

Copy And Paste Model

Philosophy

Part 1 of 5

Claudio Serdar-Petkovski

Independent Researcher, Systems Geometry Group

Version 2.0, 10 September 2026

split from the single-document Version 1.4 of 10 August 2026

Abstract. The Copy And Paste Model claims that the propagation of information across radically different media, from river deltas and vascular trees to silicon networks, immune tissue, perception and language, is governed by one least-action principle, and that the substrate changes while the governing equations do not. This document states that claim in plain language, distinguishes it from the physical unification it is sometimes mistaken for, records how the framework was built and what that origin permits it to claim, previews what it says about thought and about allergic disease, and reports the narrower claim that survived testing: the propagation machinery is substrate-invariant, and the applied apparatus built on top of it is not. It contains no equations. The formal apparatus is the Core document; its applications are the Cognitive and Biomedical documents.

Developed in collaborative dialogue. Corrections to earlier versions are recorded in the Technical Appendix, not silently removed.

The CAPM Map

One framework, five documents. Read the one addressed to you; the others are referred to by name, never assumed.

Document	Its job	What you are asked to evaluate
► P Philosophy	States the thesis (<i>the substrate changes; the equations do not</i>), the method, and the boundaries of the claim. Prose only.	Is the thesis coherent? Is the intuition-first method honest? Is the narrowed claim that survived testing the right one to make?
C Core Mathematical Apparatus	Defines the engine, the action, substrate memory and the population wave; reports what is known about their well-posedness; tests substrate invariance.	The normalisation pathology and its repair; the measure problem and the convergence bound; uniqueness of equilibrium; ρ_s independence; the invariance verdict.
G Cognitive Extension	Applies the engine to thought: geometry translated across substrates, and the detection of the substrate a pattern natively occupies.	Whether coherence across crossings predicts what is remembered, and whether substrate detection predicts anything its bridge costs did not contain.
B Biomedical Extension	Applies the engine to allergic sensitization; derives predictions for immunotherapy, polysensitization and pre-symptomatic detection.	A falsification ledger: every prediction with its status, whether contradicted, pre-registered, testable now, or open.
A Technical Appendix and Bookkeeping	Numerical methods, verification runs, the corrections log, coherence bookkeeping, master equation table, V1.4 concordance.	Reproducibility. Every “earlier versions” statement lives here and nowhere else.

Anchor symbols. Shared across the documents; defined where indicated.

Symbol	Meaning	Defined in
$R(t)$	Replication probability: Boltzmann-weighted share of path weight reaching the nominated destination	C
$S[\gamma, t]$	Informational action along path γ ; $\sigma(t)$ is its temperature, in the position $\hbar$ occupies but with a real exponent	C
χ_j	Bridge-operator cost of the j -th substrate crossing; χ_{ij} its matrix form between subsystems	C
$c(s)$	Node compatibility: geometric $\vec{p}^\top M_s \vec{p}$ and thermodynamic $\rho_s / (\rho_s + e^{\Delta G_0 / k_B T})$, one quantity in two coordinate systems	C
$V(s, t, \rho_s)$	Novelty: $\rho_s \cdot \text{JSD}(P_{\text{pattern}} \ P_{\text{substrate}})$ over feature bins	C
$P_{\text{substrate}}, P_0, N, \bar{n}$	Substrate memory and its pristine baseline; execution-engine population and $\bar{n} = N/K$	C
$\alpha, \gamma_{\text{mem}}$	Learning rate and <i>forgetting</i> rate of the substrate; low γ_{mem} means long memory	C

Symbol	Meaning	Defined in
ρ_s	Structural resolution: how finely the substrate discriminates configurations	C, grounded in B
Θ, Mu	Propagation efficiency $(Z/Z_{\text{ideal}})^{1/\bar{n}}$; mutation–adaptation ratio $\eta^2 d/\gamma_{\text{relax}}$	C
$C_{\text{total}}, \omega_r$	Coherence across a set of substrate crossings; rate of recall attempts	G
κ, P_{tol}	Resensitization coefficient; regulatory substrate that makes γ_{mem} plastic	B
P^*, ν	Observable untreated baseline; ambient encounter rate inverted from it	B

Conventions. γ_{mem} is a forgetting rate throughout. Differences are oriented as $(P_{\text{substrate}} - P_0)$. The population enters memory laws as $\bar{n} = N/K$. Equation, section and theorem numbers carry the document letter, so Core Eq. (C.1) is the engine wherever it is cited and the letter alone tells you which document to open. Gray tags [v1.4(*n*)] give the number the same equation carried in the single-document Version 1.4.

