S5 Supplementary analyses

Geometric reconstruction and temporal inference under explicit assumptions

27 September 2026

The fixed numerical model ranks Cygnus first in the declared bright-star panel and returns an epoch within Stonehenge's initial-enclosure interval without fitting archaeological dates. The metric extension preserves that result in five declared variants. Independent spatial and temporal scales leave Cygnus nondominated, alongside other frontier members. Phase comparisons quantify the preference for the early branch. These executed retrospective analyses specify the reconstruction's recoverable relationships and expose the conventions still requiring archaeological support; free words and timing scale admit alternatives.

Sections 1–8 retain the crossed selection, compiler, phase and provenance analyses. Section 9 tests losses and timing scale; section 10 separates temporal ranking from spatial requirements; section 11 reports joint-scale frontiers; section 12 tests continuous metric words. Section 13 gives the execution and file map. All references to the main bibliography use the unchanged numbering of the manuscript.

1. Protocol, inputs and selection order

PROTOCOL.md was recorded before the new executions, after the primary findings and reviewer criticisms were known. No prospective discovery claim follows from that ordering. The catalogue is the archived S4 table of 5,041 Hipparcos entries with V≤6. For each of the 88 modern constellation regions, select the five lowest V magnitudes, with HIP as the tie-breaker. There is no prior filter on shape, spatial tuple, archaeological date or visibility. The modern regional partition is an explicit analytical convenience, not an inferred prehistoric star classification. The original protocol is preserved. PROTOCOL_v126_ADDENDUM.md records that the original-compiler controls were formulated after these unrestricted results were known. PROTOCOL_v128_ADDENDUM.md records that the bounded-epoch comparisons were formulated after the fixed-word results and chronology objections were known.

The temporal graph uses five distinct role-ordered stars and three unsigned right-ascension widths: roles 1–2, 2–3 and 4–5. Widths are multiplied by four to express sidereal minutes. Positions at catalogue epoch J1991.25 are propagated using the archived proper motions and long-term precession engine. Primary and six-dimensional astrometry are kept as separate nominated-input sensitivity branches. They are not mixed with the common-catalogue panel when ranking competitors.

Each regional set supplies 120 labelled role assignments. Exchanging roles 4 and 5 leaves their unsigned interval unchanged, so 60 computations per region represent 120 assignments. The 512 ordered words from V={4,16,28,32,40,68,72,100} yield 5,406,720 labelled role-word descriptions, or 2,703,360 classes under that one symmetry. These are dependent descriptions, not independent statistical trials. First/third-role reversal with exchange of the first two word channels remains separately labelled.

2. Objective, search and numerical verification

For each word W and role graph, the absolute objective is $RMS(t,W)=\{\Sigma_i[T_i(t)-W_i]^2/3\}^{1/2}$. Search astronomical years −8000 through +2000, without chronological fitting constraints. Each sampled local minimum is refined with bounded scalar minimization; endpoints and appropriate boundary cells are also evaluated. The primary new run uses a 1-year grid and the verification run a 0.5-year grid. The 3000–2920 BCE initial-enclosure window is inspected only after full-window fitting. Section 3.3 separately imposes explicit chronological bounds and does not replace this original unconstrained search.

Because the library is a Cartesian product, its nearest word at any epoch is found by independently selecting the nearest allowed value in each channel. For this qualifying threshold, no passing word is lost by searching that lower envelope: an RMS≤0.390126 minute bounds each absolute residual by about 0.676 minute, below half the minimum library-value spacing of 4 minutes. A passing word must therefore be

componentwise nearest at its passing epoch. For every qualifying graph-word pair the best discovered qualifying minimum is retained. This shortcut is verified against direct evaluation of all 512 words for 1,000 deterministic random width vectors; scores agree exactly.

The two temporal grids return the same 336 qualifying graph-word classes and the same 39 qualifying ratio classes. Across all 5,280 role graphs, the maximum difference in best absolute score is 4.09×10^{-11} minute and in relative-dispersion score 9.99×10^{-10}. A unit-conversion check agrees within 1.12×10^{-15}. Agreement across two grids is numerical evidence, not an analytic proof against every unbracketed minimum. CSVs contain all graph summaries and all qualifying classes; they do not claim to archive every point of every score curve. Scripts and both execution logs are included.

3. Crossed stars, roles and words: results

The uniform-catalogue nominal benchmark has RMS 0.390126064 minute at 2965.425 BCE, with temporal vector (27.789129,72.605469,16.213391) minutes. The primary archived specification remains approximately 0.418975 minute at 2965.2185 BCE. These input conventions produce slightly different scores and are kept separate.

The new panel contains 336 qualifying classes, equivalent to 672 labelled descriptions, in 57 regions. Six classes have their best qualifying word-specific minima within 3000–2920 BCE. They comprise four Cygnus descriptions and two Lacerta descriptions related by first/third reversal. The six entries are not six independent astronomical systems. All appear in the table below; HIP identifiers give the complete role order. Final-pair order has been canonicalized by increasing HIP.

