

**From Ideonomic Binomials to Computational
Dimensions: Extending Patrick Gunkel's System of
Divisions for Tensor Calculus and Large Language Models**

William C Russell

Abstract

The term “dimension” is used across geometry, physics, linear algebra, tensor calculus, machine learning, and classification theory, yet its senses remain only loosely unified. This paper reconstructs Patrick Gunkel’s ideonomic system of binomial divisions as a generative methodology for proposing conceptual dimensions. A dimension is defined as a reproducible distinction along which a system can vary independently enough to support structured representation; equivalently, it is the answer space to a structured question whose poles orient a continuum of possible responses. Ideonomy supplies the generative step through systematic proposal of binomial divisions . Statistics and linear algebra supply validation via measures of meaningfulness, operationalizability, non-redundancy, and stability, aggregated in the composite utility score . Tensor calculus supplies interaction modeling by realizing selected dimensions as factor spaces of an ideonomic tensor . Large language models supply a contemporary experimental environment in which an explicit low-dimensional coordinate system can be projected from high-dimensional latent states. The reconstruction yields a hybrid architecture that layers an interpretable ideonomic superstructure onto emergent representations without replacing them. A proof-of-concept design centered on literary fiction illustrates how a modest basis of dimensions can be scored, residualized, and compared against ordinary embeddings, while higher-order tensor analysis tests whether familiar semantic categories are better understood as intersections of continuous axes. The paper advances a conceptual dimensional calculus whose elementary operation is division and whose derived objects are coordinates, vectors, and tensors, thereby offering one systematic response to the problem of organizing an effectively unlimited universe of ideas.

I. Introduction: The Problem of Conceptual Dimensions

1.1 The Ambiguity of “Dimension”

The term “dimension” appears with striking frequency across contemporary inquiry, yet its meaning shifts according to the discipline in which it is invoked. In ordinary speech the word may designate little more than a measurable aspect or a figurative “level” of complexity; in technical registers it acquires sharper, though still heterogeneous, senses. The resulting polysemy is not merely terminological inconvenience. When researchers speak of “high-dimensional data,” “conceptual dimensions,” or “the dimensions of a classification,” they risk conflating distinct formal structures under a single label, thereby obscuring both the genuine analogies that unite those structures and the precise points at which the analogies break down. A preliminary clarification of the principal senses of the term is therefore indispensable before any systematic theory of conceptual dimensions can be advanced.

In classical geometry a dimension is an independent degree of freedom. A point in the Euclidean plane may be located by two real numbers; a point in ordinary space by three. Each additional coordinate axis enlarges the space of possible positions without being reducible to the preceding axes. Independence here is geometric: no continuous motion along one axis can be reproduced by motion along the others. The cardinality of the set of such axes is the dimension of the manifold. This geometric sense supplies the most intuitive prototype: a dimension marks a direction in which variation is possible and, crucially, independent of variation in other directions.

Physics inherits the geometric notion and specializes it. Coordinates become the numerical values of measurable physical quantities (length, time, mass, temperature, electric charge) relative to chosen units and origin. A physical dimension is thus both a coordinate axis and a type of measurable phenomenon. Dimensional analysis, as formalized in the nineteenth and twentieth centuries, treats these types as elements of an abelian group under multiplication, so that the dimension of a derived quantity is a monomial in the fundamental dimensions. The physical sense therefore adds two constraints absent from pure geometry: the axes must correspond to empirically accessible magnitudes, and the algebraic relations among those magnitudes must be respected. Variation remains central, but it is now variation of observables under controlled experimental conditions.

Linear algebra abstracts still further. Here a dimension is the cardinality of a basis of a vector space. Any vector may be uniquely expressed as a linear combination of basis vectors; the coefficients are its coordinates. Two bases of the same space have the same

cardinality, so dimension is an invariant of the space itself. Independence is now linear independence: no non-trivial linear combination of the basis vectors yields the zero vector. The geometric and physical senses reappear as special cases once an inner product or a physical interpretation is supplied, yet the linear-algebraic definition stands free of metric or empirical commitments. It is this sense that underwrites the language of “feature spaces” and “coordinate representations” throughout contemporary mathematics and computation.

Tensor calculus extends the linear-algebraic framework by admitting multi-linear objects whose components transform according to prescribed rules under change of basis. A tensor T of type (r, s) lives in a tensor product of r copies of a vector space and s copies of its dual. Each factor space contributes an index whose range is the dimension of that space. The “dimensions” of a tensor are therefore the extents of its index sets together with the factor spaces themselves. Interaction among dimensions is no longer a secondary phenomenon but the primary object of study: contraction, outer product, and covariant differentiation encode how variation along one index may depend upon, or induce, variation along others. The tensorial sense thus retains the idea of independent axes while systematically tracking their mutual dependencies.

Machine learning deploys the term in at least three overlapping ways. First, a feature is a measurable or computable attribute of an instance; the collection of features defines a coordinate system in which each instance is a point. Second, latent variables are unobserved coordinates inferred by a model: principal components, factors in a factor analysis, or the continuous dimensions of a variational auto-encoder. Third, embedding coordinates are the real-valued vectors assigned to discrete tokens or graph nodes by a neural network, so that geometric relations among the vectors approximate semantic or structural relations among the original objects. In each case the operative notion is that of a coordinate along which instances vary, whether the coordinate is hand-crafted, statistically extracted, or learned end-to-end. High dimensionality is routinely treated as both a computational obstacle (the curse of dimensionality) and a representational resource (the expressive power of large embedding spaces).

Finally, within philosophy and classification theory a dimension is a distinction along which phenomena may differ. To classify organisms by “vertebrate / invertebrate,” or ethical theories by “consequentialist / deontological,” is to introduce a binary cut that partitions a domain. More refined classifications admit graded or multi-valued distinctions, yet the underlying idea remains the same: a dimension is a respect in which items can be compared or contrasted. The philosophical sense is the least formalized of the six, yet it is also the most general. Every geometric axis, every physical observable, every basis vector,

every tensor index, and every machine-learning feature can be read as a specialized instance of a conceptual distinction.

Despite their surface diversity, the six senses share a common structure. In each case a dimension is a way things can vary. Geometry isolates independent directions of spatial variation; physics isolates independent kinds of measurable change; linear algebra isolates independent directions of linear combination; tensor calculus isolates independent index ranges whose interactions are themselves objects of variation; machine learning isolates coordinates (observed, latent, or embedded) along which data points differ; classification theory isolates distinctions that render phenomena comparable.

The shared structure may be stated more formally. Let $\mathcal{U}$ be a conceptual universe, the domain of discourse under consideration. A dimension relative to $\mathcal{U}$ is a structured distinction that induces a non-trivial partition or gradation of $\mathcal{U}$ (or of some distinguished subdomain $X \subseteq \mathcal{U}$). The poles of the distinction define the extreme values of a coordinate; intermediate values, when admitted, fill the interval between the poles. The collection of such distinctions, when mutually independent in an appropriate sense, spans a space of possible configurations. What varies from discipline to discipline is the precise criterion of independence, the algebraic or statistical apparatus used to measure coordinates, and the epistemic status of the distinctions themselves, whether they are stipulated, observed, or inferred. This structural commonality raises a precise methodological question. If dimensions are ways things can vary, and if computational systems (especially those built upon high-dimensional vector spaces and tensor operations) must choose or discover such ways of variation, then the quality of the resulting representations depends upon the principled character of that choice. Ad hoc feature engineering, purely statistical extraction of latent axes, and unconstrained end-to-end learning each possess well-documented strengths and limitations. None, however, begins from an explicit theory of conceptual distinction. The question therefore becomes unavoidable: can a systematic theory of conceptual distinctions provide a principled method for defining computational dimensions?

The ensuing sections examine one historically articulated candidate for such a theory and reconstruct it as a generative methodology whose outputs may be validated by statistical criteria, modeled by tensorial interaction, and experimentally probed within contemporary large language models.

1.2 Gunkel, Research Questions, and Thesis

Ideonomy, as conceived by Patrick Gunkel, constitutes an ambitious attempt to systematically organize the possible structures of ideas. Rather than cataloguing existing

concepts or refining particular classificatory schemes, ideonomy seeks a higher-order inventory of the forms that conceptual organization itself may take. Its central gesture is the isolation of recurrent patterns of differentiation—patterns that recur across domains and that, once made explicit, can be applied to generate new distinctions. Among these patterns the binomial division occupies a privileged place. Formally, a binomial division may be written $D(X) = \{X_1, X_2\}$, or more compactly $X \rightarrow (A, B)$, where X is a conceptual domain drawn from a larger universe $\mathcal{U}$ and A and B are the two poles that result from the division. The operation does not merely split an already given set; it proposes a respect in which items belonging to X may be differentiated.

It is essential to read the binomial division as a conceptual operator of differentiation rather than as a producer of mutually exclusive empirical categories. The poles A and B need not be interpreted as exhaustive and disjoint classes that partition the world. They function instead as limiting markers that define an axis of possible variation. An entity may be located at either pole, at intermediate positions, or, in limiting cases, may resist placement altogether; the axis remains well-defined even when empirical instances cluster unevenly or leave regions unoccupied. In this respect the binomial division anticipates the coordinate structure later formalized in linear algebra and tensor calculus: the poles supply the extreme values of a coordinate $x_i \in [-1, 1]$, while the continuum between them supplies the range of admissible positions. The division is therefore generative rather than taxonomic. It does not claim that every phenomenon must fall cleanly into one or the other class; it claims only that the distinction itself constitutes a coherent dimension along which phenomena can be compared, ordered, or measured.

