

Two Hours to Gerver

An independent numerical reproduction of the moving sofa optimum to 3×10^{-5} , and the cost asymmetry it exposes between approaching an answer and certifying one.

V. Golubenko — Zeltrex Claude Opus 5 — Anthropic (computational work)

15 September 2026 Single session · consumer laptop · no GPU

NumPy 2.5.2 · SciPy 1.18.0 · Python 3.13.15

ABSTRACT

Moser's moving sofa problem (1966) asks for the largest-area rigid shape that can be moved around a right-angled corner in a corridor of unit width. Gerver (1992) constructed a candidate of area 2.2195316688 comprising eighteen analytic arcs; Baek (2024) proves it optimal. We report an independent numerical reproduction of that value obtained in a single working session from a misremembered problem statement, without recourse to Gerver's construction. Using the dual “sofa from a path” formulation, an exact contact-measure subgradient, Richardson extrapolation in the rotation angle, and an annealed log-sum-exp surrogate ascended under a Sobolev-preconditioned step, we obtain a lower bound of **2.219466141**, within 2.95×10^{-5} relative of the accepted optimum and closing 99.46% of the interval between Hammersley's 1968 record and the optimum. The converged shape independently exhibits Gerver's gross morphology. Total cost: approximately two hours, ~500 lines of NumPy, minutes of CPU on consumer hardware, no GPU. **This accuracy is not competitive with the published numerical literature:** Gibbs (2014) reached a deficit of 1.7×10^{-7} by substantially the same method, and He (2026) 2.3×10^{-8} . We rank below all of them, and every algorithmic ingredient used here is standard.

We argue the principal finding is not the agreement but the asymmetry producing it. The same session's independent *upper* bound is 2.2336 — 0.63% above the optimum and uncertified. Absent prior knowledge of the answer, the defensible claim would have been an interval two orders of magnitude wider than the accuracy actually attained. Constructing and approaching has become cheap; certifying has not. We quantify the gap, report the pipeline's own defect rate (seven defects, two retracted results), and derive six predictions for engineering practice. No mathematical contribution is claimed.

Keywords: moving sofa problem · reproduction study · non-smooth optimisation · Sobolev preconditioning · semi-infinite programming · computational verification

1 Introduction

This is a reproduction study, not a mathematical contribution. Its subject is a well-posed geometric optimisation problem that remained open for fifty-eight years, and its question is how much it now costs to arrive independently at the answer.

The distinction between *replication* and *reproduction* carries the evidential weight of what follows. Replication re-runs an original method and detects arithmetic error. Independent reproduction reaches the same value by different machinery and detects method-specific artefacts, which is why it is worth more. The work reported here is the latter: the only object shared with prior art is the problem's dual formulation, standard since Gerver and Romik. The solution route is disjoint.

We state at the outset what this study does not claim. It does not solve the moving sofa problem, which was solved constructively in 1992 and as an optimality proof in 2024. Its headline accuracy figure is measured against a value already known, and is therefore a reproduction metric carrying no independent claim of accuracy. Section 6 examines what remains when that yardstick is withdrawn, and argues that the remainder is the study's actual result.

1.1 Contributions

- An independent numerical reproduction of Gerver's value to 2.95×10^{-5} relative, from below, obtained without reference to his construction — less accurate than four prior treatments (§2.1, §4.1).
- *Path Symmetric Inflation*, a Sobolev-preconditioned annealed ascent that resolves the ridge on which a plain subgradient method and L-BFGS both stall. Both ingredients are textbook (§3.5); the assembly is what is reported, not the parts.
- An empirical finding that annealing past an intermediate smoothing level *loses* area, reproduced independently in two optimisation modes (§4.3).
- A direct experimental demonstration that the residual subgradient atom is a kink rather than an improving direction (§4.4).
- A rediscovery, and then a correction, of the standard monotonicity argument — including the two gaps the literature had already identified in it (§4.5).
- A quantified statement of the asymmetry between constructive lower bounds and certified upper bounds, with both computed in the same session — the one claim here that survives the literature (§6).
- A worked case of independent arrival without novelty: the contact-pressure formulation of §3.4, derived here from physical intuition, is the published driving force of [10] (§3.4).

2 Problem statement and prior work

A corridor of unit width turns through a right angle. Let $S \subset \mathbb{R}^2$ be a connected compact set that can be moved by a continuous rigid motion from one arm to the other, remaining inside the corridor throughout. Moser [1] asks for the supremum of $|S|$.

