V1.0.0

The Diagonal Vector: A Vector-Space Model of Interdisciplinary Knowledge, Novelty, and the Synthesis Bottleneck

Rowan Brad Quni-Gudzinas

QNFO-ULA

ABSTRACT

Interdisciplinary research is routinely praised and structurally penalized: researchers who span fields face fragmented vocabulary, siloed evaluation, and funding systems organized by discipline, while the cognitive cost of holding many fields in one mind is unquantified. This paper develops a vector-space model of interdisciplinary knowledge in which disciplines are orthogonal basis directions of an n-dimensional knowledge space, an individual's or a paper's knowledge state is a vector of expertise coefficients, and innovativeness is scored by the number of significantly loaded dimensions (the participation ratio) together with angular distance from existing work. The gap addressed is twofold: no existing formalism connects a measurable vocabulary-silo statistic to a synthesis-cost model, and the claim that large language models shift the research bottleneck from information access to synthesis bandwidth has not been given a computable form. Using the measured partitionality index of six arXiv disciplines (0.9721, from 453 of 466 technical bigrams being domain-local), we derive the concentration of random-vector overlap in that space (standard deviation $1/\sqrt{466} \approx 0.0463$), a joint-fidelity model in which per-interface reliability $p = 0.95$ (assumed) yields end-to-end fidelity $0.95^{10} \approx 0.599$ for a five-field integration and $0.95^{45} \approx 0.099$ for a ten-field integration, and a pairwise-cost law $k(k-1)/2$ giving 45 cross-pairs at $k = 10$ versus 190 at $k = 20$. Under a linear value hypothesis the model predicts an optimal integration width $k^* \approx 4.7$ fields. The framework is analytical (TRL 1-2); its principal use is to turn interdisciplinarity from a rhetorical virtue into a set of measurable quantities that corpus studies and funding taxonomies can test.

PUBLISHER CLASSIFICATION High-quality scholarly research (QNFO publisher classification for open-access, versioned research with a public revision history and reader correction channel; this is not a claim of independent peer review.)

PEER-REVIEW STATUS Not independently peer-reviewed

CITE AS Quni-Gudzinas, R. B. (n.d.). The Diagonal Vector: A Vector-Space Model of Interdisciplinary Knowledge, Novelty, and the Synthesis Bottleneck. QNFO Preprint v1.0.0.

1. Introduction

A researcher in quantum foundations, a researcher in number theory, and a researcher in machine learning can work on the same structural object and never find each other's work. This is not a hypothetical: a quantitative audit of six arXiv categories (quantum physics, number theory, machine learning, genomics, materials science, economic theory) found that 453 of 466 distinct technical bigrams occur in exactly one domain, that the partitionality index is 0.9721, and that the shared core of terms occurring in three or more domains is empty [8]. The cost of this lexical wall is paid twice. First, duplicate and parallel structures go unrecognized, so the combinatorial surface of possible cross-domain connections is left unexplored. Second, when researchers do attempt integration, they must do so across a vocabulary gap with no shared metric for what counts as a genuine synthesis rather than a surface-level juxtaposition of keywords.

The problem has become more acute, not less, with the arrival of large language models (LLMs). When the marginal cost of retrieving and restating knowledge from any field approaches zero, the scarce resource shifts from information access to the cognitive bandwidth needed to hold several fields' structures in mind at once and to check that the combination is coherent. Yet no existing framework quantifies this shift: how the value of an integrative idea scales with the number of fields it spans, where the curse of dimensionality begins to degrade synthesis quality, and how a genuine "diagonal" contribution can be distinguished computationally from a semantic collage of borrowed terms. Funding agencies and editors currently decide interdisciplinarity cases with rhetoric and review-panel intuition; a computable model would let them ask instead how many dimensions a proposal meaningfully spans, whether those dimensions are orthogonal in practice, and whether the integration compresses or merely concatenates.

This paper builds that framework. We model disciplines as orthogonal basis vectors of an n-dimensional knowledge space; an agent's knowledge state is a coefficient vector; and a research idea's innovativeness is a function of its significant components and its angular separation from the existing corpus. The model is deliberately simple, and its value lies in what it makes computable: silo statistics become geometry, synthesis cost becomes a function of integration width, and the distinction between synthesis and collage becomes a testable property of coefficient profiles.

The contributions are:

- A formal definition of the knowledge space, the expertise vector, and two novelty statistics — the participation ratio N_{eff} and the angular distance θ from prior work — with worked derivations (Sections 3 and 4).
- A concentration result for random-vector overlap in a vocabulary space of measured dimension, showing that chance lexical overlap is bounded at the level of ± 0.09 cosine similarity for the measured six-domain space, so that observed cross-domain bridge structure above that band is informative (Section 4.2).
- A synthesis-cost and fidelity model: pairwise cost $k(k-1)/2$ and joint fidelity $p^{\binom{k}{2}}$ under an assumed per-interface reliability, which together predict a finite optimal integration width $k^* \approx 4.7$ and locate the curse of dimensionality for synthesis (Sections 4.4 and 4.5).
- A computable criterion distinguishing genuine synthesis from semantic collage, based on the decay exponent of sorted squared coefficients, illustrated with worked profiles ($\alpha \approx 2.54$ for a structured profile versus $\alpha \approx 0.33$ for a collage) (Section 4.6).
- A mapping of the framework onto existing taxonomies of scientific contribution and onto a proposed corpus-validation design, with readiness honestly assessed as TRL 1-2 (Sections 6 and 7).

