. Finally, Proposition 2.1 is what makes all of this a statement about the source's own f rather than about a transcription of it: the piecewise formulas for F and G used by every later proof are proved inside the kernel to be the integrals $\int_\tau^1 f$ and $\int_0^\tau tf$ of (1). Two qualifications belong here. Of Lemma 3.1, what the development verifies is what its proofs need — the evaluation identity $\widehat{P}(s, 1-s, u, 1-u) = P(s,u)$ and the nonnegativity of every monomial on $[0,1]^2$; the identification of the coefficients of $\widehat{P}$ with the Bernstein basis is classical, is used only to say where the certificates come from, and is deliberately not part of the formal development. And Corollary 5.5 and the numerical restatements in the introduction are arithmetic rephrasings of Theorem 5.4, not separate Lean statements. Every numerical value quoted above was first produced by an independent implementation in exact rational arithmetic in Python, and the formal proof was written against it afterwards.

The source of the development is currently held in a private repository: repository reference to be added at submission.

References

[1] K. Elbassioni and S. Ray, *Threshold rounding for the standard LP relaxation of some geometric stabbing problems*, arXiv:2106.12385v1 [cs.CG], 23 June 2021. The paragraph "Limitation of the analysis" defines the density (1), the ratio (2), and contains the two quoted sentences. Checked live on 3 September 2026: one submission-history entry, no journal reference and no related DOI on the arXiv record.

[2] K. Elbassioni and S. Ray, *Geometric stabbing via threshold rounding and factor revealing LPs*, Discrete Comput. Geom. **71** (2024), no. 3, 787–822, doi:10.1007/s00454-023-00608-8. The journal descendant of [1]; its Appendix B, "Limitation of our analysis for SegStab", carries the same density and the same two sentences, with τ_x, τ_y renamed τ_1, τ_2. Received 9 December 2021, accepted 17 September 2023, published online 27 November 2023, issue dated April 2024; Crossref, dblp and the publisher's own recommended citation all give the volume year as 2024, which is why it is used here in place of the online-publication year.

[3] K. Elbassioni, R. Gajjala and S. Ray, *Tight UGC thresholds for geometric stabbing problems*, arXiv:2607.28062, 30 July 2026. The only forward citation recorded for [2]; it quotes the factors 1.935 and 29/15 and does not mention the constant 1.89.

[4] S. Kovaleva and F. C. R. Spieksma, *Approximation algorithms for rectangle stabbing and interval stabbing problems*, SIAM J. Discrete Math. **20** (2006), no. 3, 748–768, doi:10.1137/S089548010444273X. Read here in full from the publisher's typeset version; the theorem numbers cited above are those of this journal version. Its earlier conference version, *Approximation of rectangle stabbing and interval stabbing problems*, Algorithms — ESA 2004, Lecture Notes in Computer Science **3221**, Springer, 2004, 426–435, carries a different title and was not used.

[5] D. R. Gaur, T. Ibaraki and R. Krishnamurti, *Constant ratio approximation algorithms for the rectangle stabbing problem and the rectilinear partitioning problem*, J. Algorithms **43** (2002), no. 1, 138–152, doi:10.1006/jagm.2002.1221. Conference version: Algorithms — ESA 2000, Lecture Notes in Computer Science **1879**, Springer, 2000, 211–219.

[6] S. Ben-David, E. Grant, W. Ma and M. Sharpe, *The approximability and integrality gap of interval stabbing and independence problems*, in: Proceedings of the 24th Canadian Conference on Computational Geometry (CCCG 2012), Charlottetown, Prince Edward Island, 8–10 August 2012, 47–52. No DOI: the CCCG proceedings are self-published. The page range was confirmed against the official electronic proceedings.

[7] C. Muñoz and A. Narkawicz, *Formalization of Bernstein polynomials and applications to global optimization*, J. Automat. Reason. **51** (2013), no. 2, 151–196, doi:10.1007/s10817-012-9256-3.

[8] L. de Moura and S. Ullrich, *The Lean 4 theorem prover and programming language*, in: Automated Deduction — CADE 28, Lecture Notes in Computer Science **12699**, Springer, 2021, 625–635, doi:10.1007/978-3-030-79876-5_37.

[9] The mathlib Community, *The Lean mathematical library*, in: Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2020), ACM, 2020, 367–381, doi:10.1145/3372885.3373824.