VERSF Standard Model Derivation Ledger Nov 26

November assessment against the ledger of 19 October 2026

Programme assessment

About 80% mature on the programme's planning scale, compared with about 75% in October. The November range remains 76–82%.

The main change is that VERSF now has a definite candidate for the overall strengths of the three gauge interactions. In October the theory related the strengths to each other while leaving a shared multiplier undetermined. A seven-site construction now gives that multiplier as 31/59 under explicit assumptions. This turns the remaining disagreement into a numerical test for each interaction. [B, §§4–5; N14, §10; N15, §§1, 5–8]

With colour setting the comparison scale, the original weak and hypercharge inverse-squared couplings need changes of +0.795% and −1.635%. A joint Gaussian model now removes 81.7% of the hypercharge discrepancy for one placement. Across its relaxed family the weak correction is exactly zero. Both effects must be explained by a physical coupling to matter and records. [N15; A1, §§3–7]

What else changed

The public Yang–Mills supplement adds a volume-independent mass-gap bound for specified strongly coupled SU(N) models, including SU(3). A separate finite model combining matter, gauge fields and records has a unique lowest state and a positive gap. These results extend the earlier Abelian analysis. The physical non-Abelian continuum construction remains open. [N11; W1, theorem 4.2, §§10, 14]

Particle masses were substantially covered by October. The programme calculated charged-lepton and Higgs ratios, absolute mass benchmarks, all six quark Yukawa values, a W/Z/photon mass pattern and a relative neutrino spectrum. The website also contains a six-quark calculation with a down-quark input and an earlier proton calculation with an electron-scale input. Pages 5 and 11 show what each uses and predicts. [B, §6; W4–W9]

The next two milestones

To reach 90% Use one physically admissible finite model to pass independent force and flavour tests and calculate a colour observable. The four tests are on page 14.

To reach 100% Derive the remaining inputs, establish the interacting quantum limit and reproduce the claimed observables with controlled uncertainties. The completion criteria are on page 15.

The website audit and its sources are on page 13. Three follow-up calculations used to define the roadmap appear as A1–A3 in the source register.

How the assessment changed

The five layers retain October's definitions. Each November range expresses a judgement about the maturity of the work at that layer. The overall estimate is made separately, because several layers depend on the same physical construction. [B, §1]

Layer	October	November	Reason for movement
Mathematics and representation	94–96%	95–97%	Exact blocking and probability-measure results add to an already mature layer. The rank-three and census premises carry forward. [N02, N10]
Conditional construction from one action	85–89%	89–93%	Finite field–matter–record models now test both spectral positivity and sector-dependent response. The wheel gives a shared stiffness on its own assumptions. [N13–N15; W1–W2]
Numerical reproduction	55–61%	62–68%	A fixed candidate can now be tested in two sectors after one calibration. Three-sector running and explicit residual tests strengthen this layer. Mass and mixing predictions retain their October grades. [N03, N08, N14–N15]
Primitive physical admission	46–51%	50–56%	Native event and current calculations determine more of the allowed response. The physical source profiles, permitted record changes and common normalisation still affect the result. [N07–N09, N13–N15]
Continuum and scheme or pole completion	45–50%	49–55%	Exact chiral blocking, a Gaussian refinement result, the SU(N) fixed-spacing volume bound and one finite scheme conversion address distinct limits. Physical joint refinement and pole matching remain open. [N02, N11, N15; W1]
Overall planning assessment	72–77% ≈75%	76–82% ≈80%	Gauge normalisation shows the largest new gain. Earlier numerical mass results count in the baseline. Matching them to observables, confinement and the interacting limit still carry significant weight.

Why the range stays at 76–82%

The upload-based review put November at 79%. About 80% is the appropriate precision after the website check. The non-Abelian bound and the calculated response in twelve gauge directions strengthen the upper part of the assessment. The weak-force residual, failed quark-mixing tests and physical continuum limit constrain it. The older mass results belong in October's baseline. The evidence does not resolve a one-point change.

This is a judgement of research maturity using October's definitions. The percentage has no statistical interpretation and cannot be read as a count of measured parameters derived.

