A Discrete-Substrate Perspective on the Yang–Mills Mass Gap Problem

D. Elliman
Neuro-Symbolic Ltd.

May 24, 2026

Abstract

The Clay Yang–Mills mass gap problem asks for a constructive quantum field theory on $\mathbb{R}^4$ for any compact simple gauge group G, satisfying the Streater–Wightman or Osterwalder–Schrader axioms, and exhibiting a strict mass gap $\Delta > 0$. We argue that the Holographic Circlette (TCH) discrete-substrate framework based on the bipartite tensor network $\mathbb{Z}^3 \otimes Q_3$ does *not* solve this problem — and cannot, because the framework is discrete from the start, satisfies Lorentz invariance only emergently, and is built around $\mathrm{SU}(2)_\chi \times \mathrm{SU}(3)_c$ rather than arbitrary G. However, we identify three structural insights that complement the constructive–QFT programme. First, in the TCH framework a finite mass gap is the structural default and gaplessness requires a symmetry reason: the T_{1u} photon multiplet is protected by a representation-parity theorem, and only the $E = +1$ transmission resonance of the macroscopic gauge web admits a flat-band crossing. Second, confinement arises mechanistically from a 3D string-tension argument on the parity-check graph of the $[\![8, 4, 4]\!]$ code, giving a substrate-level lower bound for closed-loop excitations (the structural analogue of a glueball mass floor). Third, the framework's finite-dimensional local Hilbert space places it in the *decidable* regime of the Cubitt–Pérez-García–Wolf undecidability theorem, providing worldview-level evidence that the Yang–Mills mass gap is in principle answerable rather than blocked by Turing undecidability. We conclude that the framework offers a structural complement to constructive QFT, not a substitute, and that the open targets it identifies (§15 items 71 and 73 in the framework's anchor document) constitute concrete next steps for a substrate-level approach to the Clay question.

1 The Clay Problem: What It Actually Demands

The Clay Mathematics Institute Yang–Mills mass gap problem, as formulated by Jaffe and Witten [1], demands the conjunction of three statements:

1. A non-trivial quantum Yang–Mills theory exists on $\mathbb{R}^4$ for any compact simple gauge group G;
2. This theory satisfies axiomatic properties at least as strong as those of Streater & Wightman [2] (in the Minkowski-signature formulation) or Osterwalder & Schrader [3, 4] (in the Euclidean formulation);
3. The theory has a strict mass gap $\Delta > 0$.

The actual difficulty resides almost entirely in item (1) understood in conjunction with item (2). Wilson's lattice gauge theory [5] has demonstrated a mass gap on the lattice for nearly five decades, and modern lattice-QCD numerical simulation computes glueball masses to several-percent precision. None of this resolves the Clay problem because the open question is whether the continuum limit $a \to 0$ exists in the sense required by the Streater–Wightman or Osterwalder–Schrader axioms — in particular reflection positivity, Poincaré invariance as an *exact* symmetry, the spectrum condition, and vacuum uniqueness — and whether the resulting continuum theory retains the lattice-level gap. The lattice serves as a regulator to be removed; the Clay problem is about what remains.

2 Where the TCH Framework Falls Short of Clay

The Holographic Circlette (TCH) framework [9] treats the vacuum as a discrete bipartite tensor network $\mathbb{Z}^3 \otimes Q_3$, with Q_3 the face-adjacency graph of the oblate square bipyramid tiling the 4.8.8 Archimedean structure. For the purposes of the Clay problem, four limitations are immediate and should be stated openly.

Discrete from the start

The lattice spacing $a_0 \approx 0.594$ fm is a *physical* ultraviolet cutoff in the framework, not a mathematical regulator to be removed in a limit. The framework's worldview is that continuum field theory is an emergent long-wavelength approximation; the Clay problem demands the opposite, namely that a continuum theory exists as a fundamental object with the lattice serving only as scaffolding.

A note on lattice scales. The framework's anchor document uses multiple lattice spacings in different physics regimes (the chiral-scale $a_0 = \hbar c/\Lambda_{\mathrm{QCD}} \approx 0.594$ fm of [10, §1.4]; finer sub-scales appear in atomic-shell calculations [10, §10.6]). We use the [10, §1.4] canonical chiral-scale inscription throughout, as appropriate to the QCD mass-gap regime of the Clay problem.