This reading of the binomial division immediately raises three research questions that structure the present inquiry. First, can binomial divisions be formalized as a generative procedure for conceptual dimensions? That is, can the informal practice of proposing opposing poles be rendered precise enough to yield a reproducible method for expanding a candidate set $\mathcal{D} = \{D_1, \dots, D_n\}$ of distinctions, each $D_i = (A_i, B_i)$, and for assigning to each an ideonomic utility score $\Phi(D_i)$ that reflects meaningfulness, operationalizability, non-redundancy, and stability? Second, how do conceptual divisions relate to mathematical dimensions in tensor spaces? Once a selected set $\mathcal{D}^*$ of size k has been obtained, the corresponding factor spaces $V_{D_1}, \dots, V_{D_k}$ may be assembled into a tensor $T \in V_{D_1} \otimes \dots \otimes V_{D_k}$. The question is whether the algebraic structure of such a tensor faithfully captures the intended relations of independence and interaction among the original conceptual distinctions. Third, can explicitly structured ideonomic dimensions improve interpretability, controllability, or analysis of large-language-model

representations? Given a hidden state $\mathbf{h} \in \mathbb{R}^d$, a linear map $W \in \mathbb{R}^{n \times d}$ may project $\mathbf{h}$ onto an ideonomic coordinate vector $\mathbf{z} = W\mathbf{h}$. The empirical question is whether the resulting coordinates afford clearer semantic interpretation, more precise steering, or more diagnostic error analysis than the unconstrained latent dimensions of the model itself.

The central thesis of this paper answers the first question affirmatively and supplies a framework within which the second and third may be pursued. Patrick Gunkel’s ideonomic system of binomial divisions can be reconstructed as a generative methodology for proposing conceptual dimensions. A dimension, on this reconstruction, is a formally specified distinction whose poles define a space of possible relations. Ideonomy supplies the generative step: the systematic proposal of candidate distinctions $D_i = (A_i, B_i)$. Statistics and linear algebra supply the validation step: measures of redundancy R_{ij} , together with the components of the utility score $\Phi(D_i) = \alpha M_i + \beta O_i + \gamma N_i + \delta S_i$, permit selection of a non-redundant, operationalizable subset $\mathcal{D}^*$. Tensor calculus supplies the interaction-modeling step: once the selected dimensions are realized as factor spaces, their joint structure is expressed by an ideonomic tensor T . Finally, large language models supply a contemporary experimental environment in which the projected coordinates $\mathbf{z}$ can be tested for interpretability and control. The reconstruction does not claim that Gunkel anticipated tensor calculus or neural networks; it claims only that the conceptual operator he isolated can be embedded within a modern mathematical and computational setting without violence to its original intent. Nor does the reconstruction assert that binary distinctions imply a binary ontology; the poles of a division mark extremes of a continuum, not exhaustive metaphysical categories. Still less does it propose that designed ideonomic dimensions should replace learned embeddings; the two regimes are complementary, the former offering explicit structure and the latter offering data-driven coverage.

A brief roadmap indicates how the argument unfolds. Section 2 develops the generative methodology in detail, formalizing the production of candidate divisions, the scoring functions f_i that map entities to coordinates x_i , and the criteria that prune $\mathcal{D}$ to $\mathcal{D}^*$. Section 3 examines the passage from selected conceptual dimensions to tensorial representations, clarifying the sense in which an ideonomic tensor encodes both independence and interaction. Section 4 turns to large language models, constructing the projection $\mathbf{z} = W\mathbf{h}$ and reporting experimental probes of interpretability and controllability. Section 5 consolidates the results, acknowledges limitations, and indicates directions for further work. Throughout, the exposition remains deliberately modest in its historical claims and explicit about the boundaries of the proposed reconstruction.

The next section begins the constructive task by examining the internal logic of binomial division as a generative operator.

II. Patrick Gunkel's Ideonomic System

2.1 Ideonomy as a General Science of Ideas

Ideonomy presents itself as a general science of ideas, oriented not toward the enumeration of existing objects or the refinement of classificatory schemes, but toward the systematic study of the structures that ideas themselves may assume. Where traditional taxonomy begins with a domain of entities and asks what kinds of things exist within it, ideonomy begins with the space of possible conceptual operations and asks what kinds of distinctions are possible. The shift is decisive. Classification of objects presupposes that the relevant axes of difference have already been identified; ideonomy treats those axes as the primary objects of inquiry. Its governing ambition is therefore meta-classificatory: to isolate recurring forms of differentiation that appear across systems of thought and to render those forms available as generative instruments.

This orientation rests on a substantive conviction about the nature of conceptual possibility. Ideonomy proceeds from the working hypothesis that the structures of ideas are not indefinitely plastic or arbitrarily idiosyncratic, but exhibit recurrent patterns that can, in principle, be inventoried. Just as the natural sciences once moved from the cataloguing of specimens to the isolation of underlying regularities, ideonomy seeks to move from the collection of concepts to the isolation of the formal operations by which concepts are differentiated, combined, and elaborated. The hypothesis does not entail that every conceivable distinction has already been articulated, nor that the inventory is finite in any trivial sense. It entails only that the space of distinctions is structured, and that systematic exploration of that space is therefore a coherent research program rather than an exercise in unconstrained invention.

Central to this program is the recognition that systems of thought (scientific, philosophical, legal, aesthetic) repeatedly employ a limited repertoire of organizational devices. Binary contrasts, graded scales, hierarchical nestings, and cross-cutting partitions reappear in domains that share no obvious subject matter. Ideonomy treats such recurrence as evidence that the devices themselves constitute a higher-order subject matter. By abstracting from the contents that fill these structures, one obtains a set of formal patterns that can be reapplied, varied, and combined. The resulting catalogue is not a taxonomy of the world but a taxonomy of the ways in which a world can be carved. In this respect ideonomy stands closer to a morphology of conceptual form than to an ontology of substances.

Formal division supplies the principal method by which this morphology is explored. A division takes a conceptual domain $X \subseteq \mathcal{U}$ and proposes a structured differentiation of it. The simplest and most generative case is the binomial division $D(X) = \{X_1, X_2\}$, or $X \rightarrow (A, B)$, which isolates two poles and thereby defines an axis of possible variation. The operation is not primarily empirical; it does not assert that the poles correspond to observed clusters in data. It asserts that the distinction itself is coherent, that the poles stand in a relation of meaningful opposition or complementarity, and that intermediate positions along the resulting axis are at least thinkable. Once such a division has been proposed, it may be tested for fruitfulness: whether it organizes existing material illuminatingly, whether it suggests further distinctions, whether it remains stable under modest reformulation. The method is therefore exploratory rather than merely descriptive: each successful division expands the stock of available conceptual dimensions, and each expansion alters the space in which subsequent divisions may be framed.

The contrast with traditional taxonomy can be stated with greater precision. Taxonomy, in its classical form, begins with a population of individuals and seeks partitions that maximize intra-class similarity and inter-class difference according to some chosen metric of resemblance. The axes of comparison are typically taken as given (morphological characters, molecular sequences, functional roles) and the task is to discover the natural or most useful cuts along those axes. Ideonomy inverts the priority. It treats the axes themselves as the phenomena to be generated and assessed. The question is not which cut best respects an antecedently fixed set of characters, but which characters, or more generally which distinctions, are worth introducing in the first place. A distinction earns its place not solely by its capacity to recover existing clusters, but by its capacity to open new comparative possibilities, to interact productively with other distinctions, and to remain intelligible across shifts of domain.

This generative stance carries methodological consequences. Because the objects of ideonomy are distinctions rather than individuals, the criteria of success cannot be reduced to empirical fit alone. A proposed division may be assessed for internal coherence (whether the poles are genuinely opposed or complementary), for scope (whether the axis applies beyond the domain in which it was first noticed), for productivity (whether it suggests further divisions or illuminates anomalies), and for non-redundancy (whether it is independent of distinctions already in hand). These criteria are not algorithmic, yet they are not merely subjective; they constitute a disciplined practice of conceptual experimentation. The practice remains fallible—some proposed divisions will prove sterile, others will collapse into prior ones—but the fallibility is of a kind continuous with other forms of theoretical inquiry.

The same stance also explains why ideonomy remains largely independent of any particular metaphysical commitment concerning the ultimate furniture of the world. One may hold that the most fundamental distinctions track real joints in nature, or that they are pragmatic instruments shaped by human interests, or that they occupy some intermediate status. Ideonomy as a method does not decide among these options. It supplies a procedure for proposing and refining distinctions; the further question of their ontological status can be left open or settled by independent argument. In this respect the program is modular: its generative and evaluative machinery can be adopted by realists and anti-realists alike, provided only that both parties accept the legitimacy of systematic attention to conceptual form.

The foregoing characterization remains deliberately abstract. It isolates the governing questions and the characteristic method without yet examining the detailed apparatus of binomial division, the criteria by which candidate distinctions are scored and selected, or the mathematical structures that can be built once a set of dimensions has been secured. Those tasks belong to the sections that follow. What has been established is the distinctive orientation of ideonomy: a science of ideas that classifies distinctions rather than objects, that treats conceptual possibility as systematically explorable, and that employs formal division as its primary instrument of investigation.

The next section turns from this general orientation to the specific operator that Gunkel placed at the center of the ideonomic enterprise (the binomial division) and examines its internal logic as a generative device.

2.2 Binomial Divisions and Their Three Kinds

The binomial division constitutes the elementary generative operation of ideonomy. Given a conceptual domain $X \subseteq \mathcal{U}$, a binomial division is an operation that yields a pair of poles, written $D(X) = \{X_1, X_2\}$ or, equivalently, $X \rightarrow (A, B)$. The poles A and B are not required to be subsets of X in any set-theoretic sense; they function as limiting markers that define a respect in which elements of X (or of a larger universe containing X) may be differentiated. Once the poles are fixed, an entity may be assigned a coordinate $x_i \in [-1, 1]$ along the resulting axis, where the endpoints correspond to pure instantiation of A and of B respectively. The division therefore does more than classify; it opens a one-dimensional space of possible positions.

The same formal pattern appears across widely divergent domains. Familiar oppositions such as static/dynamic, internal/external, discrete/continuous, actual/potential, individual/collective, and deterministic/stochastic all instantiate the binomial schema. In each case two poles are isolated, a continuum of intermediate positions is at least

conceivable, and the resulting axis supplies a dimension along which phenomena may be compared or ordered. These examples illustrate the generative reach of the operator: the poles need not be drawn from a single established taxonomy, nor need they correspond to empirically observed clusters. What matters is that the opposition is intelligible, that it marks a coherent respect of variation, and that it can be applied beyond the context in which it was first noticed.

