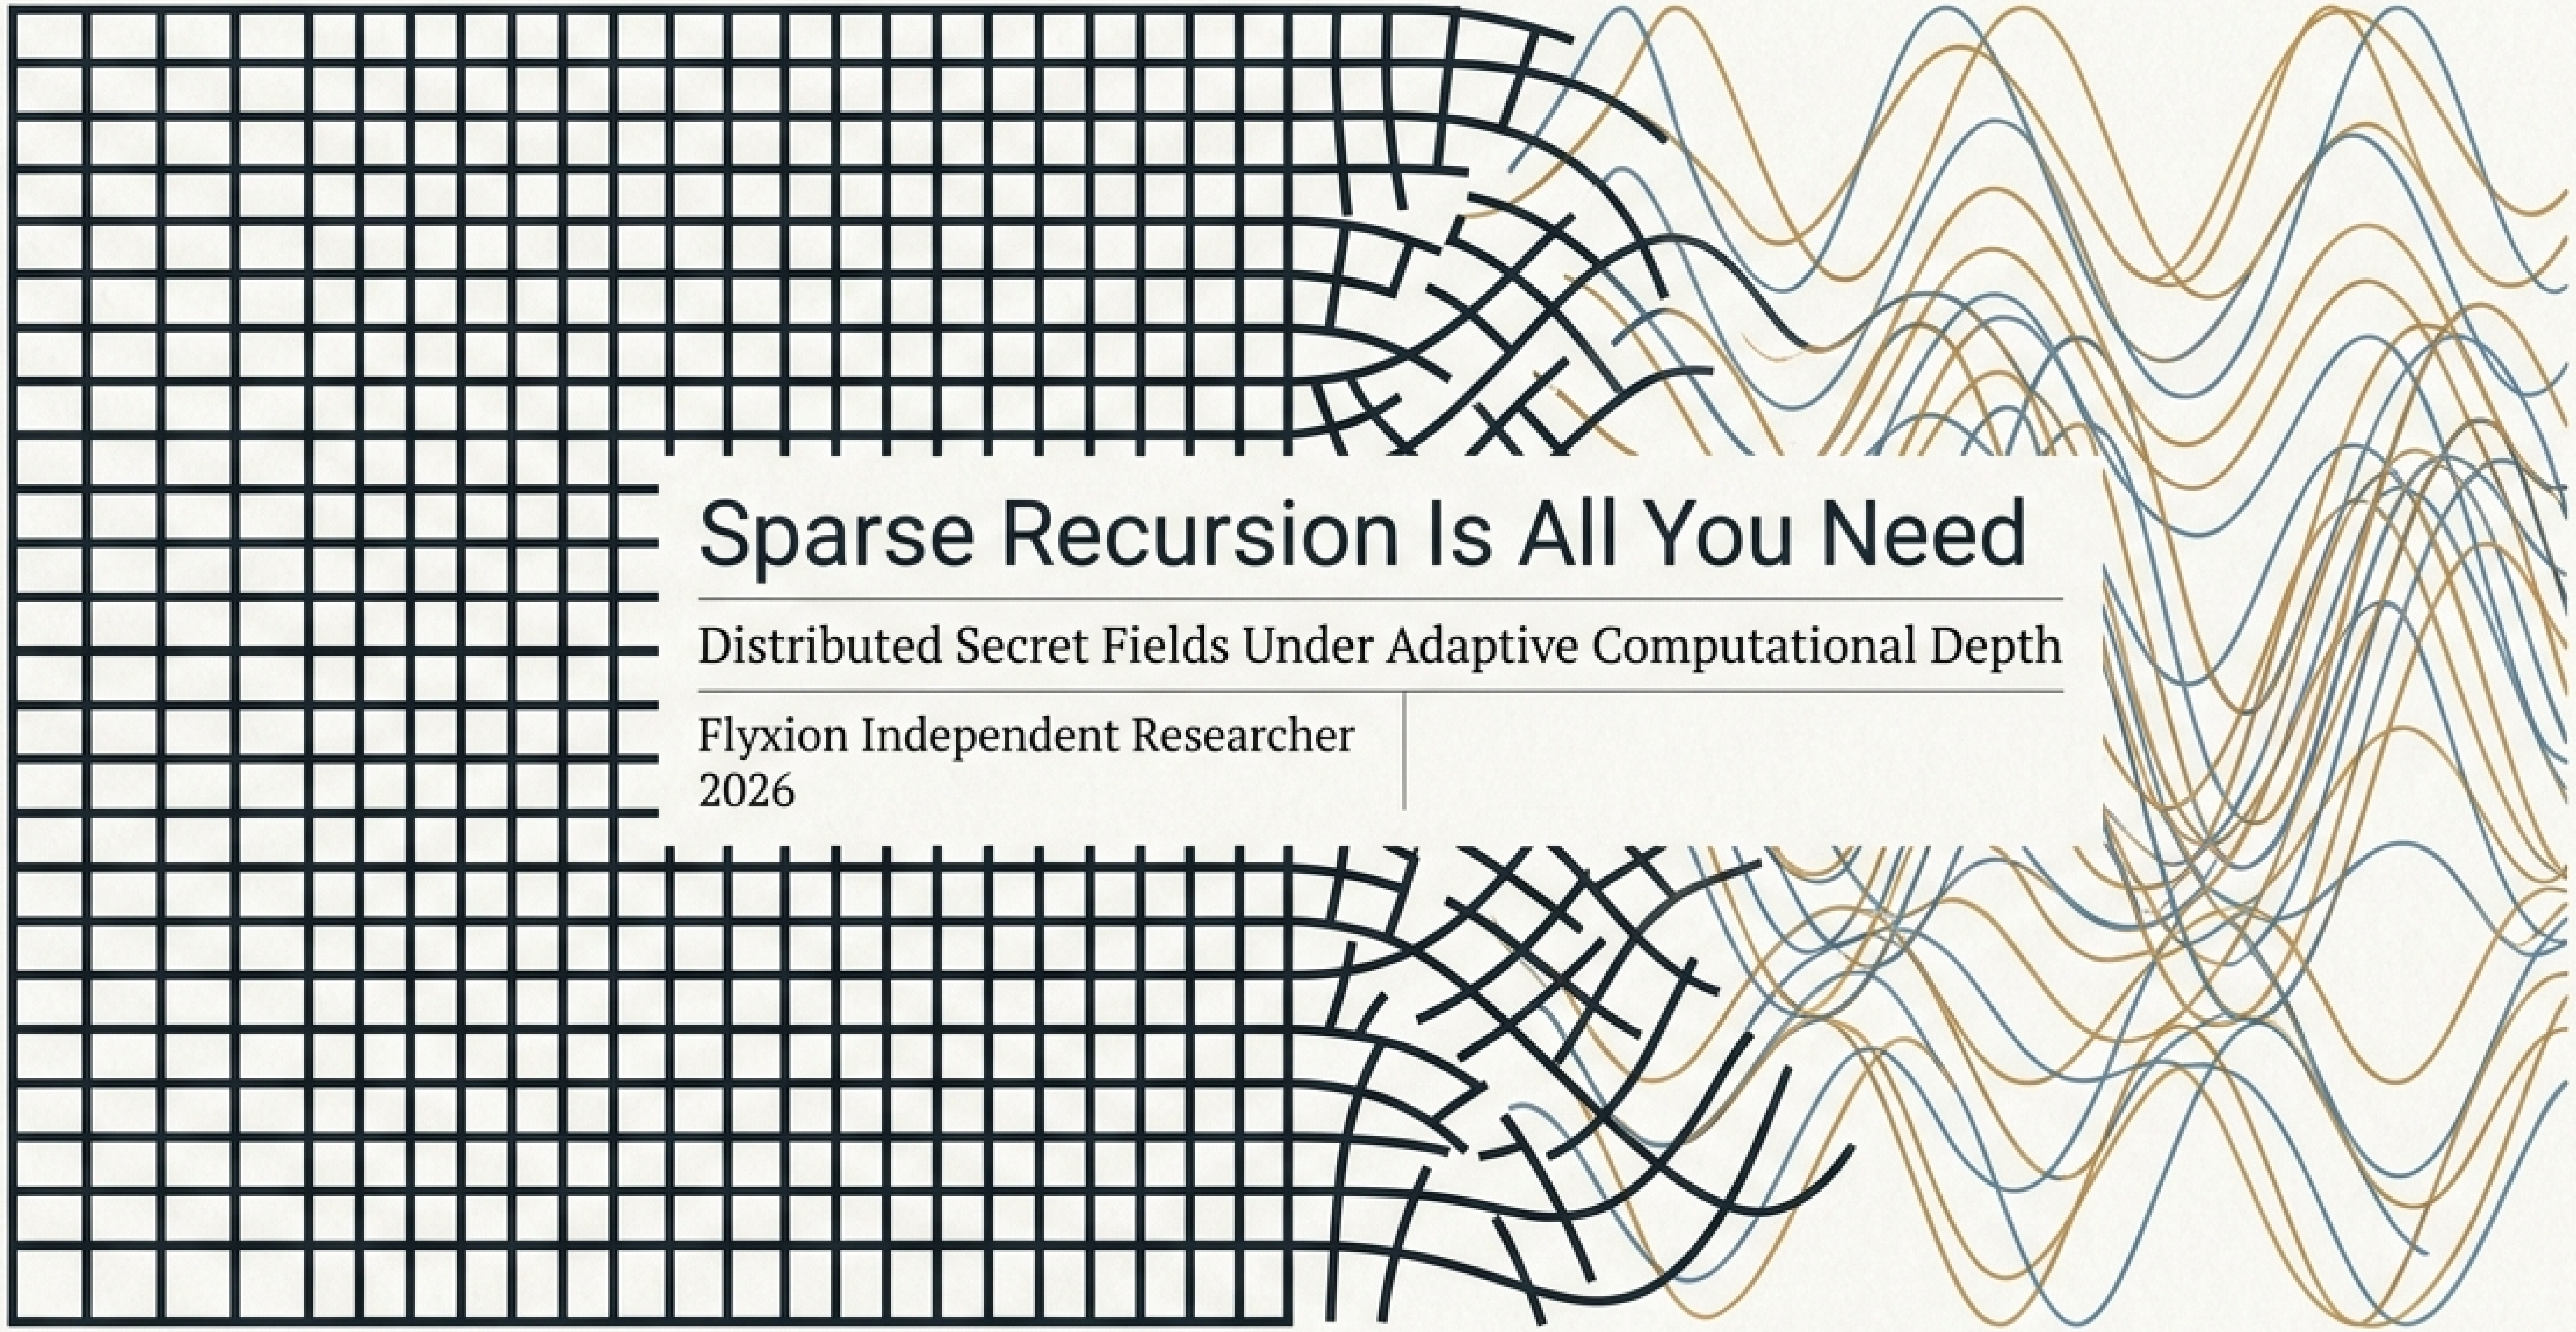

Sparse Recursion Is All You Need

Distributed Secret Fields Under Adaptive Computational Depth

Flyxion Independent Researcher
2026

The Problem With One-Shot Intelligence

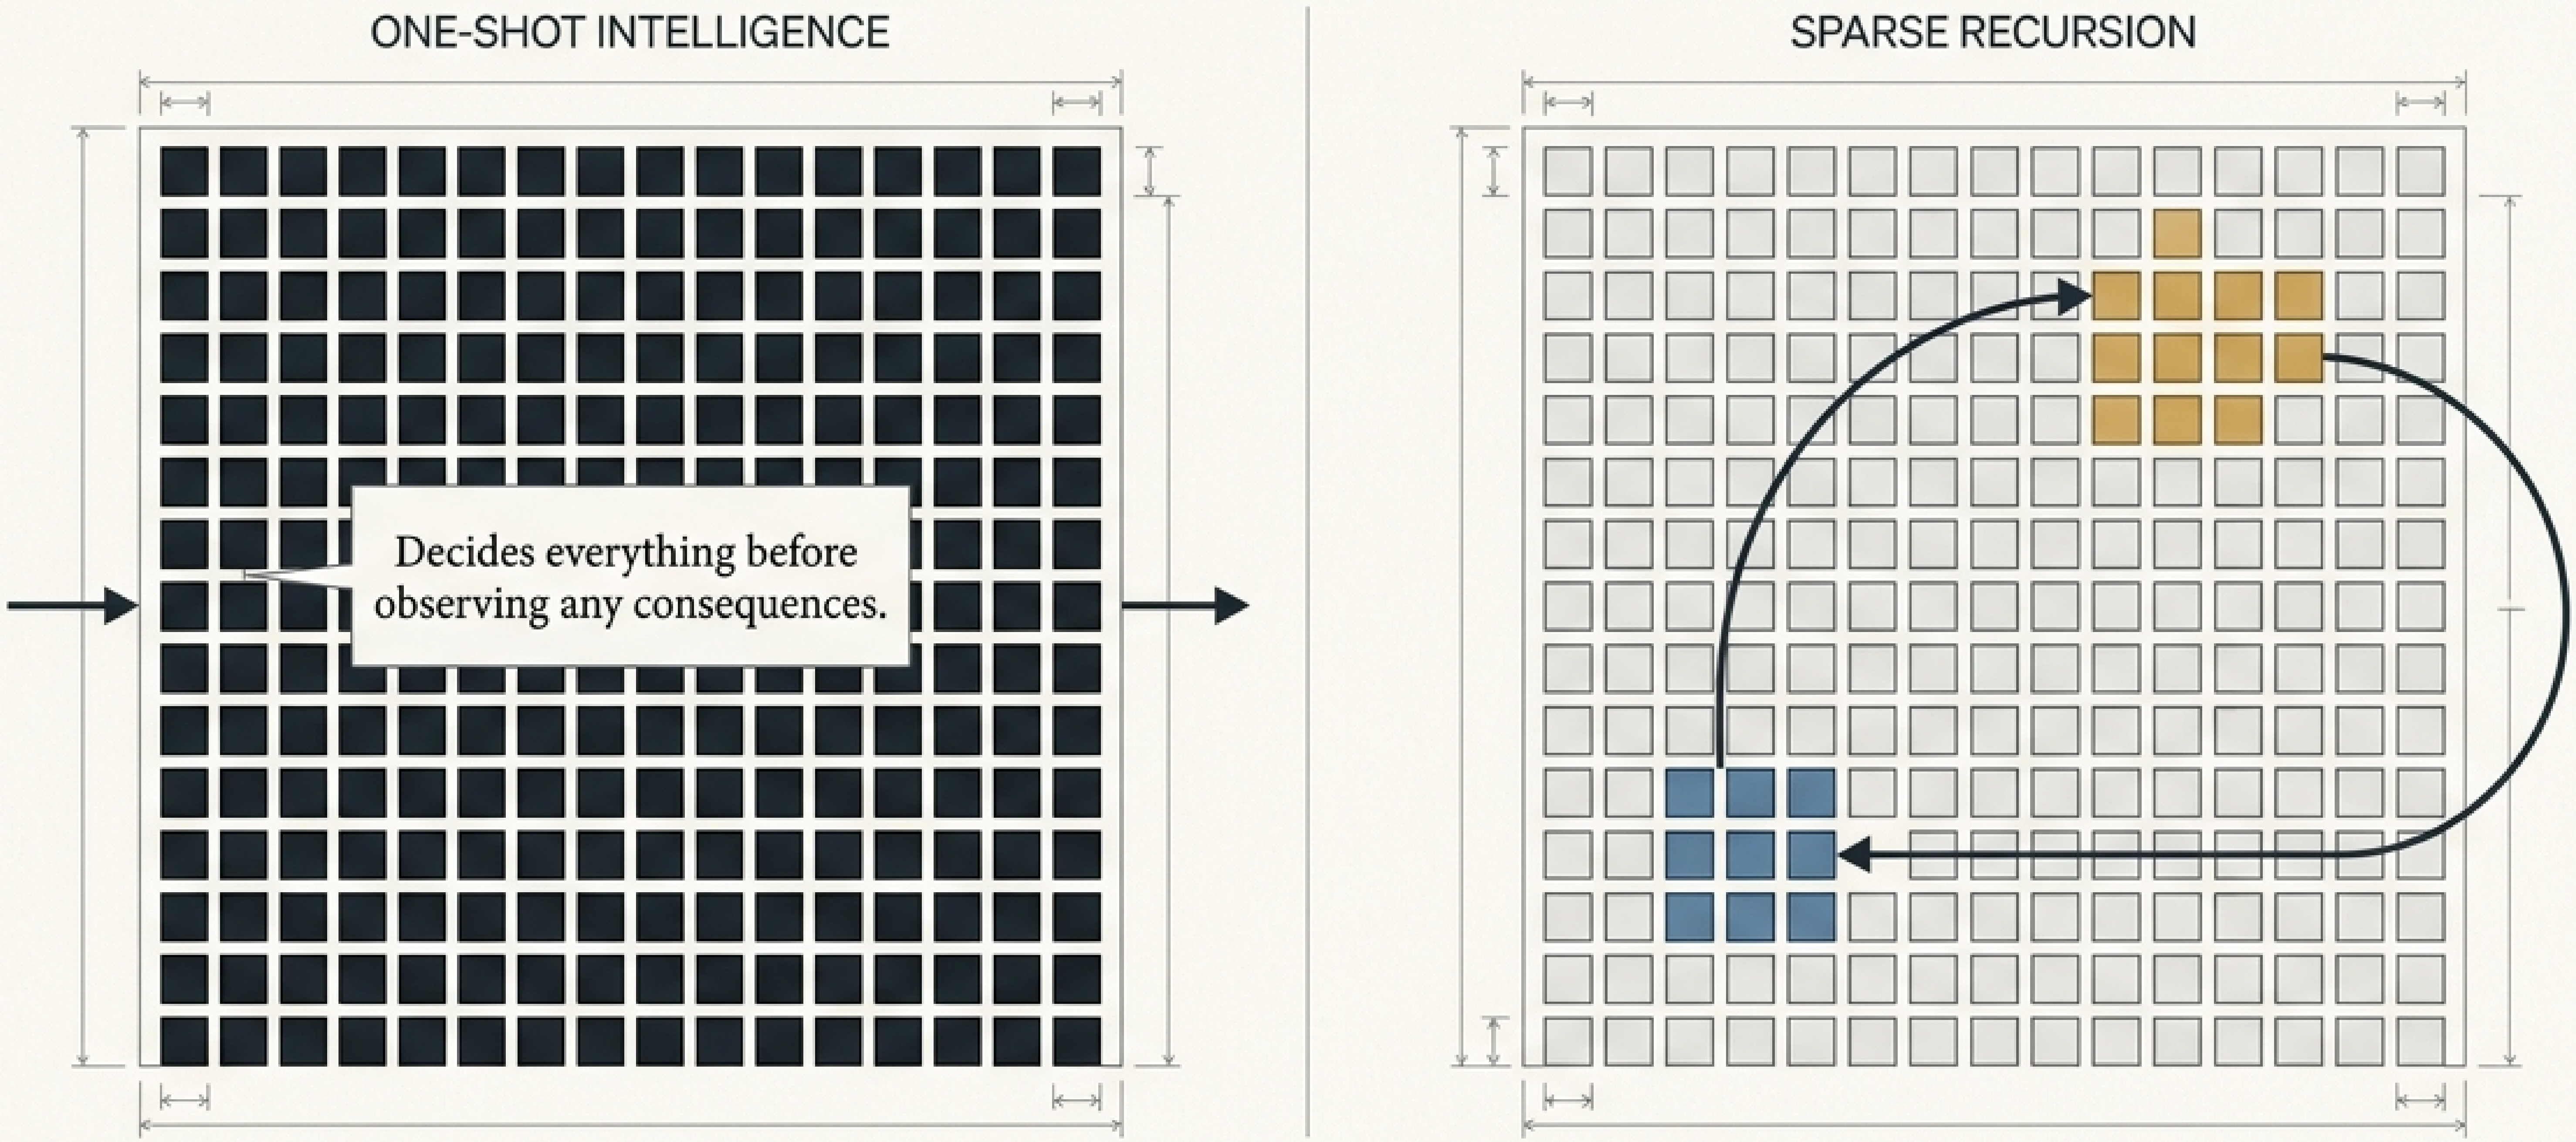

A one-pass system commits its entire computational budget up front, blindly hoping the mistakes are tolerable.

When one-pass systems fail, the default answer is simply more: more parameters, layers, and data.

Perhaps capability scales not by doing everything at once, but by repeatedly changing only the few parts that still matter.

Steganography: The Perfect Clean Test

Conventional encoders distribute a secret message across pixels, emit the image, and only afterward is the image cropped or compressed. The encoder commits its full distortion budget before learning what will survive.