Existence was never the difficulty. The Blaschke selection theorem renders the space of compact sets compact in the Hausdorff metric, and area is upper semicontinuous along such limits, so a maximiser exists. The

difficulty was identification.

YEAR	RESULT	AREA	METHOD	STATUS
–	Unit semicircle	1.570796	elementary	folklore
1966	Problem posed [1]	–	Blaschke selection	existence settled
1968	Hammersley [2]	2.207416	one-parameter ansatz	conjectured optimal
1992	Gerver [3]	2.219532	variational, 18 arcs	locally optimal
2018	Kallus & Romik [4]	≤ 2.37	certified computation	rigorous bound
2024	Baek [6]	$= 2.219532$	concave majorant	preprint, v1

Table 1. The record and the method behind each advance. Each jump followed a change of method rather than additional effort with the previous one. Baek's proof constructs a functional Q that upper-bounds any admissible sofa's area, is exact for Gerver's, and is concave — whence a local maximum is global. It is not computer-assisted.

2.1 Prior numerical work, and where this result stands

This section was written after the computational work, following a literature check. It should have been written before, and its absence is the largest methodological failure in this study. Numerical attacks on the problem go back to 1973, and the approach taken here has a direct antecedent.

WORK	METHOD	AREA	DEFICIT
Maruyama 1973 [7]	polygonal approximation	–	historical
Gibbs 2014 [8]	rotating corridor, path optimisation, $O(1/N)$ corner term, extrapolation across successive N	2.219531843	$1.7e-7$
Leng et al. 2024 [9]	differentiable area, ReLU PINN, Adam + L-BFGS; monotonicity and symmetry dropped	2.219515	$1.8e-5$
He 2026 [10]	gas-pressure dynamics; driving force = pressure $\times$ contact length	2.2195316460	$2.3e-8$
This work	contact-measure subgradient, LSE annealing, Sobolev preconditioning	2.219466141	6.6e-5
Gerver 1992 [3]	analytic construction	2.2195316688	0

Table 2. This work ranks last, below even the unrefined mean of [9] (2.219482 , $\sigma = 1.1 \times 10^{-5}$). Gibbs [8] anticipated the frame change, the path optimisation, the diagnosis of the $O(1/N)$ term to the inner-corner boundary and its removal by extrapolation across successive N — the four choices this study arrived at independently — and did so roughly four hundred times more accurately. His Table 1 runs $N = 25$ to 501 and reports 2.219531843 , “to nearly eight significant figures”. He [10] is a paper about corridor variants whose headline result is a 43.327° corner; the value quoted here is its 90° validation case. Since Gerver's sofa is explicitly constructible, a lower-bound search that fails to reach it is simply underperforming; only exceeding it would have been news, and [6] makes that impossible.

The individual ingredients are equally well established. Sobolev and H^1 gradient preconditioning in shape optimisation dates to Neuberger [11] and the H^1 gradient method of Azegami and Wu; Gaussian-kernel-smoothed gradients with an annealed kernel width are standard in the active-contour literature [12]. Entropic and log-sum-exp smoothing of semi-infinite minimax problems is textbook [13]. Applying kernel preconditioning to a one-dimensional path rather than a boundary is the textbook case, not a generalisation of it.

As of 15 September 2026, reference [6] remains at v1 with no journal reference and has not completed peer review. All corroboration reported here is therefore corroboration of an unrefereed claim.

3 Method

Five components, presented in the order they became necessary. Each was forced by the failure of its predecessor; none was planned.

3.1 Dual formulation

Hold the sofa fixed and rotate the corridor about it. At rotation angle $t \in [0, \pi/2]$ the outer walls carry unit normals

$$n_1(t) = (-\sin t, \cos t), \quad n_2(t) = (\cos t, \sin t) \quad (1)$$

$$L_t = \{\langle p, n_1 \rangle \leq b_1\} \cap \{\langle p, n_2 \rangle \leq b_2\} \setminus \{\langle p, n_1 \rangle < b_1 - 1 \wedge \langle p, n_2 \rangle < b_2 - 1\} \quad (2)$$

$$S = \bigcap_{t \in [0, \pi/2]} L_t \quad (3)$$

Two half-planes for the outer walls, less the quadrant behind the inner corner; unit width enters as the -1 . Shape ceases to be an unknown: for a given motion the maximal sofa is forced, and the entire remaining freedom is the pair of scalar offset functions $b_1(t)$, $b_2(t)$.