The geometric view developed here may also bear on problems outside research policy — any setting where agents combine knowledge from disjoint label systems, such as ontology alignment or multi-source data fusion — insofar as those settings share the orthogonal-vocabulary structure measured in [8].

2. Background and Related Work

The mathematical core of the model — treating a knowledge space as a vector space with orthogonal discipline directions — has precedents across several literatures, which we review for the specific tools each supplies.

Orthogonal decomposition in function spaces. [1] investigates an orthogonal decomposition of a Sobolev space, determining the dimension of the orthogonal difference space and showing how the spaces expand as their regularity exponent increases: the change in regularity "causes increase in length or norm, expansion of the space in terms of distance or separation between distinct elements and change in orthogonality." This is directly relevant as a formal precedent: orthogonality and distance are not fixed properties of a space but depend on the inner product chosen, and refining the space's structure changes the geometry of separation. Our model inherits this lesson — the "orthogonality" of disciplines is a property of the chosen representation (here, disjoint vocabulary), and changing the representation changes which ideas count as diagonal.

Sparse and compressible representations. [5] proposes the optimum subspace decomposition of infinite-dimensional compressible random processes in a locally convex Hausdorff space, measuring the dimension of the decomposition and proving that with a sufficient number of separating points the space admits a Banach-space structure; the work frames compressed sensing as reconstructing a sparse signal "by finding the infimum subspace." This supplies the second formal precedent: knowledge states, like compressible signals, are sparse in the right basis, and the effective dimension of a state is far smaller than the ambient dimension. Our participation-ratio statistic (Section 3) is the elementary analogue of this optimum-subspace idea.

Interpreting vector components as concepts. Whether the components of a learned vector space carry conceptual meaning is an empirical question. [4] trains a neural model to "conceptualize" word vectors — activating higher-order concepts recognized in a given vector — and reports a weighted classification F_1 score of 0.22, with performance varying sharply across concepts: high for sharply bounded concepts such as *province*, poor for vague ones such as *location*. The authors conclude that

"the more precise a concept boundary, the higher the classification performance tends to be." This is a caution for our framework: if a knowledge-space axis is to be interpreted as "the discipline," the interpretability of that axis depends on how sharply the discipline's boundary is defined, and axes corresponding to vague fields will support weaker geometric claims.

Novelty as a relational quantity. [3] analyzes sentence-level novelty detection and argues that novelty is "a combination of the partial overlap (PO...)" relations among items rather than an intrinsic property; the paper shows that clustering-based novelty methods fail because partial-overlap relations are not transitive, so sentences "need not necessarily form classes," and reports that clustering-based novelty methods fail because partial-overlap relations are not transitive. Two lessons carry over: novelty must be defined relative to a reference set (in our model, angular distance from existing work), and naive clustering of "fields" can destroy the overlap structure that novelty measurement depends on.

Empirical vector spaces of skills and knowledge states. [7] develops the Knowledge Query Network (KQN) for knowledge tracing, representing students' knowledge states and skills as vectors, and explicitly tests the sensitivity of conclusions to the embedding dimensionality: skill-distance structure is compared across dimensionalities at $0.5d_{\text{opt}}$, d_{opt}, and $2d_{\text{opt}}$, with Mantel correlations between distance matrices ranging from 0.495 to 0.825 across datasets and distance measures, and clustering quality measured by Adjusted Rand Index. The finding that distance structure is preserved but not identical across dimensionalities is exactly the robustness question our novelty metric faces: how much does N_{eff} or θ change when the vocabulary space is re-embedded at a different dimension?

Cross-dataset integration as information gain. [2] describes the SPACE mission concept, which aims to take near-IR spectra for more than half a billion galaxies, and gives a concrete instance of integration producing information that neither part contains alone: the CMB spectrum C_l depends on the combination $\Omega_m h^2$ while the matter power spectrum $P(k)$ is sensitive to $\Omega_m h$, "thus, combining these two datasets help to break degeneracies and adds tremendous scope to tighten the limits on parameters which result from the use of just one dataset in isolation." This is the physical-science prototype of the "diagonal vector" claim: the joint constraint is not the union of the individual constraints but a qualitatively new one, because the two datasets' parameter degeneracies are oriented differently — geometrically, their degeneracy directions are non-parallel, which is our orthogonality condition in a concrete form.

Quantified silo structure. [8] measures the precondition for any of this geometry to matter: whether domains share vocabulary at all. Beyond the six-domain audit (partitionality 0.9721, 13 bridge terms out of 466, of which 5 are machine-learning method vocabulary such as *machine-learning* and *foundation-models*), the paper probes the QNFO corpus itself and finds that the knowledge graph's semantic edges resolve to 40 record pairs, of which 8 are same-program and 0 are cross-program, while within-program linked pairs share title vocabulary at mean Jaccard 0.115 against 0.028 for random pairs. The within-program elevation of $+0.086$ establishes that shared vocabulary tracks semantic relatedness where links exist — which is what licenses using vocabulary overlap as a coordinate system in the first place — and the zero cross-program links show the coordinate system currently connects nothing across the wall.