Contents

P.1	The Observation	4
P.2	What This Framework Is, and Is Not	4
P.3	How the Framework Was Built	6
P.4	What It Claims About Thought	7
P.5	What It Claims About Allergy	8
P.6	What Survived Testing	9
P.7	Boundaries and the Standing of the Claim	10

P.1 The Observation

A profound structural regularity manifests across macroscopic distribution systems throughout the observable universe. The branching layouts of river deltas, the vascular networks of biological organisms, the layout of human metropolitan road systems, and the dense latent spaces of artificial neural networks all display striking geometric commonalities. These parallel structures are not coincidental; they point toward a fundamental governing principle: the continuous optimization of pattern replication across physical media.

Traditional scientific frameworks frequently analyze informational systems in isolation. Computer science models data routing via abstract network protocols; sociology monitors the spread of ideas via qualitative cultural indicators; and ecology tracks population growth via classical resource models. The Copy And Paste Model bridges these divided silos. The foundational philosophy rests upon a universal evolutionary pipeline:

$$\text{COPY} \longrightarrow \text{PASTE} \longrightarrow \text{ADAPT} \longrightarrow \text{REPEAT}$$

Every idea, piece of software, or biological blueprint behaves as an informational configuration seeking survival. To survive, it must absorb environmental energy, recognize the structural layout of its current domain, replicate its pattern onto a new substrate, and continuously optimize its internal code against ambient resistance.

That is the whole intuition. Everything else in this series is an attempt to say it precisely enough to be wrong.

P.2 What This Framework Is, and Is Not

A framework that speaks of path integrals and substrate invariance invites a reading it cannot support, and it is better to close that reading off at the start than to leave it standing.

What it is

The Copy And Paste Model is a least-action account of how information propagates through media. Its central object is a weighted sum over all the routes a pattern could take through a substrate, in which every route is penalised by the resistance it meets along the way and by the cost of each boundary it has to cross. Its second object is a substrate that the passage of the pattern rewrites, so that what propagates once propagates differently the next time. Those two objects, coupled, are the entire apparatus.

The apparatus is then deployed on three kinds of medium: network topologies made of silicon, immune tissue, and the passage of a geometry through perception into language. The claim under test is that the same equations govern all three. That claim is the reason the framework exists, and it is the thing worth attacking.

What it is not

It is not a unification of physics. The weight the model assigns to a route decays exponentially with the resistance along it. That is a Boltzmann weight, and a sum of Boltzmann weights is a partition function. What makes quantum mechanics quantum is something the model does not have: a complex phase, whose interference between routes produces the effects the theory is built to explain. The resemblance to Feynman's construction is real, and the Core document states it precisely, with the environmental noise scale sitting exactly where Planck's constant sits. But a real exponent puts this framework in statistical mechanics, alongside reaction–diffusion systems and epidemic models, and not in quantum theory. Nothing here bears on gravity, on spacetime curvature, on the Planck scale, or on the reconciliation of general relativity with quantum mechanics. A reader who came expecting that should stop here.

It is not derived from established theory. The next section records the opposite order: a geometric regularity was noticed across distribution systems, and the formalism was built afterwards to express it. That order constrains what the framework may legitimately claim, and the constraint is honoured across the series by stating the domain of validity alongside each result rather than after it fails.

It is not, in its applied parts, substrate-invariant. The Core document tests the invariance claim against two substrates the framework was not built on, and the verdict is partial in a way that matters. The propagation machinery carries across. The apparatus built on top of it for allergy does not, because that apparatus assumes a substrate which dies when its input is removed. §P.6 sets out what survived and what did not.

What is being asked of the reader

Three questions, and none of them requires the mathematics.

