A finite record reconstruction of the Standard Model representation

Why the eight-bit QEC cell is a constrained local record architecture, not a decorative geometric fit

David Elliman
Neuro-Symbolic Ltd
dave@neusym.ai

28 June 2026

Abstract

Many proposed "geometric theories of everything" begin with an attractive shape and then assign particles to its parts. Such models are easy to dismiss, because resemblance is not a derivation. This paper presents a different kind of claim. Starting from a modest record-theoretic premise—that microscopic physics is a finite, local, noise-tolerant system capable of writing stable records—we ask what binary local record architecture is needed to host the active Standard Model matter representation.

The result is a finite reconstruction statement. Stable repeatable records motivate a quantum-error-correcting substrate; the minimal balanced distance-four binary record cell is the self-dual doubly-even $[8, 4, 4]$ code. Inside that byte, the primitive local roles needed by the Standard Model require eight bits: two generation bits, two colour bits, one lepton/quark bit, one weak-isospin bit, one chirality bit, and one repeatable weak-readout bit. An explicit seven-bit compressed encoding of the Standard Model content exists, but it merges the lepton/quark distinction into the colour plane and therefore fails the local-record criterion. With the eight-bit architecture fixed, finite executable audits show that chirality, $SU(2)$ doublet completeness and exact anomaly cancellation isolate the Standard Model active sector up to colour and name relabelling.

The claim is deliberately conditional. It does not prove that nature is discrete, or that this substrate is true. It proves something narrower and sharper: if one follows the binary, local-record, QEC route, then the active Standard Model representation is not pasted onto an arbitrary geometry. It is the unique chiral anomaly-free survivor in a small natural rule class, and the code/scripts make that statement reproducible.

Contents

1 The problem with "beautiful-shape" theories

There is a large class of speculative theories that start with a geometric object—a polyhedron, lattice, knot, tiling, graph or algebraic pattern—and then place particles or forces onto its components. Sometimes the geometry is striking. Sometimes the numerical coincidences are suggestive. But the usual criticism is fair: a visual match is not a derivation. If enough labels can be moved around after the fact, almost any sufficiently rich geometry can be made to look like the Standard Model.

The present paper is meant to be read against that background. Its central question is not "does this shape look like the Standard Model?" It is:

> *What is the smallest local binary record architecture, under explicit locality and role-separation assumptions, in which the active Standard Model matter sector can appear as a chiral anomaly-free representation?*

That is a reconstruction question. A reconstruction question has a different standard of evidence from an analogy. One must state the allowed rule class, quotient away mere relabellings, enumerate the alternatives, and show where the Standard Model sits.

The paper's main result is therefore not a new particle prediction. It is a finite classification result. Under the record-local assumptions stated below, the eight-bit cell is minimal as a local architecture, and the Standard Model active matter rule is isolated by chirality plus anomaly cancellation.

What this paper is not

To avoid a common failure mode, we state several negative claims at the outset.

- It is not a claim that a decorative polyhedron explains the Standard Model.
- It is not an after-the-fact labelling of known particles onto vertices.
- It is not a replacement for quantum field theory.
- It is not a proof that space is discrete.
- It is not a proof that the wider finite-QEC substrate programme is correct.

It is instead a constrained finite-audit statement about representation space:

within a natural class of local binary record architectures, the Standard Model active matter content is the unique chiral anomaly-free survivor, up to relabelling.

2 Why take discreteness seriously?

The argument does not assume that discreteness is established. It assumes discreteness as a working hypothesis and asks what follows. That hypothesis is not arbitrary. There are several independent reasons why many physicists have suspected that the continuum description of spacetime may be effective rather than fundamental.

2.1 Operational limits

Physics is tested by finite records: detector clicks, pointer readings, memory states and correlations between them. A continuum field theory uses real numbers at every point, but no experiment reads out infinitely many real numbers. This does not prove the world is discrete; continuous mathematics is often the best effective language. It does mean that a theory built from finite records is not automatically less empirical than one written with continua.

2.2 Ultraviolet infinities and effective fields

The Standard Model is not made invalid by its ultraviolet divergences. Renormalisation is one of the great successes of twentieth-century physics: after the correct finite set of measured parameters is fixed, low-energy predictions are stable and extraordinarily accurate [15, 21]. But renormalisation also teaches a conceptual lesson. The continuum field is usually treated as an effective description whose short-distance behaviour must be controlled by a prescription, symmetry, or completion. In a continuum theory there are formally infinitely many independent field modes in any finite region, and the ultraviolet sector has to be handled with care. A finite record substrate starts from the opposite end: a finite region carries a finite information capacity, so there is no literal infinity of independent ultraviolet modes to regulate. That does not prove discreteness, but it makes a finite microscopic substrate a natural hypothesis to test.