Contribution taxonomies and unification as compression. [10] decomposes scientific progress into three functional types — Generative (G-Type, novel empirical predictions), Structural Unification (S-Type, coherence among existing theories), and Epistemic Scaffolding (E-Type, refinement of inquiry methods) — and formalizes non-trivial unification as description-length compression: Theorem 3 requires $G_U(T_U; T_1, T_2) > \delta$ with, for example, $\delta = 0.15 \times (K(T_1) + K(T_2))$, so that the unifying theory is strictly shorter than the sum of its parts. Our synthesis-versus-collage criterion (Section 4.6) is the vector-space counterpart of this compression condition: a genuine diagonal idea must be simpler than the concatenation of its source-field descriptions.

Testable cross-domain protocols. [11] specifies a Structured Cross-Domain Translation (SCDT) protocol with four executable phases and an evaluation design in which LLM agents either receive an object with a request for "interdisciplinary insights" (baseline) or execute the protocol with fixed lexicons (SCDT condition), measuring novelty, expert validity, and predictive success. Its computed search-space quantities include $N_{\text{hyp}} = 60$ pairwise hypotheses and $E_{\text{total}} = 180$ evaluation points, and its projected false-link risk falls from $P_{\text{any}} \approx 0.265$ unmitigated to $P_{\text{any}}^{\text{tested}} \approx 0.059$ with validation sensitivity $s = 0.8$. This supplies both a template for testing our metrics and a quantified warning: unstructured cross-domain generation carries a spurious-link probability on the order of one in four, which is precisely the failure mode our collage criterion is designed to catch.

Community infrastructure. [6] presents a community-wide strategy for advancing heliophysics and space-weather modeling through open science, noting that applying open-science principles to modeling "requires more careful consideration regarding open data and open software guidelines, as scientific models differ significantly" from other artifacts. For our framework, this establishes both the enabling condition and the friction: corpus-scale tests of the diagonal-vector hypothesis require open, machine-readable cross-domain data and software, and the community norms for such sharing in modeling fields are still being negotiated.

Finally, [9] illustrates the ambition and the risk of grand unification: it posits a universal Hilbert space with $H|\Psi\rangle = 0$ and formalizes cross-domain mappings as functors between information, geometric, and perceptual categories, with perceiver correlations written as mutual information $I(X_p; X_q) = \sum_{x_p, x_q} P(x_p, x_q) \log[P(x_p, x_q)/(P(x_p)P(x_q))]$. We cite it not for its physics claims but for its methodological pattern — cross-domain structure expressed as mappings between formal spaces — and as a reminder, via its own emphasis that the mapping functor "is generally not injective," that cross-domain correspondence loses information, a point our angular-distance metric makes locally and quantitatively.

3. Methods

This is an analytical paper. We define a model, derive its properties, and compute all numbers from stated inputs; no experiment, simulation, survey, or corpus analysis is performed here, and every empirical input is attributed to its source or labeled as an assumption.

Knowledge space. Let $\{e_1, \ldots, e_n\}$ be an orthonormal basis of $\mathbb{R}^n$, where each basis direction e_i represents one discipline (operationally: one domain whose technical vocabulary is disjoint from the others, per the measured structure in [8]). A knowledge state — of a person, a group, or a paper — is a vector

$$x = \sum_{i=1}^{n} x_i e_i, \qquad x_i \geq 0,$$

where x_i measures the depth of expertise in discipline i. The non-negativity constraint encodes that expertise is not signed; a paper that misuses a field's methods does not gain negative loadings but rather loadings that fail the coherence checks below.

Justification of orthogonality as an assumption. Orthogonality of the basis is an idealization justified by the measured partitionality: [8] finds 453 of 466 technical bigrams domain-local (partitionality index 0.9721) and an empty three-domain shared core. If disciplines shared substantial vocabulary, their basis vectors would have positive mutual inner products and the geometry would need a general metric $g_{ij} = \langle e_i, e_j \rangle$. We take the orthogonal model as the leading approximation under the convention that the measured 0.9721 partitionality is high enough that off-diagonal metric terms are second-order; Section 6 discusses when this fails.

Choice of novelty statistic. Two candidate statistics appeared in the source material for this framework: a threshold-count score $N = \frac{1}{n}\sum_i \mathbf{1}(|x_i| \geq \tau)$ with a heuristic cutoff, and the participation ratio. We adopt the participation ratio N_{eff} as the primary statistic under the convention that a smooth, threshold-free measure is preferable when the threshold τ has no external calibration; the threshold-count score is retained only as a secondary diagnostic whose behavior is discussed in Section 6.

Effective dimensionality. Following the sparse-reconstruction idea of finding the minimal subspace supporting a signal [5], we define the effective number of disciplines spanned by a state as the participation ratio

$$N_{\text{eff}}(x) = \frac{\left(\sum_{i=1}^{n} x_i^2\right)^2}{\sum_{i=1}^{n} x_i^4}.$$

For a state with k equal components $x_i = c$ and the rest zero, $N_{\text{eff}} = k$ exactly; for uneven profiles, N_{eff} lies in $[1, k]$ and down-weights states with one dominant component. This is the "number of dimensions meaningfully spanned" made precise.

Angular novelty. Let y be the centroid vector of a reference corpus (the existing literature in the idea's neighborhood). We define the angular novelty of idea x as

$$\theta(x, y) = \arccos \frac{\langle x, y \rangle}{\|x\| \, \|y\|} \in [0, \pi].$$

A purely intra-disciplinary advance has small θ; a diagonal idea has large θ because its mass sits in directions the corpus centroid does not occupy.