Region and HIP roles	Word	BCE	RMS min
Cyg 102098 / 100453 / 95947 / 97165 / 102488	28,72,16	2965.425	0.390126
Cyg 100453 / 102098 / 95947 / 97165 / 102488	28,100,16	2967.187	0.285972
Cyg 95947 / 102098 / 100453 / 97165 / 102488	100,28,16	2967.187	0.285972
Cyg 95947 / 100453 / 102098 / 97165 / 102488	72,28,16	2965.425	0.390126
Lac 111169 / 109754 / 110538 / 109937 / 111022	32,32,28	2921.499	0.137723
Lac 110538 / 109754 / 111169 / 109937 / 111022	32,32,28	2921.499	0.137723

The closest fit anywhere in this panel is a Scutum graph with HIP roles (92175,91117,91726,90135,90595), W=(16,4,4), RMS 0.018962 minute, near 3237.2 BCE. Reversing the first three roles and the first two word channels gives an equivalent score. This example lies outside the initial-enclosure window. The Lacerta example with W=(32,32,28), RMS 0.137723 minute and epoch near 2921.5 BCE shows directly that a better timing fit plus inclusion in that window is not exclusive to Cygnus within this panel. Its second word channel, 32, cannot equal 3q for integer q. The corresponding Scutum channel, 4, fails for the same reason. These words therefore do not arise from any integer inventory under the original compiler.

This does not show that Lacerta or Scutum reproduces the full material model. The panel releases words without requiring each word to arise from a feasible alternative count compiler, and does not require the

spatial tuple or a horizon test. Its question is the specificity of timing and chronology under joint star-word freedom. It cannot replace an equal-freedom comparison of complete archaeological pipelines. The earlier S4 censuses cover more star subsets within regions but retain spatial filters; their conditional survivors and these new alternatives are compatible results for different questions. Section 3.1 explicitly quantifies that compiler distinction using constraints already present in S1, rather than treating every free-word pass as a complete competing model.

3.1 Original word and positive-integer compiler controls

For an arbitrary positive integer word (d,m,r), the unchanged map gives q=m/3, h=(d−r)/2, B=q+h, D=t, C=t+h and A=d+q+t. A positive integer solution exists exactly when q and h are integers and B>0; every integer t>max(0,−h) then supplies one. This proof defines the test implemented in check_original_constraints.py. It is feasibility under the original map, not satisfaction of its three extra closure equations, and not archaeological feasibility.

Only 55 of the 512 words satisfy these conditions, giving 55×5,280=290,400 classes in the same regional role panel. Of the 336 unrestricted timing passes, 26 are compiler-feasible. The existing complete qualifying lists suffice for this classification because feasibility removes words without changing their timing scores. The arbitrary example inventories in compiler_classification.csv demonstrate existence; they are not proposed inventories of Stonehenge.

The fixed-word control directly minimizes the absolute-RMS objective for W=(28,72,16) on all 5,280 role classes. It reuses the archived astrometry and precession engine but implements that objective independently of audit.py's Cartesian lower envelope. The 1-year grid, refined local minima and endpoint cells rank nominal Cygnus first, with RMS 0.390126064 minute at 2965.425143 BCE. Reticulum is second, with HIP roles (19780,17440,19921,18597,18744), RMS 0.554285025 minute at 2848.633911 BCE. Reticulum has 1.42 times the nominal RMS. The “pass” threshold used in the retained files is Cygnus’s own score; one pass denotes sole first place, not an independently calibrated absolute cutoff. No chronological, spatial-tuple or visibility filter is imposed. This direct control was run on a 1-year grid; the two-grid agreement reported in section 2 concerns the original crossed and ratio searches.

3.2 Sensitivity to the archaeological comparison interval

Among the 26 compiler-feasible passes, only nominal Cygnus has its best qualifying minimum in Darvill et al.'s 3000–2920 BCE initial-enclosure window. Greaney et al.'s revised 2980–2720 cal BCE 95% ditch interval retains three classes: Cygnus and two Vela assignments. The 2940–2800 cal BCE 68% interval retains the two Vela assignments and excludes the unconstrained Cygnus optimum. These are comparisons after full-window fitting, not reruns with a date prior (main references 2 and 18).

Vela's roles are (44816,39953,45941,42913,52727), W=(32,72,68), RMS 0.364695145 minute at 2905.808369 BCE, and (44816,45941,39953,42913,52727), W=(40,72,68), RMS 0.276188419 minute at 2884.338143 BCE.

Both fail the full three-closure model because r=68 cannot equal D^2 for integer D. Requiring all three original closures uniquely restores (56,30,10,4) and hence the original word. This mathematical restriction is distinct from evidence that prehistoric builders used those closures.