Every normal in (1) has non-negative y -component on $[0, \pi/2]$, so (3) is the region between an upper and a lower profile and the area reduces to a one-dimensional integral:

$$\text{top}(x) = \min_{t,i} (b_i - n_{i,x}x)/n_{i,y}, \quad \text{bot}(x) = \max_t \min_i (b_i - 1 - n_{i,x}x)/n_{i,y}, \quad |S| = \int (\text{top} - \text{bot})^+ dx \quad (4)$$

3.2 Validation before optimisation

Two closed-form identities fix the implementation. Constant offsets $b_1 = b_2 = 1$ must return the unit semicircle, $\pi/2$; the offsets $b_1 = 1 + r \sin t$, $b_2 = 1 + r \cos t$ with $r = 2/\pi$ must return Hammersley's $\pi/2 + 2/\pi$. Both reproduce to approximately 2×10^{-6} . No optimiser was run until both held.

3.3 Richardson extrapolation in the rotation angle

Naive sampling of t inflates the area by order 10^{-3} — the same magnitude as the quantity under study, and always in the direction that flatters the result. A parabolic-vertex correction was applied first and made the error *worse*. The diagnosis: the lower envelope in (4) attains its extremum at a **corner** in t , where two wedge

lines cross, so the sampling error is $O(\Delta t)$ rather than $O(\Delta t^2)$ and a quadratic model is invalid. A single Richardson step

$$A \approx 2A(\Delta t) - A(2\Delta t) \quad (5)$$

reduces the error to approximately 2×10^{-6} , roughly five hundred times smaller than the Hammersley-Gerver interval being resolved.

3.4 Exact subgradient as a contact measure

Differentiating (4) with respect to a wall offset gives

$$\partial|S|/\partial b_i(t) = |\{x: \text{wall } i \text{ at angle } t \text{ bounds the top}\}| - |\{x: \text{the wedge at angle } t \text{ bounds the bottom}\}| \quad (6)$$

a difference of contact-set lengths, computable in closed form from the `argmin` / `argmax` indices and verified against central differences to six digits. It is a *measure*, not a function: $t = 0$ and $t = \pi/2$ each own an entire straight edge of the sofa and contribute atoms of finite mass that do not diminish under grid refinement.

THE SAME QUANTITY, ALREADY PUBLISHED AS A PHYSICAL ALGORITHM

Equation (6) was not reached from the literature. It was reached from a physical intuition raised mid-session: treat the sofa as compressible, force it through the corner, and let the pressure show where the shape still has room. Formalised, that gives the contact measure above, which we called a generic pressure force.

He [10] publishes exactly this as a driving force. A gas pressure inflates the shape against the corridor walls, and — quoting the paper — “*the magnitude of each component equals the product of the pressure and the effective contact length of that side*”, $F = p \cdot b$, with b the contact-segment length. That is equation (6) with a physical constant in front. It reaches 2.2195316460, roughly two thousand eight hundred times closer to the optimum than this study.

This is the sharpest illustration the paper has of its own thesis. The formulation was correct, reconstructible in one sentence of physical reasoning, and already in the literature at far better accuracy. Independent arrival is not novelty. The guards of §3.6 confirmed the quantity was computed correctly; nothing inside the work could report that it already had a name.

3.5 Path Symmetric Inflation

Plain ascent on (6) stalls at 2.218686; L-BFGS terminates abnormally at iteration zero. The cause is structural: (4) is a max-of-min of linear pieces, hence a polyhedral surface whose optimum lies on a **ridge**, across which a first-order method oscillates. Three modifications resolve it.

Annealed smoothing.

Both envelopes in (4) are replaced by log-sum-exp at inverse temperature β , swept from 5×10^2 to 3×10^4 . The ridge is rounded at low β and sharpened as the iterate approaches.

Sobolev preconditioning.

Rather than stepping along the pointwise gradient, ascend in u where $b = Ku$ and K is a Gaussian in the angle variable, σ annealed with β . The step becomes a smooth inflation of the whole path. A ridge is precisely a locus where the pointwise gradient alternates sign at the sampling scale; smoothing the *search direction* averages that oscillation instead of following it.

