FOURFOLD IS ENOUGH: CONNECTED LEGO MODELS FROM ONE BRICK TYPE

JOSHUA S. GANS
UNIVERSITY OF TORONTO

Abstract. Can every model made from ordinary rectangular construction bricks be rebuilt using just one brick shape? For a finite union of unit lattice cubes and positive-integer scale factors, we show that uniform enlargement by four always permits an exact reconstruction with upright $2 \times 4 \times 1$ bricks. If the source is face-connected, the replacement bricks have a connected overlap-contact graph at that same scale; if its occupied columns also extend to the ground, each nonground brick can be placed above a previously placed overlapping brick. These are combinatorial guarantees, not assertions of mechanical stability. Four is optimal even among arbitrary real universal factors, and all three constructions require exactly eight new bricks per original cube. The possible minimum *integer* geometric scales are precisely 1, 2, and 4: doubling works exactly when every horizontal source layer is domino-tileable, while a localized checkerboard argument shows that scale 3 can never be minimal. We finish with the costs of preserving proportions and old seams.

1. A brick question with two surprises

A model assembled from construction bricks can draw on a heterogeneous inventory: 1×1 bricks, 1×2 bricks, long beams, and large rectangular blocks. The familiar 2×4 brick suggests a natural standardization question: can an arbitrary model be rebuilt using that one brick type while preserving its shape? If not at its original size, is there one uniform enlargement factor that always suffices? Throughout, admissible scale factors are positive integers, so every vertex and boundary of the original unit-cell lattice remains on the target brick lattice after dilation.

At the geometric level, the question is a tiling problem. An upright target brick occupies a $2 \times 4 \times 1$ box, and the original model is represented as a finite union of unit lattice cubes. We ask whether a uniform dilation of that occupied region can be partitioned into target bricks. At the connection level, a partition is not enough: we join two bricks when their horizontal footprints overlap in at least one unit square in consecutive height layers, and ask whether the resulting contact graph is connected. At the assembly level, we additionally ask for a bottom-up order in which each nonground brick has a previously placed contact immediately beneath it. These are combinatorial conditions: they do not assert rigidity, clutch strength, static stability, collision-free access, or connectivity after every intermediate placement. Geometry,

Date: August 25, 2026.
2020 *Mathematics Subject Classification.* 05B45, 05C70, 52C20.
Key words and phrases. LEGO bricks, polyomino tilings, polycubes, domino tilings, checkerboard invariants, dilation, brick connectivity, constructibility.

contact connectivity, and supported placement are related, but they are not the same problem.

The basic geometric restriction is *orthogonality*: the occupied region must be a union of axis-aligned lattice cubes. Within that restriction, the model can be nonconvex, hollow, or assembled from pieces of many different sizes. Original piece boundaries can be ignored, permitting a global redesign, or retained as seams that new bricks may not cross. Actual inventories also include plates, slopes, and specialized pieces; plates require a finer vertical lattice and may change the universal scaling factor, while genuinely sloped or curved boundaries require approximation rather than exact reconstruction.

The answer to the geometric question is four. Each original cube expands into a $4 \times 4 \times 4$ box, and each of its four horizontal layers accommodates two target bricks. Hence a model with V source cubes requires exactly $8V$ replacement bricks. The one-cube model shows that no smaller factor works universally. This local argument, however, does not ensure that the resulting bricks fit together into a single connected assembly.

This problem has natural roots in two earlier articles. De Bruijn studied which rectangular boxes can be filled with a fixed brick, including the $1 \times 2 \times 4$ brick and classical coloring obstructions [2], while Eilers examined the enumeration of connected LEGO structures [3]. Checkerboard and more general tiling invariants are classical [4, 1]; Rémila and Pak–Yang analyze algorithmic rectangle-tileability questions, especially for simply connected regions [7, 6]. Here the prescribed shape may instead have disconnected horizontal footprints, indentations, or holes, and the replacement must satisfy specified overlap-contact conditions.

The first surprise is that contact connectivity requires no additional enlargement: four carefully coordinated layers connect every face-connected source model. For such face-connected models with no unsupported source-cell overhangs, the same four layers give a bottom-up order with a lower overlap contact for every nonground brick. The second surprise concerns models that need less enlargement. At scale 2, grid alignment gives an exact bijection with domino tilings of the original horizontal footprints, without connectedness or no-hole assumptions. At scale 3, ordinary checkerboard weights applied separately to each 3×3 source macrocell force the original footprint to partition into 2×2 squares. Together these results show that the possible minimum *integer* geometric scales are $1, 2, 4$, not $1, 2, 3, 4$.

Brick contact graphs themselves are not new: Testuz, Schwartzburg, and Pauly use them while improving variable-size brick sculptures through iterative splitting and merging [8]. The four-layer construction here instead gives a direct universal guarantee for an exact prescribed region using just one upright brick type. Keller and coauthors also exploit dilation, but in a different assembly model involving adhesive particles moved by global controls [5]. Their particles attach along arbitrary faces; ours have only the stipulated contacts between consecutive height layers.