3.3 Fixed-model phase contrasts and interval-restricted ranking

These additional diagnostics preserve the word, five roles, minute-per-integer scale, equal channel weights and propagation engine. They minimize absolute RMS separately within 3000–2920, 2980–2720, 2940–2800 and 2620–2480 BCE, using one-year samples, local refinement and both boundaries. The original full search, astronomical years −8000 to +2000, is also repeated for the nominated configuration under primary, six-dimensional and common-catalogue astrometry. None of these new interval-constrained fits is described as chronology-free.

The global optimum belongs to both the initial-enclosure and revised 95% ditch intervals, leaving the minimum unchanged there. The three input variants give the following phase contrast. All Stage 2 minima occur at its oldest boundary, 2620 BCE. Ratios compare deterministic errors within one model; they are not likelihood ratios or independent archaeological confirmations.

Astrometry	Global BCE	Global RMS min	Stage 2 RMS min	Ratio
Primary	2965.2185	0.418975	2.081039	4.967
Six dimensional	2965.8550	0.420359	2.084465	4.959
Common catalogue	2965.4251	0.390126	2.079616	5.331

Restricting the epoch to Greaney et al.'s 68% interval, 2940–2800 BCE, gives a minimum at 2940 BCE for each nominated input set. Primary RMS is 0.444750 minute (+6.15%); six-dimensional RMS is 0.447313 (+6.41%); common-catalogue RMS is 0.418210 (+7.20%). These are best fits at an interval boundary, not replacements for the nominal astronomical epoch.

The fixed-word common-catalogue ranking follows from a simple bound. For every competitor j and restricted interval I, min over I of RMS_j is at least its unrestricted minimum. All 5,279 non-nominal classes have archived unrestricted RMS ≥0.554285025 minute, whereas Cygnus attains 0.418209685 inside I. It therefore remains first. Reticulum's unrestricted optimum is at 2848.633911 BCE, inside I, so its restricted score equals 0.554285025 and it remains second, with 1.325376 times the Cygnus RMS. This numerical ranking proof uses the archived all-class global scores and the new bounded Cygnus curve; it does not claim a new full 5,280-class census.

The count of unrestricted optima inside a window answers a different question. Of the 5,280 fixed-word classes, 22 have global optima in 3000–2920 BCE, 65 in 2980–2720, 39 in 2940–2800 and 20 in 2620–2480. These are dependent panel frequencies, not historical false-positive probabilities. Neither these fractions nor a ratio of chronological window lengths calibrates the entire retrospective selection procedure.

Figure 11 in the main manuscript shows the phase contrast and constrained ranking. temporal/audit_temporal.py reproduces results.json and temporal_profiles.csv from the retained inputs; temporal/verify_temporal.py checks the endpoint scores, the archived ranking bound, phase ratios and exact input identity. The new comparison does not remove the free-word Vela cases, scale-free alternatives or lack of independent evidence for prehistoric use of the compiler.

3.4 Phase preference under arbitrary nonnegative channel weights

This retrospective test asks whether the phase comparison relies on equal weights. The word, roles, timing scale and propagation remain fixed. For each input branch, let t_0 be its previously determined unweighted early optimum and $e_i(t)=|T_i(t)-W_i|$. Each residual is minimized separately over I=[2480,2620] BCE on 1-year and 0.25-year grids, refining bracketed local minima and checking both endpoints. Both grids agree: sampled errors increase towards 2480 BCE in all channels, with minima at 2620 BCE.

For the common catalogue, t_0=2965.425144 BCE. The following table reports absolute errors in sidereal minutes.

Channel and target	Error at t_0	Stage 2 minimum	Ratio
Deneb–Sadr / 28	0.210871	0.742451	3.520874
Sadr–Albireo / 72	0.605469	0.728061	1.202475

Channel and target	Error at t_o	Stage 2 minimum	Ratio
Aljanah–Fawaris / 16	0.213392	3.448638	16.161059

Let m_i be the interval minimum of $e_i(t)$. The numerical bounds give $m_i > e_i(t_0) > 0$ in every channel. For any $w_i \geq 0$ with $\Sigma w_i > 0$, define the weighted RMS by

$$Rw(t) = [\Sigma_i\, w_i\, e_i(t)^2 \,/\, \Sigma_i\, w_i]^{1/2}$$

For $c = \min_i[m_i/e_i(t_0)] > 1$, the bound $Rw(t) \geq c\, Rw(t_0) > Rw(t_0)$ holds throughout I. This algebraic implication includes zero weights that omit channels; it does not rely on a selected list of weighting scenarios.

The lower factors c are 1.043916 for primary inputs, 1.039665 for six-dimensional inputs and 1.202475 for the common catalogue, each relative to its own early reference optimum. The common-catalogue factor 5.33 reported in section 3.3 is the equal-weight contrast, not a uniform factor for every weighting. The smallest margin occurs in the second channel of the six-dimensional branch: 0.681339 versus 0.655345 minute, approximately 4%. No additional astrometric uncertainty distribution is evaluated here.