From a relative gauge relation to a fixed candidate

What October had already established

On the declared particle census, one relation among the three inverse-squared couplings eliminates both the common conversion factor and the comparison scale at one loop. With strong and electromagnetic inputs, the subsequent gauge two-loop plus leading-top calculation gave $\sin^2\theta W = 0.227367$ against 0.23122. That two-input result is the baseline. [B, §5]

What the later calculations supply

Integrating the effective fluctuations of matter and gauge fields gives the one-loop coefficients (−7, −19/6, 41/6) for colour, weak and conventional hypercharge. The calculation explains how the strengths vary with scale on the inherited light spectrum. The derivation is insensitive to a class of 46 real matter-normalisation parameters. [N02, §§6–9; N03, §§3–7]

On a seven-site cell the complete closed-path calculation gives six positive cycle responses. Ordinary quadratic transport preserves the charge-trace proportions. Nonlinear composition can change them on a connected pair of cells. Separate history and completed-event calculations show that recording, relaxation and the choice of observed variables determine which changes can reach the field. [N01, §§3–6; N04, §§3–8; N07–N08]

With provisional fixed normalisation, an earlier quadratic model gave an inverse electromagnetic coupling of 249.88 against 127.95. A nonlinear scenario brought this to 143.12 under a common-factor assumption; changing sector weights could remove the enhancement. A separate Abelian Hamiltonian subsequently calculated current memory and an 8.31% reduction of inverse field susceptibility. These examples establish mechanisms on their own models. [N08, §§14–15; N09, §§9–14, 18; N13, §13]

How the wheel fixes a number

A least-squares response, common charge trace and spacetime isotropy leave one stiffness factor. A further assumption places the responding variable in series with the field gradient. On the specified seven-site transfer, the contrast between hub and rim has eigenvalue 31/28. The series response then gives $(31/28)/(1 + 31/28) = 31/59 = 0.525424$. The same value multiplies electric and magnetic terms in all twelve directions. [N14, §§10.2–10.4]

The result is 0.43% above the data-derived one-loop mean 0.523163 on the baseline (6, 13/2, 21/2). The target was known when the wheel assumption was formulated. The doubled locking baseline would give a one-loop target near 0.26158. Physical use therefore requires the field normalisation and clock to match the chosen baseline. One admissible realisation satisfying these conditions suffices. [N14, §§10.5–10.6]

What the residual test actually measures

The residual manuscript fixes the conventional matter vertices and takes the candidate inverse-squared couplings to be Z = (31/59) × (6, 13/2, 21/2). The strong interaction sets the scale to approximately 8.09 × 10^{13} GeV in the paper's running scheme with zero additional native finite matching. This leaves weak and hypercharge as comparisons. [N15, §§1, 5, 7.4]

Sector	Candidate Z	Comparison Z	Required change in Z	Required change in g
Colour	3.152542	3.152542	Calibration	Calibration
Weak	3.415254	3.442396	+0.7947%	−0.3950%
Hypercharge	5.516949	5.426726	−1.6354%	+0.8279%

Each change is relative to the candidate at the same scale. The last column converts via g = 1/√Z. The comparison values come from running the paper's physical inputs to that scale. They carry the scheme and matching assumptions of the calculation. [N15, §1, eqs. 1–2]

How the proposed corrections perform

Boundary tested	Weak residual ΔZ	Hypercharge residual ΔZ
Original wheel	+0.02714224	−0.09022286
Relaxed participation (D41)	+0.02714224	−0.01649434
Clamped participation (D41)	+0.09317979	+0.54666087
Finite fermion conversion alone	+0.05059707	+0.02137394

ΔZ is the empirical running coefficient minus the candidate after colour recalibration for that row. Each row is a separate test. Adding rows would reuse incompatible boundary assumptions. [N15, §7.4, table 3]

In the follow-up joint Gaussian action, four source placements remove 0%, 39.06%, 43.21% and 81.72% of the hypercharge discrepancy. Each leaves the weak residual at +0.02714224. Clamping changes the weak response in the wrong direction. A physical source and its record condition must determine the placement before comparison. [N15, §7; A1, table 1 and §7]

The compact wheel also softens both sectors at leading order after colour recalibration. Preserving an alternating record can increase stiffness through quantum fluctuations even when its direct source overlap is zero. The total tested weak correction still has the wrong sign. Progress requires an interaction or record condition whose full weak response is positive in the inverse-squared convention. [N15, §§4–8; A1, §§5–7]

Particle masses already calculated

By October, VERSF had calculated quantities in every major mass sector. Two examples show the precision reached on the stated assumptions: the muon-to-electron mass ratio is 0.17% high, and the Higgs-to-vacuum ratio is about 0.05% low against the comparisons used in the papers. An absolute mass also depends on how the calculation sets its energy scale.