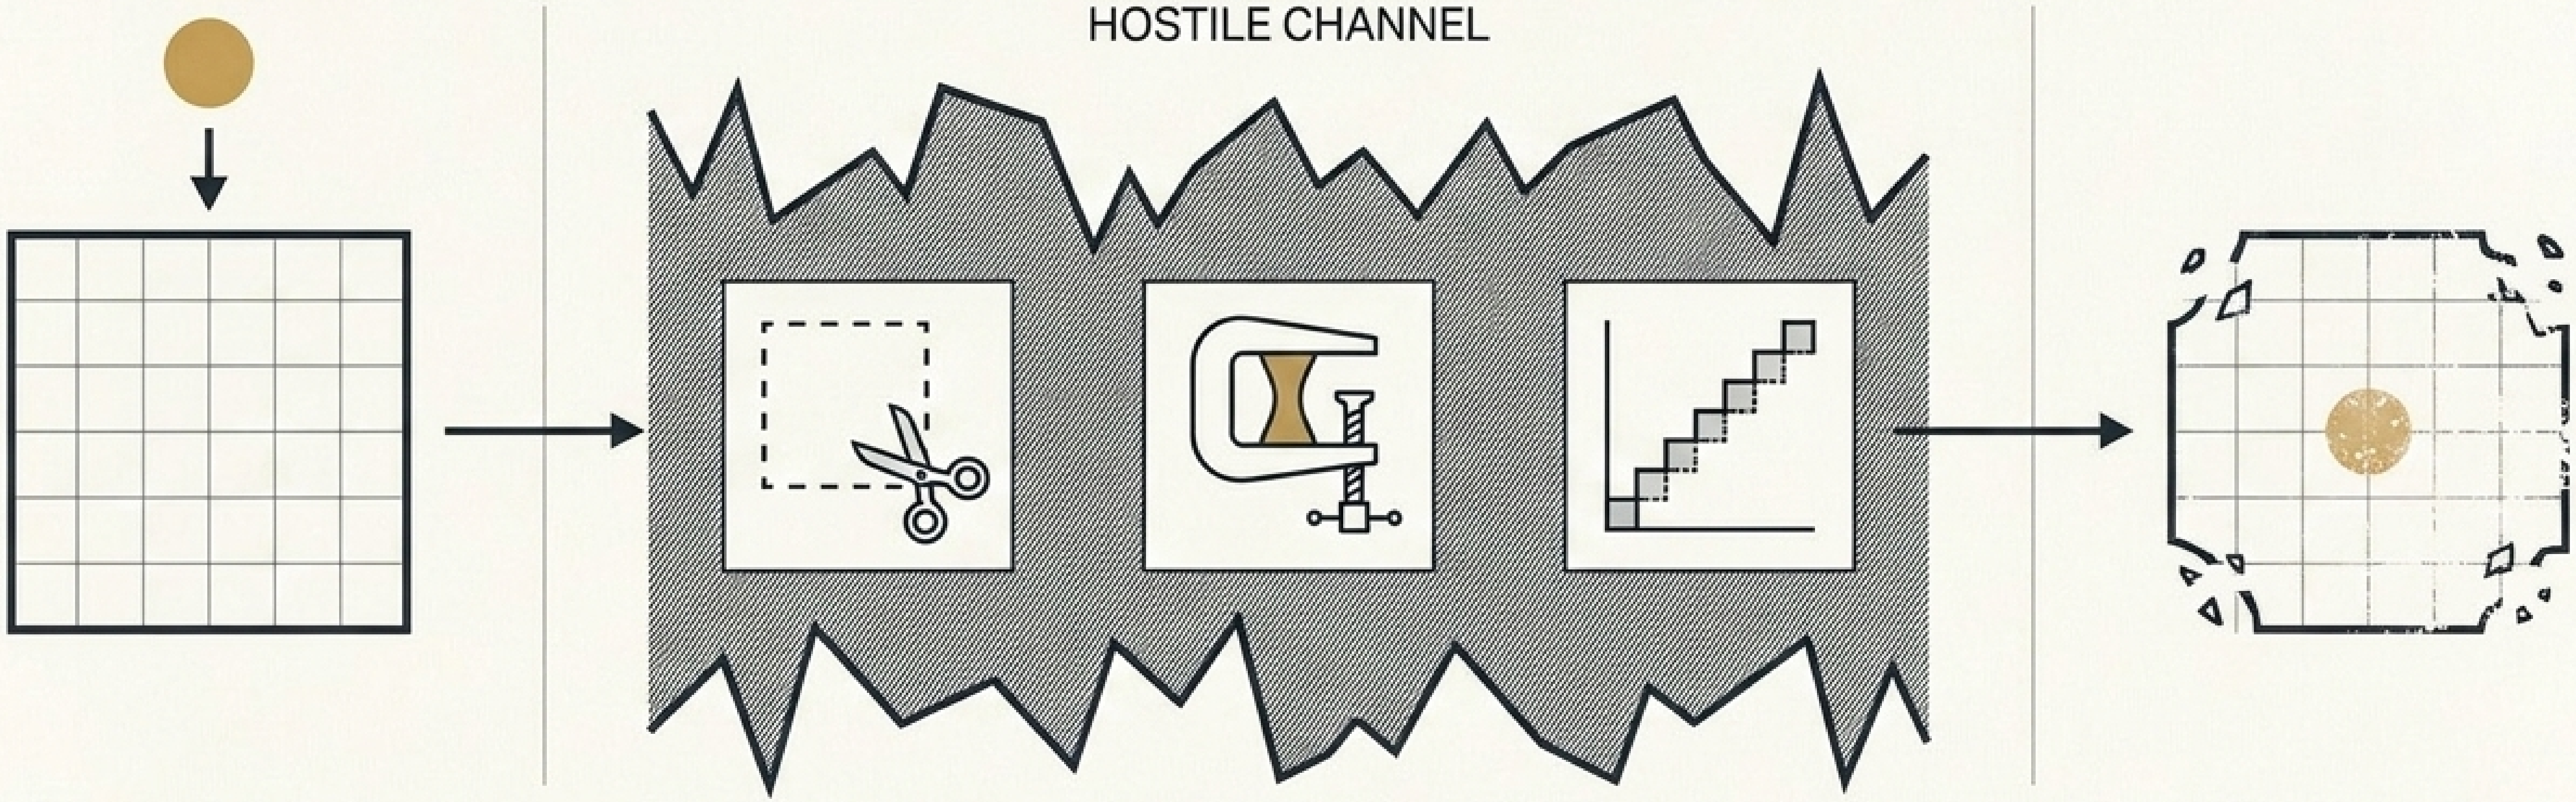

To survive uncertain channels, an encoder must stop requiring its first prediction to be final.

Spatial Assignment vs. Field Construction

	TRADITIONAL PARADIGM	SPARSE RECURSIVE HOLOGRAPHIC STEGANOGRAPHY (SRHS)
Payload Mapping	Localized fragments (Pixel/Block)	Distributed measurements (Overlapping field)
Failure Mode	Brittle boundaries (Cropping destroys payload)	Graceful degradation (Cropping reduces overall evidence)
Compute Allocation	Dense, one-pass, static	Sparse, recursive, adaptive (conditional depth)
Channel Strategy	Blind prediction (Hoping changes survive)	Active interrogation (Observing the channel iteratively)

Pillar 1: Holographic Payload Projection

LOCALIZED ASSIGNMENT

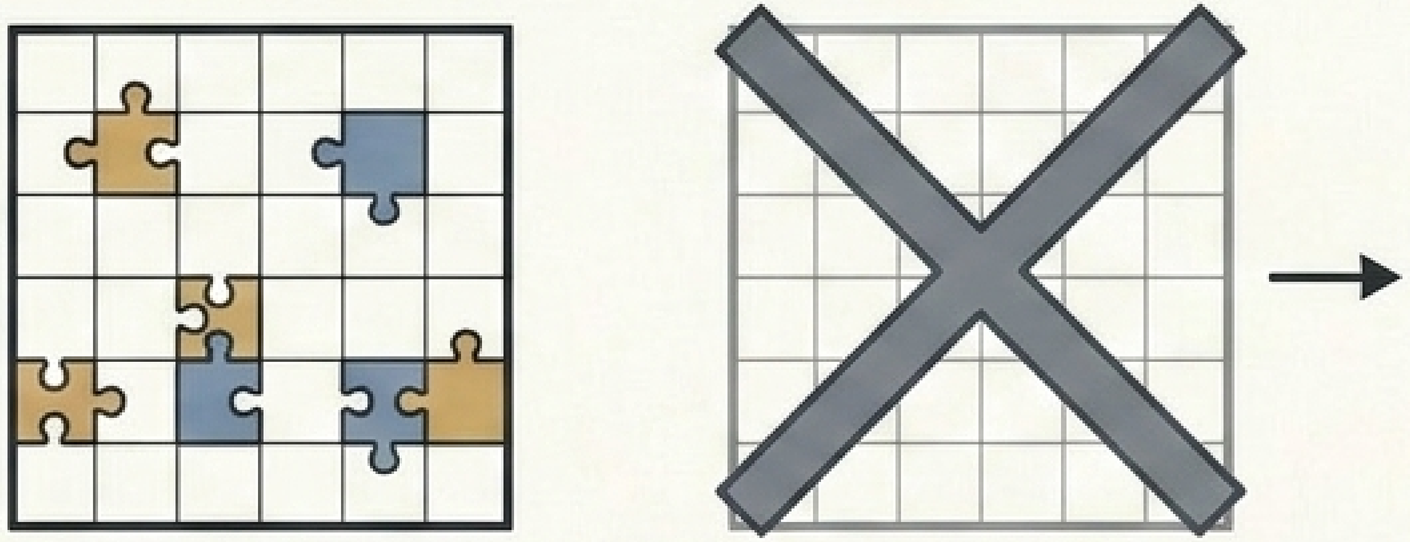

HOLOGRAPHIC PROJECTION

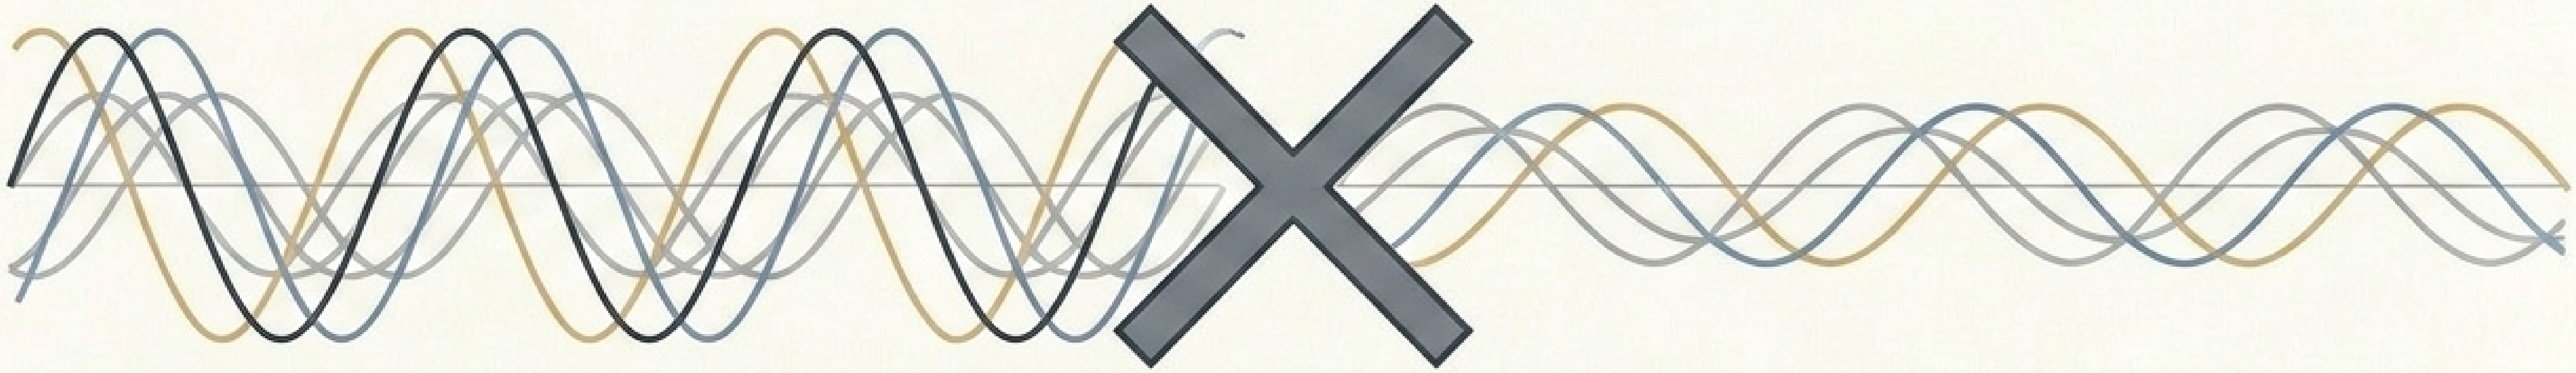

Holographic means distributed recoverability. No pixel contains bit 37. Every partial view contains degraded evidence about the whole message rather than exact evidence about one isolated fragment.