The algebra is exact conditional on the interval bounds; their application is numerical, checked on two grids with local refinement, rather than an analytic certificate for every continuous epoch of the propagation engine. This result prevents Stage 2 from outperforming t_0 within the fixed model under the stated weights. It does not fix a reweighted global optimum at 2965 BCE or within 3000–2920 BCE, and does not preserve the all-constellation ranking under changed weights. Free-word and free-scale alternatives remain. The three input branches are sensitivities of one model, not independent historical confirmations. Executed code and all branch values are in channel_dominance/test_dominance.py and channel_dominance/results.json.

4. Free temporal scale

Define $z_i(t) = T_i(t)/W_i$ for W=(28,72,16), let μ=mean(z), and set $E(t) = \{\text{mean}[(z_i/\mu - 1)^2]\}^{1/2}$. Multiplying all T by a constant unit conversion leaves E unchanged. Report μ as the common timing scale in minutes per numerical unit. This is the coefficient of variation of the three channel ratios; it changes weighting relative to absolute RMS and does not transfer the original ranks. The numerical scale changes with units, while E and the selected epoch do not.

Nominated input	Minimum BCE	E percent	Scale min/unit
Primary archival astrometry	2986.805	0.681862	1.001865
Six-dimensional variant	2987.516	0.681667	1.001907
Uniform catalogue	2985.482	0.649026	1.001529

The primary ratio curve also has a worse local minimum near 6547.1 BCE, with E about 1.588%; the complete primary and 6D minimum lists, including endpoints, are retained in the summaries. The free-scale primary result is therefore near the earlier branch without being the same objective or the same numerical epoch.

In the common-catalogue panel, 39 classes in 25 regions match or improve Cygnus, equivalent to 78 labelled role assignments under final-pair exchange. The best is a Draco graph with HIP roles (83895,94376,87833,80331,85670), E=0.0205813%, scale 2.032299 minutes per unit, near 3052.0 BCE. Its epoch falls outside 3000–2920 BCE. No archaeological filter is imposed on the ratio ranking. Alternative regions also score well; all qualifying graphs are supplied.

The comparison tests a temporal scale convention rather than just acknowledging it. A near-2987 BCE Cygnus branch survives; exclusivity does not. This is not historical justification for reading architectural numbers as angles. It also does not remove the fixed-J2000 absolute spatial tuple or the auxiliary solar degree convention. Ratios cannot identify the absolute size of W: replacing W by any positive multiple leaves E unchanged.

5. Inventory identification and construction phases

The temporal map is $W=M(A,B,C,D)^{T}$, with rows of M equal to (1,−1,1,−2), (0,3,−3,3) and (1,−1,−1,0). M has rank three and kernel span{(1,0,1,1)}. Consequently W=(28,72,16) admits exactly the affine family (56+t,30,10+t,4+t); positive integer inventories require integer $t\geq -3$. This follows by subtracting any solution from the nominated one and solving the homogeneous system. The script checks the rank and kernel and samples t from −3 to 1,000 as an arithmetic check, not as the proof of the infinite family.

For fixed W, r=16 and the additional closure $r=D^{2}$ select positive D=4, hence t=0. The original theorem under all three specified closures is preserved. The new result isolates what timing alone cannot infer. It does not argue that the full conditional compiler is underdetermined. Nor does mathematical uniqueness establish that its assumptions were available or used during the earliest archaeological phase.

With k=(10,24,26) and exterior x=(33.435,80.095,86.44) m, $u=(k\cdot x)/(k\cdot k)=3.331412721893491$ m. Applied without refitting to inner k=(4,10,12,16) and x=(13.7,33,39,54) m, the signed percentage residuals are −2.732475, +0.951901, +2.505007 and −1.291475. RMS is 2.019623% overall, 2.046037% for trilithon/sarsen and 1.992858% for Z/Y. These subsets share one scale and selected model roles; they are not independent statistical trials. This calculation uses the original outer-only scale, distinct from the all-channel descriptive fit also printed in S1 Note 2.

The earliest Aubrey circuit is compared with the initial enclosure chronology. Darvill et al. [main reference 2] (Table 3) place the major sarsen settings and Stations in Stage 2, 2620–2480 cal BCE, and Y/Z excavation in 1630–1520 cal BCE. Station dating remains insecure; assignment in that model is not an independent determination for every setting. English Heritage's broader context places central stones around 2500 BCE and late holes at 1800–1500 BCE (main reference 35). Thus the complete inputs cannot securely be assigned to the nominal early epoch. Shared approximate diameters alone do not demonstrate inheritance.

