The Soliton and the Cloud: A New Intuition for the Quantum Particle

Lee Bloomquist, with Claude (Anthropic)

ABSTRACT

The two inherited images of the quantum particle — a literal wave in some medium, alternating with a literal point; or Bohm's more sophisticated picture of a point particle steered by a "pilot wave," like a ship navigated by a radio signal — both keep the particle *and* something wave- like as separate ontological citizens. This paper proposes a third image. The quantum particle is a single point, a soliton of zero classical action, moving through a nonstandard halo of infinitesimally close possibilities. What looks, after the fact, like a "wave" or a "cloud" is nothing but the trace of one particle's jumps, each jump too small to register as a jump, so that the record of many jumps reads as a continuous spread. We give the image a formal home in five pieces of existing mathematics not usually brought together for this purpose: Aczel's non-well-founded sets and Lambek's Lemma, which let the particle's history be written as a self- generating coalgebraic process, provably the same shape at every point along its own unfolding; the Hamilton–Jacobi formulation of Bohmian mechanics, which supplies the actual zero-action surface the particle surfs on; Robinson's nonstandard analysis, which supplies the infinitesimal `dt` and the *halo* of trajectories infinitesimally close to the classical one; Barwise–Perry situation theory, which upgrades "points left behind" to "facts left behind," each one a partial, local, nested situation; and Barwise–Seligman channel theory, in which the Born rule appears not as a postulate bolted onto the formalism but as an *infomorphism* — a structure-preserving translation between the theorist's classification of outcomes and the experimenter's. We read a well-known Bohmian-trajectory figure through this new image, show where it agrees with, and where it diverges from, the pilot-wave account, and close by showing how the image's graphics attach directly to the `dt`-bearing equations of the innerProperTime (iPT) formalism.

1. Introduction: A Third Image

A NOTE ON TERMINOLOGY

A note on terminology, before anything else. This paper does not claim that the quantum particle obeys a nonlinear soliton equation — no Korteweg–de Vries, no nonlinear Schrödinger equation, no literal soliton dynamics borrowed from fluid mechanics or nonlinear optics. "Soliton" is used here only for the intuitive shape that word carries: a localized, self-sustaining disturbance that keeps its identity as it travels, as opposed to a dispersive wave that spreads out and loses its shape. We borrow that picture, not the equation that produces it elsewhere. The particle's actual

> *dynamics, throughout this paper, are ordinary linear quantum mechanics, in its Hamilton–Jacobi/guidance-equation formulation, given in full in §3.6 below.*

Two images of the quantum particle circulate. The oldest, and the one still handed to undergraduates, is *complementarity by substance swap*: sometimes the electron behaves like a wave in a medium, sometimes like a point particle, and which one shows up depends on what you measure. This image asks the electron to be two different kinds of thing at two different times, with no story connecting them.

The second, more disciplined image is Bohm's. Here the particle is always a point, with a definite position and a definite trajectory — but it is steered by a separate field, the "pilot wave," which itself obeys the Schrödinger equation. David Bohm's own analogy was a ship under autopilot, guided by a radio signal too weak to push it around but rich enough in information to steer it. This is a real improvement: it restores a continuous trajectory to the particle. But it does so by doubling the ontology — a point *and* a guiding wave, two entities, one of which (the wave) lives on configuration space and is never itself directly observed.

This paper introduces a third image, in which nothing needs to be added to the particle. **The quantum particle is a soliton of zero action, moving through a medial space, leaving behind a cloud of facts embedded in that space.** It is still, at every instant, one point. It is not "sometimes" anything else. What makes it look like a cloud, or a wave, is a matter of *resolution*, not of substance: the point is jumping among infinitesimally close possibilities so continuously that its trail, seen at ordinary (real-number) resolution, looks like a smear. Zoom the resolution back down to the nonstandard scale at which the jumps actually happen, and there is only ever the one point.

2. The Image, Stated Plainly

A soliton is a self-reinforcing wave that keeps its shape as it travels — it is a *wave*, but a wave that behaves, for many purposes, like a localized *particle*: it doesn't disperse, it has a well-defined position, it can collide with other solitons and come out the other side intact. This is already the right kind of object for what we want: something that is genuinely wave-like in its dynamics but genuinely particle-like in its localization, without needing to be two different things at two different times. As the note above insists, this is a borrowed shape, not a claimed equation.