Not every binomial division, however, functions in the same way. Three distinct kinds may be distinguished according to the logical and empirical relations that obtain between the poles. The distinctions are ideal types; actual proposals may approximate one kind more closely than another, and the same pair of terms may be interpretable under more than one reading. Nevertheless the typology clarifies the range of work that binomial division can perform and, crucially, isolates the variety that most directly licenses passage to mathematical dimension.

The first kind is the complementary division. Here the poles are required to be mutually exclusive and jointly exhaustive with respect to the domain under consideration. Formally, if A and B are interpreted as classes, one has $A \cap B = \emptyset$ and $A \cup B = X$ (or some suitable enlargement of X). The division therefore induces a strict partition: every element of the domain falls into exactly one of the two classes. A classic illustration is the opposition odd/even restricted to the integers; every integer is either odd or even, and none is both. Complementary divisions are powerful when the domain admits a clean dichotomy, yet they are correspondingly brittle. As soon as intermediate or mixed cases appear, or as soon as the domain is enlarged, the demand for exhaustiveness and exclusivity forces either ad-hoc stipulations or abandonment of the division. Complementary divisions therefore tend to function as classificatory tools rather than as open axes of variation.

The second kind is the contrasting division. The poles remain opposed in meaning (each is understood as the contrary or the negation of the other) yet they are not required to be exclusive or exhaustive. Empirical instances may exhibit features of both poles simultaneously, or may lie outside the contrast altogether. The opposition static/dynamic, for example, functions in many contexts as a contrast rather than a complement. A physical system may be described as predominantly static for certain purposes and dynamic for others; a social institution may display both inertial and transformative tendencies at once. The poles retain their oppositional force (static is still the contrary of dynamic) but the relation is one of semantic polarity rather than of set-theoretic partition. Contrasting divisions are correspondingly more flexible. They tolerate overlap, gradation, and incomplete coverage, and they therefore serve well as exploratory instruments. Their

limitation is precision: because the poles do not carve the domain into exclusive regions, the assignment of coordinates remains partly a matter of interpretive judgment.

The third kind is the orthogonal division. Here the poles are not interpreted primarily as classes that stand in relations of exclusion or opposition, but as the extreme values of an independent axis of variation. The division does not claim to partition the domain; it claims only to define a direction in which entities can differ independently of other directions already under consideration. The opposition discrete/continuous, when read orthogonally, does not assert that every phenomenon is either wholly discrete or wholly continuous. It asserts that discreteness and continuity mark the endpoints of a dimension along which models, data structures, or conceptual schemes may be located, and that this dimension may be treated as independent of, for instance, the dimension actual/potential or the dimension individual/collective. Orthogonal divisions therefore abstract from the classificatory force of the poles and retain only their capacity to orient a coordinate axis. An entity's position on the axis is expressed by a real number $x_i \in [-1, 1]$, and a collection of such axes yields a coordinate vector $\mathbf{x} \in \mathbb{R}^n$.

The three kinds are not mutually exclusive in practice, yet they differ in the constraints they impose and in the further constructions they license. Complementary divisions maximize classificatory sharpness at the cost of flexibility. Contrasting divisions maximize semantic vividness at the cost of formal precision. Orthogonal divisions maximize independence and coordinate structure at the cost of immediate classificatory utility. It is the third kind that supplies the bridge from ideonomic division to mathematical dimension. Once poles are treated as endpoints of an axis rather than as rival classes, the apparatus of linear algebra becomes available: linear independence of axes can be assessed, redundancy R_{ij} between pairs of divisions can be quantified, and a selected set $\mathcal{D}^*$ of orthogonal distinctions can be assembled into a coordinate system. The same reading also prepares the ground for tensorial constructions, in which each orthogonal dimension contributes a factor space and the joint structure of several dimensions is expressed by a tensor $T \in V_{D_1} \otimes \cdots \otimes V_{D_k}$.

The recognition that orthogonal divisions function as independent axes does not eliminate the complementary and contrasting varieties from ideonomic practice. Both remain valuable as generative heuristics; many productive orthogonal axes begin life as vivid contrasts or as attempted complements that later prove incomplete. The typology therefore serves a dual purpose. It disciplines the interpretation of proposed divisions, and it isolates the precise variety (the orthogonal) whose formal properties align most closely with the mathematical notion of dimension. In what follows, the focus narrows to the conditions under which a binomial division can be treated as orthogonal, the criteria by

which candidate orthogonal distinctions are evaluated, and the procedures by which a coherent set of such distinctions is selected and coordinated.

The next section examines the evaluative criteria that transform a raw collection of proposed divisions into a principled candidate set $\mathcal{D}$ and, ultimately, into a selected set $\mathcal{D}^*$ of ideonomic dimensions.

2.3 Recursive Division and the Binary Tree as Coordinate Structure

Binomial division becomes a systematically generative procedure once it is applied recursively. Beginning with a conceptual domain $X \subseteq \mathcal{U}$, the initial division $X \rightarrow (A, B)$ yields two poles. Each of those poles may itself be treated as a domain and subjected to a further division: $A \rightarrow (A_1, A_2)$ and $B \rightarrow (B_1, B_2)$. The process may be iterated. After one level there are two terminal categories; after two levels there are four; after n levels of exhaustive binary branching there are 2^n terminal categories. The resulting structure is a binary tree whose root is the original domain X and whose leaves are the finest distinctions obtained by successive differentiation.

The tree is more than a hierarchical classification. It functions as a coordinate-generating structure. Each root-to-leaf path corresponds to a unique sequence of pole choices. If the left-hand (or first) pole at every node is conventionally assigned the sign $-$ and the right-hand (or second) pole the sign $+$, then every path becomes a finite sign sequence of length n , for example $(+, -, +, +, -)$. Such a sequence may be read directly as a vector in $\{-1, +1\}^n$, or, more flexibly, as the skeleton of a real-valued coordinate vector $\mathbf{x} \in \mathbb{R}^n$ once intermediate positions along each division are admitted. In this way the recursive application of binomial division does not merely multiply categories; it systematically produces a set of coordinates that locate any entity relative to the entire cascade of distinctions.

The generative power of the construction is immediate. A single initial division supplies one axis. Each subsequent level doubles the number of terminal leaves while lengthening the coordinate sequences by one component. The same formal pattern can be applied selectively rather than exhaustively: only some poles need be subdivided, and different branches may be developed to different depths. The result is a (possibly unbalanced) binary tree whose paths still yield well-defined sign sequences, albeit of unequal length. Padding or hierarchical encoding can restore uniformity when a fixed-dimensional representation is required. In every case the tree supplies an explicit, inspectable procedure for expanding a modest initial distinction into a richer coordinate system.

Yet the coordinate reading must be handled with care. In a nested tree the lower-level distinctions are conditional on the upper-level ones. The division $A \rightarrow (A_1, A_2)$ is applied only to entities already located at pole A ; it is undefined, or at least differently motivated, for entities located at pole B . Consequently, the coordinates obtained from successive levels are not independent axes in the linear-algebraic sense. The second component of a path vector is meaningful only given the value of the first component; the axes are hierarchically ordered rather than freely combinable. A change in an upper-level assignment can render an entire subtree of lower-level distinctions inapplicable. The resulting representation therefore captures a nested system of conditional differentiations, not a Cartesian product of independent dimensions.

This hierarchical character constitutes both a strength and a limitation. It is a strength insofar as many natural conceptual systems are themselves nested: broad distinctions are refined by narrower ones that presuppose the broader cut. Biological taxonomy, legal categorization, and certain forms of diagnostic reasoning all exhibit this conditional structure. The binary tree faithfully mirrors such nesting and thereby preserves information about dependency that a flat list of independent axes would erase. It is a limitation, however, when the aim is to construct a space in which every dimension can vary freely with respect to every other. Orthogonal divisions, as characterized in the preceding section, are precisely those that license such free variation. A collection of orthogonal divisions yields coordinates that behave like the components of a vector in $\mathbb{R}^k$, whereas the path coordinates of a nested tree behave like addresses in a hierarchical filing system.

The tension between hierarchical path coordinates and independent axial coordinates is therefore intrinsic to the recursive use of binomial division. One and the same formal operation (repeated binary differentiation) can be read in two ways. Read as a tree, it generates nested, conditional distinctions whose natural representation is a path signature. Read as a set of orthogonal axes, it generates independent dimensions whose natural representation is a free coordinate vector $\mathbf{x} \in \mathbb{R}^n$. Both readings are legitimate; they simply serve different purposes. The hierarchical reading is indispensable for reconstructing the internal logic of many existing conceptual schemes and for generating fine-grained terminal categories. The orthogonal reading is indispensable for the subsequent mathematical constructions that treat dimensions as factor spaces of a tensor $T \in V_{D_1} \otimes \cdots \otimes V_{D_k}$ and for the projection of latent representations onto an ideonomic coordinate system.

The present section has isolated the recursive mechanism and the coordinate interpretation it immediately suggests. It has also registered the crucial caveat: path coordinates derived from a nested tree are hierarchical rather than independent. That

caveat does not invalidate the coordinate reading; it delimits its scope and indicates the work that remains to be done. The resolution of the tension, how a generative process that naturally produces nested distinctions can nonetheless yield a selected set of approximately independent axes, belongs to the formal apparatus developed in Section III. There the evaluative criteria of non-redundancy, stability, and operationalizability are introduced, and procedures are given for extracting from a larger tree (or from an unstructured collection of proposed divisions) a subset $\mathcal{D}^*$ whose members can be treated as orthogonal for mathematical purposes.

Until that apparatus is in place, the binary tree stands as a powerful but incomplete instrument. It demonstrates that repeated binomial division can systematically enlarge a coordinate system, and it supplies an explicit map from conceptual operations to sign sequences. At the same time it leaves open the question of how the hierarchical dependencies internal to the tree are to be relaxed or controlled so that the resulting dimensions approximate the independence required by linear algebra and tensor calculus. The next section begins the construction of the evaluative machinery that addresses precisely this question.