Synthesis cost and fidelity. Integrating k disciplines requires, at minimum, checking every pairwise interface: we take the pairwise cost law $C(k) = \binom{k}{2} = k(k-1)/2$. For end-to-end fidelity we assume each discipline-interface is handled with per-interface reliability p (an assumption; no source measures it), and interfaces fail independently, giving joint fidelity

$$F(k) = p^{\binom{k}{2}}.$$

We adopt this interface-fidelity model as the primary synthesis-cost model, under the convention that the binding constraint on wide integration is the number of interfaces that must each be validated, rather than a per-dimension quality penalty; Section 6 discusses the alternative per-dimension decay hypothesis and what observation would discriminate between them.

Synthesis versus collage. A genuine synthesis, following the compression condition of [10], must be structurally simpler than the concatenation of its parts. In vector terms we use the sorted squared-coefficient profile $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_k$ (with $\lambda_i = x_{(i)}^2$). We define the profile decay exponent α by fitting $\lambda_i \propto i^{-\alpha}$ over the significant components and propose, as a hypothesis to be tested rather than a derived theorem, that genuine syntheses exhibit a structured profile (a small number of load-bearing dimensions with $\alpha \gtrsim 1$) while collages exhibit near-flat profiles ($\alpha \approx 0$) — many borrowed terms with near-equal weight and no load-bearing core. The threshold is a convention; we state it as a hypothesis in Section 6.

Chance-overlap floor. To distinguish signal from chance in vocabulary-level overlap statistics, we compute the concentration of random-vector cosine similarity in dimension d (Section 4.2), which sets the floor above which observed cross-domain structure is informative.

4. Analysis

Every input is stated with its source; every arithmetic step is shown.

4.1 The measured space and its statistics

Inputs from [8]: 466 distinct technical bigrams across six arXiv domains; 453 domain-local; 13 bridge terms (two or more domains); 0 shared-core terms (three or more domains); partitionality index reported as 0.9721.

Check of the partitionality index:

$$\frac{453}{466} = 0.972103\ldots \approx 0.9721.$$

Bridge share:

$$\frac{13}{466} = 0.027897\ldots \approx 0.0279,$$

matching the reported bridge share 0.0279 [8]. The two independent ratios confirm internal consistency of the source's table.

Within-program lexical elevation, from [8]: mean title-vocabulary Jaccard of 0.115 for same-program linked pairs versus 0.028 for random pairs:

$$\Delta J = 0.115 - 0.028 = 0.087,$$

which the source reports as $+0.086$; the 0.001 difference is rounding in the source's inputs, and we use the source's stated elevation $+0.086$ when quoting it.

4.2 Chance-overlap floor in the measured space

Model the vocabulary space as d-dimensional with $d = 466$ (the number of distinct technical bigrams, from [8]). For two independent random unit vectors in $\mathbb{R}^d$ with i.i.d. zero-mean components, the cosine similarity c has

$$\mathbb{E}[c] = 0, \qquad \mathrm{Var}(c) = \frac{1}{d}.$$

(This is the standard concentration result for random directions; it follows from each pairwise product $u_i v_i$ having mean 0 and variance $1/d^2$, summed over d coordinates.) With $d = 466$:

$$\sigma_c = \frac{1}{\sqrt{466}} = \frac{1}{21.5870\ldots} = 0.046323\ldots \approx 0.0463.$$

The two-sided 95% band under approximate normality is

$$\pm 1.96\,\sigma_c = \pm 1.96 \times 0.046323 = \pm 0.090793\ldots \approx \pm 0.0908.$$

Interpretation: in a space of the measured size, two unrelated knowledge states will show cosine overlap within about ± 0.09 by chance alone. The within-program lexical elevation $+0.086$ measured in [8] sits at the edge of this band — a point we return to in Section 6, since it warns that single-pair lexical signals near 0.09 are weak evidence on their own and must be aggregated.

4.3 Worked example: participation ratio

Take a hypothetical researcher with expertise coefficients (assumed for illustration, in arbitrary expertise units) $x = (0.9, 0.8, 0.6, 0.1, 0.1)$ across five disciplines, all other components zero.

Compute the squares: $x_i^2 = (0.81, 0.64, 0.36, 0.01, 0.01)$.

$$\sum x_i^2 = 0.81 + 0.64 + 0.36 + 0.01 + 0.01 = 1.83.$$

Fourth powers: $x_i^4 = (0.6561, 0.4096, 0.1296, 0.0001, 0.0001)$.

$$\sum x_i^4 = 0.6561 + 0.4096 + 0.1296 + 0.0001 + 0.0001 = 1.1955.$$

Then

$$N_{\text{eff}} = \frac{(1.83)^2}{1.1955} = \frac{3.3489}{1.1955} = 2.8013\ldots \approx 2.80.$$

Although the vector has five nonzero components, its effective width is 2.80 disciplines: the two 0.1-loadings contribute almost nothing. This is the quantitative content of "meaningfully spans": naive nonzero counting would say 5, the participation ratio says 2.80. A uniform five-field polymath, $x_i = c$ for $i \le 5$, would give $N_{\text{eff}} = 5$ exactly, since $(5c^2)^2/(5c^4) = 25c^4/5c^4 = 5$.