The image proposed here: the quantum particle is a soliton of *zero action* — a localized disturbance that travels not through ordinary space but through the "medial space" of infinitesimally close possibilities, where the classical action functional, evaluated along any path in that disturbance's neighborhood, rounds off to the same value. It moves by jumping, at every instant, to another possibility infinitely close to the one it just actualized. Each jump deposits a fact — a *this happened, here* — at the point it just left. After enough jumps, looking back over the trail, those deposited facts appear as a continuous cloud. But the cloud is not made of many particles, or of a wave "in" anything. It is the retrospective appearance of one particle's jumping.

The sections that follow give this image the mathematics it needs to be more than a metaphor — first in the abstract, then, in §3.6–3.7, as a direct calculation inside ordinary Bohmian mechanics.

3. Mathematical Scaffolding

3.1 Non-well-founded sets

Ordinary set theory forbids a set from containing itself, directly or through any chain of memberships — this is the Axiom of Foundation. Peter Aczel's Anti-Foundation Axiom (AFA) replaces it: any graph of membership, including a circular one, denotes a unique set. This licenses equations like `Ω = {Ω}`, and more usefully, equations like

```
self = { possibility → fact, self }
```

in which a process is defined by unfolding into itself forever, with no first instant. This is exactly the shape the soliton's history needs: it is not a sequence built up from a starting point, it is a self-sustaining process, and the mathematics that can hold such a process without paradox is AFA, not ordinary Foundation. §3.7 below makes this exact.

3.2 Nonstandard analysis and the halo

Robinson's nonstandard analysis (NSA) extends the real numbers to the hyperreals, `*ℝ`, which properly contain actual infinitesimal and infinite numbers obeying all the ordinary laws of algebra. Every finite hyperreal has a unique **standard part** — the ordinary real number it is infinitely close to — and the set of all hyperreals infinitely close to a given real `r` is called the **halo** (or monad) of `r`.

This gives "zero action" a literal mechanism. The soliton doesn't travel along a single classical path; it travels through the halo of that path — a whole cloud of nonstandard trajectories, each differing from the classical one by an infinitesimal amount, every one of which has the same *standard part* of the action. "An infinitesimal wave front of zero action" is, precisely, the halo of the classical stationary-action trajectory.

3.3 The double halo

A single halo is a cloud around a point. A **double halo** is what you get when a point carries two such clouds at once, pointing in opposite directions: one generated by an infinitesimal step into the future (`dtF`) and one by an infinitesimal step into the past (`dtP`). Rather than a durationless instant, the "present" becomes a hinge — the joint between a future-halo and a past-halo, neither of which collapses into a single real-valued instant. This is the structure each of the soliton's jump-points actually has: not a bare dot, but a two-lobed hinge, with the `dtF` lobe opening toward what the particle could next become and the `dtP` lobe closing over what it just was. §3.6 shows exactly where these two lobes come from in ordinary Bohmian mechanics.

3.4 Situation theory

Barwise and Perry's situation theory replaces the idea of a single, maximal "possible world" with **situations**: partial, local, limited chunks of reality. A situation *supports* an **infon** — a discrete unit of information, $\langle\langle R, a_1, \ldots, a_n, i\rangle\rangle$ — when that piece of information holds within it. Situations can nest inside other situations.

This is what upgrades "a cloud of points" to "a cloud of facts." A bare point cannot carry information; a situation can. Each hinge the soliton passes through is not merely a location — it is a situation supporting the infon "the particle was here, with this momentum, in this phase," and the whole trail is a nested situation containing all the local ones. This nesting is why the AFA framing of §3.1 is not decorative: the cloud and the process generating it are mutually defined, and neither bottoms out at a first fact.

3.5 Channel theory and infomorphisms

Barwise and Seligman's channel theory formalizes how information flows *between* situations. A **classification** consists of tokens, types, and a relation saying which tokens are of which types. An **infomorphism** between two classifications is a pair of structure-preserving maps running in opposite directions — types pushed forward, tokens pulled back — satisfying a condition that keeps classification consistent across the translation. A **channel** connects classifications through a shared core.

3.6 The Hamilton–Jacobi route: where `dtF` and `dtP` actually come from

Sections 3.2–3.3 described the double halo abstractly, as a general nonstandard-analytic structure. It is worth showing, concretely, exactly where in ordinary Bohmian mechanics that structure lives — not by analogy, but by direct calculation, using nothing beyond the standard polar decomposition of the wave function.

