Transdimensional Painter — Sample Chapter

Chapter 1 — The Flash on the Lemniscate

By Kenneth A. Mendoza

2026-05-13

Chapter 1 — The Flash on the Lemniscate

It did not begin with an equation.

It began with a motion.

I was watching Dan Pike's lemniscate animations on Mathstodon — the kind of mathematical animation most people treat as a beautiful toy. A curve breathes. A level set folds. A polynomial field becomes visible for a second, then hides again behind the smooth authority of the screen.

But something in the animation did not feel decorative.

There was a flash on the level set.

Not a glitch, exactly. Not random jitter. It was the kind of motion a careful eye sees before the mind has language for it: the system started to misbehave, then snapped back into order. A local instability appeared, as if the orbit was about to leave the lawful track, and then the geometry seemed to correct it. The curve did not merely trace a path. It acted as if some deeper program had reached in and re-positioned the orbit back into normal behavior.

That was the first clue.

Most people would have watched the animation as a picture. I watched it as instrumentation.

That distinction matters. A picture is a surface. An instrument is a thing that lets reality push back. The moment the lemniscate pushed back, it stopped being visual ornament. It became a probe.

The question I had been asking — without quite knowing I was asking it — changed.

It changed from *What does this curve look like?* to *What computation is this curve performing?*

The answer did not arrive whole. It came as a collision of separate intuitions that had been building for years.

The first intuition was an old beat from earlier work. A system near transition does not simply move from one state to another. It hesitates. It dwells. It flickers. The warning is not in the final state; it is in the rhythm of almost changing. You do not catch the transition by staring at a static photograph. You catch it by listening to the timing. I had learned this years before in a different context entirely — the early-warning signal that lives in the timing of a heartbeat, in the dwell of an immune response, in the flicker of a climate system at the edge of a tipping point. The shape of the warning was always the same: not *what* but *when*, not *where* but *how long*.

The second intuition was a different animation altogether. Quadratic motion, drawn plainly enough that the hidden idea could be seen. What looks linear in one plane can become elliptical in another. A line, under the right lift, becomes an orbit. A simple rule, moved into the right space, becomes curvature.

That was enough.

I had already been living, for a long time, with the tension between two ways of looking at the world. Euclidean space is where the eye lives — points, distances, curves, bodies, orbits. Hilbert space is where the program lives — projections, modes, coefficients, inner products, squared error. Euclidean geometry shows you the object. Hilbert geometry shows you the computation that makes the object identifiable.

The lemniscate flash suggested that these were not two worlds.

They were two views of the same event.

In the Euclidean view, the orbit looked like a curve correcting itself. In the Hilbert view, the correction looked like a projection: a bad component being squeezed out, a mode being re-weighted, a coefficient returning to the subspace where the system could still be decoded.

The animation was no longer just a curve. It was a translator.

That is where Transdimensional Painter was born.

Not as a brand. Not as a theorem. As an instrument.

The painter was the missing object between the visible orbit and the hidden mode. It asked: if a dynamical system is painting its behavior in ordinary space, what higher-dimensional brush is actually moving? What if the apparent curve is only the shadow of a modal computation? What if the right coordinates are not the coordinates of the body, but the coordinates of the rule?

This is the first-principles hinge. A dynamical system is not only a path through space. It is a compression program whose output happens to look like a path through space. Once you see that, even the oldest problems change character.

Take the N-body problem, the problem of figuring out where three or more bodies go under their mutual gravity. The problem is more than three centuries old. It is hard. Henri Poincaré broke his teeth on it in the 1890s and discovered, in the breaking, what we now call chaos. But the question Poincaré was asking was the classical question: *Where do the bodies go?*

Transdimensional Painter asks a second question, underneath the first one. *What is the shortest modal program that makes the orbit intelligible?* Not just where do the bodies go, but what is the smallest hidden rule that explains why they go there?

This is where the project quietly touched two other lines of work I had been doing for years.

The first was a long, slow argument with the Erd■s problems. The Erd■s problems — Paul Erd■s's published list of unsolved combinatorial questions — are, on the surface, isolated puzzles about specific structures. Sidon sets. Prime gaps. Forbidden configurations. Lattice arrangements that cannot be too dense or too sparse without violating some rule. But sit with the list long enough and another structure emerges. Each Erd■s problem is a compression test in disguise. Each one asks whether a complicated combinatorial object can be described by a shorter law. Each one is a battlefield between raw enumeration and hidden structure.

The lemniscate flash made the same fight visible in motion.