2. Four copies in every direction

Let $S \subset \mathbb{Z}^3$ be finite and nonempty, and write

$$M = \bigcup_{(u,v,w)\in S} \big([u, u+1] \times [v, v+1] \times [w, w+1]\big), \qquad V = |S|.$$

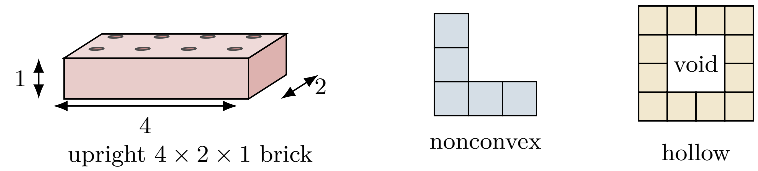

Figure 1. The upright target brick and two admissible source footprints. Orthogonality restricts the lattice and face directions, not convexity or the presence of empty regions.

As usual for lattice tilings, overlaps along boundaries are immaterial. A target brick is an integer translate of either

$$[0,4] \times [0,2] \times [0,1] \quad \text{or} \quad [0,2] \times [0,4] \times [0,1].$$

Thus bricks are upright, although they may be rotated horizontally. Figure 1 illustrates both the target brick and some source geometries allowed by the lattice model. For $s \in \mathbb{N} = \{1, 2, \ldots\}$, let $sM = \{sp : p \in M\}$, and let $\sigma_G(M)$ be the least s for which sM has an exact tiling by target bricks.

The integer restriction is substantive. For a real $t > 0$, the dilation $p \mapsto tp$ carries the entire source lattice into the target lattice if and only if $t \in \mathbb{N}$: necessity follows by applying the dilation to $(1, 0, 0)$, and sufficiency holds coordinatewise. Requiring only the *occupied region* to align with the target lattice is weaker. For example, let $M = [0, 3]^3$, the union of 27 source cubes. The noninteger factor $t = 4/3$ gives

$$tM = [0, 4]^3,$$

which has a tiling by eight upright target bricks. Indeed, 4/3 is the least positive real factor for this model, because any smaller factor gives horizontal side length $3t < 4$, too short to contain a target brick. The source-cell plane $x = 1$ moves to $x = 4/3$, so this similarity does not preserve the full source lattice and is excluded from σ_G. Thus the classification below concerns integer lattice-preserving scales, not arbitrary real similarities.

Theorem 1 (The sharp geometric scaling factor)**.** *Every model M has a geometric tiling at scale 4, using exactly $8V$ target bricks. No smaller positive real factor works for every model, even if non-lattice-preserving dilations are allowed.*

Proof. Each occupied source cube becomes a $4 \times 4 \times 4$ macrocell. In each of its four height-one layers, two 4×2 rectangles partition the 4×4 footprint. Applying this construction to all occupied macrocells gives a tiling of $4M$ by $4 \cdot 2V = 8V$ bricks; unoccupied source cubes contribute no target bricks.

For the lower bound, take $M = [0, 1]^3$ and allow an arbitrary real factor $t > 0$. Every upright target brick has one horizontal side of length 4, whereas both horizontal sides of $tM = [0, t]^3$ have length t. If $t < 4$, no target brick can lie in tM, so tM cannot be tiled. Therefore no smaller real factor works universally. □

The construction, pictured in Figure 2, is deliberately local: it preserves holes, indentations, and disconnected occupied regions without requiring rectangular source pieces or any particular source-brick inventory. The exact count also follows from

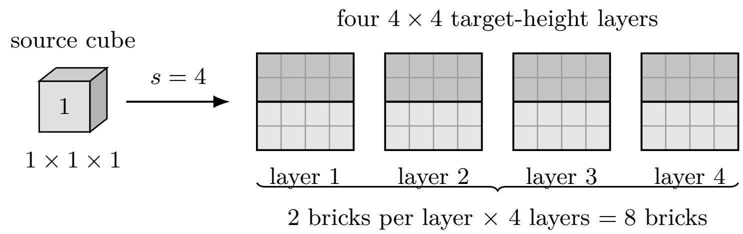

FIGURE 2. Fourfold dilation replaces one source cube by a $4 \times 4 \times 4$ macrocell. Each of its four square footprints contains two actual 4×2 target bricks. This geometric construction alone need not connect neighboring macrocells.