Sector	Result already on the record	Physical interpretation
Charged leptons	The two mass steps are exp(16/3) = 207.127 and exp(16/3)/12 = 17.2606, respectively 0.174% and 2.63% high. With α and v = 246.22 GeV as inputs, the electron, muon and tau are 0.521691, 108.056 and 1865.12 MeV. [a]	The three absolute benchmarks are 2.09%, 2.27% and 4.96% high. The paper separates the common-scale discrepancy from the two ratio discrepancies. Physical normalisation and radiative matching remain.
Higgs	The conditional ratio mH/v = 32/63 is about 0.0528% below its comparator of 0.508204766. Route A gives v = 243.723 GeV and a leading Higgs mass of 123.796 GeV. Its scalar-only one-loop pole shift is +8.889 MeV. [b, c]	The ratio and absolute mass test different things. The dimensional estimate uses the Route A scale. The physical Higgs kinetic normalisation and remaining gauge, Yukawa and threshold contributions enter the full pole.
W, Z and photon	Electroweak breaking gives a massless photon, massive W and Z directions, and the leading relation mW = mZ cos θW. [d]	The mass pattern follows on the stated vacuum. Absolute W and Z values also use the weak couplings and vacuum scale, followed by pole matching.
Quarks	The conditional charged-flavour calculation gives three up-type Yukawa values (0.00000722, 0.00371, 0.971983) and three down-type values (0.0000156, 0.000313, 0.016196). A separate projection gives a first down-sector candidate of 4.695 MeV. [e, f]	The six values specify a mass hierarchy on the chosen action. The 4.695 MeV example has its own channel-addition and state-identification conditions. A physical quark-mass comparison needs a specified scale and scheme.
Neutrinos	The projected neutral construction gives relative masses 1 : 1.046825 : 2.096684 and a mass-splitting ratio of 0.0282214. One observed splitting sets the indicative masses to 27.2, 28.5 and 57.0 meV. [B, §6]	The splitting ratio is 5.3% low against the paper's NuFIT 6.0 comparison. One splitting supplies the absolute scale. A common physical mass, kinetic and weak-current construction remains. [W7]

Page 11 gives the one-anchor six-quark results, an earlier baryon benchmark and the neutrino comparison under two uncertainty estimates. November's fold-spectrum paper addresses excitation energies and probe visibility. [N12, §§8–11]

Primary papers [a] Charged-lepton scale, §19 · [b] Numerical confrontation, §4 · [c] Higgs potential and pole · [d] Broken electroweak phase, §9 · [e] Charged flavour, p. 52 · [f] Quark current-mass projection

Quantum structure and particle spectra

Fermion handedness survives exact coarse graining

When the eliminated fermion block is invertible, the retained operator preserves the chiral relation, zero modes and index. A uniform free bound establishes invertibility for the specified two-site blocking, and the free branch retains one Weyl pole per supplied chiral representation. The determinant's fermion-only sector proportions are exactly 6 : 6 : 10. The global chiral measure and fully interacting continuum spectrum remain outstanding. [N02, §§3–9, table 11]

The probability measure has a conditional uniqueness theorem

On the stated complex inner-product carrier, nonnegative nonzero finite weights that respect global phase and add across orthogonal branches are proportional to squared norm. The theorem needs no continuity assumption and includes two-dimensional systems. A second route reconstructs the weight from reference-record filtering and grouping. Implementing the admitted record operations in VERSF is the physical premise. This closes a quantum-probability subproblem shared by several entries. [N10, §§2–4, 10]

A lowest internal energy can now be tested

A fold is the minimally distinguishable thing that can exist in reality. Its definition assigns no shape. A physical realisation can acquire an excitation threshold through comparison between neighbours, restoration towards a reference, or a twist around a loop. The seven-site example reproduces a relaxation gap of ½ per step using its native weights. A second calculation identifies the existing radial benchmark as a reference threshold. Assigning a particle mass still requires the physical inertia, clock, observable coupling and mass scale. [N12, §§3–9, 11]

Two different limits for the mass gap

In the Abelian Gaussian theory, restoring the relative record phase produces a gap that persists at large volume and fine resolution when physical coefficients are held fixed. The supplement also proves a volume-uniform gap on specified strongly coupled non-Abelian lattices at fixed spacing. The latter theorem includes SU(3). Joint physical dynamics, its vacuum, physical refinement and pure Yang–Mills identification remain open. [N11, §§4, 7–11; W1, theorem 4.2, §14]

Event histories sharpen the physical readout

The directed seven-address clock gives 0.883 bits of adjacent-address information and 18.798 bits of stopped-path distinguishability under the stated calibration. The recorded crossing sequence retains ten of twelve rate directions in the tested family. A separate aggregation model distinguishes pair binding from many-body stability and shows how locked records preserve cluster shapes. These are useful constraints on record models and binding. They supply no numerical QCD string tension. [N05, §§5–9; N06, §10; N08, §§3–6]

Standard Model entries 1 to 8

The strongest direct movement in this group concerns recorded information. The new chiral and Gaussian results also test assumptions that several entries share.