6. Context figure and measurement provenance

Fig 1A reproduces an unaltered historical aerial photograph by Philip Henry Sharpe (Royal Engineers, 1906), public domain through expired Crown copyright according to its Wikimedia Commons record. It illustrates the condition in 1906 and supplies no measurements. The full source URL and acquisition date appear in README.md. Panels B and C plot existing archaeological drawing proxies; D identifies the primary measurement boundaries.

Panel B uses 56 manually read Aubrey symbols and four Station proxies from Cleal Plan 1. The 20 m bar spans 357 pixels, or 0.0560224 m/pixel: sampling, not archaeological uncertainty. Stations 91/93 use projected footprint centres, 92/94 socket-outline proxies, with partial excavation at 94. Panel C separately uses Rees click-plan label or hatch-centroid proxies, transformed using the archived registration constants. It labels entrance 1–30 and complete inner opening pairs 51–52, 53–54 and 57–58; separate fragment labels are retained. Dotted normals describe these proxies, not securely reconstructed original mounting axes. The panels are not mutually registered or independent surveys. Neither supplies full outlines, every missing centre, or measured per-feature uncertainty. The primary dimensional fit continues to use published measurements, not distances newly extracted from this figure.

7. Observations that can contradict the reconstruction

Claim to test	Adverse observation	Present status
Geometric transfer across phases	Withheld later coordinates disagree beyond justified survey and construction tolerances under fixed registration; alternatives on identical holdouts perform better.	A homogeneous independent coordinate and uncertainty set is still required.
Opening or cord realization	An original socket or mounting axis is incompatible with the nominated position or direction.	Original-setting evidence is needed; drawing proxies are not decisive.
Early availability of all material roles	Stratigraphy establishes that required physical roles arose later.	Already constrains contemporaneity; inheritance requires distinct evidence.
Exclusive timing and early-window selection with arbitrary words	Another admitted star–word description equals or improves the benchmark within the window.	Fails in the unrestricted panel (Lacerta). Those words fail compiler integrality; fixed W and scale retain Cygnus in first place.
Independent cultural identification of the stellar convention	A documented alternative reference or material-role convention defeats the claimed correspondence under rules fixed beforehand.	The absolute-angle convention lacks independent cultural calibration.
Exclusive fixed-word selection in the declared panel	Another assignment matches the benchmark with fixed W and scale and identical astrometry.	Not observed among 5,280 classes; other subsets and compilers remain outside the control.

Absence of a portable astronomical instrument, incomplete excavation or a missing socket is not automatically a falsifier. The relevant distinction is between an observed contradiction of a specified prediction and insufficient evidence to score it. A successful geometric transfer would support continuity of geometric information; it would not independently identify Cygnus or establish timekeeping intention.

8. Reproducibility and relationship to the original archive

The README specifies commands and file meanings. The v1.25 archive contains both temporal grids and verification. The content audit reran the full S5 1-year panel and verify.py, compared against the archived 0.5-year outputs, reran the original three-closure audit, and executed the new direct fixed-word search. The complete original S2–S4 runners were not rerun. All 5,280 fixed-word results and 336 feasibility classifications are supplied, together with three-window summaries. S1 begins with a revision guide and reference crosswalk and carries notices at the relevant original notes. S2–S4 retain their original analysis records. No construction-date confidence interval, global false-positive probability or probability of intention is produced. The v1.28 addition executes the bounded-epoch diagnostics and their verifier, and adds Figure 11. make_figure_v128.py reconstructs that figure from the numerical profile. The current manifest separately records this execution and the untouched older comparison outputs. The v1.29 diagnostic separately executes the channelwise searches on two grids and derives the implication for arbitrary nonnegative weights with positive total. PROTOCOL_v129_ADDENDUM.md records its retrospective order.

8.1 Public antecedent and limits of the provenance claim

The public record https://zenodo.org/records/22103786 was created and updated on 26 August 2026. Its Giza-titled ZIP contains the subtree S3/repro/CALANAIS/PRECESSIONAL_CLOCK/. The script

calanais_chronology_audit.py fixes CLOCK_TARGET_MIN=(28,72,16), the Deneb–Sadr, Sadr–Albireo and Aljanah–Fawaris channels, and the four-minutes-per-degree conversion. observed_results.json records the clock optimum as astronomical year −2964.2184951979593, labelled 2965.2 BCE, with RMS 0.4189747781503438 minute. Its executed vector agrees with the primary Stonehenge result at the reported precision. The downloaded archive matches the MD5 checksum returned by the public record; provenance/provenance_ledger.json supplies the exact paths and SHA-256 hashes of the extracts.

The earlier routine obtains the nominated clock optimum within astronomical years −3400 to −2500; the present Stonehenge analysis tests the full −8000 to +2000 window. The earlier archive therefore documents the prior word, operator and nominated epoch, not prior execution of every Stonehenge control. The public Calanais account of 10 September 2026, https://zenodo.org/records/22691368, also states the word and approximately 2965 BCE result. These records establish documentary anteriority of the ingredients relative to their stated Stonehenge application; they do not establish ignorance of the already published archaeological chronology, a blind discovery protocol, or statistical independence between applications. Any material system yielding the same word under this operator receives the same clock epoch.