Symmetrisation and gauge projection.

Optionally impose $b_2(t) = b_1(\pi/2 - t)$; always project out the two-dimensional translation subspace, which alters no shape and consumes steps.

3.6 Guards

TWO PROPERTIES WITHOUT WHICH NO NUMBER HERE WOULD BE TRUSTWORTHY

The surrogate never scores. With normalised weights the log-sum-exp envelope biases upward, so the smoothed objective *over-estimates* area. It is used only to steer. Every reported figure is the exact Richardson area of the path actually produced.

Re-score on a finer grid than the one optimised on. Optimisation ran at $n_t = 401$; every result was re-scored at $n_t = 3201$. Without this an optimiser exploiting its own discretisation would manufacture a value above the true optimum — precisely the false positive this study was constructed to detect.

4 Results

4.1 Lower bounds

BEST AREA 2.219466 exact, $n_t = 3201$	ACCEPTED OPTIMUM 2.219532 Gerver / Baek	RELATIVE GAP 2.95e-5 0.003%, from below	INTERVAL CLOSED 99.46% from Hammersley 1968
---	--	--	--

ROTATION PATH	AREA	GAP TO OPTIMUM	PROVENANCE
Unit semicircle	1.570796	-6.49e-1	$\pi/2$
Hammersley	2.207416	-1.21e-2	$\pi/2 + 2/\pi$, 1968
Plain subgradient ascent	2.218686	-8.46e-4	this work, stalled
PSI, mirror-symmetric	2.219413	-1.19e-4	this work
PSI, unconstrained	2.219466	-6.55e-5	this work, best
Gerver's sofa	2.219532	0	1992 / proved 2024

Table 3. All figures are exact Richardson areas of the sofa each path generates, re-scored on a grid eight times finer than that used for optimisation. Across twelve starting points in two modes, no run exceeded the accepted optimum — the only outcome this method could produce as evidence in favour of [6].

Restarts cluster into three distinct basins near 2.2137, 2.2187 and 2.2192–2.2195, confirming that the landscape carries multiple local optima; none lies above Gerver's value.

4.2 Independent recovery of Gerver's morphology

No element of the pipeline encodes the expected form of the answer. Decomposing the converged boundary by the constraint generating each point yields:

lower — flat edge ($t = 0$) | **one monotone envelope arc, $81.3^\circ \rightarrow 10.7^\circ$** | flat edge ($t = 90^\circ$)

upper — envelope arc | straight edge spanning 49% of the width | envelope arc

The telephone-handset morphology: a flat back, two rounded ends, and a single smooth bite swept from the front by the inner corner. Overall width 3.2362. The envelope is sharply determined throughout — the band of angles within 10^{-4} of tangency measures 0.02° , one grid step. Gerver's finer eighteen-section classification lies below our resolution and is neither confirmed nor contradicted.

(7)

4.3 Annealing past the optimal smoothing level loses area

B	Σ	SYMMETRIC MODE	UNCONSTRAINED MODE
8×10^3	0.015	2.219407	2.219468
3×10^4	0.008	2.219420	2.219458
1.2×10^5	0.003	2.219359	2.219416
5×10^5	0.000	2.219352	2.219416

Table 4. Both modes peak at intermediate smoothing and regress when sharpened further. Every figure is the exact Richardson area, not the surrogate, so the peak is not an evaluation artefact. This is nonetheless a known phenomenon rather than a finding: it is the standard basin-loss and ill-conditioning failure of graduated non-convexity and exponential smoothing. Two modes inside one codebase is also not independent reproduction.

4.4 A subgradient atom that indicates no improvement

The converged path retains an atom of +0.81 at $t = 0$, which presents as a large unexploited gradient. Perturbing that wall in the direction the atom indicates:

$$\delta b = +0.02 \Rightarrow \Delta A = -3.63 \times 10^{-2}; \quad \delta b = -0.02 \Rightarrow \Delta A = -3.64 \times 10^{-2} \quad (8)$$

Area decreases in both directions. The point is a kink; the computed “gradient” is one arbitrary element of a subdifferential, not a direction of ascent. This is a cheap and unambiguous experimental demonstration of why first-order methods stall on this objective.