Axioms only emergent

Reflection positivity, exact Poincaré invariance, the spectrum condition, and vacuum uniqueness all hold in TCH only as long-wavelength statements with corrections of order $(a_0/\lambda)^d$ for appropriate dimensional exponents d. The framework's emergence of the Dirac equation from the chiral coin operator is established at the structural level [10, §3.5] but not in the constructive sense demanded by Streater–Wightman or Osterwalder–Schrader.

Specific gauge group, not arbitrary G

TCH is structurally a $\mathrm{SU}(2)_\chi \times \mathrm{SU}(3)_c$ theory via the bipartite Pati–Salam embedding [10, §2.13]. The Clay problem demands construction for arbitrary compact simple G, including the exceptional series F_4, G_2, E_6, E_7, E_8. Generalisation to arbitrary G is not the framework's natural mode.

Continuum limit asserted, not constructively proved

The discrete walk's long-wavelength limit is argued by emergence [10, §3.5] via the coarse-graining of the chiral coin operator. This is a structural-emergence argument, not a constructive proof in the Glimm–Jaffe or Osterwalder–Schrader sense [6].

Given these limitations, the framework cannot — and in this paper does not claim to — resolve the Clay problem. What it can offer is structural perspective: a discrete-substrate worldview in which the mass-gap question takes a different shape, and three concrete insights that complement the constructive–QFT approach.

3 Mass Gap as Structural Default

In TCH, the Q_3 unit cell carries a 256-dimensional local Hilbert space partitioned by the three Boolean $\mathbb{F}_2$ parity constraints (R1, R2, R3) into a 48-dimensional physical subspace $\mathcal{P}$ and a

208-dimensional invalid subspace $\mathcal{Q}$ [10, §2.6]. The constraint Hamiltonian assigns an energy penalty λ per constraint violation, and in the $\lambda \to \infty$ limit yields the projector $\Pi_{\mathcal{P}}$ onto the physical subspace. The spectral gap between $\mathcal{P}$ and $\mathcal{Q}$ is finite and definite:

$$\Delta = 2 \quad \text{(lattice units)}, \qquad 2\Delta = 4 \text{ activates the second-order CKM mixing.} \tag{1}$$

Because the local Hilbert space is finite-dimensional and the projector is explicit, the existence of a gap is structurally automatic. There is no continuum-limit issue, no convergence question, and no analytic subtlety. The framework starts from a gap.

The interesting question in this worldview is the opposite of the Clay question: *where do gapless modes come from?* The framework's answer is that gaplessness is restricted to symmetry-protected modes, and the protection mechanism is explicit.

The representation-parity mass protection theorem

The cleanest statement of the gaplessness-requires-symmetry principle is the framework's mass-protection theorem for the T_{1u} photon multiplet [10, §7.7]:

Theorem (representation-parity mass protection). *Any real antisymmetric matrix projected into an odd-dimensional irreducible representation of O_h possesses at least one exactly zero eigenvalue. The T_{1u} photon irrep is 3-dimensional (odd); the photon is therefore topologically forbidden from acquiring a mass gap. The E_g graviton irrep is 2-dimensional (even); the graviton is permitted a non-zero gap.*

The proof is elementary: eigenvalues of a real antisymmetric matrix come in conjugate pairs $\pm i\lambda$, so in a d-dimensional irreducible representation these pairs account for at most $2\lfloor d/2 \rfloor$ eigenvalues. For odd d at least one eigenvalue must be exactly zero; for even d no such constraint applies. The photon is generated by the antisymmetric directed-flux operator $\partial H_{\text{anti}} = H_{\text{forward}} - H_{\text{backward}}$ projected through the T_{1u} Clebsch–Gordan operator, which inherits the antisymmetric structure.

Numerical verification on the discrete substrate gives the T_{1u} projected mass to 0.000 GeV at machine precision; the E_g band carries a non-zero stiffness ~ 0.22 GeV at the period-4 stable minimum $k = \pi/2$ [10, §10.2], consistent with the $\mathcal{O}(1)$ GeV bare strong-gravity Planck scale.

The $E = +1$ transmission pole as universal substrate invariant

A second mechanism for gaplessness operates at the macroscopic gauge-web level. The walk operator $\mathcal{W} = \mathcal{S} \cdot \mathcal{C}$ on the line graph $L(\mathbb{Z}^3)$ of the simple-cubic gauge bridges has a flat-band crossing at spectral eigenvalue $\lambda = +1$ [10, §7.4]. Under the $C_{4v} \to O_h$ promotion that takes the 2D vertex-figure E_x, E_y doublet to the 3D $T_{1u} = (T_x, T_y, T_z)$ vector triplet, this pole becomes a universal substrate invariant: the fundamental flat-band crossing of the gauge sector's local connectivity graph. The same $E = +1$ pole reappears in [10, §9.7] as the divergence of the bulk Hartman cavity delay $\tau^{\text{bulk}} = -1/(E-1)^2$ and in [10, §13.5] as the band-crossing driving the velocity bifurcation. Three structurally independent appearances of the same spectral object at $E = +1$ are unified at the anchor level as the kinematic origin of the gauge sector's gapless modes.