III. From Binomial Divisions to Mathematical Dimensions

3.1 A Distinction as an Axis

The transition from ideonomic division to mathematical dimension begins with a decisive reconception of the poles. In the complementary and contrasting readings examined earlier, the poles A_i and B_i of a distinction $D_i = (A_i, B_i)$ function primarily as classes or semantic attractors. An entity is assigned to one pole, to the other, or, in the contrasting case, is recognized as mixed or unclassifiable. The orthogonal reading already weakens this classificatory demand by treating the poles as endpoints of an axis.

The present step completes the transformation: category assignment is replaced by continuous coordinate assignment. For each distinction D_i one introduces a real-valued coordinate $x_i \in [-1, 1]$ such that $x_i = -1$ registers maximal association with pole A_i and $x_i = +1$ registers maximal association with pole B_i . Intermediate values represent graded degrees of association, indifference, or mixture according to the semantic character of the particular distinction. The poles cease to be rival bins and become the extreme markers of a one-dimensional continuum.

Once this coordinate interpretation is in place, a finite collection of binomial distinctions immediately yields a finite-dimensional representation of any entity. Given n distinctions $\{D_1, \dots, D_n\}$, an entity is mapped to the vector

$$\mathbf{x} = (x_1, \dots, x_n) \in \mathbb{R}^n,$$

where each component x_i is the entity's coordinate on the corresponding axis. The integer n is thereby promoted from a mere count of distinctions to the dimension of a conceptual coordinate space. Each binomial distinction, under the continuous reading, contributes one conceptual dimension; the selected set of distinctions becomes a basis (in the informal sense) for locating entities relative to the ideonomic system. The representation is explicit, inspectable, and, in principle, comparable across entities and across occasions of measurement.

The mapping from entity to coordinate is effected by a scoring function f_i . Formally, f_i assigns to each entity a value in the closed interval $[-1, 1]$ that is interpreted as the coordinate x_i . The concrete realization of f_i may take several forms. It may be an annotator rating: human judges, presented with an entity and the polar descriptors A_i and B_i , locate the entity on a numerical scale whose endpoints are the poles. It may be the output of a supervised classifier trained to predict polar association from features of the entity. It may be a model-derived probability or logit, suitably rescaled to the interval $[-1, 1]$, obtained from a language model or other predictive system prompted or fine-tuned to assess the relevant distinction. Hybrid procedures are possible: an initial model score may be calibrated against a smaller set of human ratings, or multiple annotators may be aggregated by averaging or by more robust estimators. What matters for the formal framework is not the provenance of the score but the existence of a reproducible function that lands in the target interval and that can be applied uniformly to the entities under study.

This reconception of distinctions as continuous axes constitutes one of the strongest constructive claims of the present reconstruction. It asserts that the generative output of ideonomy—the binomial division—can be lifted, without residual classificatory commitment, into the setting of real coordinate geometry. The claim is strong because it converts an essentially qualitative operation (the proposal of opposing poles) into a quantitative representation that inherits the algebraic structure of $\mathbb{R}^n$. Once entities are points in that space, distances, angles, linear projections, and, later, tensorial constructions become available as analytic instruments. The same claim is also strong because it supplies a uniform interface between conceptual work and statistical or machine-learning work: any method capable of producing a score in $[-1, 1]$ can serve as an implementation of f_i , and any method capable of operating on vectors in $\mathbb{R}^n$ can operate on the resulting ideonomic representations.

The claim rests on several assumptions that must be stated with equal explicitness. First, it assumes that intermediate degrees of association are semantically intelligible for the distinctions in question. Not every opposition admits a natural continuum; some complementary divisions may be genuinely binary, and forcing intermediate values upon them may introduce artifactual precision. The framework therefore applies most cleanly to distinctions already regarded as orthogonal or contrasting, and only cautiously to those regarded as strictly complementary. Second, it assumes that a scoring function f_i of acceptable reliability can be obtained. If human annotators disagree sharply, or if model probabilities prove unstable under prompt variation, the numerical coordinates become noisy and the subsequent geometric constructions inherit that noise. Third, it assumes that the collection of axes can be treated, at least provisionally, as a coordinate system (i.e., that the distinctions are sufficiently independent for the vector representation to be meaningful). Full linear independence is not required at this stage, but gross redundancy would collapse the effective dimension and undermine the interpretation of $\mathbf{x}$ as a location in an n -dimensional space. The measurement of redundancy R_{ij} and the selection of a non-redundant subset $\mathcal{D}^*$ are therefore indispensable sequels to the coordinate construction; they are taken up in the sections that follow.

A further assumption concerns the relation between hierarchical and axial readings. As noted earlier, recursive division naturally produces nested trees whose path coordinates are conditional rather than independent. The continuous-axis construction abstracts from that nesting. It treats each distinction as an axis that can, in principle, receive a coordinate regardless of the values assigned on other axes. When the underlying generative process remains hierarchical, this abstraction is an idealization: the coordinate x_j on a lower-level distinction may be undefined or irrelevant for entities that fall outside the relevant upper-level pole. The idealization is nevertheless indispensable if the goal is to obtain a representation compatible with the linear-algebraic and tensorial machinery deployed later. Section III as a whole is devoted to managing the resulting tension: first by adopting the continuous-axis reading, then by introducing explicit criteria of non-redundancy and stability that allow a subset of distinctions to be promoted to approximately independent dimensions.

Under these assumptions the coordinate construction stands as a bridge. On one side lies the ideonomic practice of proposing and refining binomial divisions; on the other lies the mathematical practice of treating entities as points in a structured vector space. The scoring function f_i is the span of the bridge: it converts qualitative polar association into quantitative location. Once the bridge is crossed, the full resources of linear algebra become available for validation, and the resources of tensor calculus become available for

the modeling of interactions among dimensions. The remainder of this section develops the validation step (the assessment of redundancy, meaningfulness, operationalizability, and stability) and thereby converts a raw collection of scored distinctions into a selected set $\mathcal{D}^*$ of ideonomic dimensions suitable for subsequent mathematical and experimental work.

The next section introduces the formal measures of redundancy and independence that determine which candidate axes survive into the selected set $\mathcal{D}^*$.

3.2 Binary Distinctions Versus Continuous Dimensions

A persistent misunderstanding of binomial division equates the introduction of two poles with the imposition of a binary ontology. The misunderstanding is natural but consequential. Because the generative act isolates a pair of opposing markers, it is easy to suppose that the world itself must be partitioned into two exclusive and exhaustive classes. The continuous-axis construction developed in the preceding section dissolves this inference. A binary distinction identifies two poles; it does not force the phenomena under study into two bins. The poles function as reference extremes that orient a continuum, not as ontological ultimatums that exhaust the space of possible being.

The distinction between a binary conceptual instrument and a binary ontology can be stated with precision. Ontology, in the relevant sense, concerns what kinds of entities or properties fundamentally exist. A binary ontology would assert that reality, or some significant portion of it, consists of two mutually exclusive and jointly exhaustive categories—that every relevant item is determinately of one kind or the other, with no remainder and no mixture. A binomial distinction, by contrast, is a methodological proposal: it advances two polar descriptors as the endpoints of an axis along which entities may be located. The axis may be sparsely populated, densely graded, or even empty in intermediate regions; none of these empirical distributions alters the status of the distinction as a coordinate rather than a metaphysical decree. The same distinction can be retained whether the observed distribution is bimodal, unimodal, uniform, or irregular. What is required is only that the poles remain intelligible as extremes of a coherent respect of variation.

A concrete illustration is supplied by the opposition order $\leftrightarrow$ disorder. Let $D_{\text{order}} = (A, B)$ where A is maximal order and B is maximal disorder. Under continuous reading the associated coordinate is a real number $x_{\text{order}} \in [-1, 1]$, with -1 registering maximal association with order and $+1$ maximal association with disorder. Nothing in this construction requires that every system be classified as either ordered or disordered. A crystal near absolute zero may receive a score near -1 ; a turbulent fluid or a

high-temperature plasma may receive a score near $+1$; a partially annealed alloy, a living cell, or a social institution under moderate stress may receive intermediate values that reflect graded or domain-relative degrees of organizational coherence. The score may itself be probabilistic, representing an expected value under a distribution over micro-states, or it may be a deterministic assessment relative to a chosen scale of description. In all such cases the binary poles remain in place as reference points, while the ontology of the systems under study is left free to be as continuous, hierarchical, or multi-faceted as independent evidence indicates.

The same logic extends immediately to the multi-dimensional case. An entity is represented by a vector $\mathbf{x} = (x_1, \dots, x_n) \in \mathbb{R}^n$, each component of which is the entity's coordinate on a distinct ideonomic axis. Because the axes are coordinates rather than exclusive class labels, an entity can occupy intermediate or polar positions on many dimensions simultaneously without contradiction. A single physical sample may be high on an order–disorder axis, intermediate on a discrete–continuous modeling axis, and low on an individual–collective causal axis. A legal text may be polar on a rigid–flexible interpretive axis while remaining intermediate on a formal–substantive axis and polar on a retrospective–prospective temporal axis. The vector representation records these joint locations without forcing any one distinction to absorb or eliminate the others. Rigid dichotomous classification is thereby escaped not by abandoning binary instruments, but by refusing to treat those instruments as ontological partitions and by embedding them instead within a higher-dimensional coordinate geometry.

This escape is not merely pragmatic; it is conceptual. Once distinctions are axes, the logical space of possible descriptions expands from a discrete set of class assignments to a continuum of vector positions. The cardinality of the descriptive possibilities grows from 2^n (the number of vertices of the hypercube) to the continuum of the solid cube $[-1, 1]^n$. Intermediate positions, mixtures, and graded memberships cease to be anomalies that must be explained away or forced into residual categories; they become ordinary points in the ambient space. At the same time the original binary poles retain their normative and interpretive force: they continue to supply the meanings of the extreme values and thereby keep the coordinates semantically anchored. The system therefore preserves the generative clarity of binomial division while discarding the ontological rigidity that a naïve reading would impose.