As a corpus-level expectation, if each of the 13 bridge terms in [8] is taken to mark one significant cross-domain loading and the six sampled domains define the ambient space, the expected number of significant dimensions per randomly drawn term is

$$k_{\text{exp}} = N_{\text{share}} \times n = 0.027897 \times 6 = 0.16738\ldots \approx 0.17.$$

A typical vocabulary item in this corpus spans far less than one full cross-domain dimension on average, which is the term-level reflection of the empty three-domain shared core.

4.4 Synthesis cost

With the pairwise law $C(k) = k(k-1)/2$:

$$C(5) = \frac{5 \cdot 4}{2} = 10, \qquad C(10) = \frac{10 \cdot 9}{2} = 45, \qquad C(20) = \frac{20 \cdot 19}{2} = 190.$$

Doubling the integration width from 10 to 20 fields multiplies the interface count by $190/45 = 4.2222\ldots$ — cost grows quadratically, so the marginal discipline is increasingly expensive to integrate.

4.5 Joint fidelity and the optimal integration width

Assume per-interface reliability $p = 0.95$ (assumption; no source measures interface reliability — we state it to make the model computable and examine sensitivity). Then

$$F(k) = 0.95^{\binom{k}{2}}.$$

Compute:

$$F(5) = 0.95^{10} = 0.598737\ldots \approx 0.599,$$

$$F(10) = 0.95^{45} = e^{45 \ln 0.95} = e^{45 \times (-0.051293)} = e^{-2.30819} = 0.099457\ldots \approx 0.099,$$

$$F(20) = 0.95^{190} = e^{190 \times (-0.051293)} = e^{-9.74568} = 5.8517 \times 10^{-5} \approx 5.9 \times 10^{-5}.$$

Check $F(5)$: $0.95^2 = 0.9025$; $0.95^4 = 0.9025^2 = 0.81451$; $0.95^8 = 0.81451^2 = 0.66342$; $0.95^{10} = 0.66342 \times 0.9025 = 0.59874$. The chain confirms the exponential evaluation.

Under this model, a ten-field integration has about a 10% chance of being fully coherent end-to-end, and a twenty-field integration is essentially guaranteed to break somewhere. This is the curse of dimensionality for synthesis: not the geometry of distance concentration (which is mild, Section 4.2) but the combinatorial explosion of interfaces that must each be validated.

Value-fidelity trade-off and k^*. Suppose the value of an idea grows with its effective width as $V(k) = v\,k$ (assumption: linear value in width, the simplest hypothesis; the degeneracy-breaking gain exemplified in [2] suggests value can grow super-linearly, which would only strengthen the interior optimum). Expected net value:

$$\mathcal{V}(k) = v\,k\,p^{k(k-1)/2}.$$

Maximize over k: with $p = 0.95$, $\ln \mathcal{V} = \ln(vk) + \frac{k(k-1)}{2} \ln p$. Setting the derivative to zero:

$$\frac{1}{k} + \left(k - \tfrac{1}{2}\right) \ln p = 0 \quad \Rightarrow \quad \frac{1}{k} = \left(k - \tfrac{1}{2}\right) |\ln p|.$$

With $|\ln p| = 0.051293$:

$$0.051293\,k^2 - 0.0256465\,k - 1 = 0,$$

$$k = \frac{0.0256465 + \sqrt{0.0256465^2 + 4 \times 0.051293}}{2 \times 0.051293} = \frac{0.0256465 + \sqrt{0.00065774 + 0.205172}}{0.102586} = \frac{0.0256465 + \sqrt{0.205830}}{0.102586}.$$

$$\sqrt{0.205830} = 0.453687, \qquad k^* = \frac{0.479334}{0.102586} = 4.6725\ldots \approx 4.7.$$

Under these assumptions the value-optimal integration width is about 5 disciplines — enough for a genuine diagonal, few enough that the interface web stays checkable. Sensitivity: with $p = 0.99$ ($|\ln p| = 0.010050$), the same quadratic gives $0.010050\,k^2 - 0.005025\,k - 1 = 0$, so $k = \frac{0.005025 + \sqrt{0.0000253 + 0.0402}}{0.0201} = \frac{0.005025 + 0.20057}{0.0201} = 10.23$; with $p = 0.90$ ($|\ln p| = 0.10536$), $k^* = \frac{0.05268 + \sqrt{0.002775 + 0.42144}}{0.21072} = \frac{0.05268 + 0.65132}{0.21072} = 3.34$. The optimum is sensitive to p (roughly $k^* \propto 1/\sqrt{|\ln p|}$), which is the single most important quantity for any empirical follow-up to measure.

4.6 Synthesis versus collage: worked profiles

Construct two five-component profiles with identical ℓ_2 norm and identical total mass but different shapes.

Profile S (structured, hypothesized synthesis): squared coefficients $\lambda = (0.60, 0.25, 0.10, 0.04, 0.01)$, total 1.00.

Profile C (hypothesized collage): squared coefficients $\lambda = (0.2699, 0.2145, 0.1872, 0.1700, 0.1586)$, constructed as $\lambda_i = A\, i^{-0.33}$ with $A = 1/\sum_{i=1}^{5} i^{-0.33}$. Compute the normalizer: $i^{-0.33}$ for $i = 1, \ldots, 5$ gives $(1, 0.7946, 0.6936, 0.6297, 0.5877)$, summing to 3.7056, so $A = 1/3.7056 = 0.2699$. Both profiles have $\sum_i \lambda_i = 1.00$, hence identical ℓ_2 norm $\|x\| = \sqrt{1.00} = 1$ and identical total expertise mass; only the shape differs.