Entry	19 October	17 November	Still required
SM-1 Weak orientation admissibility	Orientation conditionally selected.	The chiral calculation preserves the handedness supplied to the free construction. It provides a downstream check. [N02]	Install the selected orientation in one physical gauge, completion and matter action.
SM-2 Full weak doublet Hessian	Full specified mixed tensor descends.	The October result carries forward. The new gauge stiffness calculations evaluate a different response.	Derive the family tensor coefficient, physical normalisation and readout on the same action.
SM-3 Leakage support trace	Neutral candidate and active-state diagnostics available.	The conditional neutral construction carries forward. The batch supplies no new physical weak-current leakage trace.	Compute the support and leakage with the physical neutral kinetic and charged-current matrices.
SM-4 Ordered chi readout	Exact conditional record ranks.	The ordered crossing record retains ten of twelve rate directions. The source clock now has exact completed-path information diagnostics. [N06, N08]	Identify the physical preparation, event grain and record that survives the readout.
SM-5 Gauge census closure	Conditional origin and completeness on the stated minimal census.	The quotient group is fixed on the stated census. The one-loop running uses this census and checks the three traces. [W14; N02–N03]	For a reduced-input origin, establish the remaining carrier and minimality identifications.
SM-6 Weak octant sign	Conditional mixing diagnostics.	The October diagnostics carry forward. No new common neutral construction chooses the physical octant.	Return neutral masses, kinetic form, charged current, octant and phase on one branch.
SM-7 Adjoint removal positivity and microcausality	Conditional quotient; free and static tests.	The Abelian family has reflection positivity. A separate finite joint positive transfer has a unique lowest state and a certified gap. [N11; W1, §10.7]	Establish dynamical non-Abelian positivity and the physical locality and causal limit.
SM-8 Kappa and localisation convention	Local clock and convention fixed on stated settings.	Route M gives a length and traversal time on a declared length input. New event and response models test their clock and units. [W10; N06, N08, N13]	Tie physical time and length, including the common field normalisation, to one branch.

October judgements and outstanding conditions trace to B, §4, SM-1–SM-8. N keys refer to the uploaded manuscripts and W keys to the public sources on page 13.

Standard Model entries 9 to 16

Matching conventions and event occupancy now have concrete tests. The Higgs radial construction gains a dynamical interpretation. The generation count retains its inherited rank premise.

Entry	19 October	17 November	Still required
SM-9 Convention and normalisation audit	Exact interfaces; physical matching open.	The branch clock gives a duration. The action unit, field baseline and finite conversions still require a physical identification. [W10; N08, N13–N15]	Match the native observable to one physical scheme, with the clock, thresholds and full finite terms.
SM-10 Substrate anomaly descent	Exact global charge map; full chiral measure open.	Zero-mode and index preservation under exact blocking strengthens regulator control. Local anomaly traces hold on the supplied census. [N02]	Complete the microscopic chiral measure, determinant-line construction and global anomaly test.
SM-11 Generation and species census	Conditional three-generation census; representation map closed.	The three-object quotient carries S_3. The census uses the rank-three premise C9 and the family-to-terminal identification. [W13; N02–N03]	Derive or retain the rank-three premise and supply the separate physical family-to-terminal map.
SM-12 World to framework occupancy	Exact conditional support results.	The event law selects a rate-biased configuration distribution within an explicit model. The physical background remains conditional. [N07–N08]	Identify the occupied support, background and actual recording law with the needed field derivatives.
SM-13 Hermitian lift and frame firewall	Local restrictions on phase and frame.	Chiral blocking supports a regulator. The physical charged-flavour phase and current frame retain their October status. [N02]	Return the phase and frame on the same action, or establish output equivalence for the alternatives.
SM-14 Strange confinement dynamics	Nonperturbative framework; coefficient undetermined.	Strange observables pass through colour singlets. Distinct wall costs share the same retained data. A positive model gap alone leaves the wall cost undetermined. [W1, W12, W16]	Determine the colour response, flux cost and connection to physical confinement observables.
SM-15 Physical half lazy gate	Retired as an independent route.	The new wheel's holding step has an explicit role in its conditional normalisation; the historical gate stays retired. [N14]	Track the chosen transfer and clock within the live calculation. No separate scalar claim.
SM-16 Higgs potential and radial mode	Conditional potential, radial benchmark and leading mass.	The spatial kinetic residue reduces to one scalar on the stated branch. The fold-spectrum paper identifies the radial reference threshold. [W15; N12, §9.3]	Supply the physical metric, potential and couplings, then compute the matched Higgs pole.

October judgements trace to B, §4, SM-9–SM-16. The conditional Higgs benchmarks $v \approx 243.723$ GeV and leading mass ≈ 123.796 GeV predate this audit. A physical pole requires the remaining matching calculation. [B, SM-16]

Standard Model entries 17 to 24

Several historical shortcuts in this group had already been superseded. The new gauge and spectral tools can serve the active mass and mixing calculations without reviving the retired entries.