Pillar 2: Adaptive Sparse Routing

Marginal routing value
Should we alter region j at step t?

Expected error reduction
Will this fix the payload?

$$v_{t,j} = \frac{\mathbb{E}[\ell_t - \ell_{t+1} \mid j]}{c_{t,j} + \lambda d_{t,j}}$$

Modification cost
Does this violate the distortion budget?

Detectability risk
Will this alert an adversarial statistical detector?

Insight: Sparsity is not forced; it is an emergent consequence. Most carrier regions have negative marginal value at a given step. The system simply leaves them alone.

Pillar 3: Recursive Residual Correction

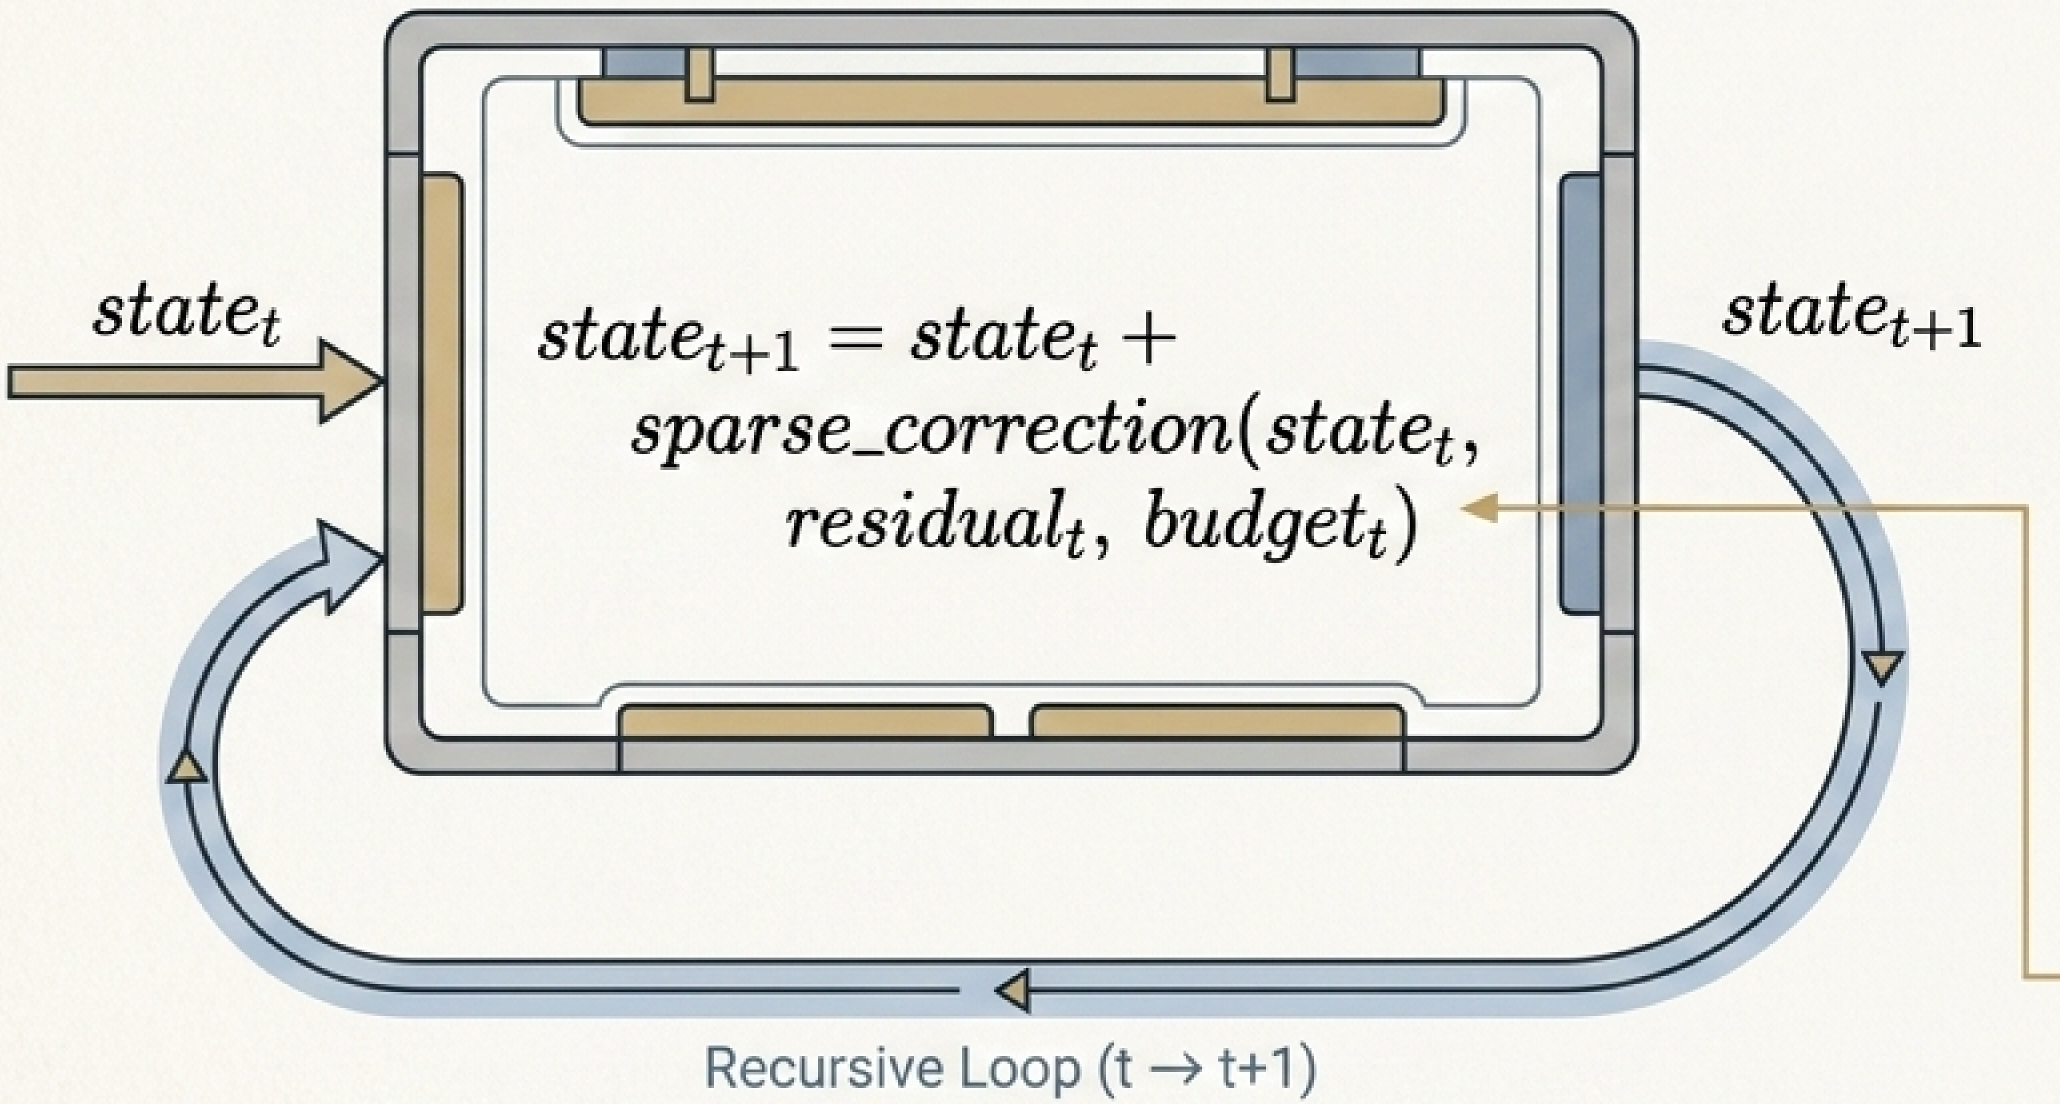

"Try again is not an architecture. Useful recursion requires state."

Increasing recursion depth increases adaptive computation without inflating the parameter count.

Easy cases stop early. Difficult cases receive more steps.

Parameters ψ, θ, ϕ are shared across all steps t.

The Joint Requirement

Sparsity Alone

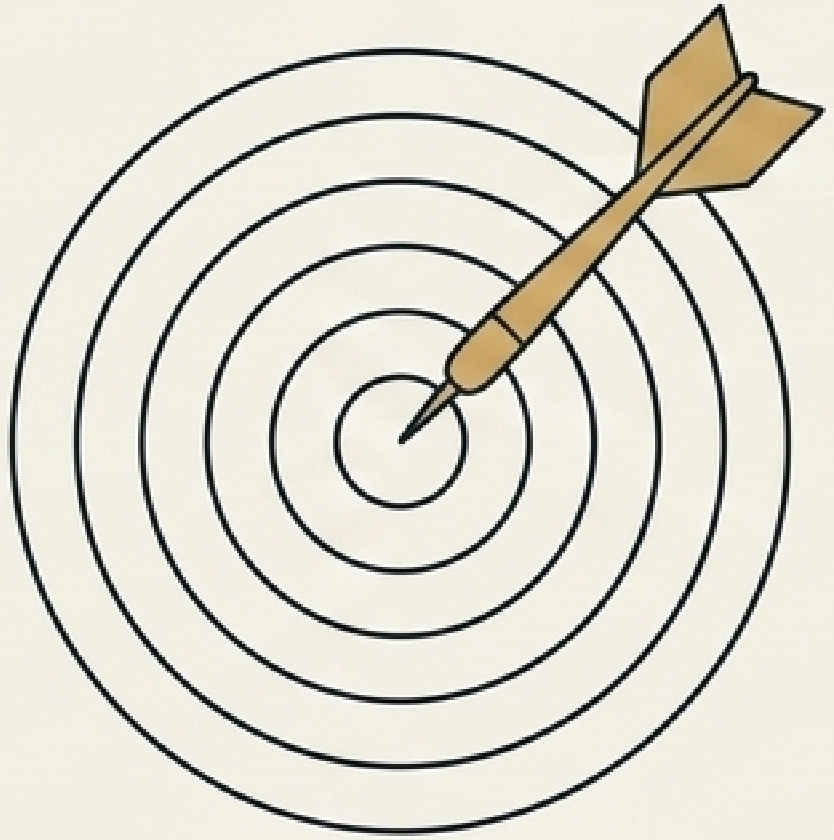

Commits prematurely to a single, possibly mistaken, active set.

