Project Tavnit

Theological Researcher Briefing

March 2026 · tavnit.ashestoaltar.com

1. What This Project Is

Scripture is dense with internal connections — cross-references, shared vocabulary, literary parallels, prophetic echoes. Generations of scholars have mapped these connections within individual books and traditions. *Tavnit* (Hebrew: תבנית, "pattern" or "blueprint") asks a different question: if you map *all* of them — across the full canon, from every available tradition — does the resulting network have a discoverable shape? Does Scripture's internal structure reveal a pattern, and if so, what does it look like?

To answer this, the project assembled **1,411,192 cross-textual connections** from **17 independent data sources** and over **2,000 teachers and commentators** spanning five major interpretive traditions. The combined network was analyzed using standard computational methods — the same tools used in genomics, social network research, and digital humanities.

The approach is Torah-positive and data-led. The project does not impose a theological framework on the text; it asks what structure the text's own connections reveal. When the data contradicted our initial hypothesis, we reported the rejection. When control tests caught overreach — and they caught it twenty-two times — we corrected it publicly. Every correction is logged; every methodological decision is documented.

2. The Approach

The text is divided at the paragraph level using the Bible's own ancient section markers: *petuchah* and *setumah* (open and closed sections in the Hebrew Masoretic text) and *kephalaia* (chapter divisions in the Greek New Testament). This yields 3,151 nodes — one per structural paragraph — whose boundaries predate modern chapter-and-verse divisions by centuries.

Connections between these paragraphs come from 17 independently constructed data layers, each reflecting a different method or interpretive tradition: curated scholarly cross-reference systems (Treasury of Scripture Knowledge, Beale & Carson's *Commentary on the NT Use of the OT*), shared-vocabulary networks (Hebrew and Greek), classical Jewish commentary via Sefaria (1,369 commentators), midrashic teachers (26), evangelical commentators (9), patristic Church Fathers (639), and several smaller specialized layers. The combined network is analyzed using standard algorithms — no bespoke or proprietary method is used. The result is a computational-scale extension of what Michael Fishbane documented in *Biblical Interpretation in Ancient Israel* (1985): the Bible's internal cross-referencing is a real structural phenomenon, and it can now be mapped exhaustively.

Every major finding is subjected to adversarial validation using a 3-LLM protocol (Grok, Claude, GPT-4o) in which at least one model operates in "blind interpret" mode — receiving only raw data with no project context.

3. What We Found

The Sinai Attractor

All 61 Old Testament books and all 27 New Testament books gravitate toward the Sinai legislation (Exodus 19–Numbers 10) as their structural center of mass. This result holds across three independent computational methods (Hebrew shared vocabulary, Greek vocabulary A and B) and two languages. The gravitational pull is not imposed by any particular interpretive lens — it emerges from the vocabulary network itself.

Bootstrap confidence intervals: all-layer gravity **22.5 (CI 21.8–23.1)**; leitwort-only gravity **25.8 (CI 25.6–26.0)**. These intervals are extremely narrow, indicating the result is robust to resampling.

13 Structural Axes

PCA of the paragraph-level connection matrix reveals 13 validated interpretive axes — distinct structural dimensions along which the text organizes itself. Each axis captures a different principle of biblical organization. These axes fall into four categories:

Category	Axes	Meaning
Substrate (5)	PC1, PC2, PC5, PC6, PC9	Recovered by 2+ independent bridges; text-level structure visible regardless of which tradition built the data
Tradition-specific (4)	PC3, PC4, PC10, PC11	Real connections, but which ones you find depends on which interpretive tradition's data you use
Tradition-engaged (3)	PC7, PC8, PC19	Axes that interpretive commentary actively engages — reinforcing or disrupting the text's baseline structure
Genre bedrock (1)	PC14	A structural boundary in the text (priestly law vs. wisdom literature) that no interpretive tradition comments across

The number "13" reflects axes that have been individually validated through loading-profile analysis, ablation testing, and blind review by independent evaluators. The Bible's paragraph-level matrix distributes its organizational structure extremely broadly — far more so than comparable ancient literature — so these 13 axes represent distinct organizing principles, not arbitrary divisions of a single pattern.