What this changes for the Clay framing

In the TCH worldview the mass gap is not a property to be proved; it is the generic state of the substrate's spectrum. The structurally non-trivial fact requiring derivation is the reverse: which modes are permitted to remain gapless, and which symmetries protect them. The framework's two explicit answers are (i) odd-dimensional irreducible representations of O_h (the T_{1u} photon) and (ii) the universal $E = +1$ band crossing of the macroscopic gauge web. Everything else is gapped by construction.

This is a worldview-level reframing, not a Clay-problem proof. The bulk mass gap of the 3D TCH octahedral voids — which expels topological flux to the 2D skeleton and produces the string-worldsheet confinement structure of the next section — is anchored explicitly as §15 item 73 of the framework's anchor document [10]: "Substrate derivation of the TCH bulk mass gap from octahedral-void structure" is an *open* computational target awaiting an explicit numerical value identifiable with the string-breaking threshold Λ_{QCD}. This is the framework-internal precision form of the Clay mass-gap question, scoped to the framework's natural energy scale rather than to arbitrary G.

4 Confinement via 3D String Tension

For the strong sector specifically the Clay problem demands a glueball mass lower bound: pure-gauge bound states must not be arbitrarily light. The framework gives a substrate-level mechanism for this, again via the geometric structure of the parity-check code rather than via continuum analytic arguments.

Error strings on the parity-check graph

An error in the $[\![8, 4, 4]\!]$ code creates a parity violation at a single Q_3 cell. Propagating that violation to a separated cell on the macroscopic $\mathbb{Z}^3$ web requires breaking a chain of $\mathbb{F}_2$ constraints — one for each gauge link the error traverses. The energy cost grows linearly with the chain length:

$$E_{\text{error}}(L) = 2\Delta + \sigma L, \tag{2}$$

where σ is the topological string tension — the energy penalty per flipped gauge link — and L is the chain length in lattice units. This is the discrete-substrate analogue of the Wilson area-law statement of confinement [5]: a quark-antiquark pair separated by distance L is connected by a flux tube whose energy grows linearly with L, forbidding their isolation at any finite energy.

2D worldsheet structure and dimensional reduction

The framework's anchor document [10, §15 item 71] establishes a stronger structural statement: under the bulk mass gap of the 3D octahedral voids, the propagating meson sweeps out a strictly 2D worldsheet embedded within the 3D TCH bulk, with transverse hopping into the orthogonal bulk energetically forbidden. The system is therefore Pólya-recurrent in 2D (recurrence theorem applies), and the chiral logarithm of §9.9 in the anchor emerges natively from the 2D worldsheet Green's function rather than from continuum chiral perturbation theory. The string-breaking energy at Λ_{QCD} is identified with the 2D-to-3D dimensional transition: below Λ_{QCD}, the system is rigorously 2D and Pólya-recurrent; above it, the string fragments and the system becomes genuinely 3D and Pólya-transient. The 2D-3D transition is the strict mathematical definition of deconfinement in the framework.

Glueball lower bound as structural consequence

A glueball in this framework is a closed flux loop in the gauge web — a closed cycle on the line graph $L(\mathbb{Z}^3)$. Its mass is bounded below by the smallest closed cycle that satisfies the parity-check structure, multiplied by the natural energy scale Λ_{QCD}. The smallest closed-cycle excitation on the cubic gauge web has spectral weight set by the eigenvalue structure of the simple-cubic line graph; for the 4-cycle (C_4) that bounds a single plaquette, this is the smallest non-trivial gauge-invariant closed-loop excitation supporting a glueball-like state. The result is a substrate-level statement of the same form as the Clay glueball lower-bound demand: glueballs cannot be arbitrarily light because the parity-check structure has a finite minimum cycle length.

A rigorous substrate-level glueball mass calculation for SU(3) is anchored as a downstream open target connected to §15 items 71 and 73; the framework provides the structural ingredients but not yet the closed numerical answer. As with the bulk mass gap, this is the framework-internal precision form of the Clay question, scoped to the specific gauge structure $\mathrm{SU}(3)_c$ rather than to arbitrary G.