2.3 Black-hole entropy and holography

Black-hole thermodynamics assigns a finite entropy to a finite horizon area,

$$S_{\mathrm{BH}} = \frac{k_B A}{4\ell_P^2}, \tag{1}$$

suggesting that a finite region has a finite information capacity [3, 9, 20, 19]. Holography then sharpens the point: the maximum number of independent degrees of freedom appears to scale like area, not volume. Again, this does not prove a literal lattice. But it is strong evidence that the continuum field-theoretic counting of independent modes is not fundamental all the way down.

2.4 Quantum gravity and singularities

General relativity treats spacetime as a smooth manifold. Quantum mechanics treats measured information as finite and probabilistic. Their direct combination sharpens both ultraviolet and gravitational singularity problems. Many quantum-gravity programmes respond by replacing the smooth continuum with more primitive combinatorial, algebraic, causal or informational structures: spin networks in loop quantum gravity [16], causal sets [5, 17], tensor networks and holographic codes [14, 8]. These are not the same proposal, but they share one lesson: continuum geometry may be an emergent large-scale description of a more discrete microscopic substrate.

2.5 Lattice field theory as a warning and a guide

Lattice gauge theory shows that much of quantum field theory can be regularised on a discrete structure without losing its low-energy continuum content [21, 10]. That does not mean nature is a lattice; in ordinary lattice QCD the lattice is a regulator, not an ontology. But it demonstrates that locality, gauge structure and continuum physics can emerge from finite local data. The present paper asks a different, representation-level question: if the microscopic data are finite local records, what record architecture is forced?

3 From stable records to QEC bytes

The first step is not geometric. It is informational.

Definition 1 (Stable local record)**.** *A stable local record is a finite physical degree of freedom whose alternatives can be read repeatably, copied into other records, and compared with other local records without destroying the fact being recorded.*

Repeatability and copyability are not innocent requirements. In quantum theory, only orthogonal alternatives can be copied without disturbance; non-orthogonal states cannot be cloned or broadcast [23, 6, 2]. Repeatable yes/no tests are projectors. Reversible record writing is an isometry, as in Stinespring/Naimark dilation [18, 11, 13]. A world of reusable records is therefore already pushed toward the language of Hilbert spaces, projectors, and error-correcting subspaces.

The finite-QEC substrate programme adds a sharper requirement: local records must survive finite noise. For binary stabilizer-like records, the relevant protection question is not "can one store bits?" but "what is the smallest balanced cell with enough distance to protect erasure/fault structure while retaining reusable local syndrome records?"

Theorem 1 (Minimal balanced record cell, canon status)**.** *Among finite binary local record systems satisfying repeatable readout, commuting read/write stabilizers, local tomography compatibility, and distance-four erasure protection, the unique minimal balanced cell is the self-dual doubly-even* $[8, 4, 4]$ *code.*

In the wider framework this theorem is carried by the reconstruction layer and the script `minimal_balanced_record_cell_theorem.py`. The present paper does not re-prove the coding-theoretic uniqueness from first principles; it uses it as the record-cell input and then asks what matter representation the cell supports. The crucial point is conceptual: the starting point is not "an octahedron is attractive." It is "stable local binary records under finite noise force a protected byte."

4 The eight local roles

The active Standard Model matter content has several logically different labels: generation, colour, lepton/quark type, weak isospin, chirality, and weak participation. In a local record architecture these should not be hidden inside arbitrary lookup tables. They should be local record predicates.

The canonical eight-bit register is

$$c = (G_0, G_1, \mathrm{LQ}, C_0, C_1, I_3, \chi, W) \in \mathbb{F}_2^8. \tag{2}$$

The roles are:

Record role	Bits	Meaning
Generation	G_0, G_1	punctured two-bit rule gives three generations
Colour plane	C_0, C_1	00 is singlet; three nonzero states are colours
Lepton/quark	LQ	independent one-bit matter type
Weak isospin	I_3	two components of a weak doublet
Chirality	χ	left/right record
Weak readout	W	repeatable weak-participation record

Proposition 1 (Primitive local-record lower bound)**.** *If these six roles are primitive disjoint local binary records, then at least eight bits are required:*

$$2_{\text{gen}} + 2_{\text{colour}} + 1_{\text{LQ}} + 1_{I_3} + 1_{\chi} + 1_{W} = 8. \tag{3}$$

This lower bound is deliberately modest. It does not say that no smaller encoding of the particle list exists. It says that no smaller *comparable local-record architecture* exists if lepton/quark type, colour, chirality and weak readout are independent primitive records.