Two further consequences follow. First, the continuous multi-dimensional representation supplies a natural setting in which questions of independence and redundancy can be posed quantitatively. If two axes are highly correlated across a population of entities, the corresponding coordinates convey overlapping information; the redundancy measure R_{ij}

becomes well-defined, and decisions about retention or elimination of dimensions can be evidence-based rather than purely intuitive. Second, the same representation prepares the ground for interaction modeling. Once entities are points in $\mathbb{R}^n$, the joint variation of coordinates can be studied by the methods of multivariate statistics and, at a further stage of abstraction, by the constructions of tensor calculus. Neither step would be available if the distinctions remained locked in a dichotomous classificatory mode.

The argument of this section may be summarized in a single contrast. Binary ontology asserts that the world (or a domain) divides into two exclusive and exhaustive kinds. Binary distinction asserts only that two poles have been identified as reference extremes of a respect of variation. The continuous coordinate $x_i \in [-1, 1]$ and the vector $\mathbf{x} \in \mathbb{R}^n$ implement the second assertion while remaining neutral toward the first. In so doing they convert the ideonomic instrument into a mathematical dimension without smuggling in an unexamined metaphysics of dichotomy.

The next section turns from the justification of continuous coordinates to the formal assessment of redundancy among them, introducing quantitative criteria that determine which candidate axes are retained in the selected set $\mathcal{D}^*$.

3.3 Independence, Orthogonality, and the Ideonomic Dimensionality Criterion

Not every conceptually intelligible distinction deserves elevation to the status of a computational dimension. The generative phase of ideonomy can produce a large candidate set $\mathcal{D} = \{D_1, \dots, D_n\}$ of binomial divisions, each equipped with a scoring function f_i that maps entities to coordinates $x_i \in [-1, 1]$. If every such division were retained without further scrutiny, the resulting representation $\mathbf{x} \in \mathbb{R}^n$ would rapidly become redundant, unstable, and difficult to interpret. The transition from candidate distinctions to selected dimensions therefore requires explicit criteria of independence. Only those distinctions that vary with sufficient autonomy relative to the others can contribute new axes to the coordinate system without merely duplicating information already captured elsewhere.

Independence for ideonomic purposes is defined relationally. Two distinctions $D_i = (A_i, B_i)$ and $D_j = (A_j, B_j)$ are said to be orthogonal, written $D_i \perp D_j$, when variation along one does not necessarily determine variation along the other. The modal character of the definition is deliberate. Perfect statistical independence across every conceivable population is neither expected nor required; what is required is that there exist admissible ways of varying the coordinate x_i while holding x_j fixed, and vice versa. When this condition holds, the two axes span a genuine two-dimensional region rather than collapsing onto a lower-dimensional manifold. Conversely, when variation on one axis is

largely predictable from variation on the other, the pair is redundant: the second distinction adds little descriptive power once the first is known. The quantity R_{ij} is introduced as a generic measure of this redundancy or dependency; its precise numerical form will depend on the estimation method adopted, but its conceptual role is uniform.

Several families of estimation methods are available once entities have been scored on both axes. The simplest is linear correlation: the Pearson coefficient (or a robust alternative) computed between the paired coordinates (x_i, x_j) across a suitable sample. High absolute correlation signals redundancy; near-zero correlation is consistent with orthogonality, though it does not entail it. Mutual information supplies a more general, non-linear alternative. By estimating the reduction in uncertainty about x_i afforded by knowledge of x_j , mutual information detects dependencies that correlation may miss. Conditional dependence measures refine the analysis further: one may examine whether x_i remains informative about an external target once x_j is controlled for, or whether the conditional distribution of one coordinate given the other is effectively degenerate. Principal-angle analysis, applied to subspaces spanned by groups of axes, extends the assessment from pairs to higher-order collections and is especially useful when the candidate set is large. Finally, representational similarity measures (comparing the geometry of distances or kernels induced by different subsets of coordinates) provide a global diagnostic of whether two distinctions induce approximately the same relational structure among entities. Each method has characteristic strengths and limitations; in practice a battery of complementary estimators is preferable to any single statistic.

These estimation procedures must be applied with awareness of the hierarchical-dependence problem identified in the recursive construction of binomial trees. In a nested tree, lower-level distinctions are defined only conditionally upon the poles of higher-level distinctions. The coordinate associated with a subordinate division is therefore not free to vary independently of the superordinate coordinate; certain combinations are undefined or semantically incoherent by construction. Path coordinates derived from such trees are consequently hierarchical rather than axial. When the continuous-coordinate representation is imposed upon a set of distinctions that originated in nested fashion, residual conditional dependencies are likely to appear in the estimated R_{ij} values. The hierarchical origin does not automatically disqualify a distinction, but it does raise the evidentiary bar: a subordinate distinction earns independent dimensional status only if, across the population of interest, its coordinate continues to exhibit substantial residual variation after the superordinate coordinates are considered. In short, the statistical criteria of independence serve as a filter that can promote a conditionally generated

distinction into an approximately orthogonal axis, or that can demote it back to the status of a nested refinement.

The cumulative force of these considerations is expressed in the ideonomic dimensionality criterion. A conceptual division D_i becomes a useful computational dimension when it is sufficiently meaningful, sufficiently operationalizable, and sufficiently non-redundant relative to the other divisions under consideration. Meaningfulness (M_i) concerns the semantic clarity and theoretical or practical significance of the polar contrast; a distinction whose poles are obscure or whose axis organizes nothing of interest fails this test.

Operationalizability (O_i) concerns the existence and reliability of a scoring function f_i ; a distinction that cannot be mapped reproducibly onto $[-1, 1]$ remains a purely verbal proposal. Non-redundancy (N_i) is the complement of aggregate dependency: it rises as the collection of pairwise (and higher-order) redundancies R_{ij} involving D_i falls. A fourth component, stability (S_i), will be examined in due course; it concerns robustness under changes of sample, scoring method, or minor reformulation of the poles. Together these components determine whether a candidate survives into the selected set $\mathcal{D}^*$ of size k .

The criterion is deliberately conjunctive. High meaningfulness alone does not suffice if the distinction cannot be scored or if it merely recapitulates axes already present. High operationalizability is of limited value if the resulting coordinates are almost linearly dependent on existing ones. Non-redundancy without meaningfulness yields statistically independent but conceptually empty dimensions. Only the joint satisfaction of the three primary demands—later augmented by stability—warrants promotion from candidate distinction to computational dimension. The criterion thereby converts the open-ended generative output of ideonomy into a disciplined, evidence-sensitive selection procedure. It does not claim that the surviving axes are metaphysically privileged; it claims only that they meet explicit thresholds of clarity, measurability, and informational novelty relative to one another.

In operational terms the criterion is implemented by first scoring a broad candidate set, estimating the matrix of redundancies R_{ij} , evaluating the component scores M_i , O_i , and N_i (and later S_i), and then retaining those distinctions whose composite utility $\Phi(D_i)$ exceeds a chosen threshold or ranks within the top k . The precise weighting of components and the choice of threshold remain matters of methodological judgment informed by the aims of the inquiry. What is fixed is the architecture: generation produces candidates; measurement of meaning, operationalizability, and independence filters them; the survivors become the axes of the ideonomic coordinate system.

The hierarchical-dependence problem is managed rather than eliminated by this architecture. Distinctions that remain stubbornly conditional after residualization are simply assigned low non-redundancy scores and are excluded from $\mathcal{D}^*$. Distinctions that, despite nested origins, exhibit substantial independent variation are retained and thereafter treated as ordinary axial coordinates. The continuous multi-dimensional representation introduced earlier supplies the common space in which these statistical decisions are executed. In this way the tension between the generative tree and the independent axial system is resolved pragmatically: the tree remains a valuable source of candidate distinctions, while the dimensionality criterion determines which of those candidates function as free dimensions for subsequent mathematical and computational work.

The next section completes the evaluative apparatus by formalizing the composite utility score $\Phi(D_i)$ and specifying the selection rule that extracts the final set $\mathcal{D}^*$ from the scored and residualized candidates.

IV. An Ideonomic Basis for Tensor Calculus

4.1 Tensor Dimension and Ideonomic Tensor Construction

Tensors furnish the natural mathematical setting in which multiple dimensions interact. In their most elementary presentation they appear as multidimensional arrays: a scalar is a tensor of order zero, a vector a tensor of order one, a matrix a tensor of order two, and so on. Abstractly, a tensor T of order k is an element of a tensor product of vector spaces,

$$T \in V_1 \otimes V_2 \otimes \cdots \otimes V_k.$$

Each factor space V_i contributes an index whose range is the dimension of that space; the components of T encode multilinear relations among the indices. The algebraic operations of contraction, outer product, and change of basis then become available for the systematic study of how variation along one collection of axes covaries with variation along others. The question that organizes the present section is whether the ideonomic distinctions already developed can guide the selection or the interpretation of these factor spaces.

The construction begins from the apparatus assembled in preceding sections. Let $\mathcal{D} = \{D_1, \dots, D_n\}$ be a candidate set of binomial distinctions, each $D_i = (A_i, B_i)$. For a given entity or textual unit X the scoring functions supply coordinates

$$x_i = f_i(X) \in [-1, 1],$$

and the entity is thereby represented by the ideonomic coordinate vector

$$\mathbf{x} = (x_1, \dots, x_n) \in \mathbb{R}^n.$$

When a subset $\mathcal{D}^*$ of size k has been selected according to the ideonomic dimensionality criterion, the corresponding coordinates span a k -dimensional space. Each selected distinction $D_i \in \mathcal{D}^*$ may now be associated with a one-dimensional factor space V_{D_i} , which may be identified with $\mathbb{R}$ (or with a suitable subspace) whose ordered basis is oriented from the negative pole A_i to the positive pole B_i . The ideonomic tensor is then the element

$$T \in V_{D_1} \otimes \dots \otimes V_{D_k}$$

whose components record measured interactions among the selected dimensions.