Decay exponent of Profile S. Fit $\lambda_i \propto i^{-\alpha}$ using the endpoints $i = 1$ and $i = 5$:

$$\alpha = \frac{\ln(\lambda_1/\lambda_5)}{\ln 5} = \frac{\ln(0.60/0.01)}{\ln 5} = \frac{\ln 60}{\ln 5} = \frac{4.09433}{1.60944} = 2.5440 \approx 2.54.$$

Decay exponent of Profile C.

$$\alpha = \frac{\ln(0.2699/0.1586)}{\ln 5} = \frac{\ln 1.7018}{1.60944} = \frac{0.53175}{1.60944} = 0.3304 \approx 0.33.$$

Participation ratios of the two profiles. Since $\sum \lambda_i = 1$ for both, $N_{\text{eff}} = 1/\sum_i \lambda_i^2$.

Profile S: $\sum \lambda_i^2 = 0.36 + 0.0625 + 0.01 + 0.0016 + 0.0001 = 0.4342$, so

$$N_{\text{eff}}^{\text{S}} = \frac{1}{0.4342} = 2.3029 \approx 2.30.$$

Profile C: $\sum \lambda_i^2 = 0.07285 + 0.04601 + 0.03504 + 0.02890 + 0.02515 = 0.20795$, so

$$N_{\text{eff}}^{\text{C}} = \frac{1}{0.20795} = 4.8088 \approx 4.81.$$

The two statistics separate the profiles in opposite directions, which is what makes the pair diagnostic. The collage spreads mass almost evenly over five dimensions ($N_{\text{eff}} = 4.81$ of a possible 5) with a flat decay exponent $\alpha \approx 0.33$; the synthesis concentrates on a load-bearing core ($N_{\text{eff}} = 2.30$) with a steep exponent $\alpha \approx 2.54$. Under the compression reading of [10], the structured profile is the one a unifying description would produce: a few dimensions carry the integrated structure and the remainder are refinements, whereas the collage profile is what concatenating five field-vocabularies with equal weight yields. We emphasize the epistemic status: the $\alpha \gtrsim 1$ threshold and the association of flat profiles with collage are hypotheses stated for testing, not derived results; the derivation here is only that the two profile shapes are separable by (N_{eff}, α) jointly, which the arithmetic above establishes for these constructed cases.

5. Results

All numbers below are computed in Section 4 from the stated inputs; none is measured by this paper.

Measured space (from [8], checked in Section 4.1). Partitionality index 0.9721 $(453/466)$; bridge share 0.0279 $(13/466)$; shared core at the three-domain threshold: 0 terms; within-program lexical elevation $+0.086$ (source-stated; recomputed as 0.087 from the source's inputs).

Chance-overlap floor (Section 4.2). With $d = 466$, the standard deviation of chance cosine similarity is $\sigma_c = 1/\sqrt{466} = 0.0463$, and the two-sided 95% band is $\pm 1.96\,\sigma_c = \pm 0.0908$. Observed cross-domain cosine overlap above about 0.09 is therefore informative at the 5% level for single pairs; the measured within-program elevation $+0.086$ lies inside this band, so single-pair lexical signals of that size are not individually significant.

Participation ratio (Section 4.3). The illustrative five-component vector $(0.9, 0.8, 0.6, 0.1, 0.1)$ has $N_{\text{eff}} = 2.80$ against a naive nonzero count of 5. The expected significant width per vocabulary item under the bridge-share reading is $k_{\text{exp}} = 0.0279 \times 6 = 0.17$.

Synthesis cost and fidelity (Sections 4.4–4.5). Interface counts: $C(5) = 10$, $C(10) = 45$, $C(20) = 190$. Under the assumed per-interface reliability $p = 0.95$: $F(5) = 0.599$, $F(10) = 0.099$, $F(20) = 5.9 \times 10^{-5}$. Under the linear value hypothesis $V(k) = vk$, the optimal integration width is $k^* = 4.67 \approx 4.7$ fields, with sensitivity $k^* = 10.23$ at $p = 0.99$ and $k^* = 3.34$ at $p = 0.90$.

Synthesis-versus-collage separation (Section 4.6). For the two constructed profiles with identical norm and mass: structured profile, $(N_{\text{eff}}, \alpha) = (2.30, 2.54)$; collage profile, $(N_{\text{eff}}, \alpha) = (4.81, 0.33)$. The pair of statistics separates the shapes by construction; whether real syntheses and real collages occupy these regions is the hypothesis handed to Section 7's proposed test.

6. Discussion

What the model claims and what it does not. The paper derives three things: the chance-overlap floor ± 0.0908 for the measured six-domain vocabulary space; the interface-count and fidelity consequences of wide integration under an assumed p; and the separability of constructed synthesis-like and collage-like profiles by (N_{eff}, α). It does not measure any real researcher, paper, or corpus, and it does not establish that the orthogonal-basis idealization holds beyond the lexical evidence of [8].

Limitations, with sensitivity.