4.1 The seven-bit counterexample

There is a useful counterexample to any overstrong claim. Merge the lepton/quark predicate into the colour plane:

$$(G_0, G_1, CL_0, CL_1, I_3, \chi, W),$$

where $CL = 00$ means lepton and the three nonzero values are colours. This seven-bit scheme has

$$3 \times 4 \times 2 \times 2 = 48$$

states before removing the three right-handed neutrinos, and 45 active states afterward. It can reproduce exact anomaly cancellation and chirality.

But it is not comparable to the eight-bit local-record architecture. There is no primitive one-bit LQ observable: "lepton" means $CL = 00$, while "quark" means $CL \neq 00$. The sterile-neutrino exclusion must inspect the colour/lepton plane. Thus the seven-bit object is a compact *encoding*, not a compact *local record architecture*. The distinction is the point.

5 The four record rules

The finite rule system uses four simple constraints:

$$R1: \quad \neg(G_0 \wedge G_1), \qquad \text{three generations, not four,} \tag{4}$$

$$R2: \quad W = \chi, \qquad \text{weak readout locked to chirality,} \tag{5}$$

$$R3: \quad \text{LQ} = 0 \Leftrightarrow (C_0, C_1) = (0, 0), \qquad \text{leptons colourless, quarks coloured,} \tag{6}$$

$$R4: \quad \neg(\text{LQ} = 0 \wedge I_3 = 0 \wedge \chi = 1), \qquad \text{remove active right-handed neutrino.} \tag{7}$$

The resulting active set has 45 states: three generations of the usual chiral Standard Model matter without an active right-handed neutrino.

Charges are read using the standard weak relation

$$Q = T_3 + Y,$$

with the colour-blind finite readout used in the executable audits. The exact anomaly sums tested are the usual six consistency conditions: $[SU(3)]^2U(1)$, $[SU(2)]^2U(1)$, $[U(1)]^3$, mixed gravitational-$U(1)$, $[SU(3)]^3$, and the Witten $SU(2)$ parity condition [1, 4, 22, 15].

6 Finite classification: where the Standard Model sits

The central computation is not numerical fitting. It is a finite enumeration over Boolean sector rules on the three matter/colour bits (LQ, C_0, C_1).

Definition 2 (Natural Target-A rule class)**.** *The Target-A rule class consists of finite Boolean sector masks on* (LQ, C_0, C_1)*, quotiented by colour* $GL(2, \mathbb{F}_2)$ *relabelling, with fixed* $R1$ *and* $R2$*, and with either:*

(a) the canonical colour-blind right-singlet exclusion $R4$*, or*

(b) any bit-local colour-blind right-singlet exclusion of the form $(\mathrm{LQ} = \ell, I_3 = i, \chi = 1)$*.*

Allowed survivors must be chiral, $SU(2)$*-complete, and exactly anomaly-free.*

The executable classification gives the following ledger.

Rule class	Total	Anomaly-free	Chiral anomaly-free	SM-like survivors
Fixed canonical $R4$	255	2	1	1
Fixed $R4$, colour quotient	79	2	1	1
Wider bit-local $R4$	1020	6	2	2
Wider $R4$, colour quotient	316	6	2	2

The two anomaly-free fixed-$R4$ rules are not two chiral alternatives. One is the Standard Model active rule. The other is vectorlike. In the wider bit-local sterile-exclusion class, the two chiral survivors are the Standard Model rule and its LQ-name flip. Finally, all 3360 role maps tested by the audit are Standard-Model-isomorphic: they are relabellings, not new physical structures.

Theorem 2 (Target-A finite classification)**.** *Within the natural Target-A local rule class, chirality,* $SU(2)$ *doublet completeness and exact anomaly cancellation isolate the Standard Model active matter rule, up to colour relabelling and the* LQ*-name flip.*

7 A small grammar for "simplicity"

One might still ask whether the Standard Model rule is merely one survivor among many complex rules. To test this, the grammar audit defines a small Boolean language over the sector bits. Each atom and each connective costs one abstract-syntax-tree node. Two grammars are used:

Bit-literal grammar: atoms are equality predicates on LQ, C_0, and C_1.

Record-local grammar: the same atoms, plus primitive colour-plane predicates $C = 00$ and $C \neq 00$.