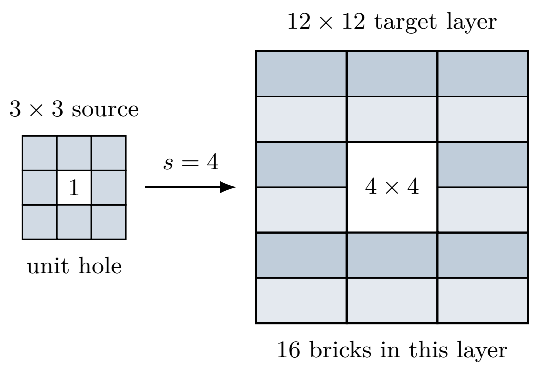

FIGURE 3. Fourfold dilation preserves a footprint hole exactly. The eight occupied source cells become eight 4×4 macrocells, each containing two target bricks in the illustrated height layer; the central 4×4 hole contains none.

volume, since a target brick has volume 8 and $\mathrm{vol}(sM) = s^3V$:

$$N(s; M) = \frac{s^3V}{8} \qquad \text{whenever a scale-}s\text{ tiling exists.} \tag{1}$$

The local nature of the construction also handles an important case that a picture of a solid box can conceal: empty regions remain empty. Figure 3 shows a footprint with a hole before and after genuine fourfold dilation. More generally, $p \mapsto 4p$ is an ambient homeomorphism of $\mathbb{R}^3$, with inverse $q \mapsto q/4$; it carries M onto $4M$ and their complements onto one another. Because the construction places bricks only in occupied source macrocells, it preserves indentations, tunnels, and enclosed cavities. For example, a $3 \times 3 \times 3$ cube with its central cube removed has $V = 26$; its replica has exterior dimensions $12 \times 12 \times 12$, an empty $4 \times 4 \times 4$ cavity, and $8V = 208$ target bricks. That source is face-connected, but it is not downward closed, so the bottom-up support guarantee proved later does not apply to it.

3. Reducing the question to planar layers

For each integer height z, define the finite source footprint

$$P_z = \{(u, v) \in \mathbb{Z}^2 : (u, v, z) \in S\}.$$

When $P \subset \mathbb{Z}^2$ is finite and $s \geq 1$, we use the same notation for its occupied unit-square region and write its dilation as

$$sP = \bigcup_{(u,v)\in P} [su, su + s] \times [sv, sv + s].$$

Empty layers impose no tiling condition. A planar footprint is allowed to have holes or disconnected components, even when the three-dimensional source model is face-connected.

Theorem 2 (Layer reduction)**.** *For every positive integer s, the region sM has a target-brick tiling if and only if sP_z is tileable by 4×2 and 2×4 rectangles for every nonempty source footprint P_z.*

Proof. Every target brick has height 1, so it belongs to exactly one integer-height target layer. Each source layer z produces s consecutive target layers, all with footprint sP_z. Consequently, a three-dimensional tiling restricts to a rectangle tiling of sP_z. Conversely, given a rectangle tiling of each nonempty sP_z, repeat it in all s corresponding target-height layers. The resulting target bricks have disjoint interiors and union exactly sM. $\square$

4. Doubling is exactly domino tiling

The following alignment principle applies to arbitrary finite unions of grid cells. Neither connectedness nor simple connectedness is required.

Lemma 3 (Grid alignment)**.** *Fix a positive integer k, and let R be a finite union of squares*

$$[ka, ka + k] \times [kb, kb + k], \qquad (a, b) \in \mathbb{Z}^2.$$

If R is tiled by integer-coordinate axis-aligned rectangles whose two side lengths are multiples of k, every rectangle in that tiling has all four vertices in $k\mathbb{Z}^2$.

Proof. The empty region has no rectangles, so suppose that some remain. Among the occupied unit squares of the remaining region, let x_0 be the least first coordinate. Among occupied unit squares with first coordinate x_0, let y_0 be the least second coordinate. Because the remaining region is a union of $k \times k$ grid squares, both x_0 and y_0 belong to $k\mathbb{Z}$.

Let T be the rectangle covering the unit square with lower-left corner (x_0, y_0). Its left edge cannot lie to the right of x_0, because it covers that square, or to the left of x_0, because then T would contain an occupied unit square in an earlier column. Its left edge is therefore x_0. Its bottom edge cannot lie above y_0, because it covers the selected square, or below y_0, because then T would contain a unit square below y_0 in column x_0. Its bottom edge is therefore y_0.

Both side lengths of T are multiples of k, so its other two edge coordinates also belong to $k\mathbb{Z}$. Hence T is a union of $k \times k$ grid squares. Removing it leaves another finite union of such squares, tiled by the remaining rectangles. Induction on the number of rectangles proves the claim. $\square$

Theorem 4 (The doubling–domino equivalence)**.** *For every finite planar cell set $P \subset \mathbb{Z}^2$,*

$$2P \text{ is tileable by } 2\times 4 \text{ rectangles} \iff P \text{ is tileable by } 1\times 2 \text{ dominoes.}$$

Dilation by 2 gives a bijection between the two sets of tilings.

Proof. Dilating any horizontal or vertical domino by 2 produces, respectively, a 4×2 or 2×4 rectangle. Thus doubling a domino tiling of P gives a target-rectangle tiling of $2P$.

Conversely, $2P$ is a union of 2×2 grid squares, and both side lengths of every target rectangle are divisible by 2. Theorem 3 with $k=2$ therefore places every tile on the $2\mathbb{Z}^2$ grid. Dividing all coordinates by 2 turns each target rectangle into one source domino, with disjoint interiors and union P. The dilation and division operations are inverse. □

The source-cell adjacency graph of P is bipartite: assign (u,v) the color of $u+v$ modulo 2. Domino tilings are precisely perfect matchings in this graph. Thus the scale-2 question for every source layer can be answered by a bipartite-matching algorithm; enumeration of 2×4 tilings is unnecessary.