The Synthesis

Controlled exploration followed by targeted correction.

Recursion Alone

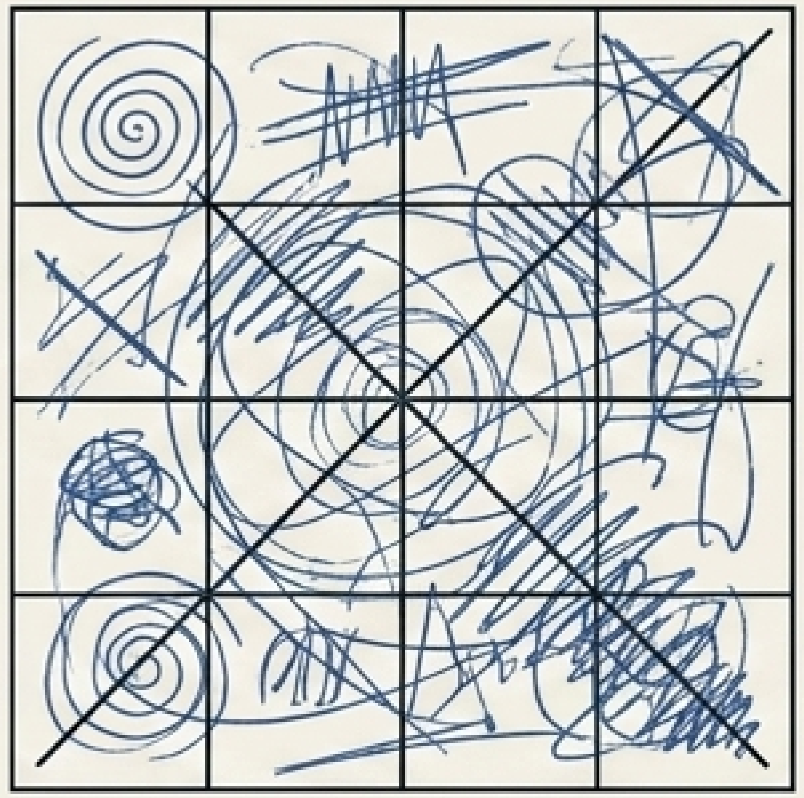

Revisits everything. Perturbs the entire carrier, compounding detectable residue across steps.

Sparsity prevents recursion from disturbing what is already correct.
Recursion prevents sparsity from stranding the system with a bad initial guess.