The result is deliberately not overclaimed.

Grammar	SM sector cost	SM sector formula
Bit-literal	11	$((C_0 = 0 \wedge C_1 = 0) \vee \mathrm{LQ} = 1) \wedge ((C_0 = 1 \vee C_1 = 1) \vee \mathrm{LQ} = 0)$
Record-local	7	$(C \neq 00 \vee \mathrm{LQ} = 0) \wedge (C = 00 \vee \mathrm{LQ} = 1)$

Complexity alone does not select the Standard Model. A simpler anomaly-free rule exists at cost 1, but it is vectorlike. The stronger statement is:

> *among the chiral,* $SU(2)$*-complete, anomaly-free survivors of the natural local rule class, the Standard Model rule is the unique fixed-*$R4$ *survivor and has the minimal record-local grammar cost of the surviving wider-*$R4$ *pair.*

This is the right kind of simplicity claim. It is not "the shortest Boolean formula in the universe." It is minimality inside a physics-motivated local rule class.

8 Relation to standard anomaly results

Nothing in this paper changes the known continuum anomaly mathematics. The Standard Model charge assignments are constrained by anomaly cancellation, and classical work by Geng–Marshak and by Minahan–Ramond–Warner showed how strongly hypercharge quantisation follows once the fermion representation content is given [7, 12].

The present result sits upstream. It asks why that chiral representation content is available in the first place. The answer proposed here is: because the $[8, 4, 4]$ local record cell has a small natural Boolean rule class, and the chiral anomaly-free survivor is the Standard Model active sector.

Thus the claim is not that anomaly cancellation is new. The claim is that the finite record cell compresses the input data to which anomaly cancellation applies.

9 Reproducibility

The central claims are executable finite audits. The current scripts are:

Script	Purpose
`target_a_genericity.py`	Exact-rational anomaly arithmetic, colour-weight forcing and genericity scan.
`target_a_classification.py`	Finite classification of natural sector rules, colour quotients and role maps.
`target_a_minimal_record_architecture.py`	Seven-bit encoding counterexample and eight-bit local-record minimality audit.
`target_a_rule_grammar_complexity.py`	Small Boolean grammar complexity audit for sector rules.
`target_a_colour_count.py`	Check that anomaly cancellation forces $N_c = 3$ in the representation pattern.
`target_a_q1q2_cpt.py`	CPT and weak-target audit for the colour-blind charge readout.
`z3_colour_resolution.py`	Z_3 colour-centre/triality audit for the two-bit colour sector.

The public code repository is

`https://github.com/dgedge/itfrombit`.

The working canon repository carries the same scripts under `python_code/`. The classification scripts are self-asserting: exit status zero means the stated finite counts and anomaly sums reproduce.

10 Status ledger

The table below separates what is being claimed from what is not.

Status	Claim
Finite audit	The Target-A rule class has the survivor counts stated above.
Finite audit	The seven-bit compressed encoding exists but fails primitive LQ locality.
Finite audit	The eight-bit primitive local-record lower bound is attained.
Finite audit	The grammar-complexity statement is minimal only after chirality and anomaly-freedom are imposed.
Standard theorem	Gauge anomaly cancellation is required for quantum consistency.
Standard theorem	Stable repeatable records in quantum theory are orthogonal/projective.
Conditional reconstruction	Stable noisy local records motivate a QEC substrate.
Canon-derived input	The minimal balanced distance-four binary record cell is $[8, 4, 4]$.
Open wider question	Whether nature is in fact a binary local-record/QEC substrate.
Open wider question	Whether all dynamics and constants follow from the same substrate.

11 Discussion

The strongest honest sentence is not "this is the simplest possible theory of the Standard Model." That is too broad. A nonlocal encoding, a non-binary alphabet, a continuum group construction, or a different reconstruction programme can evade the premises.

The strongest honest sentence is:

> *If microscopic physics is a finite binary local-record system with stable QEC records, then the eight-bit* $[8, 4, 4]$ *cell is a minimal local architecture for the Standard Model roles, and the active Standard Model matter sector is isolated by chirality and anomaly cancellation inside the natural local rule class.*

That sentence is still strong. It says that the model is not a free geometric labelling. The geometry is downstream of record requirements; the matter content is downstream of a finite classification; the scripts are short enough to inspect; and the caveats are explicit.

For readers wary of speculative "theories of everything", this is the intended point of entry. Do not begin by believing the whole finite-QEC substrate. Begin by checking the finite record reconstruction problem. If the rule class is accepted, the Standard Model active sector is not inserted by hand. It is the anomaly-free chiral occupant of the minimal local binary record architecture.

