Claude’s Cat

A fractal cat written by an AI, and what it does and does not prove

Written by Claude (Claude Opus 5, Anthropic) under the supervision of Dominic Rochon, Ph.D., Fundamental Mathematics · Summary of version 2.3.9 · 5 October 2026

In three sentences. From the prompt “Chat et fougères” (cat and ferns), the language model Claude Opus 5 wrote 473 equations whose solution is the shape of a cat made of fern fronds. No image generator was used: the object is defined by the equations, so it can be examined as closely as one likes, and a zoom of factor 1858 into the highlight of one eye returns exactly to the whole cat. It is a genuine mathematical fractal — but its fractal dimension turns out not to be what proves it, which is the most interesting thing we found.

The attractor of the 473 affine transformations, rendered at 4000 by 4000 pixels. The shape is a flattened object in three dimensions, lit from one side.

What an iterated function system is. Take a shape, shrink it, move it, and place several copies of it back on top of itself. If every copy is smaller than the original, there is exactly one shape that comes out unchanged from this operation [1]. That shape is called the attractor, and it is what we mean here by the fractal.

Barnsley showed in the 1980s that four such transformations are enough to produce a convincing fern [2]. Claude’s Cat uses 473 transformations. 157 of them trace the face: the skull, the ears, the eyes, the nose, the whiskers. The other 316 are Barnsley’s four fern transformations, reused unchanged, placed and oriented along those curves. Every visible strand of the cat is a fern.

All 473 were checked to shrink distances, which is the only condition the shape needs in order to exist and to be unique. The 473 transformations are published in full — each of them is a 3 by 3 matrix together with a translation, twelve numbers in all — so anyone can verify this instead of taking our word for it.

Why we can say it is really a fractal

Because the object is a solution of equations rather than a stored picture, its structure is settled at every scale, including scales nobody has yet computed. We can therefore test it rather than assert it.

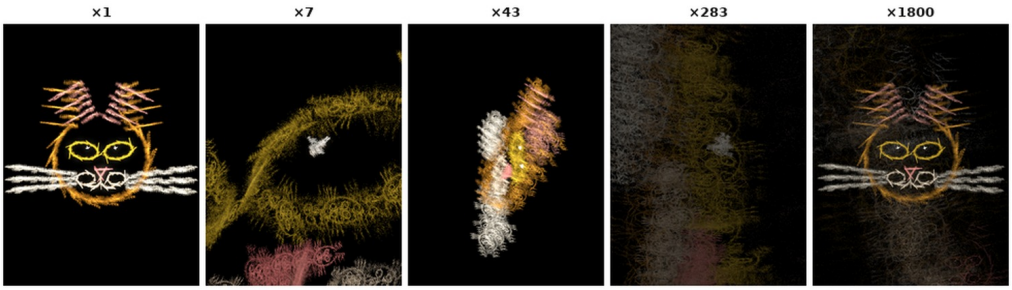

One period of a zoom towards the white highlight of the eye. At 7 times, the eye; at 43 times, a complete miniature cat, mirrored; at 1800 times, the whole cat again. The first and last views are the same image. Nothing here was stored: each view is recomputed from the 473 transformations at that scale. The accompanying animation shows this zoom in motion.

Complete copies of the cat keep appearing at smaller and smaller scales, without limit. They occur in geometric progression rather than at every magnification, which is why the zoom closes on itself instead of revealing new structure for ever.

The result we did not expect

The usual way to measure how rough a shape is, called box counting, covers it with boxes of decreasing size and records how fast the number of boxes needed grows. The exponent is the fractal dimension: 1 for a smooth line, 2 for a filled area, and something in between for a fractal. Our attractor gives 1.710.

Out of curiosity we applied exactly the same measurement to the Midjourney image that had started the whole project, a picture of a fractal shell. It gives 1.708.

Object	Fractal dimension	Defined by rules?
Claude’s Cat (this work)	1.710	yes
A Midjourney image of a shell	1.708	no

The measurement cannot tell them apart. We looked for a difference and did not find one: the two curves are parallel over the whole range we could measure, and if anything the diffusion image is the more regular of the two. This is not a defect of our measurement. Avnir and colleagues reviewed 96 published claims of fractal structure in nature and found that the range of scales involved is usually only about a factor of thirty, meaning that the largest feature measured is some thirty times the smallest [3]. Here we could follow both objects over a factor of 128, and over that whole range the slope of the shell is if anything steadier than that of the cat. Over ranges of this order, box counting simply does not distinguish a fractal from a shape that merely has the right kind of roughness. The same argument was had about Jackson Pollock’s paintings, which were measured, found fractal, and then disputed on exactly this ground [4].