9. Timing scale and loss function

This retrospective sensitivity keeps the Cygnus word, role order and input branches fixed. Besides absolute RMS, mean absolute error and largest absolute error, it evaluates the existing scale-free E, RMS with a profiled common scale $k=(T\cdot W)/(W\cdot W)$, and the minimum largest fractional error. For positive $z_i=T_i/W_i$, the last criterion is $[\max(z)-\min(z)]/[\max(z)+\min(z)]$, attained at $k=[\max(z)+\min(z)]/2$. The descriptive scale mean(z) used by E is not the least-squares scale.

Each of six losses is searched globally and in the declared intervals under primary, six-dimensional and common-catalogue inputs. One-year and half-year grids refine bracketed local minima, include endpoints and, for nonsmooth losses, relevant cusp roots. Their scores agree within the recorded tolerances. Analytic scale formulas are checked against independent scalar optimization. These are numerical searches, not analytic global certificates.

Common-catalogue loss	Global BCE	Global error	Stage 2 error	Ratio
Relative dispersion E	2985.482	0.649026%	10.422720%	16.059
Profiled RMS, min	2976.475	0.240015	2.075426	8.647
Minimax fractional error	2997.929	0.793826%	11.057715%	13.930

All Stage 2 minima occur at 2620 BCE. For E, the later-to-global error ratios are 15.312 and 15.339 in the primary and six-dimensional branches and 16.059 in the common catalogue. Even the optimal later-phase minimax fit leaves one channel discrepant by 11.06% in the common catalogue, whereas the early optimum permits all three within 0.794%. These are ratios of deterministic errors, not historical likelihood ratios.

Across six losses and three input branches, global epochs span 2933.467–3000.628 BCE. The six-dimensional minimax optimum extends about 0.63 year beyond the initial-enclosure boundary. The free-E optima precede the 2980 BCE boundary of the revised 95% interval. The epoch range is procedural sensitivity, not chronological uncertainty. The scale estimates of E remain within 0.2% of 1 in all three branches; that agreement does not establish an ancient time unit.

The componentwise fixed-scale inequalities in section 3.4 also imply the early preference for weighted positive-power sums of channel errors and the weighted maximum. A merely nondecreasing saturating loss can tie. Those implications keep scale fixed. They cannot be combined with the finite free-scale checks into a claim covering arbitrary weights and scales simultaneously: with one positively weighted channel and freely fitted scale, that channel can be fitted exactly at any epoch.

10. Interval-constrained competitors and spatial requirements

Filtering the global best epochs and refitting every candidate inside a window answer different questions. Complete constrained free-E censuses evaluate all 5,280 classes on annual and half-year grids, with local refinement and endpoints. An earlier bound-based re-evaluation reaches the same ranking; both records are included. No class is excluded because its best global epoch lies outside the window.

BCE window	Best region	Best E %	Cygnus E %	Unfiltered rank	Rank with j>g
3000–2920	Her	0.523001	0.649026	2	1
2980–2720	Her	0.579465	0.669776	2	1
2940–2800	Col	0.693772	1.509824	4	2

Her denotes Hercules and Col denotes Columba. The Hercules boundary solution has HIP roles (84345,80816,84380,81693,84379); at 3000 BCE its E is 0.523001% and its descriptive scale is 1.258438, against Cygnus's 0.649026% and 1.001529. Profiling scale by least squares or minimax also preserves that competitor's advantage in the initial window. The complete tables identify all role assignments. These are numerical comparators, not proposed archaeological identifications.

The material map requires $q=j-g>0$ for a positive second temporal channel. In the fixed-J2000 spatial descriptors, nominal Cygnus gives g=10.487708, j=34.077570 and v=62.160104. The above temporal competitor gives g=15.060675 and j=4.084509, so it fails the sign condition; multiplying by a positive spatial scale cannot reverse that ordering. This is a necessary condition from the existing map, applied retrospectively, not a newly preregistered test.

The weaker continuous condition j>g retains 2,640 classes. Cygnus is globally 24th in their free-E ranking and ranks first, first and second within the initial, 95% and 68% windows. The initial runner-up is Mic at E=1.484582%, 2.2874 times Cygnus; in the 95% window Lyn is next at 1.244650%, 1.8583 times Cygnus. In the 68% window Lyn improves on Cygnus (1.398290% versus 1.509824%). Thus the sign condition does not produce unconditional temporal exclusivity.

Requiring the full rounded spatial tuple (10,34,62), with both assignments of the final pair to latitude examined, retains the nominal class alone in this five-brightest panel before timing. A subsequent temporal success under that restriction cannot be counted as a second independent act of stellar selection. The different restricted results do not erase the unrestricted comparisons.