- (i) Is *the substrate changes; the governing equations do not* a coherent claim, or does it dissolve under pressure into a claim about analogy?
- (ii) Is the intuition-first method of the next section an honest epistemology for a framework of this kind, given the failure mode that §P.6 documents?
- (iii) Is the narrowed claim that survived testing the right thing to be claiming, and is it interesting enough to justify the apparatus?

The formal apparatus is the Core document. Its application to thought is the Cognitive document, and to allergic disease the Biomedical document. Nothing in this document depends on reading any of them.

P.3 How the Framework Was Built

The engine at the centre of this framework did not emerge through conventional axiomatic derivation or systematic literature review. Its origin and development methodology are documented here not as biography, but as epistemology: the nature of a framework’s origin constrains the kind of claims it can legitimately make, and transparency on this point is a prerequisite for honest scientific evaluation.

The core geometric observation

The framework originates in a structural observation made across macroscopic distribution systems of radically different physical composition. Terrestrial road networks, biological vascular architectures, neural pathway formations, colony foraging routes, and large-scale human migration patterns all exhibit a striking geometric regularity: each system evolves pathways that minimize cumulative resistance along routes of propagation. Street networks converge toward layouts that minimize travel cost. Neural circuits form to minimize signal latency. Biological colonies establish routes requiring minimum energetic expenditure. The observation is domain-independent.

This regularity was recognized as an instance of a deeper organizing principle: *least-action optimization governs the geometry of information propagation across physical media*. The engine is, at its core, a formal statement of that principle. Its weighted sum over candidate routes is the mathematical encoding of least-action selection, and the quantity it weights each route by is the cumulative resistance a propagating configuration encounters along it.

Intuition first, formalism second

The development methodology inverts the conventional sequence of mathematical inquiry. Rather than deriving equations from established axioms and subsequently seeking empirical domains of application, the geometric intuition was observed first and the mathematical framework was constructed afterward to encode it with precision. The formalism is therefore best understood as *a language for expressing an underlying geometric regularity*, not as a proof structure descended from prior theory.

This approach has notable historical precedent. Feynman’s path-integral formulation of quantum mechanics originated in a geometric intuition about least-action trajectories, with rigorous measure-theoretic foundations constructed subsequently. The present framework makes no claim to equivalent mathematical depth. It does claim an equivalent epistemological structure: geometry first, axioms second.

Development process

The framework crystallized through an extended period of iterative cross-domain examination, spanning many months of structured dialogue across multiple analytical systems and research

contexts, during which the geometric regularity described above was stress-tested against progressively more distant domains. The central question driving each iteration was whether the least-action propagation principle retained predictive structure when the substrate changed entirely: from silicon network routing, to sociological information diffusion, to immunological memory dynamics. The mathematical formalism was refined iteratively to preserve structural invariance across these crossings, with each substrate translation serving as a progressively stronger theoretical stress test of the underlying geometric claim.

The development process is itself an instance of the framework. The geometric pattern propagated across successive analytical substrates, paying a bridge cost at each domain crossing, and its invariant core, least-action optimization, survived every translation because its compatibility with each new domain never fell to zero. Whether that is an elegant observation or a warning about unfalsifiable self-description is a fair question, and one this document does not settle.

Epistemological boundaries

The strength of the framework rests not on mathematical elegance alone but on whether the geometric principle it encodes has predictive validity across novel domains not used in its construction. This is the appropriate standard for an intuition-first theoretical framework: not internal consistency, which is necessary but insufficient, but *external predictive reach*.

The framework invites falsification. That openness is not a concession but a commitment, and the Biomedical document keeps it in a form that can be checked: a ledger of every prediction with its current status, one of which is already recorded as contradicted.

P.4 What It Claims About Thought

The framework's cognitive claim can be stated without any of its machinery, and it is worth stating here because it is the claim most people find either obvious or absurd on first hearing.

Thought is geometry, translated through successive substrates.

Consider recognising a fan. The object's physical shape enters the visual system as structured light. The visual cortex does not receive a label. It receives shape: edges, rotational symmetry, spatial periodicity. That geometric signal propagates inward, is matched against stored geometric templates, crosses the boundary into linguistic substrate, and produces the word "fan". The word is not the primary event. It is the compressed label the final substrate assigns to the geometry that arrived before it. Thought, on this account, is what geometry looks like after passing through enough substrate crossings to become conscious.