4.5 Non-monotone rotation cannot improve the optimum

A natural route to a larger sofa is to question the assumption that the corridor's rotation proceeds monotonically from 0 to $\pi/2$. The assumption closes rather than opens the search. For any admissible motion, $S \subseteq \bigcap_{\tau} L(\tau)$. The rotation angle is continuous and attains every value in $[0, \pi/2]$ by the intermediate value theorem, so selecting one placement per angle yields $S \subseteq \bigcap_{t \in [0, \pi/2]} L_t$ — a *subset* of the constraints and therefore a *superset* region. Additional traversal only adds constraints. The monotone family already attains the supremum.

This argument is not new and is not sufficient. Gerver [3, Thm. 1] proves the monotonicity reduction; Baek [6, Prop. 1.2.2] gives the identical intermediate-value sentence and then states explicitly what it does *not* establish: that the selected hallways vary continuously in t , and that the resulting intersection is connected — which Moser's formulation requires, since a sofa is by definition connected. Repairing those two gaps occupies a chapter of [6]. Reproducing a known argument together with both of its known holes is the clearest single illustration of what a literature check is for. Leng et al. [9] additionally tested non-monotone motions numerically and found no improvement.

5 Defects and retractions

A reproduction study that omits its own error rate has not reported its method. In one session the pipeline produced seven defects and two retracted results. The seventh is instructive because it lies in the publication toolchain rather than the mathematics, and would have shipped a document whose typography was silently not its own.

#	DEFECT	DETECTED BY	CONSEQUENCE
1	Inverted mask selecting the second wall family in the gradient	runtime IndexError	none
2	Parabolic refinement applied to an $O(\Delta t)$ corner	accuracy degraded	method replaced
3	Boolean mask applied before rather than after indexing	shape mismatch	none
4	Best path tracked in latent coordinate u while K varied between stages	reload disagreed with run	two results retracted
5	Undefined name after a partially applied patch	NameError	run lost, relaunched
6	Segment detector over-split one arc into nine; monotonicity test checked only the increasing direction	contradiction with tangency-band measurement	two readings corrected
7	PDF build embedded fallback system fonts: <code>--print-to-pdf</code> did not await the webfont fetch, and once inlined the document's own CDN <code><link></code> re-won the cascade with a variable font Chrome will not embed	build gate 4	caught pre-publication

Table 5. Every defect was caught by a guard, a validation identity, an internal contradiction or a build gate — none by code review.

RETRACTED

Defect 4 produced the figures 2.175752 and 2.061435, briefly reported as phase-two results; they are artefacts of re-expanding a stored path under the wrong kernel and carry no meaning. An interim statement of “72% of the interval closed” was also incorrect; the value at completion was 93%. Defect 6 produced a transient claim of nine distinct arcs along the lower boundary and a verdict of non-monotonicity, both withdrawn in §4.2.

The pattern merits a name: **defect generation scales with development speed and detection does not**. Five of six defects were caught because a guard existed, not because anyone inspected the code. The load-bearing component of this work was never the solver.

6 The asymmetry between approaching and certifying

The search of §3 constructs motions and measures the sofas they admit. Every such number is a **lower bound**: a shape that demonstrably fits. Bounding from above requires the opposite operation — retain finitely many angles and *discard* the remaining constraints. Every admissible sofa still satisfies those retained, so

$$|S| \leq \max_b |\cap_k L_{t_k}(b)| =: U(T) \quad (9)$$

valid for every sofa and tightening as angles are added — **but only when the inner maximisation is global**. A local numerical maximum is a *lower* bound on $U(T)$, and a lower bound on an upper bound constrains nothing. The numbers below are therefore not bounds on the optimum; they are estimates of a quantity that would bound it if maximised globally in exact arithmetic, which is the entire content of [4].

ANGLES T	3	4	5	7	9	13	17	25
U(T)	2.8018	2.5154	2.4148	2.3361	2.3027	2.2725	2.2584	2.2336

Table 6. Finite-angle relaxation values. At $|T| = 2$ the region is genuinely unbounded and the implementation returns 1.7×10^{17} rather than a plausible finite value — a useful negative control. At $|T| = 3$ the bound already improves on Hammersley's $2\sqrt{2} \approx 2.8284$.