The Encoder Watches Itself Encode

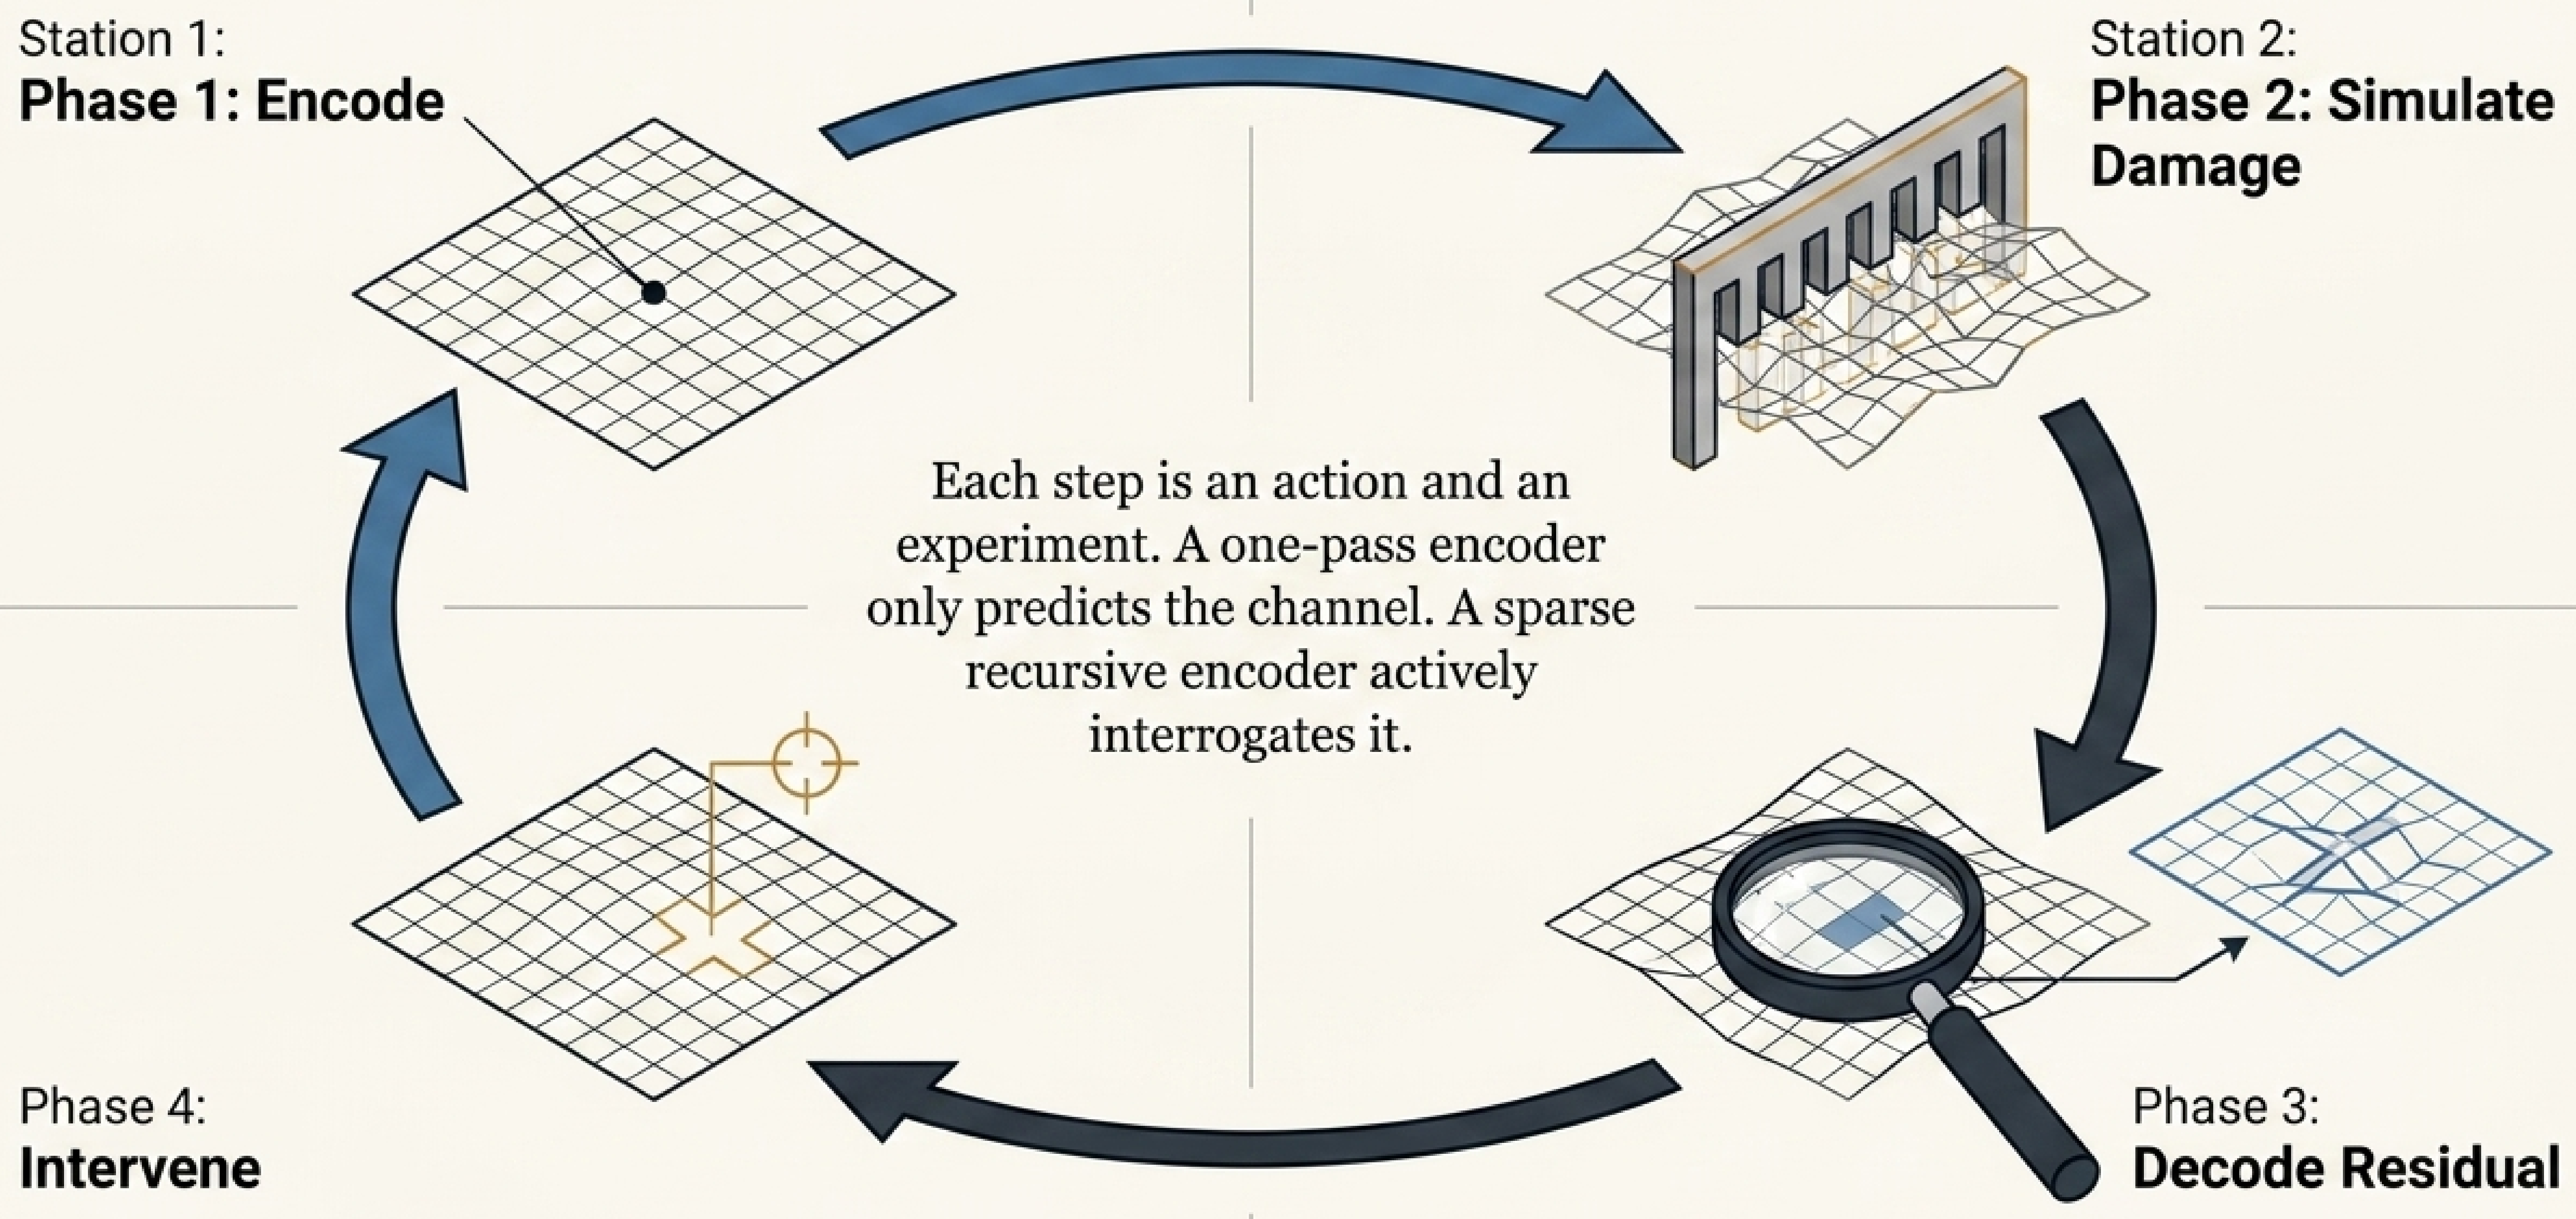

The Capacity Bound: No Free Lunch

$$L \leq C_{\mathcal{T}}(B)$$

Absolute Channel Capacity

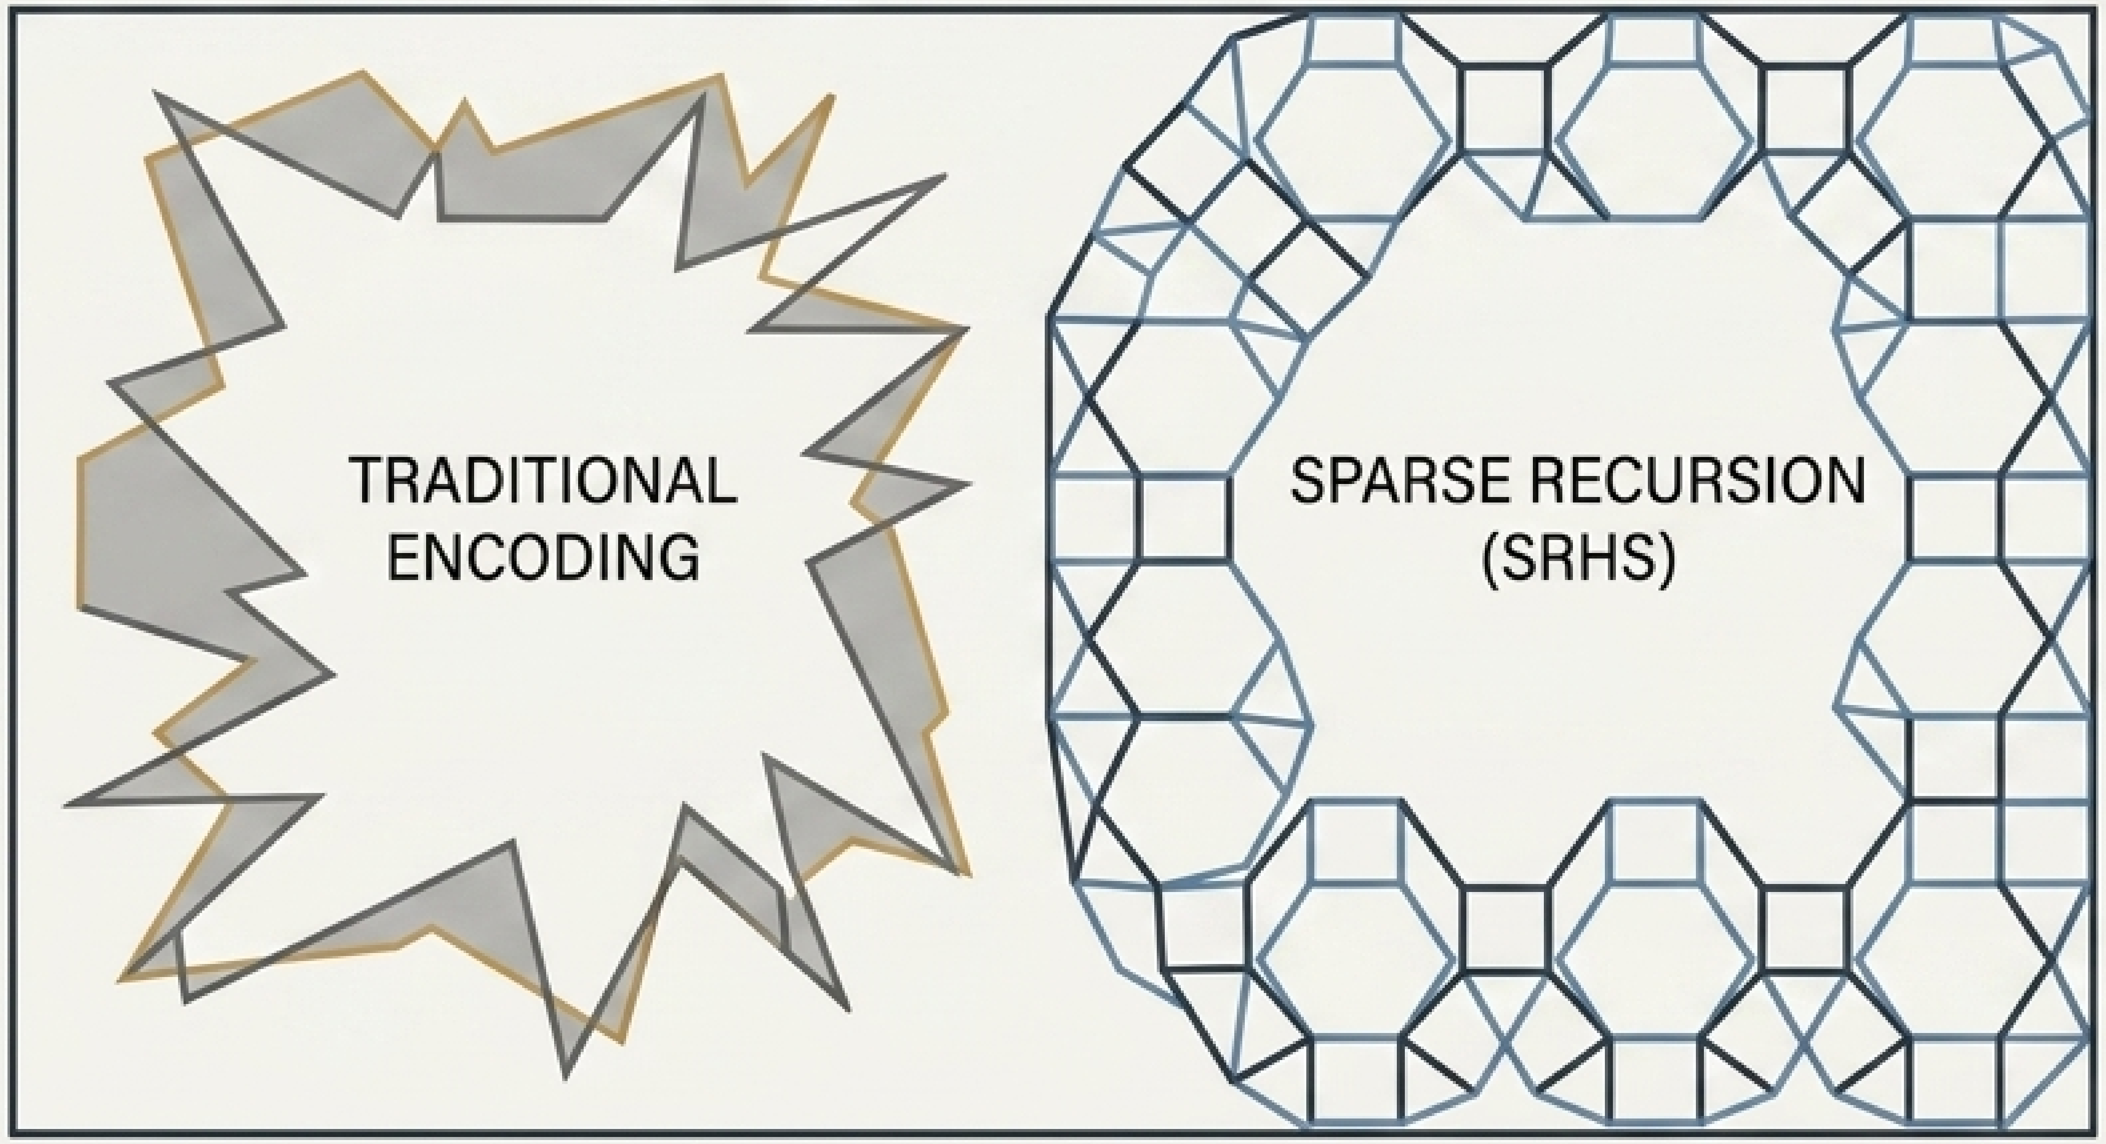

Projection, recursion, and routing cannot conjure capacity from nothing. If a transformed carrier can only support C hidden bits under distortion B, SRHS cannot reliably transmit more.

SRHS simply reduces the amount of existing capacity wasted through poor, blind allocation. It pushes the achievable operating point closer to the absolute theoretical limit.