Three consequences follow, and each is meant to be arguable rather than decorative.

Mathematical intuition is geometry recognising itself. When a mathematician sees why a proof is true before being able to say why, the geometric pattern has propagated far enough to activate well-matched pathways in the cognitive substrate but has not yet completed its translation into language. The feeling of insight is the system catching a pattern mid-propagation,

before the last crossing has been paid for.

Language is not the origin of thought but its final substrate. It does not generate thought; it is the last medium a geometric thought passes through before becoming communicable. Some thoughts resist full linguistic encoding: musical intuition, spatial reasoning, aesthetic perception. These are patterns whose cost at the geometry-to-language boundary is high. The translation is resisted, and the pattern stays partly encoded in an earlier substrate.

What you remember is what translated cleanly. A pattern that survives every crossing intact is one you cannot forget; a pattern that loses coherence at each boundary is gone within the hour. On this account the stickiness of something you learned is not a matter of how emotionally charged it was, or how often you met it, but of how cleanly it moved between the substrates it had to pass through. Which substrate a pattern natively occupies is not a matter of choice either: it depends on which substrates a particular observer has available, so a deaf reader and a hearing reader encode the same word differently, and that difference is predicted rather than excused.

The Cognitive document turns those claims into something measurable, and is candid about where the measurement currently stops: the cost of a substrate crossing is at present an input to the procedure rather than something the framework derives, which means the worked examples illustrate the machinery without yet testing it.

P.5 What It Claims About Allergy

The framework's clinical claim starts from a distinction that sounds pedantic and turns out to do real work.

An allergen does not replicate. A grass pollen protein has no machinery for copying itself, and nothing about it changes as an allergic response builds. So if the framework is about the propagation of patterns, and allergy is to be one of its substrates, something has to be doing the propagating, and it is not the allergen.

The answer is that the allergen is a *blueprint* and the immune network is the *execution engine*. The allergen supplies a molecular configuration. Dendritic cells, T cells, B cells and mast cells are what copy, amplify and maintain a response to that configuration, writing it into the body's memory a little more deeply on each encounter. What propagates is not the allergen. It is the information encoded from meeting it.

That reframing makes allergy a memory problem rather than an overreaction problem. The immune system has written a pattern into its substrate marking a harmless molecule as worth attacking, and the pattern is accurate about what the molecule is and wrong about what it means. Two parameters then govern everything clinical: how fast the system writes such a memory, and how fast it lets one decay. Aggressive learners who also forget quickly are the patients who outgrow childhood allergies and respond well to treatment. Aggressive learners who never forget are the severe cases, and every accidental exposure deepens the encoding that the treatment is trying to erase.

The framework's most useful consequence is an unglamorous one. If a memory can be written it can in principle be overwritten, and the shape of the overwriting, meaning how much and how often and for how long, is not a matter of clinical taste but something the equations constrain. The Biomedical document works that out, and also states plainly what the model cannot do: it can describe how an established sensitization is maintained and modified, and it contains no mechanism by which a person becomes allergic in the first place. That is the largest hole in the applied side of the framework, and it is named rather than hidden.

There is one further prediction worth mentioning here, because it is the kind that could be checked without any new instrument. If the immune execution engine begins working before any symptom appears, it should be changing the mechanical properties of the tissue it is working in: stiffness, density, the way sound scatters through it. Ultrasound already measures all three. The framework says what pattern to look for and roughly how far ahead of symptoms to look. Whether anything is there is an empirical question, and nobody has yet looked.

P.6 What Survived Testing

The framework's deepest claim is that the equations do not change when the substrate does. The Core document puts that claim to two tests it was not built for, and the result is a partition rather than a verdict.

The partition

The population dynamics transfer exactly. Run on silicon parameters, with clock rates seven orders of magnitude away from tissue, the equilibrium density reproduces the predicted value across three decades. Run on a multi-agent substrate, meaning messages spreading through a population, the model reproduces standard epidemic theory to three decimal places.