- *Orthogonality of the basis.* Rests on partitionality 0.9721 from a 40-abstract-per-domain sample [8]. If off-diagonal metric terms g_{ij} are non-negligible, angular distances shrink and the chance-overlap floor must be recomputed in the true metric; the floor ± 0.0908 is sensitive to d as $1/\sqrt{d}$, so a vocabulary extraction twice as large would halve it to about ± 0.065.
- *Per-interface reliability p.* Purely assumed. The headline optimum $k^* \approx 4.7$ moves to 10.23 at $p = 0.99$ and 3.34 at $p = 0.90$; the result is therefore a family of predictions indexed by p, not a single number. Measuring p is the single most valuable empirical input the framework lacks.
- *Linear value.* $V(k) = vk$ is the simplest hypothesis; the degeneracy-breaking example of [2] suggests super-linear gains are possible, which would push k^* upward, and diminishing returns would push it down. The interior optimum survives any value function that is not so steeply super-linear that it overwhelms the quadratic interface cost, but its location does not.
- *Profile-shape hypothesis.* The $\alpha \gtrsim 1$ threshold is a convention. The constructed profiles show separability, not that the separation occurs in nature.
- *Vocabulary as coordinate system.* The within-program elevation $+0.086$ sits inside the chance band ± 0.0908, so the licensing argument for vocabulary-as-geometry rests on aggregation over pairs, not on any single pair.
- *Threshold-count statistic.* The rejected alternative $N = \frac{1}{n}\sum_i \mathbf{1}(|x_i| \geq \tau)$ behaves discontinuously near τ; we retained the participation ratio partly to avoid an uncalibrated cutoff, but if external calibration of τ later becomes available, the two statistics could disagree and the choice would need revisiting.

Falsification. The main quantitative result most exposed to falsification is the fidelity-collapse prediction: under the model, ten-field integrations should be end-to-end coherent at rate about 0.099. A corpus study that identified ten-field integrative papers and found their internal cross-field consistency at rates materially above 0.2 would falsify the assumed $p = 0.95$ and, with it, the narrow optimum $k^* \approx 4.7$ under linear value. Independently, if a re-embedding of the vocabulary space at a different dimension (the robustness question posed by [7]'s Mantel-correlation finding of 0.495–0.825 across dimensionalities) changed N_{eff} rankings of real papers substantially, the statistic would fail its purpose. And if expert-rated genuine syntheses showed flat sorted-coefficient profiles (α near 0.33) at the same rate as keyword collages, the collage criterion would be dead.

Open questions. What is the correct unit of "discipline" — arXiv category, subfield, or laboratory? How should negative or borrowed loadings (methods imported without understanding) be represented given the non-negativity constraint? Does the interface-failure independence assumption hold when interfaces share a common method vocabulary, as the 5-of-13 machine-learning bridge terms in [8] suggest they often do?

7. Practical Implications and Engineering Pathways

(a) Who would use the result and what changes. Three user groups, if the framework survives testing. Funding agencies that currently judge interdisciplinarity by panel intuition could require proposals to state an intended integration width k and check it against the fidelity curve $F(k)$: a ten-field proposal under the assumed $p = 0.95$ carries end-to-end coherence below 0.1 and would need either a staged integration plan or evidence of elevated p. Editors and reviewers could use (N_{eff}, θ) as a screening statistic to distinguish diagonal contributions from keyword collages before peer review. Corpus researchers gain a concrete

target list: measure p, recompute k^*, and test the profile-shape hypothesis. If no user adopts it, the work remains foundational: it converts a rhetorical distinction (synthesis versus collage) into quantities that a later measurement program can confirm or kill.

(b) Readiness. This is analytical work that has not been peer reviewed: TRL 1–2 under the NASA definitions. No component of the framework has been validated on a real corpus; the chance-overlap floor and the profile separation are derived or constructed, not observed. Nothing here supports claims of deployment readiness, and no manufacturing readiness applies.

(c) Smallest next test. A corpus study would take a few hundred papers tagged as integrative across the six arXiv domains of [8], compute each paper's coefficient profile over a shared vocabulary embedding, and measure the sorted-coefficient decay exponent α. The outcome that refutes the framework: genuine syntheses (identified by expert rating or citation-based breakthrough labels) and keyword collages showing indistinguishable α distributions. A complementary study would estimate p directly by auditing a sample of multi-field papers for internal cross-field consistency errors, replacing the assumption $p = 0.95$ with a measurement.

(d) Binding constraints. First, data: the test requires open, machine-readable cross-domain corpora with reliable field labels; the open-science norms for such sharing in modeling fields are still being negotiated [6]. Second, the embedding choice: [7]'s dimensionality-sensitivity results (Mantel correlations 0.495–0.825 across dimensionalities) warn that N_{eff} and α may shift under re-embedding, so the pipeline must fix and report the embedding dimension. Third, validation cost: expert validity ratings are the expensive ingredient, as the SCDT protocol's own budget projection shows (20.0 h unfiltered reduced to 11.0 h with a θ-pre-filter at $q = 0.4$) [11]. Tolerance and materials constraints do not arise; the binding inputs are data, annotation, and compute for embedding.