The passage from coordinate vectors to tensors is motivated by the recognition that pairwise independence, however carefully enforced, does not exhaust the relational structure of a conceptual system. Two axes may be approximately orthogonal in the sense that R_{ij} is small, yet their joint variation with a third axis may exhibit systematic higher-order dependence. A matrix of pairwise correlations captures only second-order structure; a tensor of order three or higher can register the residual interactions that remain after lower-order effects have been accounted for. In this respect the ideonomic tensor functions as a quantitative repository of the multilinear patterns that the generative and selective phases of ideonomy have made available for measurement.

A concrete illustration is supplied by the third-order tensor whose modes correspond to the distinctions individual/collective, static/dynamic, and concrete/abstract. Denote the associated coordinates by x_{ind} , x_{stat} , and x_{conc} , each ranging over $[-1, 1]$. The component T_{ijk} is defined by reference to a corpus of entities—sentences, paragraphs, documents, or other suitably delimited units—each of which has been scored on the three axes. One natural filling rule proceeds by discretization or kernel density estimation. The interval $[-1, 1]$ on each axis is partitioned into bins (or smoothed by a continuous kernel); the value of T_{ijk} is then proportional to the number of corpus units whose coordinates fall simultaneously into the i -th bin of the first axis, the j -th bin of the second, and the k -th bin of the third, possibly normalized by the product of the marginal frequencies so that the tensor registers excess co-occurrence relative to independence. Alternative fillings are possible: T_{ijk} may encode the average value of an external response variable conditional on the three coordinates, or the strength of a predictive relationship among the three scores themselves. What is essential is that the tensor components are empirically determined quantities, not stipulative assignments. They measure how the three ideonomic dimensions co-vary across the chosen corpus.

The same construction extends immediately to tensors of higher order and to selections of different size. Once the factor spaces have been identified with the selected distinctions, any multilinear statistic computable from the joint distribution of the corresponding coordinates can serve as a component of T . The ideonomic contribution lies in the principled choice and semantic interpretation of the modes: each index of the tensor corresponds to a distinction that has already satisfied the criteria of meaningfulness, operationalizability, and non-redundancy. The tensor does not invent new dimensions; it organizes the interactions among dimensions that the earlier phases have certified as worthy of retention.

Several clarifications protect the construction from over-interpretation. First, the factor spaces V_{D_i} are deliberately minimal—one-dimensional real spaces oriented by the poles. Nothing prevents subsequent refinement in which a single ideonomic distinction is itself decomposed into a higher-dimensional subspace, but the basic construction does not require such refinement. Second, the empirical filling of T is relative to a corpus and to a choice of scoring functions. Different corpora or different operationalizations of the f_i will yield different numerical tensors; the ideonomic claim is not that a unique tensor exists, but that the modes of any such tensor can be given a transparent conceptual interpretation grounded in the selected distinctions. Third, the construction remains neutral toward the ontological status of the interactions it records. The tensor may be read as revealing real higher-order structure in the domain, or as summarizing patterns of human or model judgment, or as a pragmatic tool for prediction and control. The mathematical object itself does not decide among these readings.

The link to the continuous-axis representation developed earlier is direct. Because each entity is already a point $\mathbf{x}$ in the ideonomic coordinate space, the joint distribution of any subset of coordinates is well-defined once a population of entities is fixed. The tensor simply organizes the higher-order moments or co-occurrence statistics of that distribution according to the modes supplied by $\mathcal{D}^*$. In this sense the ideonomic tensor is not an independent invention but the multilinear completion of the coordinate geometry already in place. Where the vector $\mathbf{x}$ locates an individual entity, the tensor T characterizes the relational architecture of the entire ensemble with respect to the chosen dimensions.

The hierarchical residue of recursive division continues to require attention. If some of the selected distinctions originated as nested refinements, residual conditional dependencies may still appear in the higher-order components of T . The tensor does not erase those dependencies; it makes them visible as systematic patterns of co-variation that deviate from the independence baseline. Analysts may then decide whether to retain the implicated modes, to residualize them further, or to interpret the observed interactions as

theoretically informative rather than as statistical artifacts. The tensor thereby supplies both a diagnostic of remaining non-independence and a positive instrument for the study of genuine multilinear conceptual structure.

In summary, the ideonomic approach to tensor construction proceeds by aligning factor spaces with selected binomial distinctions, scoring entities on the corresponding axes, and filling the resulting multilinear array with measured co-variation statistics drawn from a corpus. The third-order example involving individual/collective, static/dynamic, and concrete/abstract axes illustrates the procedure without exhausting its range. The same template applies to any subset of $\mathcal{D}^*$ and to any order of interaction one wishes to examine. What distinguishes the construction is the prior ideonomic labor that has rendered each mode semantically transparent and statistically non-redundant. The tensor inherits that transparency and that selectivity; it does not replace them.

The next section examines how such ideonomic tensors may be manipulated (through contraction, decomposition, and projection) in order to extract interpretable lower-order structure and to interface with the latent representations of large language models.

4.2 Tensor Products as Combinations of Distinctions

The algebraic structure of the tensor product supplies a precise language for the combination of distinctions. Given two factor spaces V_A and V_B associated with distinctions D_A and D_B , their tensor product $V_A \otimes V_B$ is the space of bilinear maps from the dual pair into the scalar field, or, equivalently, the space spanned by formal products of basis vectors drawn one from each factor. Crucially, $V_A \otimes V_B$ does not merely juxtapose the properties already present in V_A and V_B . It contains, in addition, the full set of interactions between those properties. A pure tensor $v \otimes w$ encodes a separable combination; a general element of the product encodes inseparable correlations, interferences, and higher-order dependencies that cannot be reduced to independent contributions from each factor. The passage from the Cartesian product of the underlying sets to the tensor product of the vector spaces is therefore a passage from separate listings to structured relational possibility.

This algebraic fact licenses an analogical reading within the ideonomic framework. Just as $V_A \otimes V_B$ comprises both the individual characters of the two spaces and the ways those characters can jointly vary, the conceptual combination of two distinctions may be understood as generating not only the poles and intermediate values of each distinction taken separately, but also the structured space of their mutual determinations. The expression

$$\text{distinction}_1 \otimes \text{distinction}_2$$

is offered as a compact notation for that generated space. It does not claim that conceptual combination is literally a tensor product; it claims that the tensor product furnishes a mathematically controlled model of the combinatorial surplus that arises when distinctions are allowed to interact. The model is faithful to the extent that the higher-order components of an ideonomic tensor $T \in V_{D_1} \otimes \cdots \otimes V_{D_k}$ capture empirically observable patterns of co-variation that are invisible at the level of single axes or pairwise correlations alone.

The reading resonates with the combinatorial ambition that characterizes ideonomy as a general science of ideas. Ideonomy does not merely inventory isolated distinctions; it seeks to explore the systematic possibilities that arise when distinctions are composed, refined, and cross-applied. The tensor product supplies one formal image of that compositional process. Each new factor space adjoined to the product enlarges the space of admissible multilinear relations, just as each new ideonomic distinction, once admitted into $\mathcal{D}^*$, enlarges the space of conceptual configurations that can be located and compared. The analogy is therefore not ornamental. It translates the generative impulse of ideonomy into a setting where the results of generation can be measured, decomposed, and manipulated by the standard operations of multilinear algebra.

The proposed principle may now be stated directly. Ideonomic tensor spaces model not only what distinctions exist, but how distinctions interact. The first half of the principle is secured by the earlier construction: each mode of the tensor corresponds to a distinction that has already satisfied the ideonomic dimensionality criterion and has therefore been retained in $\mathcal{D}^*$. The second half is secured by the multilinear character of the tensor product itself. Once the factor spaces are fixed, the components of T become available as quantitative records of interaction: records that can be inspected for excess co-occurrence, conditional dependence, or predictive synergy. In this way the tensor does not merely list the selected distinctions; it organizes the relational architecture that those distinctions jointly determine across a corpus of scored entities.

It is essential to mark the status of this principle with precision. The tensor-product reading is an analogy and a modeling choice, not a derivation. Nothing in the internal logic of binomial division entails that conceptual combination must be represented by a tensor product, nor does the algebraic definition of $\otimes$ entail any thesis about the ontology of ideas. The justification for the reading is pragmatic and structural. Pragmatically, the tensor product offers a mature, computationally tractable formalism in which higher-order co-variation can be stored, decomposed (by CP, Tucker, or related methods), and contracted against vectors or lower-order tensors. Structurally, it preserves the distinction between separable and entangled contributions, thereby mirroring the conceptual difference

between independent polar associations and genuine interaction effects. These advantages are sufficient to warrant the modeling choice; they do not convert the choice into a necessary truth.

Several consequences follow once the principle is adopted. First, the ideonomic tensor becomes a natural site for the detection of conceptual synergy. If the three-way interaction term among a particular triple of distinctions is large relative to the product of the marginals, the data indicate that those distinctions do not merely coexist but modulate one another. Such synergy may be theoretically illuminating even when the pairwise redundancies R_{ij} remain modest. Second, contraction of the tensor against a fixed vector on one or more modes yields lower-order objects that inherit semantic interpretability from the contracted modes. Contracting onto a particular polar extreme, for example, produces a residual tensor or matrix whose remaining modes describe how the other distinctions behave when the contracted distinction is held at that extreme. Third, the same framework prepares an interface with latent representations of large language models. A projection matrix W that maps a hidden state $\mathbf{h}$ onto ideonomic coordinates can be viewed as extracting the first-order marginals of an implicit higher-order structure; the full ideonomic tensor supplies the natural target for any subsequent attempt to recover or regularize higher-order interactions among those coordinates.

The hierarchical residue of recursive division continues to appear inside the tensor product, yet it appears in a manageable form. Conditional dependencies that survive residualization manifest as systematic patterns of higher-order co-variation rather than as undefined entries. Analysts may therefore choose to interpret those patterns as theoretically significant interactions or to apply further residualization or mode-pruning before downstream use. In either case the tensor product does not conceal the hierarchy; it renders the residual dependence visible as multilinear structure.