11. Joint comparison with independent positive scales

This diagnostic retains the specified spatial target S=(10,34,62), temporal word W=(28,72,16), J2000 spatial convention and common catalogue. Each of 5,280 temporal classes receives both assignments of the final pair to the latitude role, giving 10,560 typed assignments. Spatial and temporal scales are fitted separately, with no requirement that they equal each other or remain near 1. No weighted sum between space and time is chosen.

For spatial descriptor vector X, let z=X/S. The first spatial score is sd(z)/mean(z) for positive mean scale. The second is the minimax fractional residual under positive scale: $[\max(z)-\min(z)]/[\max(z)+\min(z)]$ for positive triples. If a descriptor is nonpositive, its minimax infimum is 1 as the scale grows without bound; it is recorded, not omitted. A nonpositive mean scale is marked undefined for the first score. There are 4,440 classes with defined spatial CV, while all 5,280 have a minimax value or infimum. The better final-pair spatial orientation is retained for each criterion; timing is invariant to that pair exchange.

Dominance requires no greater spatial or temporal error and a strictly smaller error in at least one component. Temporal E is taken from the archived global search or from the complete constrained

censuses. The global comparison applies no archaeological chronology filter. The complete frontiers and all class scores are distributed, including alternatives that improve either component.

Spatial criterion	Cygnus spatial %	Cygnus temporal %	Dominators	Frontier classes
Spatial relative dispersion	2.146105	0.649026	0	8
Spatial minimax fractional error	2.266610	0.649026	0	6

The one spatially better class in both criteria is Vulpecula, HIP roles (95771,97886,103004,98543,99874), with spatial CV 1.063655% and minimax 1.242936%, but temporal E=4.074892%. Cygnus is therefore second on the spatial component alone. Candidates improving the temporal component worsen its spatial score. This trade-off explains nondominance without selecting a favourable cross-component weight.

BCE window	CV dominators	CV frontier	Minimax dominators	Minimax frontier
3000–2920	0	3	0	3
2980–2720	0	3	0	3
2940–2800	0	4	0	5

Cygnus remains nondominated in all three windows under both spatial criteria. It is not the only frontier member and need not win every aggregate objective. These retrospective comparisons do not release the target word, stellar definitions, reference frame or regional star selection, and they are not an equal-freedom comparison of complete archaeological models.

The material closure $r=D^2$ fixes a common rescaling of the nominated inventory: scaling it by $s>0$ gives $16s=16s^2$, hence $s=1$. It does not fix the independent conversion $T\approx kW$. A temporal scale change can leave the inventory unchanged. The closure therefore cannot be used to prohibit the alternative temporal scales admitted here, and the ordinary degree-to-sidereal-minute conversion does not establish a minute-per-architectural-number convention.

12. Continuous metric words

This diagnostic asks whether fixed-scale selection persists when the pre-existing constructive recipe is applied to the exterior measurements without snapping the temporal target to integers. For short side x, long side y and diagonal variable z, set g=x, J=x+y, V=2z+x and r=z−x. Then $W_{metric}=(V-J,3(J-g),r)/u=(2z-y,3y,z-x)/u$, with the original exterior calibration $u=(10x+24y+26z)/1352$. The ideal (10,24,26) returns W=(28,72,16). The recipe and ideal coefficients were known before this diagnostic and remain assumptions.

Use x=33.435 m and y=80.095 m from the published side means. Four declared Aubrey diameters serve as diagonal proxies, plus a separately labelled Euclidean diagonal computed from those means. These are not five independently surveyed diagonals or independent replications. Every variant is reported. Fixed-scale absolute RMS and free-scale E are evaluated globally and within all three comparison windows in the common catalogue only.

Exterior convention	Wmetric	Global BCE	RMS min	Rank
Aubrey 86.44 m	27.8515, 72.1271, 15.9107	2975.885	0.322947	1
Aubrey 86.87 m	28.0401, 71.9485, 16.0000	2964.122	0.424256	1

Exterior convention	Wmetric	Global BCE	RMS min	Rank
Aubrey 87.00 m	28.0969, 71.8946, 16.0270	2960.577	0.459085	1
Aubrey 85.83168 m	27.5832, 72.3812, 15.7835	2992.625	0.257887	1
Computed diagonal 86.79348 m	28.0066, 71.9802, 15.9842	2966.211	0.404460	1

All five fixed-scale global minima lie within 3000–2920 BCE, with no chronological filter in their selection. All lie outside 2940–2800 BCE; the MacKie minimum also precedes 2980 BCE. A constrained ranking does not move a global optimum into an interval. The measured range is sensitivity across conventions, not a confidence interval. Decimal precision records computation, not archaeological accuracy.