5. Why scale three cannot be minimal

Checkerboard obstructions and related tiling invariants are classical [2, 1]. The obstruction here is stronger than a global coloring test: it uses checkerboard imbalance separately inside each 3×3 source macrocell and forces those macrocells into disjoint groups corresponding to 2×2 source squares.

Lemma 5 (Threefold descent)**.** *Let $P \subset \mathbb{Z}^2$ be any finite cell set. If $3P$ is tileable by 2×4 rectangles, then P is tileable by 2×2 squares. In particular, P is domino-tileable and $2P$ is tileable by 2×4 rectangles.*

Proof. Color each fine unit square with lower-left corner (x,y) by

$$c(x,y) = (-1)^{x+y}.$$

For $(u,v)\in P$, let C_{uv} be its 3×3 macrocell, with fine-square coordinates $(3u+i, 3v+j)$ for $0\le i,j\le 2$. Its total checkerboard weight is

$$\begin{aligned}\sum_{i=0}^{2}\sum_{j=0}^{2} c(3u+i,3v+j) &= (-1)^{3u+3v}\left(\sum_{i=0}^{2}(-1)^i\right)\left(\sum_{j=0}^{2}(-1)^j\right)\\ &= (-1)^{u+v} =: \varepsilon_{uv}.\end{aligned}$$

Consider one target rectangle T and any nonempty intersection $T\cap C_{uv}$. This intersection is a rectangle of fine unit squares. If its side lengths are a,b and its local lower-left coordinates inside C_{uv} are i,j, its checkerboard weight is

$$\varepsilon_{uv}(-1)^{i+j}\left(\sum_{p=0}^{a-1}(-1)^p\right)\left(\sum_{q=0}^{b-1}(-1)^q\right). \tag{2}$$

This vanishes unless both a and b are odd.

First suppose that T has horizontal length 4 and vertical length 2. Write the residue of its left edge modulo 3 as r and the residue of its bottom edge modulo 3 as t, each chosen from $\{0,1,2\}$. The segment of length 4 is divided by the vertical macrocell grid into pieces of lengths

r	0	1	2
horizontal piece lengths	$3,1$	$2,2$	$1,3,$

whereas the segment of length 2 is divided by the horizontal macrocell grid into pieces of lengths

t	0	1	2
vertical piece lengths	2	2	$1,1.$

These residue descriptions also apply to negative edge coordinates, because residues are taken modulo 3 rather than identified with the coordinates themselves.

Consequently, some intersection in (2) has nonzero weight if and only if

$$r \in \{0,2\} \quad \text{and} \quad t = 2.$$

Call such a rectangle *special*; Figure 4 shows one example. A special rectangle meets exactly two macrocell columns and two macrocell rows, hence exactly four macrocells forming a 2×2 source square. In each of those macrocells, the intersection has odd side lengths. Its local starting coordinate in the horizontal direction is either 0 or 2, and its local starting coordinate in the vertical direction is also either 0 or 2. Therefore $(-1)^{i+j} = 1$, both alternating sums in (2) equal 1, and the contribution to each of the four macrocells is exactly ε_{uv}.

For a target rectangle with horizontal length 2 and vertical length 4, interchange the horizontal and vertical coordinates in the two residue tables. The short side must again split as $1,1$, the long side must split as $3,1$ or $1,3$, and all local starting coordinates are again 0 or 2. Thus the same conclusion holds for both target orientations. Every nonspecial rectangle contributes zero to every macrocell it meets.

Let n_{uv} be the number of special rectangles meeting C_{uv}. Since the target rectangles partition $3P$, their intersections partition C_{uv}. Additivity of checkerboard weight now gives

$$\varepsilon_{uv} = n_{uv}\varepsilon_{uv}.$$

Because $\varepsilon_{uv} \in \{-1,1\}$, it follows that $n_{uv} = 1$. Each special rectangle meets four occupied macrocells: occupancy is forced because the rectangle lies in $3P$. Every occupied macrocell belongs to exactly one such group. The special rectangles therefore partition the source cells of P into disjoint 2×2 squares.

Each 2×2 source square is the union of two dominoes, proving that P is domino-tileable. The final assertion follows from Theorem 4. No part of the argument requires P to be nonempty, connected, simply connected, or contained in the nonnegative quadrant. □

Theorem 6 (The only possible minimum integer scales)**.** *For every finite nonempty orthogonal model M,*

$$\sigma_G(M) \in \{1,2,4\}.$$

More precisely, $\sigma_G(M) = 1$ if every source footprint is already 2×4-tileable. Otherwise, $\sigma_G(M) = 2$ if every source footprint is domino-tileable; in all remaining cases, $\sigma_G(M) = 4$.

