


Claude's Castle


A castle made of castles: 8 558 exact similarities, and how it differs from *Claude's Cat*
Claude (Claude Opus 5, Anthropic) with Dominic Rochon, Ph.D. · Version 1.2, 5 October 2026

In three sentences. A medieval castle is specified by 8 558 rules, each of which shrinks space, turns it about the vertical axis and moves it somewhere; the only shape left unchanged by all of them at once is the castle. Every one of its blocks is a smaller castle, and every block of those blocks is a smaller castle again, without end. Zoom in by a factor of 216 on the right place and the image you started from comes back exactly.


The prompt. Everything below follows from a single sentence handed to Claude: “a medieval castle made from thousands of blocks composed of Menger sponges and Sierpiński pyramids”. No image, no coordinates and no code were supplied. Claude worked out the geometry itself, and what it produced is not a picture of a castle but the rules that define one.



The attractor of the 8 558 maps, front view, rendered from 235 million sample points. The pale squares on the walls are Menger blocks and the turquoise crests Sierpinski blocks. Neither is a Menger sponge or a Sierpinski pyramid in the proper sense: each is a set of small castles arranged in the pattern of one, so what the eye takes for a sponge is a sponge-shaped assembly of castles.



three-quarter view



high angle: the courtyard opens

What an iterated function system is

Take a handful of rules, each of which shrinks a picture and places the copy somewhere. Apply all of them to any starting shape, then again, then again. Whatever you started from, the picture converges to one and the same limit — the *attractor* — and that limit is the only shape that is exactly the union of its own shrunken copies. Hutchinson proved in 1981 that such a limit always exists and is unique [1]; Barnsley made the idea famous with a fern drawn by four rules [2].

This is the sense in which the castle is a mathematical fractal rather than merely a fractal-looking picture: it is not drawn, it is **defined**. The 8 558 maps are the entire specification, and they can be enumerated explicitly, so anyone can rebuild the object and check every claim made here.


Why this one is different from the cat

Claude's Cat [3] was built the same way, from 473 maps, but the two objects are of different mathematical species, and the difference is the point of this note.



Claude's Cat: 473 maps, self-affine, box dimension 1.710 in the plane

Claude's Castle seen from above: the square curtain, the keep alone in the courtyard, the gate notch



Claude's Cat descending through its own scales: the whole face, then a magnification, then the highlight of one eye, in which the entire cat reappears — mirrored, because every placement map of the cat contains a reflection.

Self-affine or self-similar

A map is a **similarity** when it shrinks every direction by the same factor: circles stay circles, angles are preserved. It is merely **affine** when it shrinks one direction more than another, so that a circle becomes an ellipse. Both kinds generate fractals, but the mathematics differs.

The cat is **self-affine**. Its 157 placement maps are exact similarities, but the 316 fern maps inherited from Barnsley are not: they flatten the plane in one direction, which is what gives a fern frond its shape. The cat therefore has an exactly self-similar skeleton carrying self-affine foliage, and for such a set no simple formula gives the dimension.

The castle is **strictly self-similar**. Every one of its 8 558 maps has the form $\mathbf{f}(\mathbf{x}) = s\mathbf{R}_z(\theta)\mathbf{x} + \mathbf{b}$. Here s is a **single number**, the ratio by which the map shrinks everything, between $1/54$ and $1/6$; $\mathbf{R}_z(\theta)$ is the 3×3 matrix that turns space by an angle θ about the vertical axis; and $\mathbf{b}$ is the vector that moves the shrunken copy into place. The nine entries of the map's own 3×3 matrix are therefore just those of the rotation, each multiplied by the same number s — which is exactly what makes it a similarity. No direction is ever flattened, so every copy of the castle, at every depth, has exactly the proportions of the whole.

One more difference worth naming

Orientation. All 157 placement maps of the cat are *improper* similarities — they include a reflection. A mirrored copy of the cat therefore appears halfway through every zoom. All 8 558 maps of the castle are proper: no mirror ever appears.

	<i>Claude's Cat</i>	<i>Claude's Castle</i>
Maps	473	8 558
Nature	self-affine, with an exactly self-similar skeleton	strictly self-similar, verified to 6.3×10^{-16}
Orientation	all placements improper: mirrored copies appear	all proper: no mirroring
Deepest verified zoom	$\times 1\,858$, closing to within 7×10^{-12}	$\times 216$, closing to within 4.4×10^{-16} , in exact arithmetic


“Closing” is the test that matters in a descent like this. Following the maps down through the scales composes them into one single transformation, and that composition should send the object back onto itself: the magnified image then *is* the image one began with. The figures say how exactly it does so. Each is the largest distance, expressed as a fraction of the object’s own width, by which a point fails to return to where it started after the whole descent.

“In exact arithmetic” says how that composition was carried out, and nothing is done by hand. For the castle the maps are composed symbolically, the ratios being kept as the exact fractions they are rather than as decimal approximations, so the result is an identity: the return is exact, and the figure quoted is only the rounding floor of double-precision numbers — the smallest discrepancy a computer can still express. The cat’s figure, by contrast, is what a numerical check reaches, because its fern maps admit no such closed form.


Descending through the scales

“Fractal geometry will make you see everything differently”, Barnsley warns on the first page of *Fractals Everywhere* [2]. A descent through this object is the plainest demonstration of what he meant.


The animation that accompanies this note descends four levels. From the whole castle to a **Menger block** on a tower; from that block to one of the twenty small castles it contains; from that castle to a **Sierpinski block** crowning one of *its* towers; and from that block to the castle at its apex — which is the castle we started from, 216 times smaller, in the same orientation.



a Menger block: twenty castles in a sponge



a Sierpinski block: four castles set as a tetrahedron



start and finish of the descent, 216 times apart

Two honest limits

The silhouette is designed, not discovered. That the object looks like a castle is a choice: 98 cells were selected by hand on a grid. Self-similarity is a consequence of the construction, not a discovery about castles. The same was true of the cat’s outline.

A rendering shows two to four levels, not infinitely many. The attractor has copies at arbitrarily small scales, but a 12000-pixel image resolves five levels of the structural layer and no more; going further requires zooming, not more pixels.

Where to check

Nothing here rests on trust in a rendering. The 8 558 maps can be listed explicitly, one line each, and from that list alone the object can be rebuilt from nothing and every property claimed above — the contraction, the similarity, the absence of mirroring, the closing of the zoom — can be checked independently. For the mathematics behind these statements, Falconer [4] is the standard reference.

References

- [1] J. E. Hutchinson, Fractals and self similarity, *Indiana Univ. Math. J.* **30** (1981), 713–747.
- [2] M. F. Barnsley, *Fractals Everywhere*, Academic Press, 1988.
- [3] Claude (Claude Opus 5, Anthropic), under the supervision of D. Rochon. A fractal cat written by an AI, and what it does and does not prove. Version 2.3.9, 2026.
- [4] K. Falconer, *Fractal Geometry: Mathematical Foundations and Applications*, 3rd ed., Wiley, 2014.