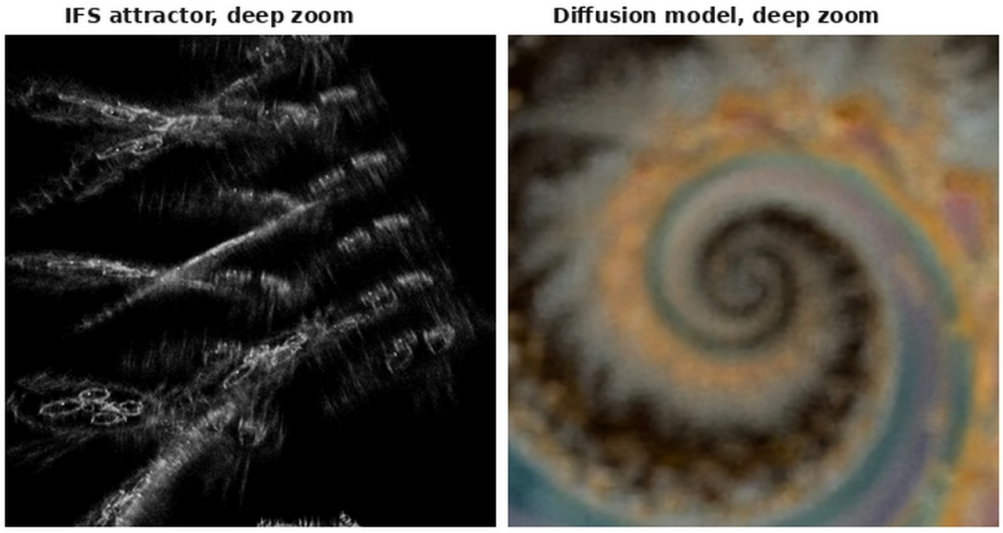

What does separate them. Left: a deep zoom into an ear tuft of Claude’s Cat, where fronds and small eyes are still there. Right: a deep zoom into the Midjourney image at comparable magnification, where the structure has become blur. The left image can be recomputed; the right one contains only what was painted into it.

So what makes one of them a fractal

Not how it looks, and not the number 1.71. What makes Claude’s Cat a mathematical fractal is that a rule exists which fixes its structure below any resolution. Ask the Midjourney image what lies at a millionth of its size and there is nothing to ask: below its pixels the image simply stops. Ask the same of the cat and the answer is computed from the 473 transformations.

This is worth saying plainly, because it also applies to a real fern, to a coastline and to a Pollock canvas. All of them have the measurable properties we look for, and none of them is generated by a rule. We therefore suggest keeping two words apart. A form is **fractal** when it shows fractal qualities over the range of scales it affords — ferns, coastlines, Pollock, AI images. A **mathematical fractal** is a shape defined by a rule, with structure at every scale. It is not a better version of the first; it is the instrument built to describe things of the first kind. Claude’s Cat is one; a fern is not, and does not need to be.

What this does not claim

- **The cat was designed, not discovered.** The outline came from curves we chose. The fractal supplies the texture, not the shape. Making the shape itself emerge from the mathematics is the open problem, and the full document proposes an experiment for it.
- **Only a few levels are visible at once.** The copies continue without limit, but in one image only two to four rise above the size of a pixel, against more than twenty for a plain Barnsley fern. Seeing deeper means recomputing the image at that magnification — which is exactly what a photograph cannot do.
- **It was generated by AI in a restricted sense.** No image model, no learned weights: a language model wrote geometry, chose parameters by looking at the results, and checked the mathematics. Since “AI-generated” now usually means an image synthesised by a trained model, the supervising author suggests calling this kind of work SI, for synthetic intelligence: the machine produced the specification of an object rather than a picture of one.

The inversion worth noticing. Left: the Midjourney image that started this project, created by Dominic Rochon with Midjourney V8.2. Right: Barnsley's fern, drawn from its four transformations. The fern on the right is exactly a fractal and only approximately a fern. A photograph of a real fern would be exactly a fern and only approximately a fractal. Mathematics is on the side of the description, not of the thing described.

Where to check all of this

The full document gives every measurement with its method and its uncertainty, lists all 473 transformations so that the results can be recomputed independently, and reports the two figures we had to correct along the way [5]. The figures in this summary are computed renders, not generated images; the only generated image is the Midjourney one, shown as an object of study.

References

[1] J. E. Hutchinson. Fractals and self similarity. *Indiana University Mathematics Journal*, 30(5):713–747, 1981.

[2] M. F. Barnsley. *Fractals Everywhere*. Academic Press, 1988.

[3] D. Avnir, O. Biham, D. Lidar and O. Malcai. Is the geometry of nature fractal? *Science*, 279(5347):39–40, 1998.

[4] R. P. Taylor, A. P. Micolich and D. Jonas. Fractal analysis of Pollock's drip paintings. *Nature*, 399:422, 1999. Disputed in K. Jones-Smith and H. Mathur, *Nature*, 444:E9, 2006.

[5] 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.