Entry	19 October	17 November	Still required
SM-17 Charm hinge magnitude	Superseded by the charged tensor route.	The audited quark mixing matrix gives |Vtd| 31% low and |J| 41% low. The closer curvature branch still needs physical selection. [W8, §3; N12]	Obtain the charged spectrum, phase and current frame, then compute the physical CKM matrix.
SM-18 Coloured participation fraction	Retired independently; exact participation geometry retained.	Collective reductions and the clamped and relaxed readings are evaluated as part of SM-26. [N01, N07, N15]	Use representation-resolved response in the common action. Count its progress under SM-26.
SM-19 Tau suppression projection	Superseded; conditional lepton ratios retained.	The earlier tau-to-muon ratio gives 17.2606, 2.63% high on its comparison. See the charged-lepton scale in SM-21. [B, §6]	Derive activation, localisation and the absolute lepton scale on a matched branch.
SM-20 Variational suppression identity	Retired historical route.	The later elimination and normalisation identities concern active response calculations. The historical row stays retired. [N13–N15]	No independent calculation required for this row.
SM-21 Charged lepton absolute scale	Conditional scale relation evaluated.	The scale and two ratios give 0.522, 108.06 and 1865.12 MeV for the three charged leptons using stated inputs. Page 5 separates scale and ratio residuals. [B, §6; a]	Return the independent common scale and physical charged-lepton pole, including thresholds.
SM-22 Full Yukawa matrices	Specified tensor descent; charged and neutral candidates.	The charged construction already returns six quark Yukawa singular values. Exact chiral blocking strengthens its kinetic setting. [e; N02]	Derive the remaining coefficients with the required addresses, kinetic forms and currents.
SM-23 Quark current mass projection	Conditional running, thresholds and transport.	The one-anchor spectrum inherits two ratios. A later operator conditionally explains s/d = 20. A first-down projection gives 4.695 MeV before matching. [W4–W5, W19; f]	Match the native scheme, thresholds and absolute quark masses to declared observables.
SM-24 PMNS closure Hamiltonian	Executable neutral candidates; physical mixing open.	Neutral masses and mixing angles have separate conditional constructions. A common weak-current readout is still needed. [W7, W17; N10]	Derive neutral masses, kinetics and charged-current mixing together, including branch and leakage tests.

October judgements trace to B, §4, SM-17–SM-24. CKM and PMNS denote the quark and lepton charged-current mixing matrices, respectively. Matching converts the theory's regulated quantities into a specified observable or scheme.

Standard Model entries 25 to 30

These entries contain the main November gains and the largest outstanding physical obligations. The distinctions between gauge running, a mass gap and colour confinement matter here.

Entry	19 October	17 November	Still required
SM-25 Neutrino mass and charge conjugation	Conditional mass-spacing ratio and local charge selection.	The projected mass-spacing ratio is 0.0282214. The faithful-bundle index separately fixes the baryon/lepton selection rule on the declared census. [W7, W17; B, §§4, 6–7]	Establish the projection, amplitude, physical kinetic and current frame, and absolute scale.
SM-26 Gauge action and coupling values	Conditional relative relation; absolute values open.	One joint Gaussian model reproduces the 81.7% hypercharge improvement. Its relaxed weak correction vanishes for all placements. [N14–N15; W2; A1]	On one admissible three-factor action, return both derivative coefficients, sector corrections, baseline mapping and finite matching.
SM-27 QCD confinement and running	Transport advanced; confinement open.	The VBF branch gives Λ_3 = 225.923 MeV. A separate strong-coupling SU(3) model has a volume-uniform gap. One-loop running follows on the inherited census. [W1, W11–W12; N03]	Determine the physical colour boundary and a nonperturbative flux or string-breaking observable in the stated scheme.
SM-28 Strong CP and theta sector	Zero strong-CP angle in a retained positive completion.	The faithful-quotient measure gives $\bar{\theta}$ = 0 in the retained positive completion. The global continuum observable still needs a measure and matching. [W12; N02]	Complete the global continuum chiral measure and observable matching for strong CP.
SM-29 Quantum gauge completion	Strong finite and static checks; global gate open.	An 84-state history tests joint assembly. A 588-direction current block matches conditionally, but the tested bond response still disagrees. Full physical joint evolution remains open. [W1; A2–A3]	Establish the interacting dynamical theory, its global chiral measure and a common continuum observable.
SM-30 Final derivation audit	October interim audit complete.	All thirty entries are reconciled. The roadmap on pages 14–15 sets distinct tests for a working finite realisation and for completion of the programme.	Combine the gauge, Higgs, charged, neutral and strong-sector outputs in one consistent physical description.

October judgements trace to B, §4, SM-25–SM-30. Its charge-selection result and conditional neutrino ratio remain credited to October. The November source set adds no independent charge-selection or neutrino-ratio result.

What the website adds to the mass ledger

The programme has explicit numerical work in every major mass sector. The public papers make it possible to distinguish an input, a ratio prediction and an absolute mass on a specified branch.