Variant	Loss	Global rank	3000–2920	2980–2720	2940–2800
Ideal reference	Fixed RMS	1	1	1	1
Ideal reference	Free E	39	2	2	4
Aubrey 86.44 m	Fixed RMS	1	1	1	1
Aubrey 86.44 m	Free E	15	1	2	6
Aubrey 86.87 m	Fixed RMS	1	1	1	1
Aubrey 86.87 m	Free E	43	2	3	4
Aubrey 87.00 m	Fixed RMS	1	1	1	1
Aubrey 87.00 m	Free E	49	2	3	4
Aubrey 85.83168 m	Fixed RMS	1	1	1	2
Aubrey 85.83168 m	Free E	10	1	2	6
Computed diagonal 86.79348 m	Fixed RMS	1	1	1	1
Computed diagonal 86.79348 m	Free E	39	2	3	4

The free-scale global ranks range from 10 to 49 across the five metric variants. The fixed-scale MacKie variant ranks second, behind Reticulum, when restricted to the 68% interval. These adverse outcomes are preserved with all fitted configurations and eligibility decisions. The result supports fixed-scale stability under this continuous extension; it does not independently validate the construction, the timing convention, or ancient intention.

For each target and domain, archived original-target global minima identify every class that could equal or beat the newly optimized Cygnus score b. For absolute RMS, a potential competitor must have old RMS≤b+RMS(Wnew−Wold), by the triangle inequality. For free E, write a=Wnew/Wold. Normalizing the new channel ratios as 1+e with mean(e)=0 and RMS(e)≤b gives old E≤[sd(a)+max(a)b]/[mean(a)−sd(a)b], provided the denominator is positive. The denominator is positive in every application. The numerator follows from sd[a(1+e)]≤sd(a)+max(a)b; the denominator follows from the covariance bound on mean(ae).

Every potentially rank-changing class is refitted on 1-year and 0.5-year grids, with all bracketed minima and endpoints considered. A local epoch offset in scalar refinement avoids a tolerance that grows with the

absolute millennial epoch. Scores agree within relative tolerance 1e−7 and absolute tolerance 1e−8. The ideal reference reproduces the original fixed-RMS first place and free-E rank 39. The eligibility files record all 5,280 classes for every target and domain. This is bounded numerical re-evaluation conditional on the archived global minima, not a new full exhaustive continuous optimization of every class under every target.

verify_results.py checks all 48 target–metric–domain combinations, the target algebra, eligibility sets, rankings and reported scores. Its positive-vector check validates the orientation of the analytic bound; it is not a statistical null. The five metric variants share measurements and the same construction. Their positive outcomes cannot be multiplied into a chance probability or relabelled as prospective confirmation.

13. Execution and file map

The v1.30 archive preserves the original S5 inputs, outputs and protocols and integrates four executed robustness analyses. The commands below are run from the S5 code root. Python, NumPy and SciPy are required. Each branch retains its original protocol, pinned inputs or explicit dependency paths, results and execution log. Portable dependency paths were updated for packaging without changing numerical algorithms; the build verifies unchanged numerical outputs. The original S2–S4 archives and S1 record remain versioned sources, not newly executed analyses.

robustness/temporal_loss: test_losses.py, test_competitors.py and verify_scale_formulas.py; loss_results.json contains all six losses, three input sets and intervals; competitor_results.json contains the bound-based constrained comparisons.

robustness/role_constraints: audit_roles.py; spatial_all_classes.csv, census_initial.csv, census_Greaney95.csv and census_Greaney68.csv contain complete 5,280-class tables. results.json retains the typed and ordinal controls and direct comparisons.

robustness/joint_scales: audit_joint_scales.py; all_classes.json and results.json contain both spatial orientations, scale estimates, sign handling, all frontiers and unit-invariance checks.

robustness/metric_word: audit_metric_word.py and verify_results.py; measurements.json, results.json and eligibility_*.csv retain all targets, bounded re-evaluations and checks. The ideal reference reproduces the original scores and ranks.

verify_v130.py checks package inputs, the integrated summary values and numerical output identities. SHA256SUMS_v130.txt covers all distributed code files other than its own manifest.

The protocols were written after the relevant previous findings were known and before their own executions; their historical timestamps are preserved. No independent survey, archaeological date distribution, probability of prehistoric intention or global false-positive rate is generated by these scripts. The contribution is reproducible conditional epoch inference and explicit reconstruction tests, with both positive and adverse controls available to reviewers.

The dependency diagram is generated by make_figure12.py, requiring Matplotlib and Pillow. The v1.33 editorial revision clarifies the main narrative and shortens the cover letter. Scientific algorithms, numerical results, tables and figures are unchanged. Full-precision results remain in these tables and the executable archives. The README retains execution history and editorial records; editorial_changes_v133.json records the new text changes and paragraph order. SHA256SUMS_v133.txt is the current code integrity manifest. Earlier protocols and manifests retain their historical scope.