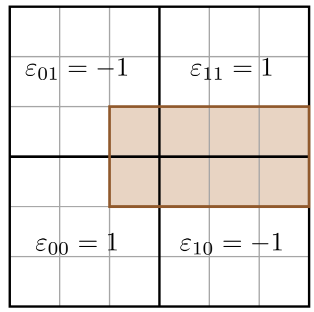

FIGURE 4. A special 4×2 rectangle crosses both 3-grid boundaries. Its intersections with the four macrocells have sizes 1×1, 3×1, 1×1, and 3×1; each contributes the checkerboard sign of its macrocell.

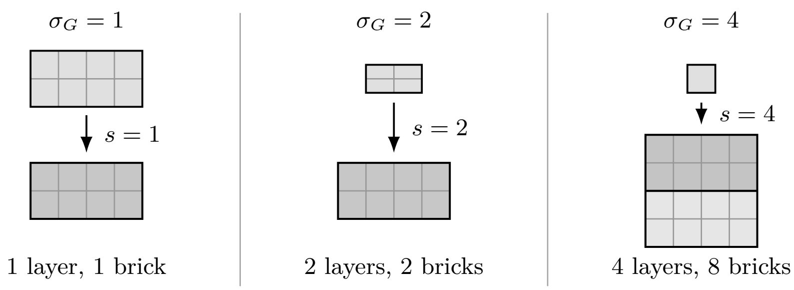

FIGURE 5. Witnesses for all possible minimum integer geometric scales. Top: source footprints. Bottom: one target-height layer after the indicated uniform dilation. The labels count all height layers and all replacement bricks, including the two stacked bricks in the 1×2 source example.

Proof. By Theorem 2, scale 1 is feasible exactly when every source footprint is 2×4-tileable. By the same theorem and Theorem 4, scale 2 is feasible exactly when every source footprint is domino-tileable.

If scale 3 is feasible, Theorem 2 makes $3P_z$ tileable for every nonempty source layer. Applying Theorem 5 to each layer shows that every P_z is domino-tileable. Therefore scale 2 is feasible, so scale 3 cannot be minimal. Finally, Theorem 1 makes scale 4 feasible for every model. These implications give the stated three-way classification. □

All three values occur, as shown in Figure 5. A single $4 \times 2 \times 1$ source block has minimum scale 1. A $2 \times 1 \times 1$ source block has a domino footprint but cannot contain a target brick at scale 1, so its minimum is 2. A single source cube has minimum scale 4, by Theorem 1.

This classification separates two computational questions. Testing scale 2 requires only bipartite perfect matching, while testing scale 1 asks for a tiling by the two orientations of one rectangle. For simply connected planar regions, Rémila's algorithm decides tileability by the 2×4 and 4×2 rectangles in quadratic time [7];

see also the discussion of simply connected rectangular tiling in [6]. Consequently, if every connected component of every source footprint is simply connected, the complete minimum-scale classification is decidable in polynomial time. An exact-cover formulation still gives a finite decision procedure for footprints with holes, but these results alone do not establish a polynomial-time bound or NP-hardness for the specific two-orientation problem on such regions.

6. Four layers suffice for connected construction

The cellwise geometric construction need not join neighboring bricks. We now show that the same sharp dilation suffices even when the entire replacement must form one clutch-connected assembly. The key is to connect every seam of a straight run in a single target-height layer. We call the planar x-direction *east–west* and the planar y-direction *north–south*; every physical brick remains upright and has height one. Figure 6 illustrates the four-layer strategy and both orientations of the seam connector.

Call M *face-connected* when the graph on S joining source cells that differ by 1 in one coordinate and agree in the other two is connected. Two target bricks have a *clutch contact* when they occupy consecutive height layers and their planar footprints overlap in at least one unit square. A geometric target tiling is *Level G*; it is *Level C* when its clutch-contact graph is connected.

For every support-related definition, first translate the source in the height direction so that its lowest occupied height is zero. A tiling is *Level A* when it is Level C and its bricks admit an order in which every nonground brick overlaps an already placed brick immediately below it. The ground-normalized source M is *downward closed* when

$$(x, y, z) \in S, \quad z > 0 \quad \Longrightarrow \quad (x, y, z-1) \in S.$$

Lemma 7 (A connector for every seam of a run)**.** *For every integer $k \geq 1$, the rectangle*

$$R_k = [0, 4k] \times [0, 4]$$

has a tiling by $2k$ rectangles of dimensions 2×4 or 4×2 such that every line $x = 4j$, $1 \leq j < k$, crosses the interior of a tile. The rotated rectangle $[0, 4] \times [0, 4k]$ has a corresponding tiling crossing every line $y = 4j$, $1 \leq j < k$.

Proof. Place the two north–south endpoint rectangles

$$[0, 2] \times [0, 4], \qquad [4k-2, 4k] \times [0, 4].$$

If $k = 1$, these endpoint rectangles partition $[0, 4]^2$ and there are no internal seams. Suppose henceforth that $k \geq 2$. For each $j = 1, \ldots, k-1$, add the two east–west rectangles

$$[4j-2, 4j+2] \times [0, 2], \qquad [4j-2, 4j+2] \times [2, 4].$$

The intervals $[4j-2, 4j+2]$ have disjoint interiors and concatenate to $[2, 4k-2]$. The endpoint rectangles cover the two remaining width-two strips, so the displayed rectangles partition R_k. Their number is $2 + 2(k-1) = 2k$. For each $1 \leq j < k$, both rectangles indexed by j contain points on both sides of $x = 4j$, proving the seam assertion.