5 The Decidability Regime

A separate and recent result casts the Clay mass-gap problem in a new light. Cubitt, Pérez-García, and Wolf [7,8] proved that the spectral gap problem is *undecidable in the Turing sense* for translation-invariant 2D lattice Hamiltonians with finite-dimensional local Hilbert space. There is no algorithm taking the local Hamiltonian as input and returning "gapped" or "gapless" in finite time. Crucially, this is an in-principle obstruction, not merely a computational-complexity one.

This does not imply the Clay problem is unanswerable. The undecidability result concerns the *general class* of finite-dimensional translation-invariant lattice Hamiltonians; the Clay problem concerns a *specific family* (Yang–Mills with compact simple G). Specific instances within an undecidable class can be decidable. Compare: integer factorisation of arbitrary numbers is not known to be in P, but factoring a *specific* number is always finite. The Cubitt–Pérez-García–Wolf result establishes that the spectral-gap question is hard in a precise sense; whether the Yang–Mills instance specifically is decidable is a separate question.

TCH lies in the decidable regime

The TCH framework escapes the Cubitt–Pérez-García–Wolf undecidability for an explicit structural reason: it operates on a *fixed* 256-dimensional local Hilbert space with a *specified* parity-check algebra and an explicit walk operator $\mathcal{W} = \mathcal{S} \cdot \mathcal{C}$. The spectral gap is computable from finite spectral algebra on this fixed structure. There is no Turing-machine encoding of an arbitrary Hamiltonian into the framework; the framework is the Hamiltonian, and the gap is a finite computation.

The structural lesson is worldview-level: by working with a discrete-substrate theory whose local algebra is fixed in advance rather than parameterised over an open class, the framework places the spectral-gap question in a regime where the undecidability does not apply. This is not a proof of Clay decidability for the continuum Yang–Mills theory — but it is structural evidence that the *specific* Yang–Mills instance, viewed as the long-wavelength limit of a discrete substrate with fixed local Hilbert space, is in principle a finite computation.

Precedent in the framework: the Strong CP problem

The framework has a working precedent for substrate-level engagement with Clay-adjacent problems. The Strong CP problem ($\bar{\theta} \approx 0$) is treated in [10, §15 item 93] via a substrate-level structural argument: "The Laplacian of a finite undirected graph is strictly a real symmetric matrix — natively possessing exactly zero complex phase degrees of freedom. You mathematically cannot continuously wind a discrete binary graph." On the discrete substrate the topological winding degeneracy that underlies the continuum θ-vacuum does not exist: the bare UV lattice has $\theta_{\mathrm{UV}} \equiv 0$ identically. The consequence is a substantive falsifiable prediction — the neutron EDM bound $d_n \sim 10^{-31}\,\mathrm{e}\cdot\mathrm{cm}$, with any experimental detection of $d_n > 10^{-30}\,\mathrm{e}\cdot\mathrm{cm}$ (SNS nEDM, n2EDM) falsifying the discrete-substrate origin of QCD.

The Yang–Mills mass gap discussion of this paper follows the same shape: a structural property of the discrete algebra (in this case, finite local Hilbert space + fixed parity-check structure) yields a worldview-level consequence for the continuum problem (in this case, decidability in principle). The framework's pattern of converting Clay-adjacent questions into substrate-level structural statements with falsifiable downstream consequences is established and replicable.

6 Scope: What This Paper Does and Does Not Claim

For clarity to a working physicist reader, the explicit list:

What this paper does NOT claim

- This paper does *not* solve the Clay Yang–Mills mass gap problem. The prize criteria demand a constructive continuum quantum field theory on $\mathbb{R}^4$; we offer a discrete-substrate perspective complementing that programme, not substituting for it.
- We do *not* establish Streater–Wightman or Osterwalder–Schrader axioms for any continuum limit of the framework. Those axioms hold only emergently in the long-wavelength regime, with corrections at scale $a_0 \approx 0.594$ fm.
- We do *not* address arbitrary compact simple gauge group G. The framework is structurally $\mathrm{SU}(2)_\chi \times \mathrm{SU}(3)_c$; generalisation to exceptional groups is not its natural mode.
- We do *not* constructively prove the continuum limit of the framework. The long-wavelength emergence of the Dirac equation is established at the structural level, not in the Glimm–Jaffe sense.