All six quarks from one down-quark anchor

The earlier grid shows how the light split, strange baseline, bottom step and heavy-family split enter. The successor takes m_d = 4.70 MeV as an illustration. It inherits u/d = 6/13 and s/d = 20. Its bottom step and heavy split carry structural assumptions, including a proposed 2/3 factor. A successor derives s/d = 20 on an explicit confinement operator with twenty balanced, oriented supports. Physical selection of its invariant projector remains open. The table uses the one-anchor paper's conventions. [W4–W5, W19]

Quark	One-anchor value	Paper's comparator	Role or difference
Down	4.70 MeV	4.70 MeV	Input setting the scale
Up	2.17 MeV	≈2.16 MeV	Inherited ratio, +0.5%
Strange	94.0 MeV	≈93.4 MeV	Inherited here, +0.6%. See W19.
Charm	1,113.5 MeV	≈1,105 MeV	Conditional split, +0.8%
Bottom	4,867.5 MeV	≈5,076 MeV	Conditional step, −4.1%
Top	292.1 GeV	≈300.8 GeV	Dependent on earlier steps, −2.9%

A direct experimental test needs all masses evolved to the same scale and scheme. The top values here are the paper's illustrative running comparison. The later six Yukawa singular values and 4.695 MeV first-down projection on page 5 use different constructions.

A composite particle with an explicit scale

The baryon companion obtains 17 closure layers from four stated hypotheses, parity and channel conditions. With 108 structural states and a 51-macrobit electron block at 10 keV per macrobit, the calculation gives 936.36 MeV for the proton, about 0.20% low. The electron scale enters the block. This is a conditional hadron benchmark distinct from the fundamental mass entries and from a QCD string-tension calculation. [W6]

What the comparison tests

The projected neutrino ratio Rv = 0.0282214 is about 2.0 local standard deviations low against the paper's NuFIT 6.0 comparator. The October ledger also records about 3.9 low using its tighter NuFIT 6.1 uncertainty. Both use a local Gaussian approximation with zero covariance. The projection was chosen after earlier constructions missed. Quark mixing has a separate failure: the controlling matrix gives |Vtd| 31% low and the CP invariant |J| 41% low. [W7–W9; B, §6]

What the website adds beyond masses

Yang–Mills dynamics and the limits of the gap theorem

Theorem 4.2 covers SU(N), N ≥ 2, with the stated product Haar measure, metric, bounded plaquette incidence and a Wilson potential plus a specified correction. On the cubic SU(3) lattice the lower bound is D(3/4 − 8βg/3), positive for 0 < βg < 9/32 and D > 0. It is independent of the number of links at fixed lattice spacing. This is a stronger result than the Abelian Gaussian theorem credited in the first review. [W1, §4.6; N11]

Section 10.6 constructs a gauge-covariant recorded operation with one normaliser. Section 10.7 constructs a separate finite positive transfer with a unique lowest state and a gap of at least 1/14,810,250 in model units. The paper keeps four physical obligations open: a joint evolution, its interacting vacuum, control under volume and physical refinement, and identification of pure Yang–Mills observables. [W1, §§10, 14]

A three-sector response and a null test

The 8 November paper proves an exact reduction for its source-independent positive quadratic families. It calculates leading fractional softening slopes of 0.08316, 0.07547 and 0.04822 for colour, weak and hypercharge in twelve directions. Its separate scalar example softens by about 27%. The 26 October background study finds corrections of only 0.013 ± 0.009 parts per million for colour and weak, and 0.012 ± 0.009 for hypercharge on its tested attachment. Applying either calculation to the 31/59 wheel requires a common physical source and normalisation. [W2–W3; N14–N15]

Strong scale, symmetry, clock and selection rules

A predeclared saturated branch gives a three-flavour scale of 225.923 MeV in the native VBF scheme without a measured strong-sector input. It still requires conversion to a conventional scale. A separate faithful-quotient construction gives a retained strong phase of zero and proves that its retained data allow multiple wall costs. Calculating a physical string tension therefore needs additional dynamics. [W11–W12]

The faithful gauge group is $[SU(3) \times SU(2) \times U(1)]/\mathbb{Z}_6$ on the stated census. On the stated bundle and three-generation premises, the charge-selection calculation requires baryon change in multiples of three and identifies the dimension-15 ΔB = 3 and dimension-18 sphaleron minima. Rates require further dynamics. The three-generation census retains its rank and family-identification premises. [W13–W14, W17; B, §7]

Route M gives ξ = 84.0718 μm after importing a length of twice the Planck length. The oriented traversal lasts 3.3562×10^{-13} s. The paper shows why these relations alone leave the action unit and Standard Model completion scale free. [W10]

How the public record was reviewed

The public site index listed 13 pages, 556 posts and 723 PDF records. All 723 files were retrieved. Byte-for-byte deduplication left 695 distinct PDFs, including short test files. A topic search covered the searchable text for all thirty ledger entries, the posts, the subject collections and the linked legacy word-processing document. An image-only overview was checked separately. The paper count includes versions with changed content.