Write the guiding wave in polar form:

```
(1)   ψ = R e^(iS/ℏ)
```

`R` is the amplitude, `S` the phase. The particle's velocity at any point is given by the guidance equation,

```
(2)   v = ∇S / m
```

— it moves in the direction the phase `S` is climbing fastest. A surface where `S` holds constant is therefore, at every point, exactly perpendicular to the trajectory passing through it, so that moving along that surface spends nothing:

```
(3)   dS = ∇S · dℓ = 0
```

This is not a metaphor for "zero action" — it is zero action, in the precise Hamilton–Jacobi sense. It is also the one direction available to move between branches of a single guiding field without doing any dynamical work. Two moves, orthogonal to each other, at every tick:

> **`dtP`** — move along the trajectory, parallel to `v`, accumulating real `S`. This lays down fact.
>
> **`dtF`** — move along the level surface of `S`, orthogonal to `v`, reaching every other trajectory the wave is sustaining at that instant — then advance each of those by `dt`. This is the situation of possibility: not a value, but a set.

This is exactly the pairing named abstractly in §3.3: `dtP` is the past lobe of the double halo, the spent, one-directional record; `dtF` is the future lobe, the whole sheaf of neighboring trajectories the wave holds open, reachable at no cost in `S`.

Why the fan is well-defined at all. This only works because Bohmian trajectories cannot cross. The velocity field `v(x,t)` is single-valued away from the nodes of `ψ`, so trajectories are integral curves of an ordinary differential equation with a uniqueness theorem behind it — two of them can never touch. Figure 1 below is what this looks like directly: not noise, but a clean, orderable stack. A vertical cut through the plot at any point along the beam enumerates the fan directly — a finite, countable list of neighbors, not an ambiguous tangle.

One thing this picture is not, and the caveat is worth stating plainly: the field genuinely supports a continuum of trajectories — one for every point in the support of $|\psi_0|^2$ — not the hundred-odd sampled curves any plot draws. Treating the forward fan as a finite set, as Figure 1 must, is a discretization imposed by the rendering, not a fact the physics hands over on its own. It is a defensible move, but it should be named as a choice, not mistaken for a derivation.

Remark: openness beyond ignorance. Left alone, this picture is fully deterministic — the particle's whole future is fixed the moment its starting point is; possibility survives only as an observer's ignorance of that starting point, which is thinner than it looks. Reinstating a real, irreducible openness inside the record itself means perturbing the forward step with a genuine stochastic term — a real Wiener increment layered onto the drift, in the manner of Nelson's stochastic mechanics:

```
(4)   dX = v(X,t) dt + √(ℏ/m) dW
```

The noise term acts transversally to the mean drift — suggestively, though not identically, the same direction as the zero-action hop across the wavefront. The resemblance is worth sitting with rather than overclaiming: Nelson's noise is isotropic across the full configuration space, not confined to a single isophase surface. But both are saying the same underlying thing — that whatever "possibility" adds beyond the one realized streamline has to live in the direction orthogonal to where the particle is already headed.

3.7 The coalgebra behind the hinge: Lambek's Lemma

Section 3.1 characterized the soliton's history as a non-well-founded, self-generating process. That process is more than merely well-defined — by a specific and unglamorous theorem, it is the same shape at every point along its own unfolding.

The stream of proper time is a **coalgebra**: a rule that always produces the same shape when you unfold it once more:

```
(5)   self = { possibility → fact, self } = F(self)
```

If `self` is the *terminal* coalgebra of that rule — the largest, least-constrained stream satisfying it — then a fact from category theory applies, not a metaphor: **Lambek's Lemma**. It says the map from `self` to one unfolding of itself,

```
(6)   out : self → F(self)
```

is an **isomorphism**. Not an approximation, not a resemblance — the whole infinite stream and one tick of that stream are structurally the same object. There is no later copy of "now" that differs in kind from this one. Every tick, examined from inside, has the identical shape: a fan ahead, a fact laid down, and the same `self` again. What changes is content — which fact, which fan. What never changes is the type.