The Brutal Test: Matched Ablations

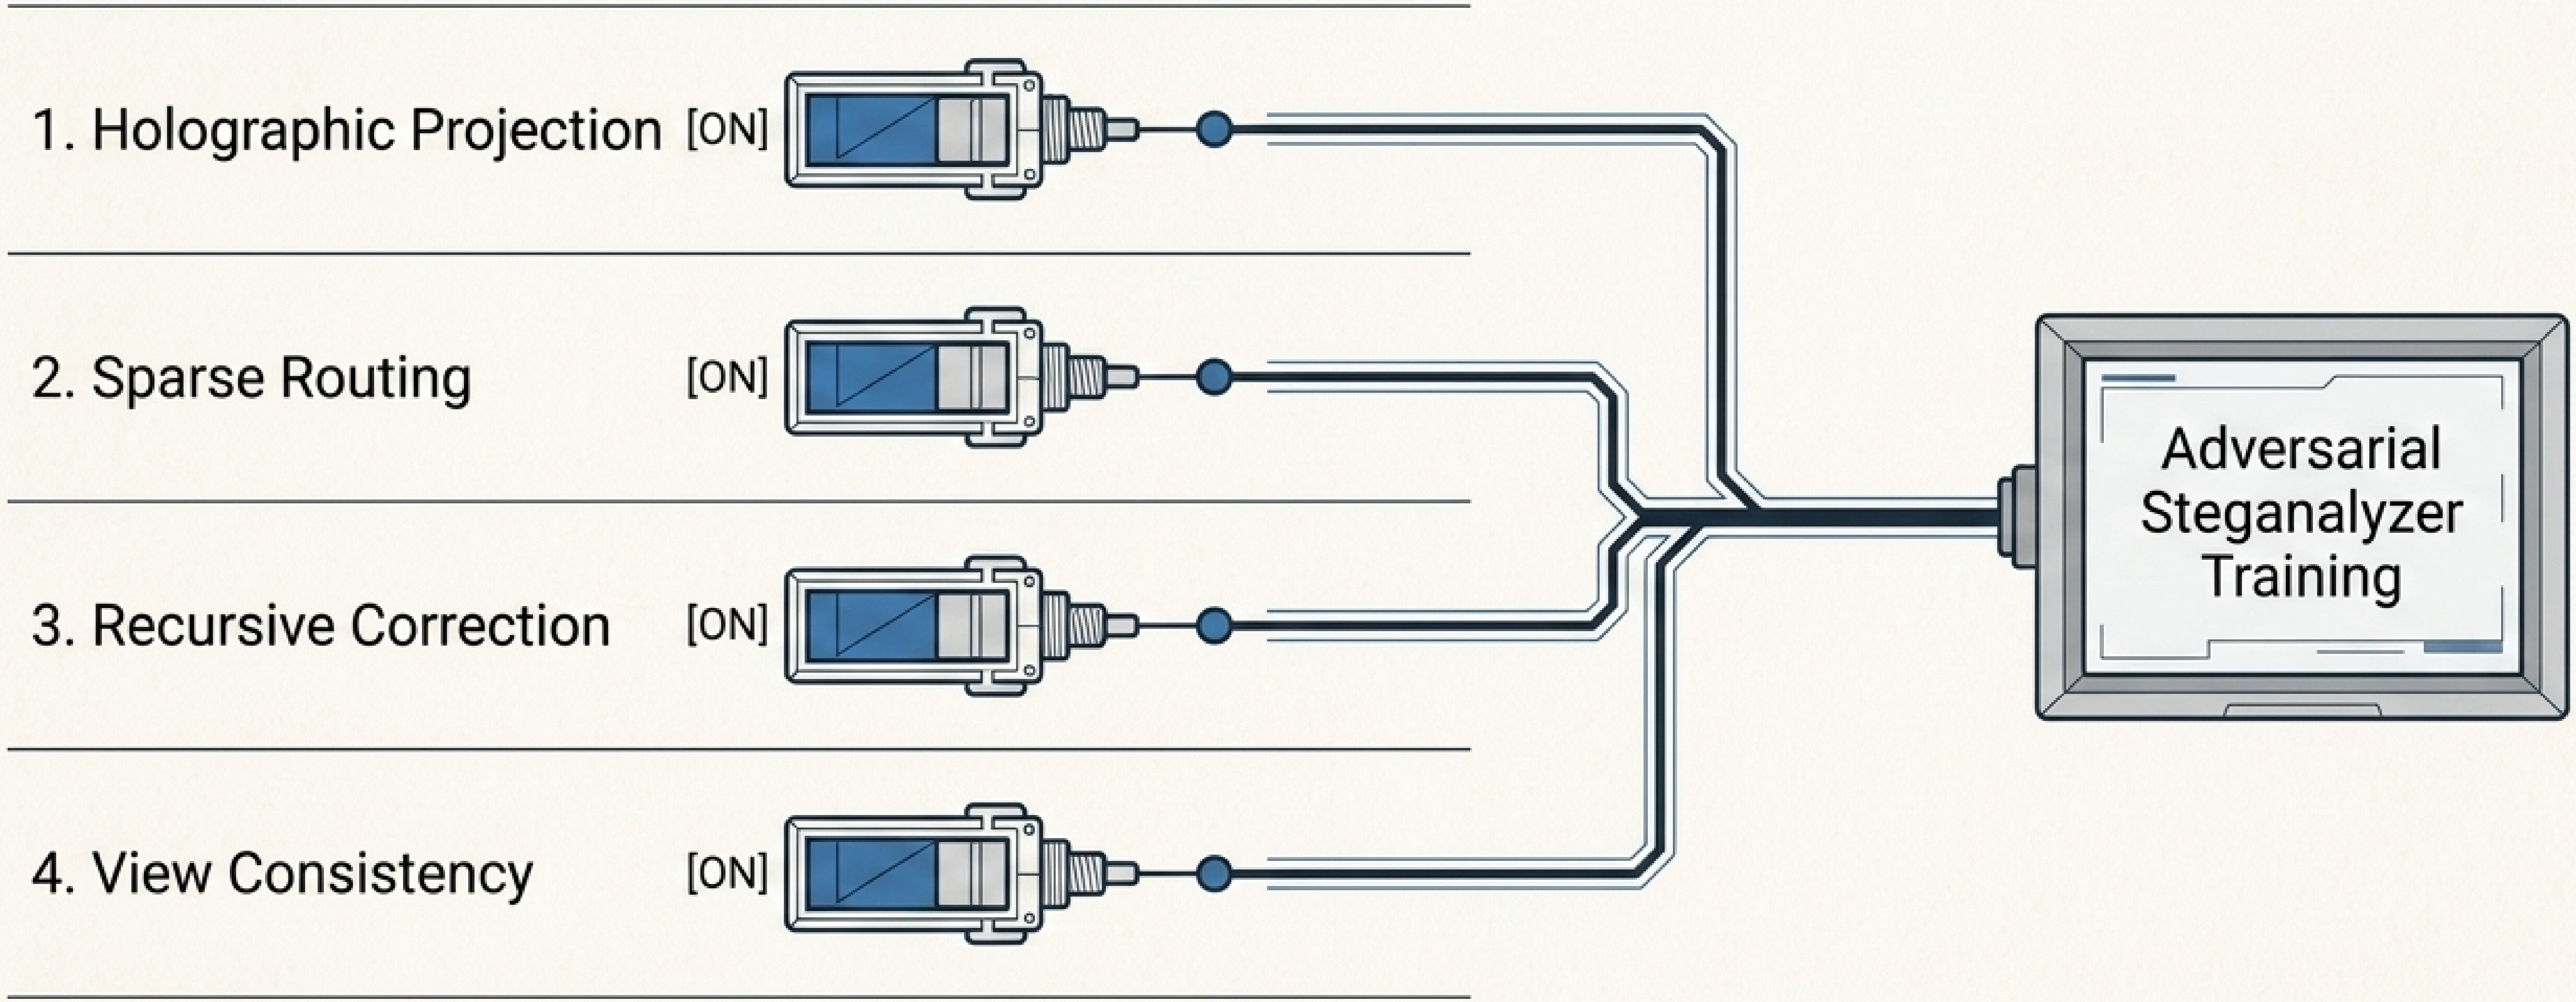

We do not just test the system; we dismantle it. We compare SRHS against strong one-pass encoders, dense recurrent networks, and transform-domain systems—giving them identical carriers, distortion limits, and security tests. Then, we flip the toggles off one by one.