(e) Order-of-magnitude bounds (projections). Using only numbers derived above or stated in [11]: if a validation study follows the SCDT evaluation design with $N_{\text{hyp}} = 60$ pairwise hypotheses and $E_{\text{total}} = 180$ evaluation points, and expert rating time scales as in its projection (11.0 h with pre-filtering), a single benchmark cycle is an eleven-hour expert-time commitment — small enough for a single-lab pilot. The unmitigated spurious-link probability $P_{\text{any}} \approx 0.265$ from [11], reduced to ≈ 0.059 at validation sensitivity $s = 0.8$, is consistent with our chance-band picture: unstructured cross-domain generation produces spurious links at a rate several times the 5% chance floor, and structured validation brings it back to the floor's neighborhood. These are projections under [11]'s stated assumptions, not measurements.

(f) Risks and misuse. The most direct misuse is metric gaming: if funding or review adopts N_{eff} or α as a criterion, authors can optimize coefficient profiles rhetorically — padding a paper with a controlled spread of field keywords to manufacture a steep-looking profile — without any genuine integration. A second risk is the opposite policy error: using the fidelity collapse $F(10) \approx 0.099$ to justify defunding all wide integration, when the model's own sensitivity shows the collapse is an artifact of the assumed p and could vanish at $p = 0.99$. Both risks argue for publishing the model with its assumptions explicit, as done here.

8. Conclusion

This paper turned three claims about interdisciplinary research into computable form. Disciplines whose vocabularies are measured to be almost perfectly disjoint (partitionality 0.9721) can be modeled as orthogonal axes; chance lexical overlap in the resulting 466-dimensional vocabulary space is bounded at ± 0.0908 cosine similarity, so cross-domain structure above that band is informative; and the cost of integrating k fields grows quadratically in interfaces while end-to-end fidelity falls as $p^{\binom{k}{2}}$, which under an assumed per-interface reliability of 0.95 and linear value yields an optimal integration width of about 4.7 fields. A pair of statistics — participation ratio and sorted-coefficient decay exponent — separates constructed synthesis-like profiles $(2.30, 2.54)$ from collage-like profiles $(4.81, 0.33)$, giving the synthesis-versus-collage distinction a testable form. The framework's value is not any single number but the conversion of interdisciplinarity from rhetoric into a small set of measurable quantities, each with a stated assumption and a stated falsifier. The immediate next steps are empirical and cheap relative to the question they answer: measure the per-interface reliability p, and test whether real syntheses and real collages occupy the profile regions the model assigns them.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.55.14). Language models were tools in that process; they are not authors. The named author is accountable for the content.
- Models: drafting by openai/gpt-oss-120b, zai-org/glm-5.3-flash; reconciliation, completion, review, revision and verification by zai-org/glm-5.3-flash, openai/gpt-oss-120b, moonshotai/kimi-k2.6, with the QNFO model router as a fallback.
- Sources: every reference is drawn from a bibliography retrieved from arXiv and the QNFO corpus before drafting; every cited work's full text was retrieved (arXiv HTML, or the QNFO paper's body) and the writers were given excerpts of its own sections; statements about a cited work are restricted to its retrieved title, abstract and those excerpts.
- Checks passed before publication: required sections present and finished, cross-references resolve, reference list rendered from the retrieved bibliography, no assistant residue or placeholders, adversarial model review with its blocking findings resolved, and a citation existence check at publication.
- Independent recomputation: the execution sandbox was unavailable, so a second language model recomputed 16 quantitative claims step by step in text from the paper's stated inputs (no code was executed); 0 differed from the stated value at the stated precision.
- Not checked: whether each cited source supports the sentence that cites it beyond the retrieved full-text excerpts; derivation steps other than the recomputed claims. This paper has not been peer reviewed, and this version was published by the pipeline without a prior manual review by the author.

References

[1] On Orthogonal Decomposition of a Sobolev Space. arXiv:1611.04249v1. https://arxiv.org/abs/1611.04249v1 [2] SPACE: the SPectroscopic All-sky Cosmic Explorer. arXiv:0804.4433v1. https://arxiv.org/abs/0804.4433v1 [3] The Nature of Novelty Detection. arXiv:cs/0510054v1. https://arxiv.org/abs/cs/0510054v1 [4] Neural Vector Conceptualization for Word Vector Space Interpretation. arXiv:1904.01500v1. https://arxiv.org/abs/1904.01500v1 [5] Compressible Topological Vector Spaces. arXiv:2208.13728v1. https://arxiv.org/abs/2208.13728v1 [6] Advancing Heliophysics and Space Weather Modeling through Open Science. arXiv:2605.30626v1. https://arxiv.org/abs/2605.30626v1 [7] Knowledge Query Network: How Knowledge Interacts with Skills. arXiv:1908.02146v2. https://arxiv.org/abs/1908.02146v2 [8] https://papers.qnfo.org/papers/terminology-silos-consilience. QNFO: Terminology Silos and the Consilience Gap: A Quantitative Audit of Cross-Domain Vocabulary. [9] https://papers.qnfo.org/papers/converging-reality. QNFO: Monistic Reality. [10] https://papers.qnfo.org/papers/quantifying-structural-unification-and-epistemic-scaffolding-in-scientific-progress. QNFO: Quantifying Structural Unification and Epistemic Scaffolding in Scientific Progress. [11] https://papers.qnfo.org/papers/structured-cross-domain-translation-as-a-testable-consilience-mechanism-protocol. QNFO: Structured Cross-Domain Translation as a Testable Consilience Mechanism: Protocol Design, Truth-Preservation Conditions, and Analytical Evaluation.