That 2.3361 at seven angles undercuts the rigorous 2.37 of [4] initially seemed implausible from forty random restarts. It is not: Leng et al. [9] report 2.3337 from a five-angle relaxation by direct numerical optimisation, so sub-2.37 values are expected once certification is dropped. We nonetheless tested the ceiling property, since a genuine global maximum cannot be raised by additional restarts. Raising restarts from 40 to 200 to 600 altered the value by +0.0000 at $|T| = 5, 7$ and 9 alike (2.414822, 2.336057, 2.302697). That is evidence of a stable local maximum and nothing stronger; float64, no interval arithmetic, no branch-and-bound.

LOWER BOUND, OURS	TRUTH, PER [6]	RELAXATION ESTIMATE
2.219466	2.219532	2.233625
constructive; a shape that fits	0.003% above our floor	0.63% above truth — not a bound

This is the study's one surviving finding. One session bought a genuine lower bound accurate to three decimal places and, in the other direction, a number that is not a bound at all — only an estimate of one, two orders of magnitude weaker, whose validity rests on an unverified global-optimality assumption. The gap is not an artefact of effort allocation: the constructive side is cheap because any feasible configuration certifies itself, and the other side is expensive because it is a statement about everything. Had Gerver's value not existed, the defensible claim would have been that the answer lies somewhere below roughly 2.2336, with no means of knowing we were already situated on it. Approaching became cheap; certifying did not.

7 Cost

RESOURCE	QUANTITY	NOTE
Wall-clock	~2 hours	one session, including three false starts
Compute	minutes of CPU	consumer laptop; no GPU used at any stage
Code, final pipeline	~500 lines	1,010 including superseded iterations
Dependencies	2	NumPy, SciPy
Prior domain knowledge	none	problem statement initially misremembered
Literature consulted	2 items	abstract and introduction of [6]

Table 7. For calibration: [3] was a research contribution; [4] was an *Advances in Mathematics* paper resting on certified computation; [6] runs to 119 pages and took seven years. None of that was replicated. What was replicated, in an afternoon, is the *value* — the component that previously required the paper in order to know.

8 Predictions for engineering practice

If the reproduction cost of a published numerical result has fallen to an afternoon, several consequences follow that are not primarily mathematical.

8.1 Independent verification becomes a step, not a project.

Checking a vendor's performance claim, a paper's benchmark or a supplier's tolerance analysis ceases to be a quarter of work. Expect “we reproduced it ourselves” to become a normal precondition for adopting a dependency.

8.2 The bottleneck migrates to specification and guards.

Nothing here was rate-limited by writing code. The decisive work was selecting the dual formulation and designing the two guards of §3.6. That is architecture, and it does not transfer to whoever writes the fastest solver.

8.3 Feasible-side answers get cheap; certificates do not.

“We found a configuration that works” and “we proved none works better” were always distinct problems. They now differ by orders of magnitude in cost. Expect systematic over-trust of the former where the latter is required.

8.4 Defect rate scales with speed; detection does not.

Seven defects and two retractions in a single session, six caught by guards rather than review. Organisations adopting this pace without adopting validation identities and independent re-scoring will publish artefacts and not notice.

8.5 Geometric feasibility problems transfer directly.

The structure exploited here — a parametrised motion, a feasible region as an intersection of half-planes, a contact measure as the gradient — is the structure of slew-envelope, sensor-coverage, mechanism-clearance and packing problems. The same five hundred lines port with the constraint set exchanged.

8.6 Safety cases are unaffected.

Where the deliverable is a certificate rather than a number — a clearance envelope that must be guaranteed rather than merely observed to hold — none of this cost collapse applies. That work still requires interval arithmetic, exhaustive argument or proof, and it is precisely the work most likely to be presumed complete because a fast numerical result appeared convincing.

9 Threats to validity

- **No result is rigorous.** All arithmetic is float64 without interval certification. The bounds of §6 are valid only if the inner maximisation attained its global optimum, supported by a stability test at three angle counts and never proved.
- **Unverified assumptions in the implementation.** A monotone quarter-turn is assumed, and connectedness of the intersection — required by Moser's formulation — is never verified.
- **The accuracy figure is a reproduction metric.** It is agreement with a published value, measured against that value, and carries no independent claim of accuracy.