This isn't only bookkeeping in category theory. Minkowski space says it too. At proper time τ on a worldline, the light cone splits into a past region $C^-(\tau)$ — accumulated fact — and a future region $C^+(\tau)$ — open possibility — meeting at a vertex. Spacetime translation invariance carries this whole decomposition, unchanged in form, to every other point on the line:

```
(7)   T_Δτ · (C⁻(τ), vertex, C⁺(τ)) = (C⁻(τ+Δτ), vertex, C⁺(τ+Δτ))
```

Nothing about the shape of "now" is special to any one τ. That is the physical sentence for what equation (6) says mathematically, and it is the exact backbone underneath the claim, made informally in §2, that looking back over the soliton's trail is always looking back from a hinge structurally identical to every other hinge on the same trail — the cloud has no privileged point, no first fact, because the process generating it never bottoms out and never changes shape.

This is the formal weight behind what §2 called "looking back over the trail": there is no privileged first hinge, and no hinge structurally unlike any other. With the scaffolding now concrete as well as abstract, we can return to how two different observers classify what that scaffolding produces.

4. The Born Rule as an Infomorphism

This machinery gives a precise reading of the oldest sore point in quantum foundations. Consider two classifications of the same underlying process: the theorist's, whose tokens are runs of the equations and whose types are amplitudes and phase relations; and the experimenter's, whose tokens are runs of the actual apparatus and whose types are detector-click bins and classical trajectory counts. Both classifications share a core — the physical setup itself.

The Born rule is the infomorphism connecting them: it pushes the theorist's types (squared amplitudes) forward into the experimenter's types (relative frequencies), while an experimenter's token (a given click) pulls back to a theorist's token (a given branch of the soliton's history) — and the two agree exactly where the classification condition says they must. On this reading, there is no collapse. There are two classifications of one and the same process, correctly related by a structure-preserving translation. "Why does the wavefunction become a classical count on measurement" is answered the way any channel-theoretic question is answered: it doesn't become anything; it is *translated*.

5. Reading a Bohmian Figure Through the New Image

The accompanying figure (Fig. 1) is a standard rendering of Bohmian trajectories through a two-slit setup: many individual paths, computed from the guidance equation (2), fanning out from each slit, weaving around each other, and settling into the familiar banded interference pattern downstream. It is usually read as showing many *possible* particles, each with its own determinate trajectory, one of which is "the" real one on any given run.

Read through the soliton image, the figure changes what it depicts without changing a single line on the page. It is not a family of possible particles. It is the halo — in the §3.2 sense — of a single soliton's jump-history, rendered at real-number resolution. Every line in the fan is a nonstandard path infinitesimally close to every other; the bunching and weaving is exactly what a double-haloed hinge-chain (§3.3, §3.6) looks like once you can no longer resolve the individual `dtF/dtP` steps that produced it. The turbulence near the source and the settled bands farther out are the same field at two stages of the same story: close in, the quantum potential is still shoving the bundle around; farther out, the trajectories — guaranteed non-crossing by the uniqueness theorem noted in §3.6 — have sorted themselves into the distinct streams that will register, one particle at a time, as an interference pattern. Where Bohm needs a separate pilot wave to explain why the trajectories bend the way they do, the new image needs nothing added: the bending *is* the shape of the halo, because the halo is defined by what keeps the action's standard part constant, and the guidance equation is one way of writing down that constraint.