Primary sections were read for claims that could change a grade. Priority went to the edition that identified the model, inputs, observable and comparison. An earlier hexagonal-model Higgs estimate of 125.8 GeV, for example, uses W and Z masses in its formula. The later 32/63 ratio and 123.796 GeV Route A estimate have different inputs and are reported separately. [W18; W9; b–d]

Public primary papers

W1 Yang–Mills supplement Thm 4.2, §§10.6–10.7, 14

W2 Physical matter, records and gauge response Props 16.1–16.2

W3 Nonlinear history and background selection 30-amplitude test

W4 Conditional quark mass-ratio grid ratio ledger

W5 Quark masses from one anchor calculator and input audit

W6 Baryon closure depth §§2, 10–11

W7 Source-fixed neutrino mass ratio §11.5, App H

W8 Full-carrier quark mixing §§3, 26

W9 SM numerical confrontation §§4–9

W10 Absolute scale and closure clock Route M and no-go results

W11 Non-perturbative QCD scale minimal branch and scheme

W12 Faithful-quotient retained measure strong phase and wall cost

W13 From Fold to three generations C9 and family map

W14 Faithful gauge quotient quotient and charge lattice

W15 Higgs spatial propagation spatial and temporal residue

W16 Strange confinement dynamics colour-singlet observability

W17 Questions a derivation must answer questions 13–22

W18 Standard Model from hexagonal geometry Apps C–H, legacy benchmarks

W19 Strange confinement operator §14, conditional strange/down ratio

W keys link to the public manuscripts. B and N01–N15 identify the editions used for the website audit. A1–A3 are the additional working papers used for the roadmap.

The review covers the publicly indexed VERSF corpus and the supplied manuscripts. The wheel eigenvalue, 31/59, the gauge residuals, the relaxed hypercharge reduction and selected mass comparisons were checked arithmetically. Computational results are credited to their papers. This review did not independently rerun the simulation archives, validate every proof or reproduce the experimental likelihoods.

From about 80% to 90%

90% would mean that one physically admissible VERSF realisation works across forces, particle masses and mixing, and the strong sector at finite resolution. One successful realisation is sufficient.

1 Specify the common physical evolution

Write down the joint evolution of fields, matter, currents and completed records. Specify the clock, measure, source couplings and which records stay fixed. Calculate the event probabilities and the history response when the physical variables change. Verify that the same evolution gives the action's quadratic coefficients and passes the seven checks for a physical joint evolution (M1.a–g). A2 supplies the assembly formulas. A3 matches the 588 current directions conditionally and identifies the bond response that still fails. [SM-1, 2, 4, 8, 9, 12, 29; W1; A2–A3]

2 Predict weak and hypercharge together

Calculate electric and magnetic response from that evolution, match the current and run both couplings to a common comparison scale. Colour may set the scale at this stage. The present inverse-squared targets are +0.795% for weak and −1.635% for hypercharge relative to the original wheel. Fix the interaction before using those values. Both predictions must agree at the accuracy supported by the calculation. A1 already exhausts the relaxed Gaussian placement test. [SM-8, 9, 26; N14–N15; A1]

3 Carry the mass calculations through to mixing

Use the same vacuum, kinetic normalisations and weak current to calculate the Higgs and charged-lepton benchmarks, the quark mass hierarchy and leading quark and neutrino mixing. Two quark tests are particularly sensitive. The current calculation puts |Vtd| 31% low and the CP invariant |J| 41% low. A viable result must resolve both discrepancies within a justified error budget. In the neutral sector, jointly predict the splitting ratio and mixing angles, including the atmospheric octant, with one declared scale anchor if needed. [SM-3, 6, 13, 16–17, 19, 21–25; W7–W9]

4 Calculate a strong-interaction observable

On the interacting finite model, calculate a gauge-invariant static colour potential or a string-breaking signal. Show how it changes with volume and lattice spacing, with uncertainties. Check it against an independent benchmark in the stated units. The 225.923 MeV native scale and fixed-spacing SU(3) gap already provide starting points. [SM-7, 14, 27–29; W1, W11–W12]

90% decision. Credit the milestone when all four tests succeed together. Absolute scales set by declared calibration, higher-order pole matching and the interacting continuum limit can still be open. A force-only success advances SM-26.

From 90% to 100%

100% would mean completion of the derivation claimed in this ledger, including its neutrino sector. The endpoint is a reproducible route from declared VERSF postulates to physical observables.