What this paper DOES claim

- The discrete bipartite parity-check structure of the $[\![8,4,4]\!]$ code on $\mathbb{Z}^3 \otimes Q_3$ makes a finite mass gap the *structural default* of the spectrum; the explicit gap is $\Delta = 2$ in lattice units between the 48D physical and 208D invalid subspaces.
- Gaplessness in the framework is restricted to symmetry-protected modes. Specifically: (i) odd-dimensional irreducible representations of O_h (the T_{1u} photon, by the representation-parity theorem of [10, §7.7]), and (ii) the $E = +1$ flat-band crossing of the macroscopic gauge web (universal substrate invariant, [10, §7.4]).
- Confinement is a substrate-level structural consequence of the parity-check algebra: errors propagate as 1D strings with linear-in-length energy $E_{\text{error}}(L) = 2\Delta + \sigma L$, giving Wilson area-law confinement and a structural lower bound for closed-cycle (glueball-like) excitations.
- The framework operates in the decidable regime of the Cubitt–Pérez-García–Wolf undecidability theorem: working with a fixed finite-dimensional local Hilbert space and explicit walk operator places the spectral-gap question in a finite computation rather than in the undecidable open class.
- The framework's anchor document already names the bulk mass gap derivation as open target [10, §15 item 73] and the string-worldsheet closure as open target [10, §15 item 71]; this paper is a perspective-level framing of these existing open targets, not a closure of them.

7 Conclusion

The Clay Yang–Mills mass gap problem and the TCH discrete-substrate framework are asking different questions. The Clay problem asks whether a continuum quantum field theory exists with specified axiomatic properties and a strict gap. The framework starts from a discrete substrate whose gap is structurally automatic and asks instead which modes are protected from acquiring it. These are complementary, not competing, programmes.

Three structural insights from the discrete substrate are worth contributing to the broader conversation around the mass-gap problem. First, the worldview inversion: gaplessness, not gappedness, is the spectrally non-trivial phenomenon, requiring explicit symmetry protection. Second, the substrate-level confinement mechanism: parity-check error strings give Wilson area-law confinement directly from the code structure, and the 2D worldsheet topological dimensional reduction provides a natural framework for glueball lower bounds. Third, the decidability lesson: by fixing the local Hilbert-space dimension and the parity-check structure in advance, the framework lies in the regime where Cubitt–Pérez-García–Wolf undecidability does not bite, supplying structural evidence that the Yang–Mills instance specifically is in principle a finite computation.

None of this is a Clay-prize-quality proof. The framework will not produce one, because it is built to answer a different question — the structural origin of the Standard Model's specific gauge group, mass spectrum, and confinement properties from a fixed discrete substrate, not to constructively prove continuum existence for arbitrary G. But the framework's pattern of converting Clay-adjacent questions into substrate-level structural statements with falsifiable downstream consequences (Strong CP / neutron EDM bound being the working precedent) suggests that the Yang–Mills mass gap question may yield to similar treatment when the open targets at §15 items 71 and 73 are closed.

References

[1] Jaffe, A. & Witten, E. *Quantum Yang–Mills Theory.* Clay Mathematics Institute Millennium Prize Problem statement, `https://www.claymath.org/millennium-problems/`.

[2] Streater, R. F. & Wightman, A. S. *PCT, Spin and Statistics, and All That.* Benjamin, New York (1964); Princeton University Press reprint (2000).

[3] Osterwalder, K. & Schrader, R. *Axioms for Euclidean Green's functions.* Communications in Mathematical Physics 31, 83–112 (1973).

[4] Osterwalder, K. & Schrader, R. *Axioms for Euclidean Green's functions II.* Communications in Mathematical Physics 42, 281–305 (1975).

[5] Wilson, K. G. *Confinement of quarks.* Physical Review D 10, 2445 (1974).

[6] Glimm, J. & Jaffe, A. *Quantum Physics: A Functional Integral Point of View*, 2nd edition. Springer (1987).

[7] Cubitt, T. S., Pérez-García, D. & Wolf, M. M. *Undecidability of the spectral gap.* Nature 528, 207–211 (2015); arXiv:1502.04573.

[8] Cubitt, T. S., Pérez-García, D. & Wolf, M. M. *Undecidability of the spectral gap (full version).* Forum of Mathematics, Pi 10, e14 (2022).

[9] Elliman, D. G. *A Model for John Wheeler's It from Bit: An Eight-Bit Discrete Origin for the Laws of Nature.* Neuro-Symbolic Ltd., Amazon Publishing (2026). ISBN 978-1919558875.

[10] Elliman, D. G. *Holographic Circlette Framework: Anchor Document*. Internal canonical reference; section numbers cited refer to the working anchor (ANCHOR.md, 2026 revision).