- **The reference itself is unrefereed.** Reference [6] is a v1 preprint with no journal reference; the corroboration offered here is corroboration of an unverified claim.
- **Single operator, single session.** No inter-operator variance was measured, and the wall-clock figure is not a controlled measurement.
- **The literature check followed the work rather than preceding it.** Had it come first, the study would have started from [8] and aimed at its accuracy. Every novelty-adjacent claim in the first draft was withdrawn on contact with prior art.
- **Prior-art figures were verified first-hand.** [8] from its own Table 1 in the source PDF, [9] from its abstract, [10] from its full text. An earlier draft cited them second-hand and mis-stated the scope of [10]; that is corrected here.

10 Conclusion

An independent numerical route reached the moving sofa optimum to 2.95×10^{-5} relative in a single session on consumer hardware, recovering the known morphology without being supplied it, and produced in the same session a relaxation estimate two orders of magnitude weaker that is not a bound. No part of the mathematics is new; the accuracy is beaten by four prior treatments, one of which used substantially the same method twelve years earlier and far better. The reproduction of a known value has become inexpensive; the manufacture of a certificate has not.

For engineering organisations the operative consequence is not that hard problems became easy, but that the *appearance* of having solved one became cheap, while the evidence that one has actually done so did not. This document is its own example. It reached three correct decimal places, recovered the right shape, passed every internal guard – and still required an external literature check to establish that it had contributed nothing mathematically, that its upper bound pointed the wrong way, and that its central argument had been published, with both of its holes, in 1992 and again in 2024. The guards caught what was wrong *inside* the work. Only the literature caught what was wrong *about* it.

References

- [1] L. Moser. Problem posed 1966; see the *SIAM Review* problem sections of the period.
- [2] J. M. Hammersley. On the enfeeblement of mathematical skills by modern mathematics. *Bull. Inst. Math. Appl.* 4 (1968). Area $\pi/2 + 2/\pi$; upper bound $2\sqrt{2}$.
- [3] J. L. Gerver. On moving a sofa around a corner. *Geometriae Dedicata* 42 (1992), 267–283.
- [4] Y. Kallus and D. Romik. Improved upper bounds in the moving sofa problem. *Advances in Mathematics* 340 (2018), 960–982.
- [5] D. Romik. Differential equations and exact solutions in the moving sofa problem. *Experimental Mathematics* 27 (2018), 316–330.
- [6] J. Baek. Optimality of Gerver's Sofa. arXiv:2411.19826 [math.MG], submitted 29 November 2024. 119 pp.; v1, no journal reference as of 15 September 2026.
- [7] K. Maruyama. An approximation method for solving the sofa problem. *Int. J. Computer and Information Sciences* 2 (1973), 29–48.
- [8] P. Gibbs. A computational study of sofas and cars. Preprint, 2014. Reaches 2.219531843 (deficit 1.7×10^{-7}) by substantially the method used here.

- [9] K. Leng, J. Bi, M. Cha, S. Pinilla, J. Thiyagalingam. Deep learning evidence for global optimality of Gerver's sofa. arXiv:2407.11106 (2024).
- [10] X. He. A gas-driven algorithm for variants of the moving sofa problem. arXiv:2608.11206 (2026). Driving force = pressure $\times$ contact length; reaches 2.2195316460 at $N = 500$ on the 90° corridor, reported there as a validation case.
- [11] J. W. Neuberger. *Sobolev Gradients and Differential Equations*. Springer LNM 1670, 2nd ed. 2010.
- [12] G. Sundaramoorthi, A. Yezzi, A. Mennucci. Sobolev active contours. *Int. J. Computer Vision* (2007).
- [13] S.-H. Pee, J. O. Royset. On solving large-scale finite minimax problems using exponential smoothing. *J. Optimization Theory and Applications* (2011).

A Appendix: reproduction

The pipeline comprises five modules. `core.py` implements equations (1)–(6) with Richardson extrapolation and the exact subgradient. `soft.py` supplies the log-sum-exp surrogate with shifted exponentials. `psi.py` and `psi3.py` implement Path Symmetric Inflation and the annealing schedule. `ub.py` computes the finite-angle bounds of equation (9).

Two checks should precede any use. The semicircle path must return $\pi/2$ and the Hammersley path $\pi/2 + 2/\pi$, both to approximately 2×10^{-6} ; and any result must be re-scored on a strictly finer angular grid than the one on which it was optimised. A pipeline failing either check will produce values above the true optimum and present them as discoveries.

ZTX-TR-2026-09 · Companion interactive figure: the corridor, the sofa, and the contact-pressure field with rotation angle as the third axis. All computational results are reproducible from the five modules listed in Appendix A.