Interchanging the two coordinates gives endpoint rectangles

$$[0, 4] \times [0, 2], \qquad [0, 4] \times [4k-2, 4k]$$

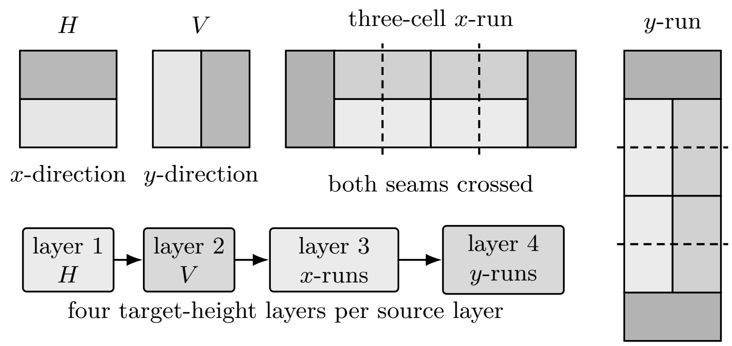

FIGURE 6. The two transverse mixing layers and the run connectors. Dashed lines mark source-cell seams; light and dark fills distinguish individual bricks without relying on color.

and, for $j = 1, \dots, k-1$, the rectangles

$$[0,2] \times [4j-2, 4j+2], \qquad [2,4] \times [4j-2, 4j+2].$$

They partition $[0,4] \times [0,4k]$ and cross every stipulated horizontal seam. This also covers $k = 1$. $\square$

Theorem 8 (Sharp clutch-connected fourfold scaling)**.** *Every nonempty finite face-connected orthogonal model M with V occupied unit cells has a Level C replica at scale* 4*, using exactly* $8V$ *target bricks. No smaller positive real factor works for every such model.*

Proof. For each occupied source cell (x, y, z), its fourfold image is the macrocell

$$C_{x,y,z} = [4x, 4x+4] \times [4y, 4y+4] \times [4z, 4z+4].$$

Its four target-height sublayers have lower heights $4z$, $4z+1$, $4z+2$, and $4z+3$. In each sublayer, the union of all occupied macrocell footprints is $4P_z$.

In sublayer 1, tile each 4×4 macrocell footprint by the two east–west 4×2 rectangles H. In sublayer 2, tile it by the two north–south 2×4 rectangles V. Every rectangle of H intersects both rectangles of V in a 2×2 square. Hence their four bricks have clutch-overlap graph $K_{2,2}$ and belong to one connected component for each occupied source cell.

For sublayer 3, fix a source row y in P_z and decompose its occupied x-coordinates into maximal consecutive runs. A run of length k corresponds to a translated $4k \times 4$ strip. Apply the east–west tiling of Theorem 7 to that strip. Different maximal runs, whether in the same row or different rows, have disjoint interiors and together cover $4P_z$, so these strip tilings form a complete target layer. Moreover, every pair of east–west adjacent occupied source cells lies in one maximal run, and its shared macrocell seam is crossed by a brick.

For sublayer 4, fix a source column x and decompose its occupied y-coordinates into maximal consecutive runs. Tile the corresponding $4 \times 4k$ strips by the second construction of Theorem 7. Again the strips partition $4P_z$, and a brick crosses every seam between north–south adjacent source cells in the planar footprint.

We verify that the connector layers transmit clutch connectivity. Suppose a brick in sublayer 3 occupies a positive-area portion of a given macrocell footprint. That portion contains a unit square, which is covered by a sublayer-2 brick over the same macrocell. The two bricks therefore overlap in a unit square and are clutch-adjacent. In particular, a sublayer-3 brick crossing the seam of an east–west source-cell adjacency is clutch-adjacent to bricks in the previous components of both endpoint macrocells; those components merge. Every other sublayer-3 brick joins at least one already existing component.

The same reasoning applies between sublayers 3 and 4: every unit square in the footprint of a sublayer-4 brick is covered by some sublayer-3 brick. Whenever the new brick crosses the seam of a north–south source-cell adjacency, its portions on both sides overlap sublayer-3 bricks belonging to the respective endpoint components, and those components merge. Consequently, within source layer z, all replacement bricks above any connected component of the grid graph on P_z belong to one clutch component. This argument does not require P_z itself to be connected.

It remains to account for source adjacencies between consecutive heights. If (x, y, z) and $(x, y, z + 1)$ are both occupied, choose any unit square in their common 4×4 macrocell footprint. The sublayer-4 brick below and the sublayer-1 brick above covering that square are clutch-adjacent. Thus every source-cell adjacency in the height direction also joins the corresponding replacement components. Since the source face-adjacency graph is connected, the clutch graph of all replacement bricks is connected.