Full Axis Table

#	Label	Category	Description
PC1	Hub Degree / Connectivity	Substrate	How densely connected each passage is to the rest of Scripture. Genesis paragraphs are the most connected — the vocabulary foundation for the entire canon.
PC2	Aramaic Island	Substrate	The Aramaic sections of Daniel and Ezra form a structural island, linguistically separated from the surrounding Hebrew text.
PC3	Consequences / Wisdom	Tradition-specific	Prophetic "therefore" passages (judgment for disobedience) vs. proverbial wisdom and instruction. Which connections you find here depends on which tradition built the data.
PC4	Warning / Execution	Tradition-specific	"Thus says the Lord" prophetic warning vs. the narrative moments where events unfold — threat vs. fulfillment.
PC5	Priestly / Journey	Substrate	Tabernacle construction, priestly regulations, and sacred space vs. patriarchal migration and journey narratives. Visible from multiple independent data sources.
PC6	Destruction / Restoration	Substrate	The fall of Jerusalem, exile, and judgment vs. prophetic promises of return and restoration. The text's deepest structural tension.
PC7	Prediction / Witness	Tradition-engaged	The arc from covenant prediction to covenant witness — "God said it would happen" to "it happened." All tested traditions reinforce this axis.
PC8	Legal / Metaphor	Tradition-engaged	Legal prescription (Levitical code, Deuteronomic law) vs. prophetic metaphor (Hosea's marriage, Ezekiel's visions). Universal Jewish axis: all five medieval commentators tested reinforce it.
PC9	Tabernacle / Sacrifice	Substrate	Worship space and structure vs. sacrificial action. The strongest cross-testament bridge: recovered by 3 of 4 independent methods. Where Hebrews meets Leviticus.
PC10	Space / Time	Tradition-specific	Sacred space vs. sacred time — the tradition divider. Kabbalistic commentary disrupts this axis; peshat and evangelical commentary reinforce it.
PC11	Temple / Eschatology	Tradition-specific	Temple-centered worship and architecture vs. eschatological expectation. Present sanctuary vs. future hope.
PC14	Legal-Priestly / Wisdom	Genre bedrock	Priestly legislation (Leviticus, Numbers) vs. wisdom literature (Job, Ecclesiastes, Proverbs). A genre boundary no interpretive tradition comments across — invisible to commentary.

PC19	Historical / Prophetic-Poetic	Tradition-engaged	Kings-era prose narrative vs. prophetic poetry (Isaiah, Habakkuk). Evangelical and midrashic commentary bridge this divide; Sefaria reinforces it.

The Covenant Curses Keystone

Deuteronomy 28:15–68 (the *Tokhekhah*, or covenant curses) scores in the extreme range on 9 of 11 originally tested axes — more than any other passage in the Bible. It is a structural keystone: a passage that participates in nearly every organizing dimension of the text. The two sequential-discontinuity boundaries flanking it are the #1 and #2 sharpest structural breaks in the entire Old Testament. A retest on the full 13-axis framework is pending.

Genesis as Vocabulary Foundation

When the Treasury of Scripture Knowledge (47.8% of all connections) is removed, the network's hub structure shifts from "Genesis + Deuteronomy" to "Genesis only." Deuteronomy's hub status depends on cross-reference scholarship; Genesis's does not. Genesis provides the vocabulary foundation that all other books draw on.

The Kabbalistic Fingerprint

Testing five medieval Jewish commentators individually (Rashi, Ibn Ezra, Ramban, Rabbeinu Bahya, Abarbanel) across all 13 axes revealed that *hermeneutical method — not historical period — determines how commentary interacts with the text's structure.*

Rabbeinu Bahya (14th century, kabbalistic-ethical method) produces the most extreme PC10 perturbation of any individual commentator tested ($z = -11.39$). This single author drives much of the aggregate Sefaria layer's $z = -58$ on PC10. By contrast, peshat-oriented commentators (Rashi, Abarbanel) *reinforce* PC10. Ramban shows mild perturbation consistent with his partially kabbalistic methodology.

Rashi (11th century) engages 9 of 13 axes — more than double any other individual commentator — suggesting his verse-by-verse peshat approach produces the broadest structural engagement. PC8 (Legal/Metaphor) is reinforced by all five commentators: a universal axis for Jewish biblical commentary. The profiles cluster by hermeneutical method (peshat vs. kabbalistic), not by century. The separation of sacred space and sacred time as independent categories of holiness — a distinction Abraham Joshua Heschel explored philosophically in *The Sabbath* (1951) — is now computationally visible as the axis (PC10) that most sharply distinguishes kabbalistic from peshat hermeneutics.