The cumulative effect of the construction is to convert the generative and selective phases of ideonomy into a multilinear observational frame. Distinctions that have been proposed, scored, and filtered become the modes of a tensor; the empirical co-variation of their coordinates becomes the content of that tensor; and the algebraic operations native to the tensor product become instruments for exploring how the distinctions jointly organize a domain. The principle that ideonomic tensor spaces model both the existence and the interaction of distinctions is the compact statement of this conversion. It remains an interpretive and methodological commitment rather than a formal derivation, yet it is a commitment that renders the combinatorial ambition of ideonomy computationally and statistically tractable.

The next section turns from the construction and interpretation of ideonomic tensors to their systematic manipulation (contraction, decomposition, and projection) and to the concrete prospects these operations open for the analysis of large-language-model representations.

4.3 From Binary Trees to Tensor Networks

The recursive application of binomial division produces a binary tree whose paths yield hierarchical coordinate sequences. The same recursive architecture finds a formal counterpart in a family of multilinear decompositions that represent high-order tensors by hierarchically organized lower-order factors. Hierarchical tensors, tensor-train (or matrix-product) decompositions, and tree tensor networks all exploit successive low-rank factorizations along a tree or path topology in order to combat the exponential growth of parameters that accompanies a full tensor product. The structural parallel is immediate: just as each successive binomial division refines a polar region produced by prior divisions, each successive core in a hierarchical tensor factorization refines the multilinear relations encoded by the cores above it. The present section develops this parallel as an explicit analogy, marks its speculative status, and identifies the empirical and theoretical conditions under which the analogy could be elevated to a working methodological correspondence.

Begin with the ideonomic side. A recursive division process starts with a root domain X and repeatedly replaces a selected pole by a new binomial pair. The resulting tree encodes conditional dependencies by construction: a lower-level distinction is defined only inside the region carved out by its ancestors. When the continuous-coordinate representation is imposed, these conditional dependencies appear as systematic patterns of higher-order co-variation inside the ideonomic tensor $T \in V_{D_1} \otimes \cdots \otimes V_{D_k}$. The full tensor still treats every mode as formally free, yet the numerical content of T retains the imprint of the generative hierarchy. A natural question is whether that imprint can be recovered, compressed, or exploited by a factorization whose own topology mirrors the original tree.

Hierarchical tensor formats supply one affirmative answer. In a hierarchical Tucker representation, for example, the ambient tensor is expressed by a collection of nested transfer tensors arranged according to a dimension tree; each node of the tree corresponds to a subspace that aggregates a subset of the original modes. Tensor-train decompositions linearize the same idea into a chain of third-order cores, while tree tensor networks allow more general branching topologies. In each case the storage cost scales linearly with order once the ranks are fixed, and contraction against vectors or matrices can be performed by successive local operations that follow the tree edges. The computational

advantages are well-documented; the conceptual interest for ideonomy lies in the shared recursive skeleton.

The analogy may be stated cleanly. The recursive binomial architecture generates a nested system of distinctions whose natural index structure is a binary tree; hierarchical factorization represents a high-order tensor by a nested system of low-rank cores whose natural index structure is likewise a tree. In both settings the global object (whether a conceptual coordinate system or a numerical array) is recovered by composing local binary (or low-arity) operations along the edges of the tree. The analogy therefore maps generative refinement of poles onto successive low-rank approximation of multilinear maps. Under this mapping, an ideonomic path coordinate becomes analogous to a sequence of core contractions, and residual hierarchical dependence inside T becomes analogous to the rank structure that the factorization is designed to capture.

The parallel is promising for several reasons. First, it offers a principled route from the generative phase of ideonomy to a compressed computational representation. Rather than materializing the full k -mode tensor and only afterward discovering its low-rank structure, one could in principle let the original division tree dictate the topology of the tensor network, thereby building compression into the representation from the start. Second, the analogy supplies a diagnostic for the hierarchical residue discussed in earlier sections. If a tensor-train or tree-network approximation of an ideonomic tensor achieves low error with ranks that respect the original nesting, the numerical evidence corroborates the claim that the dominant interactions remain hierarchically organized. Conversely, if low-rank hierarchical approximation fails while a more entangled factorization succeeds, the data indicate that the selected distinctions have escaped their generative hierarchy and now interact in a flatter, more orthogonal fashion. Third, the same network structure interfaces naturally with the contraction and projection operations already contemplated for linking ideonomic coordinates to latent states of language models: successive core contractions can implement the map from a high-dimensional hidden representation onto a hierarchically organized set of ideonomic scores.

Nevertheless, the analogy remains speculative. Several non-trivial steps would have to be secured before it could be treated as more than a suggestive structural correspondence. Empirically, one would need to demonstrate that ideonomic tensors constructed from real corpora of scored entities do in fact admit accurate hierarchical low-rank approximations whose dimension trees align with the generative division trees that produced the distinctions. Theoretically, one would need a precise translation between the semantic notion of conditional polar refinement and the algebraic notion of successive subspace nesting; without such a translation the ranks and cores remain purely numerical objects

stripped of ideonomic interpretation. Algorithmically, one would need procedures that jointly optimize the choice of distinctions, the scoring functions, and the tensor-network topology rather than treating the network as a post-hoc compression of an already fixed tensor. Until these demonstrations are in hand, the recursive-binomial / hierarchical-tensor parallel functions as a modeling heuristic: it guides the search for efficient representations and supplies a vocabulary for discussing residual hierarchy, yet it does not constitute a derived equivalence.

Even as a heuristic the parallel is already productive. It reframes the tension between hierarchical path coordinates and independent axial coordinates as a question of rank structure inside an ideonomic tensor. Perfect orthogonality among the selected distinctions would correspond to a tensor that factors into a pure outer product of vectors (one vector per mode) while strong residual nesting would correspond to a tensor whose optimal network topology recovers the original division tree. Intermediate cases, which are the cases of greatest practical interest, correspond to networks of modest rank that capture the dominant conditional interactions while still permitting approximately free variation along each retained axis. The ideonomic dimensionality criterion developed earlier can therefore be read, in this light, as a rank-control device: by penalizing redundancy one encourages the tensor toward lower interaction ranks, thereby easing subsequent hierarchical compression.

The same perspective also clarifies the status of the continuous-coordinate idealization. When lower-level distinctions remain strictly conditional, the corresponding modes of the full tensor product contain large regions of structural zeros or undefined values; a hierarchical network whose topology matches the generative tree automatically respects those regions by never attempting to form the forbidden combinations. In this limited but important sense the hierarchical factorization is more faithful to the generative process than the flat tensor product. The price of that fidelity is reduced flexibility: once the topology is fixed, certain cross-branch interactions become more expensive to represent. The modeling choice between a flat ideonomic tensor and a tree-structured network is therefore a choice between two different compromises with the hierarchical residue: one that privileges free axial combination, the other that privileges generative conditional structure.

In sum, the recursive architecture of binomial division and the hierarchical architecture of tensor-train and tree tensor networks share a common pattern of successive local refinement. The shared pattern licenses an analogy that is computationally suggestive and diagnostically useful, yet still requires empirical validation, theoretical translation, and joint optimization procedures before it can be promoted beyond the status of a promising

modeling heuristic. The analogy does not dissolve the earlier tension between hierarchical and independent readings; it relocates that tension inside the rank structure of an ideonomic tensor and thereby makes it accessible to the quantitative tools of multilinear approximation.

The next section leaves the internal architecture of ideonomic tensors and turns to their external interface with the latent spaces of large language models, examining how projection, contraction, and network evaluation can convert hidden representations into interpretable ideonomic coordinates and interaction profiles.

V. Large-Language Models and Emergent Dimensions

5.1 Latent Representation and the Emergent/Designed Contrast

Contemporary large language models encode linguistic and conceptual information in high-dimensional hidden states $\mathbf{h} \in \mathbb{R}^d$. These vectors are the product of successive non-linear transformations applied to token sequences; by the later layers of a deep network the original discrete symbols have been dissolved into a continuous geometry whose metric relations approximate a wide range of semantic, syntactic, and pragmatic distinctions. The representational capacity of such spaces is extraordinary: distances, directions, and linear subspaces have been shown to support analogical reasoning, syntactic reconfiguration, and even rudimentary world-model updates. Yet the individual coordinates of $\mathbf{h}$ are rarely interpretable in isolation. Meaning is distributed across thousands of dimensions, and the contribution of any single axis is typically entangled with many others. The result is a tension that has become central to the interpretability literature: the same distributed geometry that confers high capacity simultaneously resists direct human inspection and control.

Two distinct pipelines for obtaining dimensional structure may be contrasted. The emergent pipeline begins with data. A large corpus is used to optimize a parametric model, ordinarily by next-token prediction or a related self-supervised objective. The optimization process discovers latent dimensions as a byproduct of minimizing loss; the resulting coordinates are those that proved useful for the training task, not those that were stipulated in advance by human conceptual analysis. Post-hoc methods (probing classifiers, sparse auto-encoders, causal tracing, or representation similarity analysis) may later attempt to recover approximate semantic axes, but the axes themselves remain emergent. Their existence and alignment with human concepts are empirical discoveries rather than design choices.

The ideonomic pipeline proceeds in the opposite direction. It begins with distinctions. Binomial divisions are proposed, scored, filtered by the ideonomic dimensionality criterion,

and assembled into an explicit schema $\mathcal{D}^*$. Entities or textual units are then mapped onto the corresponding coordinates $x_i = f_i(\cdot)$, yielding an ideonomic vector $\mathbf{x}$ or, after linear projection of a hidden state, a vector $\mathbf{z} = W\mathbf{h}$. In this pipeline the dimensions are designed first and the representation is constructed to realize them. Interpretability is obtained by construction: each axis carries a publicly inspectable polar semantics, and the scoring functions f_i can be audited, recalibrated, or replaced.

The contrast is not merely chronological. Emergent dimensions maximize task performance under the inductive biases of gradient descent and architectural constraints; they are under no obligation to be sparse, disentangled, or lexically nameable. Ideonomic dimensions maximize semantic transparency and controlled variation; they are under no obligation to capture every statistical regularity present in a training corpus. Each pipeline therefore optimizes a different objective, and each inherits characteristic strengths and failure modes. Emergent representations excel at coverage and predictive accuracy yet frustrate explanation and targeted intervention. Ideonomic representations excel at explanation and intervention yet risk incompleteness or rigidity if the initial set of distinctions is poorly chosen.