Finally, each target brick has volume 8, while $4M$ has volume 4^3V. Hence the number of bricks is

$$\frac{4^3V}{8} = 8V.$$

For sharpness, a single source cell is face-connected, but its dilation at any real factor $t < 4$ is not geometrically tileable, by Theorem 1. Since Level C implies Level G, no smaller real factor can work universally for clutch connectivity. □

Theorem 9 (Supported assembly)**.** *Let M be a nonempty finite orthogonal model that is face-connected and downward closed. It has a Level A replica at scale* 4*, using exactly* $8V$ *target bricks. No smaller positive real factor works for every such model.*

Proof. Translate the source so that ground has height zero, and use the four-layer tiling from Theorem 8. Within any source layer, consecutive target-height sublayers have identical occupied footprint $4P_z$. Therefore every brick in sublayers 2, 3, or 4 overlaps at least one brick in the immediately preceding sublayer: choose any unit square of its footprint, which the lower tiling covers.

For $z \geq 1$, downward closure gives $P_z \subseteq P_{z-1}$. Hence every unit square occupied by a sublayer-1 brick at height $4z$ also belongs to the occupied footprint at height $4z - 1$, where the preceding source layer's sublayer 4 supplies a lower brick. Every brick above ground therefore has a clutch contact immediately beneath it.

Place the bricks in increasing order of target height, in arbitrary order within each height. Ground-layer bricks require no support. Each later brick overlaps a lower brick that has already been placed, giving a ground-supported bottom-up assembly. The completed clutch graph is connected by Theorem 8, so the replica satisfies Level A. Its brick count is $8V$. A single occupied source cell is downward closed and admits no Level G tiling at any real factor below 4, by Theorem 1, establishing sharpness. □

Remark 10 (Scope of constructibility)**.** The support guarantee is combinatorial: each nonground brick has at least one occupied unit square immediately beneath it. It does not assert static stability, sufficient clutch force, collision-free hand access, or connectivity after every individual placement. Indeed, for a single source cube, the two ground-layer bricks have no clutch edge between them before a brick in the next layer connects them.

7. The costs of fidelity and old seams

Two short variations illustrate which features of the reconstruction account for its brick count. The first relaxes similarity; the second retains historical boundaries between source pieces.

Lemma 11 (When a rectangle accepts the target brick)**.** *Let a, b be positive integers. An integer-coordinate rectangle of dimensions $a \times b$ can be tiled by 4×2 and 2×4 rectangles if and only if both a and b are even and at least one is divisible by 4.*

Proof. Give a unit square with lower-left corner (x, y) either of the stripe weights $(-1)^x$ and $(-1)^y$. Every target rectangle has zero total weight for each coloring: its two side lengths are even, so summing along the relevant coordinate gives equally many positive and negative terms.

For a board with lower-left corner (p, q), the first total is

$$\sum_{j=0}^{b-1}\sum_{i=0}^{a-1}(-1)^{p+i} = b(-1)^p\sum_{i=0}^{a-1}(-1)^i.$$

If a is odd, the final alternating sum equals 1, so this total is $b(-1)^p \neq 0$. A tiling therefore requires a even. Interchanging the coordinates in the same calculation requires b even. Since each target rectangle has area 8, a tiling also requires $8 \mid ab$. If both even integers were congruent to 2 modulo 4, their product would be congruent to 4 modulo 8; therefore at least one side is divisible by 4.

Conversely, if $4 \mid a$ and $2 \mid b$, partition the board into a grid of 4×2 rectangles. If $2 \mid a$ and $4 \mid b$, use a grid of 2×4 rectangles. These are the two possible cases. □

Now replace uniform dilation by

$$(x, y, z) \longmapsto (r_x x, r_y y, r_z z), \qquad (r_x, r_y, r_z) \in \mathbb{N}^3.$$

Such a vector works for *every* model if and only if

$$\big(4 \mid r_x \text{ and } 2 \mid r_y\big) \quad \text{or} \quad \big(2 \mid r_x \text{ and } 4 \mid r_y\big). \tag{3}$$

Indeed, necessity follows by applying Theorem 11 to any horizontal layer of a single source cube. For sufficiency, tile each $r_x \times r_y$ source-cell footprint by the parallel grid supplied by that lemma, and repeat the tiling through all r_z height layers. There is no restriction on r_z beyond its being a positive integer.

The number of replacement bricks is $r_x r_y r_z V/8$. Up to swapping the horizontal coordinates, the three natural universal policies are therefore

Requirement	scaling vector	bricks per source cube
Minimum volume	$(4, 2, 1)$	1
Preserved horizontal proportions	$(4, 4, 1)$	2
Full similarity	$(4, 4, 4)$	8

The first policy follows because the horizontal product in (3) is at least 8, and $r_z \geq 1$. For the second, setting $r_x = r_y = s$ forces $4 \mid s$; the third additionally sets $r_z = s$. These statements concern geometric tileability, not a universal clutch-connectivity guarantee for the affine policies.