A Appendix: exact finite statements

A.1 The active content

The canonical active sector contains 45 states. The right-handed neutrino is absent from the active content by $R4$, while the full pre-$R4$ content contains 48 states and is vectorlike. The audit confirms:

$$\sum Q = 0, \qquad \sum Q^2 = 16,$$

and all six anomaly ledgers vanish in the all-left convention.

A.2 Why the seven-bit encoding is not comparable

The seven-bit compressed encoding uses

$$CL = 00 \quad \text{for leptons}, \qquad CL \neq 00 \quad \text{for quarks}.$$

No single bit separates $CL = 00$ from the three nonzero colour states. Therefore a local sterile exclusion cannot read a primitive LQ bit; it must inspect the colour/lepton plane. This is why the encoding is compact but not a comparable local-record architecture.

A.3 Grammar-cost caveat

The grammar audit finds many sector masks cheaper than the SM sector. This is not a problem: most cheap masks are not chiral anomaly-free Standard-Model-like sectors. The conclusion is not global Boolean minimality. The conclusion is minimality inside the physically filtered local rule class.

References

[1] Stephen L. Adler. Axial-vector vertex in spinor electrodynamics. *Physical Review*, 177:2426–2438, 1969.

[2] Howard Barnum, Carlton M. Caves, Christopher A. Fuchs, Richard Jozsa, and Benjamin Schumacher. Noncommuting mixed states cannot be broadcast. *Physical Review Letters*, 76:2818–2821, 1996.

[3] Jacob D. Bekenstein. Black holes and entropy. *Physical Review D*, 7:2333–2346, 1973.

[4] J. S. Bell and R. Jackiw. A pcac puzzle: pi0 to gamma gamma in the sigma model. *Il Nuovo Cimento A*, 60:47–61, 1969.

[5] Luca Bombelli, Joohan Lee, David Meyer, and Rafael D. Sorkin. Space-time as a causal set. *Physical Review Letters*, 59:521–524, 1987.

[6] Dennis Dieks. Communication by epr devices. *Physics Letters A*, 92:271–272, 1982.

[7] Chao-Qiang Geng and Robert E. Marshak. Uniqueness of quark and lepton representations in the standard model from the anomalies viewpoint. *Physical Review D*, 39:693–696, 1989.

[8] Daniel Harlow. The ryu–takayanagi formula from quantum error correction. *Communications in Mathematical Physics*, 354:865–912, 2017.

[9] S. W. Hawking. Particle creation by black holes. *Communications in Mathematical Physics*, 43:199–220, 1975.

[10] John B. Kogut. An introduction to lattice gauge theory and spin systems. *Reviews of Modern Physics*, 51:659–713, 1979.

[11] Karl Kraus. *States, Effects, and Operations*. Springer, 1983.

[12] Joseph A. Minahan, Pierre Ramond, and Nicholas P. Warner. A comment on anomaly cancellation in the standard model. *Physical Review D*, 41:715–716, 1990.

[13] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.

[14] Fernando Pastawski, Beni Yoshida, Daniel Harlow, and John Preskill. Holographic quantum error-correcting codes: Toy models for the bulk/boundary correspondence. *Journal of High Energy Physics*, 2015(6):149, 2015.

[15] Michael E. Peskin and Daniel V. Schroeder. *An Introduction to Quantum Field Theory*. Addison-Wesley, 1995.

[16] Carlo Rovelli. *Quantum Gravity.* Cambridge University Press, 2004.

[17] Rafael D. Sorkin. Causal sets: Discrete gravity. *Lectures on Quantum Gravity*, pages 305–327, 2003.

[18] W. Forrest Stinespring. Positive functions on c*-algebras. *Proceedings of the American Mathematical Society*, 6:211–216, 1955.

[19] Leonard Susskind. The world as a hologram. *Journal of Mathematical Physics*, 36:6377–6396, 1995.

[20] Gerard 't Hooft. Dimensional reduction in quantum gravity. *Conf. Proc. C*, 930308:284–296, 1993.

[21] Kenneth G. Wilson. Confinement of quarks. *Physical Review D*, 10:2445–2459, 1974.

[22] Edward Witten. An su(2) anomaly. *Physics Letters B*, 117:324–328, 1982.

[23] William K. Wootters and Wojciech H. Zurek. A single quantum cannot be cloned. *Nature*, 299:802–803, 1982.