The Structure Is Unique to Scripture

Comparing the Bible's paragraph-level network against 31 classical Greek literary works — Homer, Plato, Aristotle, the tragedians — using identical analytical methods, the Bible's structural complexity stands apart. The Bible distributes its organizational patterns across three times more independent dimensions than comparable Greek literature. This is not a generic property of ancient multi-author texts; it is specific to this corpus.

The same result holds when the Hebrew Bible is translated into Greek (the Septuagint shows the same structural profile), confirming the effect is a property of the text's content, not its language. From a different mathematical angle, the Bible's connection network contains persistent topological loops — closed circuits of cross-referencing passages — that Greek literature does not produce at any scale tested. The Bible has 129 such persistent loops; the Greek corpus has 6.

This gap persists under every control tested: density matching, lemmatization, resolution scaling, per-genre comparison. No tested Greek work — not even Plato's *Republic*, which shows genuine internal structure of its own — approaches the Bible's multi-dimensional complexity.

Metric	Bible	Greek Literature (31 works)
Independent structural dimensions	18.1	6.0 (3× fewer)
Persistent topological loops	129	6 (21× fewer)
Septuagint (Greek Bible)	Shows Bible-like profile, not Greek-like — the effect transcends language	

Patristic Phase Transition

Decomposing the patristic data layer by era reveals an abrupt structural shift at approximately 325 CE. Pre-Nicene Church Fathers (Clement, Origen, etc.) perturb 3 axes and reinforce 0 — their cross-references disrupt the text's baseline organization. Post-Nicene Fathers (Augustine, Chrysostom, etc.) reinforce 6 axes and perturb 0 — their cross-references strengthen it. The qualitative asymmetry (0 reinforced vs. 0 perturbed) is density-independent: it persists even when the larger post-Nicene dataset is subsampled to match pre-Nicene size.

This represents a computationally detectable phase transition in how the Church engaged the Hebrew Bible's structure — from disruption to reinforcement — coinciding with the consolidation of the institutional Church in the 4th century.

Bootstrap Confidence Intervals

Key metrics with 95% confidence intervals:

Metric	Estimate	95% CI
Midrashic convergence rate (26 teachers)	9.5%	7.4–11.5%
Midrashic-leitwort Jaccard similarity	0.50	0.48–0.52
Sinai gravity, all layers	22.5	21.8–23.1
Sinai gravity, leitwort only	25.8	25.6–26.0

4. How the New Testament Connects to the Old

The Language Boundary

The Hebrew Old Testament and Greek New Testament form a single connected network, but the connection is selective, not uniform. The NT engages with 3 of 13 structural axes at statistically significant levels ($p < 0.001$) — specifically the vocabulary fault lines around priestly/sacrifice, nations/covenant, and legal/prophetic metaphor. The remaining axes survive in the combined-canon PCA only because the OT's numerical dominance (85% of nodes) mechanically preserves them.

Four Independent Bridges

Four independent cross-testament data sources were tested: the Treasury of Scripture Knowledge, Beale & Carson (11,475 connections from 25 scholars), Greek leitwort A, and Greek leitwort B. Their convergence pattern produced the four-category axis partition described above.

Worship as the Strongest Bridge

PC9 (Tabernacle-space / Sacrifice-action) is the only axis recovered by three of four independent bridges — a triple convergence. Worship and sacrifice represent the strongest structural link between testaments, visible through both vocabulary and interpretive data regardless of tradition. Harold Attridge's Hermeneia commentary

on *Hebrews* (1989) documented how that epistle builds its Christological case through Tabernacle and sacrificial vocabulary; the computational finding shows this is not unique to Hebrews but reflects the broadest vocabulary bridge between testaments.

The Tradition-Engaged Reclassification

PC7 (Prediction/Witness) and PC8 (Legal/Metaphor) were originally classified as "unengaged" by cross-testament bridges, suggesting they existed in covenant territory that the NT did not touch. Subsequent testing with the classical Jewish commentary layer and the patristic layer showed both axes are *actively engaged* by interpretive traditions — all tested traditions reinforce them. They are "tradition-engaged," not unengaged: the text's commentary traditions all find these axes and strengthen them.