That agreement is weaker evidence than it first appears, and the Core document says so plainly rather than banking it. Logistic growth with linear death *is* the standard epidemic equation, written in different letters, so the match had to hold. It is a calibration result, not a confirmation. What did not have to hold is where the boundary between the two regimes falls, and that turns out to be the informative part.

The apparatus built on top does not transfer. The step that makes the clinical work possible, which is replacing an unobservable environmental exposure rate with the observable severity a patient shows when nobody is treating them, works in tissue and fails on silicon. The reason is structural rather than incidental. That inversion assumes a substrate which dies without external input. Remove the allergen and the immune response subsides; a processor keeps its processes alive with no input at all. Substrate memory transfers formally and then saturates, discriminating nothing. The path sum transfers only with a route-length discount two orders of magnitude away from the biomedical value, because switch fabrics branch far more widely than tissue does.

The honest form of the claim is therefore narrower than the one this document

opened with: *the propagation machinery is substrate-invariant; the apparatus built on top of it for allergy is not, and assumes a driven substrate*. That is a smaller claim and a better one, because it states the condition under which the clinical work applies instead of leaving it implicit.

A recurring failure mode

The corrections recorded across this series share a shape, and naming it is more useful than listing them again.

Almost none of the claims that failed under testing were simply false. They were *conditionally* true, with the conditions unstated. The speed at which an immune response spreads through tissue was correctly derived for the case where the tissue starts empty and the response can sustain itself, and was inapplicable at the framework's own reference settings, where neither condition holds. The boundary separating tolerance from persistent attack was correct when read at the forgetting rate a patient currently has, and substantially wrong when read at the rate they started with. The requirement that three ultrasound channels agree is informative when the three are genuinely different measurements and an empty identity when they are not. The law governing how fast a mutating pattern escapes recognition describes an average over many runs and says nothing useful about any single one. In each case the mathematics was sound and the scope was missing.

This is what one should expect of a framework built by carrying a geometric intuition across substrates, and it suggests where the remaining errors will be. They are unlikely to be in the algebra. They are likely to be assumptions so natural to the original setting that nobody wrote them down, and which fail silently once the same expression is used somewhere else. Stating the domain of validity alongside each result is therefore not editorial tidiness in a framework of this kind. It is the principal defence available.

P.7 Boundaries and the Standing of the Claim

Two limits are worth stating in a document that contains no mathematics, because neither of them is mathematical.

Measurement disturbs what it measures

The framework cannot fully predict its own behaviour when the substrate is altered by the act of measuring it. The pre-symptomatic detection it proposes is mechanical: pressing sound through tissue to read its stiffness. Doing so may itself shift, slightly, the substrate state being read. This is not a defect peculiar to this framework, since it is a version of a limit that turns up wherever an instrument and its subject share a medium, but it is a real constraint on experimental design rather than a philosophical flourish, and it should be built into any protocol derived from this work.

What the origin permits the framework to claim

A framework built intuition-first cannot claim the authority of one derived from axioms, and this series does not try to. What it can claim is that the intuition was made precise enough to generate predictions that could fail, and that when they failed this was recorded rather than quietly repaired. The Technical Appendix keeps that record. It lists, for every superseded claim, what was originally said, what went wrong, and where the corrected form now lives. One prediction in the Biomedical document is marked as contradicted by published data. Several others are marked as untestable for want of a bridge between the model's units and a clinical endpoint, which is a limitation and is listed as one.

A reader is entitled to weigh that record against the framework's ambition. The ambition is large: one governing principle across river deltas, immune memory and the passage of a thought into speech. The evidence is partial, the applied apparatus is regime-bound, and the cognitive extension currently supplies its own inputs to the procedure it uses to rank them.

A lens, not a final answer

The framework proposed here is a lens, not a final answer. Its value lies not in replacing existing immunology, or existing cognitive science, but in providing a unified mathematical language, one that connects the geometry of information propagation to the dynamics of substrate memory and in doing so reveals structure that neither field could see alone.

Whether it does reveal such structure, or merely redescribes what each field already knew in unfamiliar notation, is the question this series exists to have settled by someone other than its author.

References