1 Account for the remaining physical inputs

Derive the particle representations, chiral charges and three-family structure from the stated starting assumptions. In particular, establish the rank-three premise and identify the three physical families with their records. Justify the occupied background, record operations, clock and normalisation on the common construction. One reference quantity may set the units. Derive the remaining dimensionless scale ratios. List any additional measured mass or coupling used as a calibration and identify the prediction that still depends on it. [SM-1–5, 8–13; W10, W13–W14]

2 Deliver the complete numerical spectrum and forces

Calculate all three gauge strengths, including the strong coupling that set the comparison scale at 90%. Match the Higgs, W, Z and charged-lepton pole masses, preserve the massless photon and give all six quark masses at named scales in a specified scheme. Derive the absolute neutrino scale, the mass splittings and the full quark and neutrino mixing matrices, including the observable CP phases claimed by the model. Include threshold and radiative corrections and propagate their uncertainties. Compare measured quantities with data and unmeasured outputs with available bounds. [SM-3, 6, 9, 13, 16, 21–27; B, §§5–8]

3 Establish the interacting quantum theory and confinement

Identify the physical interacting vacuum (M2). Control the relevant gap as volume grows and the regulator is removed at fixed physical units (M3). For the ledger's pure Yang–Mills claim, identify the limiting theory and its observables (M4). Separately demonstrate physical colour confinement or string breaking with matter. Establish the global chiral measure, anomaly consistency, positivity and causal locality, and connect the retained strong-CP result to a continuum observable or bound. [SM-7, 10, 14, 27–29; N02, N11; W1, W11–W12]

4 Reproduce the full derivation

Use one set of inputs and conventions to rerun the construction from its postulates through to the reported observables. Record the source, units, uncertainty and comparison for each output. Resolve material discrepancies on independently fixed inputs. Trace all thirty entries to their evidence, including which active result supersedes each retired route. [SM-15, 17–20, 30; B, §4]

100% decision. Every claimed result must follow on at least one physically justified realisation. A measured quantity used to fix the calculation earns credit only as an input. An unexplained physical parameter, unresolved discrepancy or uncontrolled limit keeps its claim open.

Source register

B is the October baseline. N01–N15 are the fifteen manuscripts used in the November website audit. A1–A3 are follow-up calculations consulted for the 18 November roadmap.

Key	Manuscript	Evidence used
B	VERSF Standard Model Derivation Ledger October 2026	§§1, 4–8. Baseline grades, thirty entries and existing predictions.
N01	Charge Participation and Coarse Graining in VERSF	§§3–6. Derived history reductions and conditional gauge interpretation.
N02	Coarse Graining and Chiral Fermion Survival in VERSF	§§3–9, table 11. Exact blocking, finite tests and fermion response.
N03	Coarse Grained Gauge Running in VERSF	§§3–7. Complete one-loop running on the inherited census.
N04	Coarse Grained History Response and Gauge Exchange in VERSF	§§3–8, 10. Six-cycle response and connected-cell tests.
N05	Fact Momentum, Record Overlap and the Emergence of Stable Aggregates in VERSF	§§5–9, 14. Pair and many-body thresholds; record locking.
N06	Information Sharing, Record Overlap, and Aggregation in VERSF	§§8–11, 13. Information retained by completed address histories.
N07	Source History Weighting and the Coarse Grained Gauge Action in VERSF	§§12–20. Coupled source tests, field stiffness and record memory.
N08	Completed Event Weighting and Force Response in VERSF	§§3–8, 13–16. Event weighting, gauge coefficients and benchmarks.
N09	Nonlinear Record Dynamics and the Force Strength Gap in VERSF	§§9–18, apps. A–E. Record relaxation, current stiffness and noise.
N10	Further VERSF Deliberations on the Born Rule	§§2–4, 10. Conditional uniqueness and operational readout. Version 7.
N11	Commitment Constraints and the Yang–Mills Mass Gap in VERSF	§§4, 7–11, conclusion. Conditional Gaussian gap and refinement.
N12	Bounded Fold Dynamics and the Particle Spectrum in VERSF	§§3–11, props. 12–13. Thresholds and two native tests.
N13	From Primitive Dynamics to Physical Force Strengths in VERSF	§§5–10, 13. Twelve-source reference and separate Abelian dynamics.
N14	Stationary Joint Response and the Force Normalisation Test in VERSF	§10, props. 5–7. Conditional common factor and baseline mapping.
N15	Nonlinear Corrections and the Force Strength Residual in VERSF	§§1, 4–9, tables 1–3. Calibrated residuals and correction tests.
A1	Joint Response and Record Constraints in VERSF	§§3–7, table 1. Joint Gaussian placements and compact quantum weak-sign test.
A2	Conditional Joint History Assembly for M1 in VERSF	§§2–6. Conditional assembly from one marked transfer and an 84-state test.
A3	Current magnitude and the joint history law: an M1 closure attempt	§§7–12. Current-block match on 588 directions and the surviving bond mismatch.