Torah-Positive Implications

These structural findings carry a Torah-positive implication worth stating directly. The priestly and sacrificial vocabulary of the Torah — the very system that replacement theology claims was made obsolete — is computationally the thickest linguistic bridge between the testaments. PC9 (Tabernacle/Sacrifice) is the only axis recovered by three of four independent methods. The Tokhekhah — the covenant curses of Deuteronomy 28, the most structurally central passage in Scripture — anchors the tradition-engaged axes: the dimensions that interpretive commentary engages but that no cross-testament bridge reaches independently. The structure does not treat the Torah as scaffolding to be discarded. It treats it as the foundation on which everything else rests.

5. What We Got Wrong

Tavnit maintains a corrections log (22 entries) and a methodology-decisions log (51 entries). Selected corrections that affected published briefing materials:

Claim	Correction
Helical geometry (H3)	Rejected. Twist correlation r = 0.009, p > 0.37. Structure is an asymmetric cone ("tree"), not a helix.
"11 axes for 80% of variance"	Software bug (COR-21). True 80% threshold requires 200+ components. 13 axes are interpretively validated, not variance-counted.
"All 11 NT axes significant"	Majority-dominance artifact. Only 3 of 13 axes show genuine NT-specific fit ($p < 0.001$).
Evangelical PC10 reinforced (z = +4.0)	Small-sample artifact (N = 1 teacher). Full 9-teacher layer: PC10 PERTURBED z = −14.59. Driven by Henry co-citation density.
Co-citation "inherently perturbs" PC9/PC10	Method artifact — density-dependent, not method-inherent. Rashi co-citation reinforces both axes. Henry's perturbation disappears at matched density.
"Single engineered system"	Replaced by: multiple distinct structural properties that different analytical lenses reveal differently.
Torah reading cycle is cyclical	The text's own connection network is not cyclical (sinusoid R^2 max = 0.31 vs. 0.7 threshold). The reading cycle is a liturgical framework, not a property of the text.
Moedim (feasts) as density hubs	Moedim provide directionality, not density ("spin, not mass"). Spring feasts → narrative fulfillment; fall feasts → interpretive theology.
"One-directional engagement" (NT hooks into OT but not vice versa)	Density artifact. At equal connection density, both testaments show similar structural fragility (~1 surviving axis).
Compact NT architecture	The NT is 7.1× denser in cross-references than the OT. "Compactness" was a density effect, not an architectural difference.
3 universal axes (PC7, PC8, PC9)	Only PC9 is triple-convergent. PC7 and PC8 are tradition-engaged (not substrate). The "universal" label was overstated.
TSK as neutral scaffolding	TSK is 47.8% of all edges and perturbs all 11 axes when tested in isolation. It is a dominant structural force, not a neutral baseline.

6. Open Questions

- **Tokhekhah on 13 axes:** The keystone passage was tested on the original 11 axes (scoring extreme on 9). A retest on the full 13-axis framework is pending.
- **Peshat vs. kabbalistic beyond the medieval period:** The kabbalistic fingerprint is confirmed for Bahya and partially for Ramban. Does this pattern extend to later kabbalistic commentary (Zohar-influenced, Hasidic)?
- **Spectral uniqueness vs. other corpora:** The Bible is dramatically flatter than Greek literature. How does it compare to other long-edited multi-author corpora (Talmud, Quran + tafsir, Hindu epics)?
- **Patristic phase-transition drivers:** The pre-Nicene/post-Nicene structural shift is real but its causes are underdetermined. Canon stabilization, genre shifts, citation practice changes, and theological development are all plausible factors.
- **Cross-corpus null:** Would Homer plus its commentary tradition produce the same "universal axis" pattern? This is the most dangerous untested threat to the finding's specificity.
- **Density-matched multi-tradition comparison:** Now unblocked after method-artifact resolution. Will compare all five traditions at equal connection density.

7. Project Details

Item	Value
Total connections	1,411,192
Data layers	17
Paragraph nodes	3,151
Library files	87,532
Library chunks	2,694,361
Teachers / commentators	2,000+ across 5 traditions
Corrections logged	22
Methodology decisions	51
Database	SQLite (single-user, local-first)
Embedding model	BGE-M3 (1024-dim)
Analysis scripts	Python (NumPy, SciPy, NetworkX, scikit-learn)
Validation	3-LLM adversarial protocol (Grok, Claude, GPT-4o)