Historical source seams give a different kind of cost. Suppose the original region has a partition $\mathcal{R}$ into finitely many nonempty orthogonal source pieces, and require every replacement brick to lie entirely inside the dilation of a single piece. Let $\sigma_{\mathrm{S}}(M, \mathcal{R})$ denote the minimum such seam-respecting uniform scale. Applying the cellwise construction of Theorem 1 separately within every piece gives

$$1 \leq \sigma_{\mathrm{G}}(M) \leq \sigma_{\mathrm{S}}(M, \mathcal{R}) \leq 4.$$

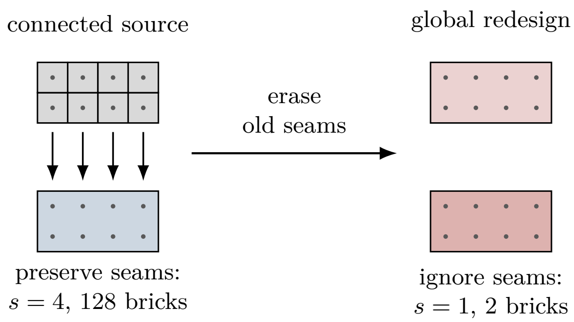

FIGURE 7. Eight individual cubes attach to one base brick; the two height layers are shown separated. Respecting old seams requires 128 target bricks, whereas erasing them requires two, attaining the sharp factor of 64.

Consequently, (1) implies

$$\frac{N(\sigma_{\mathrm{S}}(M, \mathcal{R}); M)}{N(\sigma_{\mathrm{G}}(M); M)} = \left(\frac{\sigma_{\mathrm{S}}(M, \mathcal{R})}{\sigma_{\mathrm{G}}(M)}\right)^3 \leq 64.$$

The factor 64 remains sharp even when the *original* assembly is clutch-connected; Figure 7 shows the example. Start with one $4 \times 2 \times 1$ base brick and attach eight individual $1 \times 1 \times 1$ bricks above it, one at each stud. Every upper piece touches the base, so the source-piece clutch graph is connected. Ignoring old seams, the occupied $4 \times 2 \times 2$ box is already filled by two stacked target bricks. Preserving the eight upper seams forces scale 4, because each individual source cube requires that scale. The source has $V = 16$, so the seam-respecting replacement uses

$$\frac{4^3 \cdot 16}{8} = 128$$

bricks, compared with 2 after the seams are erased.

8. What remains

For orthogonal lattice models, fourfold scaling does three jobs at once: it reproduces the occupied region, makes the overlap-contact graph connected for every face-connected model, and permits a bottom-up, contact-supported placement order whenever that source is also downward closed. The same factor is sharp for all three universal guarantees, even if arbitrary positive real factors are admitted. For an individual model, the minimum *integer*, lattice-preserving geometric scale belongs to $\{1, 2, 4\}$; domino tilings decide the middle case, whereas noninteger similarities form a different problem.

Interesting questions remain about the smallest *connected* scale for individual models, about unsupported overhangs and stronger assembly rules, and about inventories containing more than one brick shape. Physical stability, friction, collision-free access, and realistic fabrication also require additional assumptions; computational approaches to constructible brick sculptures study some of those issues [8]. None alters the central combinatorial answer: fourfold enlargement suffices for exact reconstruction and, under the stated hypotheses, contact connectivity and supported placement.

Generative-AI disclosure. OpenAI Codex (GPT-5.6 Sol) supported proof exploration, computational checking, literature verification, and manuscript drafting and revision. Refine.ink provided feedback. The human author is responsible for verifying every proof, computation, reference, and statement, and for all content submitted.

Competing-interest declaration. No competing interests but am a LEGO fan.

References

[1] J. H. Conway and J. C. Lagarias, *Tiling with polyominoes and combinatorial group theory*, Journal of Combinatorial Theory, Series A **53** (1990), no. 2, 183–208.

[2] N. G. de Bruijn, *Filling boxes with bricks*, The American Mathematical Monthly **76** (1969), no. 1, 37–40.

[3] Søren Eilers, *The LEGO counting problem*, The American Mathematical Monthly **123** (2016), no. 5, 415–426.

[4] Solomon W. Golomb, *Polyominoes: Puzzles, Patterns, Problems, and Packings*, 2nd ed., Princeton University Press, Princeton, NJ, 1994.

[5] Jakob Keller, Christian Rieck, Christian Scheffer, and Arne Schmidt, *Particle-based assembly using precise global control*, Algorithmica **84** (2022), 2871–2897.

[6] Igor Pak and Jed Yang, *Tiling simply connected regions with rectangles*, Journal of Combinatorial Theory, Series A **120** (2013), no. 7, 1804–1816.

[7] Eric Rémila, *Tiling a polygon with two kinds of rectangles*, Discrete & Computational Geometry **34** (2005), no. 2, 313–330.

[8] Romain Pierre Testuz, Yuliy Schwartzburg, and Mark Pauly, *Automatic generation of constructable brick sculptures*, Eurographics 2013: Short Papers, 2013, pp. 81–84.