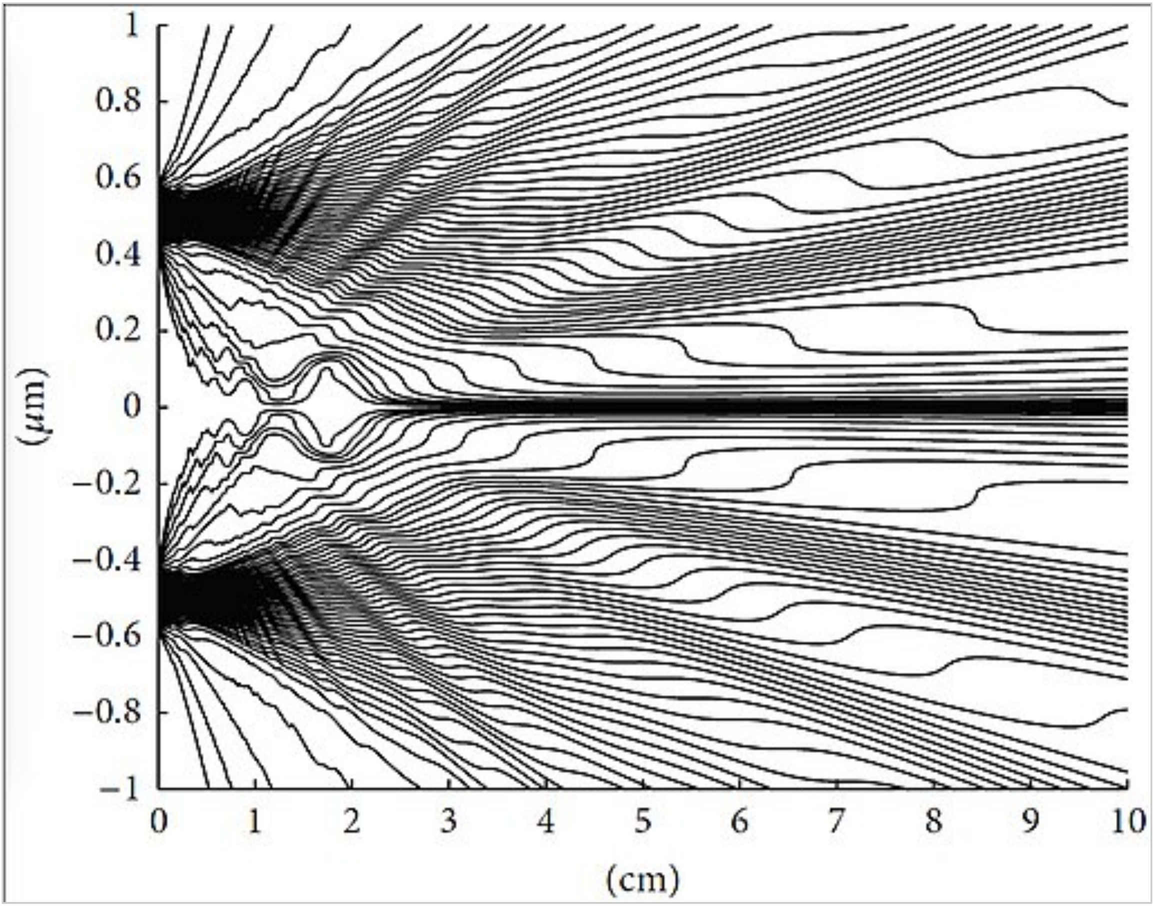

Figure 1. Bohmian trajectories through a double slit, reproduced here as the halo-cloud of a single soliton's jump-history rather than a family of distinct possible particles.

6. From Graphics to Equations: Where the `dt`s Live

The consolidated iPT equation is

```
innerProperTime = { possibilitiesInTheNonStandardFuture →
factsInTheNonstandardPast, innerProperTime }
```

Each hinge in the soliton's trail is one unfolding of this equation: the `dtF` lobe of its double halo *is* `possibilitiesInTheNonStandardFuture` — concretely, the level-surface fan of equation (3); the `dtP` lobe *is* `factsInTheNonstandardPast` — concretely, the accumulated real `S` of equation (2) — and the arrow between them is the jump itself: the deposit of a fact where a possibility just was. The recursive `innerProperTime` on the right-hand side is what keeps the process non-well-founded (§3.1, §3.7): there is no first hinge, only the ongoing generation of the next one from the last, guaranteed by Lambek's Lemma (6) to have exactly the same shape every time.

The proposed figure for the paper's centerpiece is therefore a beaded chain rather than a smooth curve: each bead a double halo labeled `dtF/dtP`, an arrow from the bead into the equation above with the two lobes labeled to match, and — beside it, rendered at coarser resolution — the same chain collapsed into the cloud that Fig. 1 shows. The caption should make explicit that the second rendering is the *standard part* of the first: it is the only one an experimenter, working at real-number resolution, will ever see.

7. Outlook

The image this paper proposes does not compete with the Schrödinger equation, the path integral, or the Bohmian guiding equation as calculational machinery — all of it survives untouched, and §3.6 shows the new image using that same machinery directly, without modification. What it replaces is the *intuition pump* underneath the machinery: not a substance that is sometimes a wave and sometimes a particle, not a point steered by a second entity, but a single point whose infinitesimal jumping, viewed after the fact, is mistaken for a cloud — a jumping that is, by Lambek's Lemma, provably the same shape at every point along its own history. The mathematics that makes this more than a slogan — non-well-founded sets and coalgebras, nonstandard halos, the Hamilton– Jacobi zero-action surface, situation theory, channel theory — already exists, in each case for reasons having nothing to do with quantum mechanics, or with each other. Bringing them together here is the paper's only novel move.