Preregistered Failure Conditions

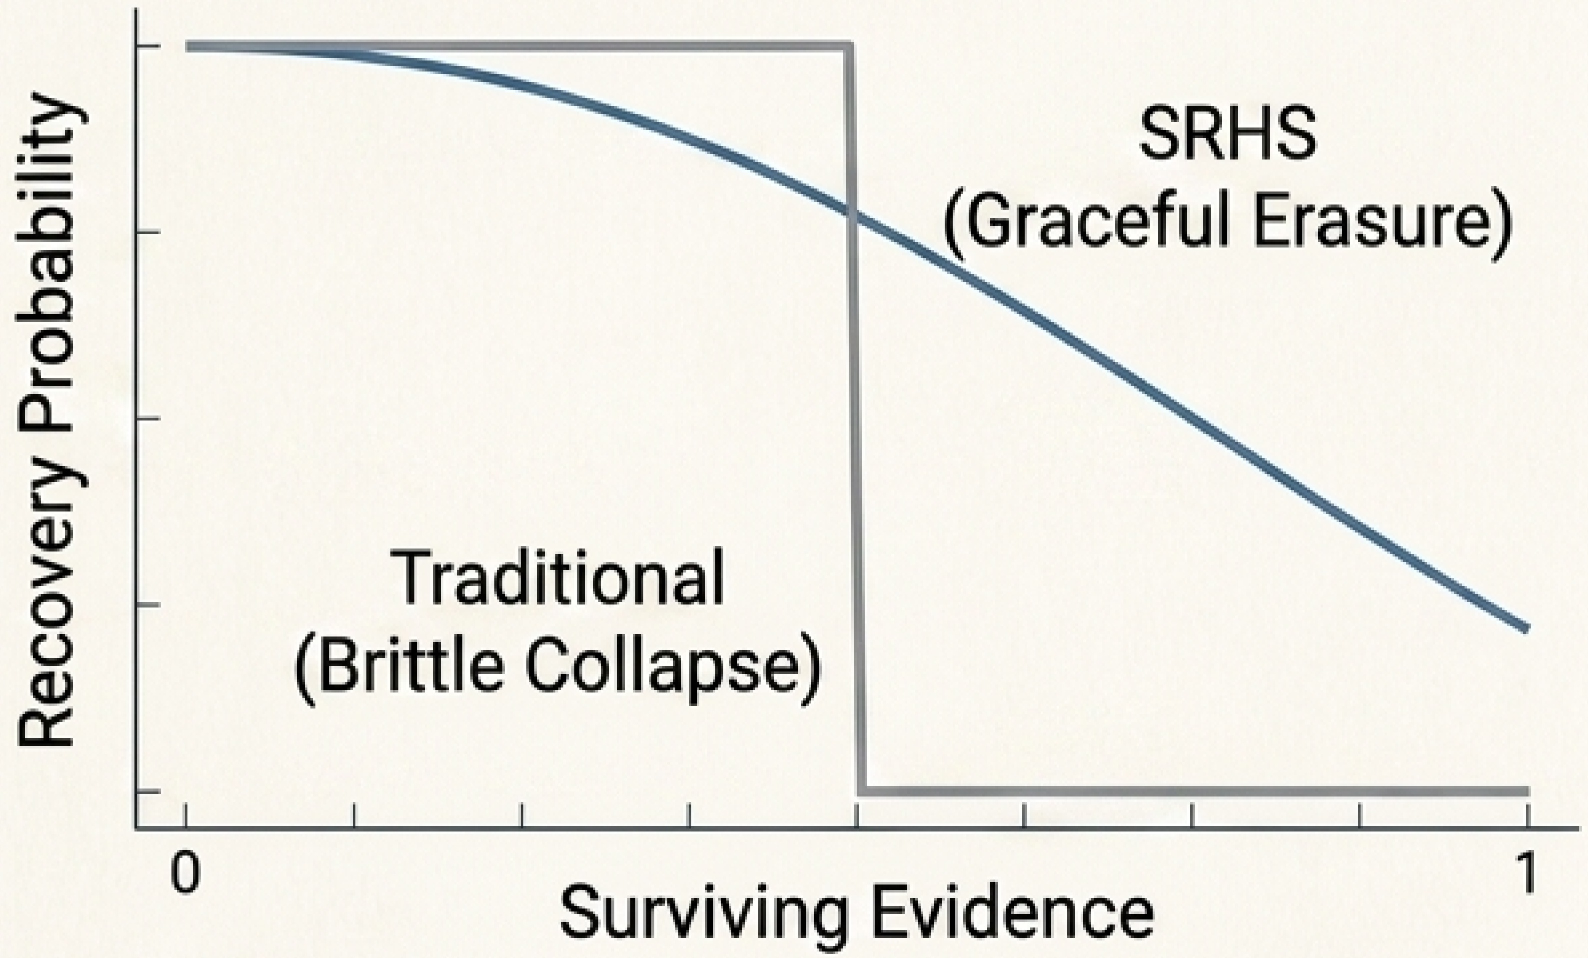

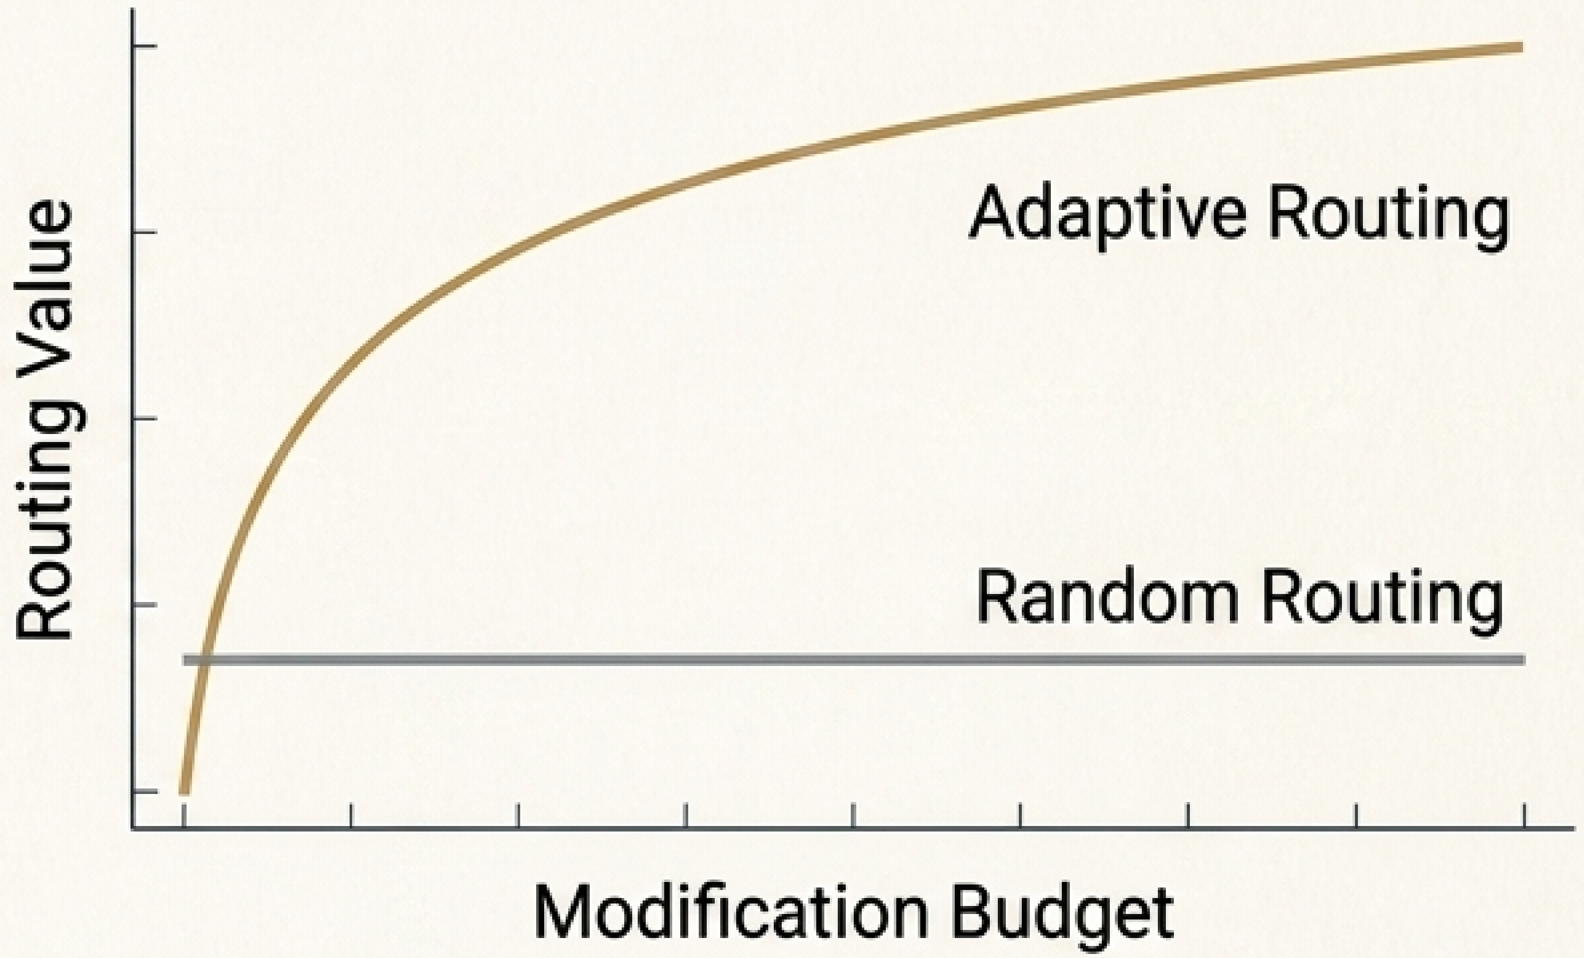

- If partial recovery collapses at fixed boundaries rather than degrading smoothly with evidence... the projection failed.
- If adaptive sparse routing cannot outperform random routing at the same modification budget... the routing failed.
- If additional recursive steps yield no gain after compute matching... SRHS is not something you need.

Beyond Steganography: The Macrocosm

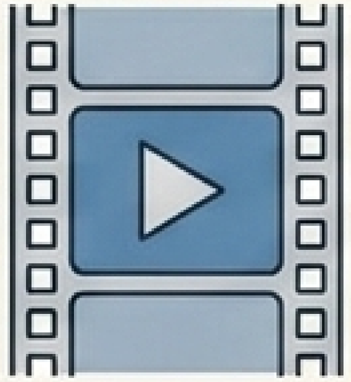

Video Generation

Revisits only objects whose identities drift between frames.

Large Language Models

Reconsiders only the specific claims lacking evidential support.

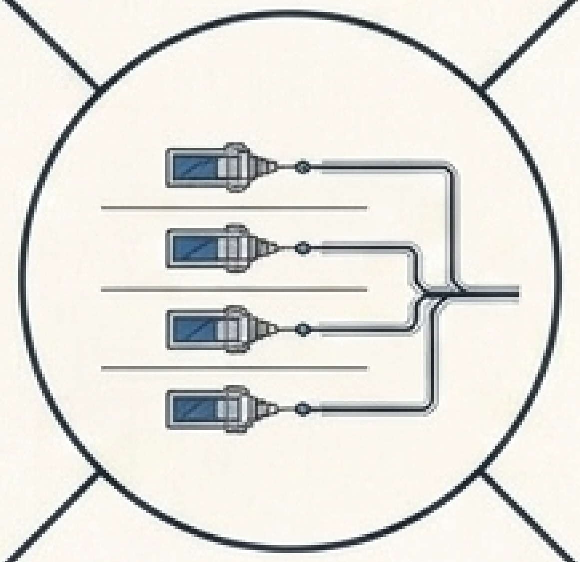

Robotics

Revises only the part of a plan contradicted by immediate physical action.

Rendering

Refines only regions with visible residual error.

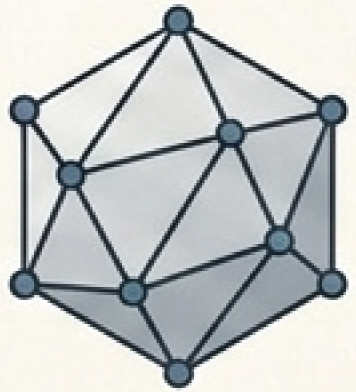

Steganography was only the first clean test. The shared architecture applies to any closed loop where sparse action produces evidence, and that evidence determines the next sparse action.

Scale gave us systems that can attempt almost anything once.

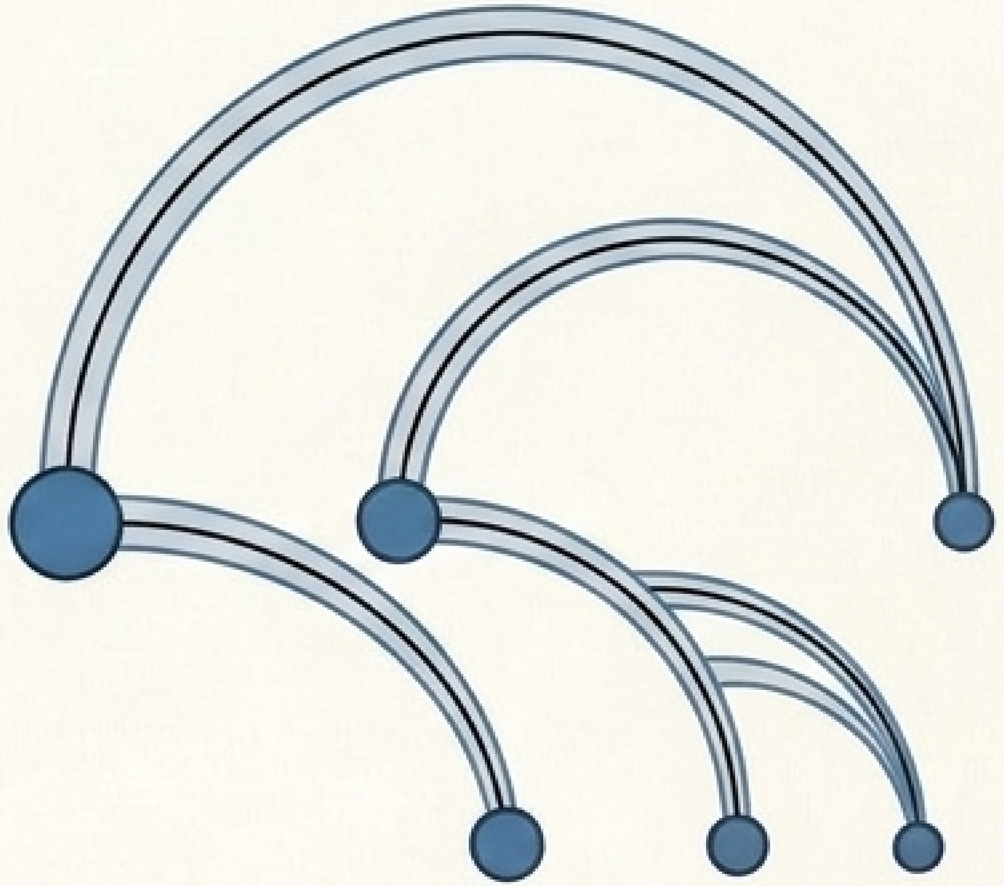

Sparse recursion gives us systems that know exactly where to try again.

Distributed recoverability · Marginal value routing · Residual correction