Raw coordinates say one thing: *Here are the positions, here are the velocities, track everything*. The painter says something different: *Find the modes. Find the projection. Find the smaller program.*

The second line of work the painter touched was the strange repetition I had noticed across domains. The same shape of problem kept appearing in completely unrelated fields. In cardiac signal processing, it appeared as a problem of false alarms and dwell time — when does a transient become a true event, and when is it just noise that almost crossed the threshold? In climate-system early warning, it appeared as a problem of tipping signals and complexity drops — what statistical signature precedes a system's collapse into a new regime? In the finite Hilbert-space approximations I was building for the formal-mathematics work, it appeared as coefficient error and projection stability — when does a truncation tell you the truth about an infinite operator, and when does it lie? In the N-body work, it appeared as known orbits whose visible choreography might be a low-dimensional modal sentence written in a higher-dimensional language.

The outsider advantage in this kind of work is not what people sometimes think it is. It is not ignorance of the field. It is freedom from the field's default eyesight.

An insider sees the accepted categories first: numerical integration, orbit classification, Koopman modes, Hilbert spaces, minimum description length, formal proof, cardiac signal processing. Each category has its own gatekeepers and its own polite vocabulary. The polite vocabulary is the field's working language, and it is the field's working language because it has been refined by generations of careful thinkers. Most of the time the polite vocabulary is the right thing to reach for.

But sometimes it is the wrong thing to reach for. Sometimes the polite vocabulary smuggles assumptions that prevent you from seeing what is actually happening in front of you. My instinct, raised on a combination of laboratory training and self-direction, is to look before the vocabulary hardens. If the animation behaves like a correction, treat it as a correction. If a line becomes an ellipse under a quadratic lift, treat that as a clue about hidden coordinates. If a heartbeat warning lives in the dwell time, do not throw away timing as noise.

This is not an argument against expertise. The framework this book describes is built directly on top of work by mathematicians and physicists whose expertise is unimpeachable — Koopman, von Neumann, Riemann, Hardy, Bourbaki, the entire chain. The outsider advantage is narrower than that. It is the freedom to ask whether a problem belongs to the category the field has put it in. It is the freedom to suggest a different category. It is, in the language a graduate advisor of mine once used about a different question entirely, the freedom to *use your coat hangers wisely* — to not insist on the category just because the category is what you were taught.

The method was crude at first only in the sense that a blade is crude before sharpening.

The discipline was the opposite of crude.

Do not fool yourself. Bind the claim to an artifact. Let the failed rows stay failed. Use the trusted reference before claiming speed. Keep the theorem ceiling lower than the intuition.

That discipline is what turned the flash into steel.

The first runs of the N-body machinery did not announce a discovery. They did something more useful. They made the instrument behave in public-readable conditions while staying local and honest. Three known three-body orbits became a smoke test against the standard physics-community reference integrator. Eighteen source-verifiable rows became a catalog gate. Then sixty rows. Then one hundred and one, with a trajectory tensor that could be loaded in a browser and inspected by any reader with a network connection.

The result was not a sentence I had been wanting to write.

The result was not *We found new orbits*.

The result was something more careful, and more useful. *We can turn the known orbit literature into a reproducible compute instrument*. Not a discovery. An instrument. A machine that takes published orbits and turns them into something the next experiment can build on, with every claim bound to a public artifact that anyone can audit.

That is a different kind of power. It is the slower kind. It is also the kind that compounds.

It means the project can now ask sharper questions. Not vague questions about beauty. Not grand claims about solving the N-body problem. Concrete, narrower, more dangerous questions. Which known orbits survive stricter invariant gates? Which fail, and why? Which modal coordinates compress the trajectories — and which do not? Does the framework's prediction match the standard reference integrator at machine precision? Where does the simple coordinate-based description fail? Where does the mode-aware description recover structure?

The book enters here.

The book is not just the story of a working researcher watching curves on a screen and having an idea. It is the story of an idea becoming an instrument. The chapter cannot end with inspiration. It has to end with the reader able to interrogate the evidence. That is why the book carries with it, as a parallel structure, the methodological discipline that turned the inspiration into something an outsider could check.

That discipline is the through-line.

Run the orbit. Save the tensor. Hash the result. Keep the failed rows. Compare to the reference. Then, and only then, let the story speak.

Transdimensional Painter began as a flash on a lemniscate.

It became a question about whether nature paints in one space while computing in another.

Now it is becoming a machine that can be inspected.

That is where this book starts.

Not with proof.

With the moment the curve moved wrong, snapped right, and refused to be just a picture.

Sample chapter ends. The book continues with Chapter 2 — Two Geometries, One Event, which develops the dual-view idea introduced here into the operational language the rest of the book uses.

For the full manuscript or for further sample chapters, please contact ken@kenmendoza.com.