Appendix: The Equations, Taught

Every equation used above, unpacked in plain terms, in the order it first appeared.

1. The wave function in polar form `ψ = R e^(iS/ℏ)` Any complex number can be split into a size and an angle. Do that to the wave function at every point in space: `R` is the size — how much amplitude is there, which becomes the probability once squared. `S` is the angle, scaled by `ℏ` — the phase. Writing `ψ` this way doesn't add new physics; it separates two questions that were tangled together: how much (`R`) and which way (`S`). Everything about the particle's actual motion, in this framework, turns out to depend only on the second question.

2. The guidance equation `v = ∇S / m` `∇S` — "gradient of `S`" — is a vector pointing in whatever direction the phase is increasing fastest at that point, with a length equal to how fast it's increasing. Divide by mass and you get a velocity. This is the one rule that tells the particle which way to go: not Newton's force law, but "follow the direction the wave's phase is climbing." It's the entire content of how the wave steers the particle.

3. Zero action along a level surface $dS = \nabla S \cdot d\ell = 0$ A "level surface" of S is the set of points where the phase has one fixed value — like a contour line on a map, but for phase instead of elevation. Walking along a contour line, by definition, doesn't change your elevation; walking along a level surface of S doesn't change S. The dot product being zero is just the statement that the gradient (which always points straight uphill, perpendicular to the contour) has no component along the contour itself — a basic fact of vector calculus, true of any scalar field. Here it says the sideways move between trajectories costs nothing in S, which is why it was called "free."

4. The stochastic modification (Nelson noise) $dX = v(X,t)\ dt + \sqrt{\hbar/m}\ dW$ The first term, $v(X,t)\ dt$, is the ordinary guided step from equation (2). The second term adds a random nudge: dW is a Wiener increment — a random number, mean zero, whose typical size grows with the square root of the time step rather than the time step itself (that's why real Brownian motion looks jagged at every scale, not smooth). Multiplying it by $\sqrt{\hbar/m}$ sets its size correctly in physical units. In plain terms: take the deterministic Bohmian step, then add a small, genuinely unpredictable jitter on top, the way a dust mote gets kicked around by unseen molecular collisions.

5. The coalgebra (self as a self-referential rule) $self = \{ possibility \rightarrow fact, self \} = F(self)$ This defines something not by building it up from smaller pieces, but by saying what one step of unwinding it looks like: a possibility, an arrow turning that possibility into a fact, and then — the same rule again. F is just a name for "the shape one step takes." An infinite stream defined this way is fully specified even though the definition mentions itself, the same way "a natural number is zero, or one more than a natural number" fully specifies the natural numbers without listing them.

6. Lambek's Lemma $out : self \rightarrow F(self)$ **is an isomorphism** An equals sign says two things have the same value. An isomorphism says two things have exactly the same structure — a perfect two-way translation between them that loses nothing either direction. Lambek's Lemma is a theorem, not an assumption: for the terminal coalgebra of a self-referential rule like equation (5), the map from the whole stream to "one unfolding of it" is guaranteed to be this kind of perfect structural match. That's the technical backbone of "it's always today" — unfolding once doesn't get you closer to some different, later structure. It gets you back to the same structure, exactly.

7. Translation invariance along a worldline $T\_\Delta\tau \cdot (C^-(\tau), vertex, C^+(\tau)) = (C^-(\tau+\Delta\tau), vertex, C^+(\tau+\Delta\tau))$ $C^-(\tau)$ and $C^+(\tau)$ are the past and future light cones at proper time τ on a worldline — everything that could have influenced this point, and everything this point could go on to influence. $T\_\Delta\tau$ is a shift: slide everything along the worldline by an amount $\Delta\tau$. The equation says that shifted picture at $\tau+\Delta\tau$ is identical in shape to the original one at τ — just relabeled. This is a basic symmetry of flat spacetime (no point on an unaccelerated worldline is special), and it's the physics-side twin of equation (6): both say the relation between "behind" and "ahead" never changes shape as you move along.

Authorship and AI Disclosure

This paper was developed collaboratively between Lee Bloomquist and Claude (Anthropic). Claude contributed to the exposition of the mathematical background (§3–4) and to drafting text under the author's direction; the central image, the iPT formalism it is grounded in, and the interpretive claims of the paper are the author's own. This disclosure follows the author's standing practice of explicit AI co-authorship attribution.