A moderate position avoids the false dichotomy that would force a choice between the two pipelines. Explicit ideonomic dimensions are best understood as an interpretable superstructure that can be imposed upon, or learned alongside, emergent representations rather than as a wholesale replacement for them. The linear map $W \in \mathbb{R}^{n \times d}$ that produces $\mathbf{z} = W\mathbf{h}$ is the simplest formal expression of this superposition: the emergent geometry is left intact while a set of designed directions is read out from it. More elaborate schemes are possible. One may train a lightweight auxiliary head to predict ideonomic coordinates from $\mathbf{h}$, thereby encouraging the underlying network to align some of its internal dimensions with the target schema. One may also use ideonomic scores as auxiliary supervision or as regularizers that penalize entanglement among the designed axes. In all such cases the emergent representation remains the primary carrier of statistical knowledge; the ideonomic layer supplies an explicit, human-readable coordinate system that can be inspected, edited, or constrained without requiring a full reinterpretation of the latent space.

This moderate stance is consistent with the non-replacement claim stated at the outset of the paper. Designed dimensions do not pretend to exhaust the geometric structure discovered by large-scale optimization, nor do they claim superior predictive power on arbitrary downstream tasks. Their value lies in complementarity. Where an analyst needs to know whether a hidden state has shifted along a particular conceptual axis (say, from more individual to more collective framing, or from more static to more dynamic description) the

ideonomic coordinate supplies a direct scalar answer. Where an intervention is required (steering generation toward one polar extreme while holding others constant) the same coordinate offers a controllable handle. The emergent geometry continues to support the bulk of fluent prediction; the ideonomic superstructure renders selected directions of variation legible and actionable.

Several technical observations follow. First, the quality of the projection $\mathbf{z} = W\mathbf{h}$ depends on the degree to which the desired ideonomic axes are linearly recoverable from the hidden state. If a distinction is encoded in a highly non-linear or highly distributed fashion, a linear readout will capture only a fraction of the available signal; non-linear probes or sparse feature extractors may then be required. Second, the stability of the extracted coordinates across model checkpoints, prompt formats, and fine-tuning runs becomes an empirical question that directly engages the stability component $\mathcal{S}_i$ of the ideonomic utility score. Third, once a set of ideonomic coordinates is reliably extracted, the higher-order co-variation among them can be assembled into an ideonomic tensor whose modes remain semantically transparent even while their numerical content is derived from model-internal geometry. In this way the tensorial constructions of the preceding section acquire an experimental realization inside contemporary language models.

The tension between capacity and interpretability is therefore not resolved by abandoning either pipeline. It is managed by layering an explicit dimensional schema onto a high-capacity emergent substrate. The ideonomic contribution is the disciplined generation and selection of the schema; the model contribution is the rich geometry from which the schema can be read out. The remainder of this section develops the concrete mechanisms of that readout—projection, probing, and controlled steering—and examines the conditions under which the resulting ideonomic coordinates improve the interpretability, controllability, or diagnostic analysis of large-language-model behavior.

The next section turns to the technical realization of the projection $\mathbf{z} = W\mathbf{h}$ and to the empirical protocols by which the fidelity and utility of the extracted ideonomic coordinates can be assessed.

5.2 The Ideonomic Projection Layer

The interface between an emergent latent geometry and an explicit ideonomic schema is realized most directly by a linear projection. Let $\mathbf{h} \in \mathbb{R}^d$ be a hidden representation extracted from a chosen layer of a language model, and let n be the number of selected ideonomic distinctions, with $n \ll d$. A matrix $W \in \mathbb{R}^{n \times d}$ defines the map

$$\mathbf{z} = W\mathbf{h},$$

where the output $\mathbf{z} \in \mathbb{R}^n$ is interpreted as an ideonomic coordinate vector. Each component z_i is understood as the model's estimated location of the input on the corresponding distinction D_i , scaled so that the theoretical range remains $[-1, 1]$ or is subsequently calibrated to that interval. The projection therefore extracts a low-dimensional, semantically labeled shadow of the high-dimensional latent state without altering the latent state itself.

Three illustrative coordinates make the construction concrete. Let z_1 record position on the concrete $\leftrightarrow$ abstract axis, z_2 on the individual $\leftrightarrow$ collective axis, and z_3 on the static $\leftrightarrow$ dynamic axis. A sentence describing a particular physical mechanism might yield a vector whose first component is strongly negative (concrete), whose second is near zero or mildly negative (individual), and whose third is positive (dynamic). A sentence advancing a general social theory might reverse the first two signs while remaining intermediate on the third. Because each axis carries an explicit polar semantics, the resulting triple is immediately legible to a human analyst; the same information, distributed across thousands of latent dimensions, would be far harder to recover by direct inspection of $\mathbf{h}$.

The matrix W itself may be obtained by any of several procedures, each of which embodies distinct assumptions about the relationship between latent geometry and ideonomic structure. The simplest procedure is supervised probing. A corpus of textual units is annotated (by human raters or by a previously validated scoring function f_i) with target coordinates on each of the n distinctions. Ordinary least-squares regression, or a regularized variant, then estimates the rows of W so as to minimize the discrepancy between $W\mathbf{h}$ and the target vector. The central assumption is linear recoverability: the desired ideonomic axes are presumed to lie approximately within the span of the latent coordinates, so that a linear readout suffices. When the assumption holds, probing is efficient and transparent; when the axes are encoded non-linearly, the residual error of the linear probe quantifies the shortfall and may motivate non-linear alternatives.

A second family of methods estimates pole directions contrastively. One constructs pairs or sets of examples that differ maximally along a single target distinction while remaining roughly matched on others. The difference of their mean latent vectors supplies a candidate direction for the corresponding row of W ; subsequent orthogonalization or joint optimization can reduce crosstalk among the estimated directions. This approach assumes that polar extremes are linearly separable in the latent space and that a sufficiently clean contrast set can be assembled. It requires less exhaustive annotation than full supervised probing (only polar exemplars are needed) yet it inherits the risk that the observed difference vector conflates the target distinction with correlated but irrelevant factors.

A third route is regularized joint training. The language-model parameters and the projection matrix W are optimized simultaneously, or in alternation, under a composite loss that combines the original predictive objective with an auxiliary term encouraging $\mathbf{z}$ to match ideonomic targets or to satisfy geometric constraints (for example, low mutual information among the z_i , or reconstruction of known polar contrasts). The assumption here is more ambitious: the latent geometry itself is plastic enough to align partially with the ideonomic schema without catastrophic loss of predictive performance. Joint training can in principle produce tighter coupling between emergent and designed dimensions, yet it also risks distorting the very geometry that made the model useful, and it demands careful weighting of the auxiliary losses to avoid either under- or over-regularization.

Each of these procedures yields a hypothetical architecture rather than a claim about existing systems. The architecture consists of an ordinary transformer (or other sequence model) that continues to maintain its high-dimensional residual stream, together with an auxiliary linear (or lightly non-linear) readout that continuously emits an ideonomic coordinate vector $\mathbf{z}$. The two representations coexist: $\mathbf{h}$ supports the fluent prediction and broad generalization for which the model was trained, while $\mathbf{z}$ supplies a compact, human-interpretable summary of selected conceptual dimensions. Because the projection is cheap to compute and does not modify the forward pass used for generation, the architecture preserves the original model’s capabilities while adding an explicit dimensional interface.

The dual-representation design realizes the moderate position articulated earlier. The emergent latent space is not replaced; it is instrumented. Analysts may consult $\mathbf{z}$ when they need to know where a hidden state sits on a particular conceptual axis, or when they wish to monitor drift along that axis across a conversation or a fine-tuning run. Interventions become correspondingly more precise: one may adjust generation by adding a steered component in the direction of a chosen row of W , or by constraining the decoder to keep selected coordinates inside prescribed ranges. At the same time, the bulk of the model’s statistical knowledge remains distributed in $\mathbf{h}$, so that the inevitable incompleteness of any finite ideonomic schema does not cripple performance on tasks outside the schema’s scope.

Several practical considerations follow. The numerical stability of $\mathbf{z}$ across random seeds, prompt framings, and model scales must be assessed; high variance would undermine the utility of the coordinates for longitudinal or comparative analysis. Calibration techniques may be required to map the raw projections onto the canonical interval $[-1, 1]$ and to correct for systematic biases introduced by the training distribution. Finally, the same projection matrix that produces $\mathbf{z}$ can be reused to populate higher-order ideonomic

tensors: once a corpus of hidden states has been mapped to coordinate vectors, the empirical co-variation among the z_i supplies the components of T . In this way the projection layer serves as the empirical bridge between the latent geometry of the language model and the multilinear constructions developed in the preceding section.

The ideonomic projection layer is therefore both a technical device and a methodological commitment. Technically, it is a thin linear map that extracts a low-dimensional labeled representation from a high-dimensional unlabeled one. Methodologically, it embodies the claim that designed conceptual dimensions can productively coexist with emergent latent dimensions, each compensating for the characteristic limitations of the other. The architecture remains hypothetical until concrete instantiations are trained and evaluated, yet the formal outline is already sufficient to guide experimental work and to clarify the sense in which ideonomy can serve as an interpretive and control layer for contemporary language models.

The next section examines the empirical protocols by which the fidelity of the extracted coordinates can be measured and by which the utility of the dual representation for interpretability and controlled generation can be assessed.

VI. The Ideonomic Dimension-Generation Algorithm

6.1 Define, Generate, Test, Prune

The reconstruction of ideonomy as a generative methodology for conceptual dimensions remains incomplete until the successive operations of proposal, evaluation, and selection are organized into an explicit procedure. The algorithm that follows converts the theoretical apparatus developed in earlier sections into a sequence of reproducible steps. It begins with the delimitation of a universe of discourse, proceeds through the generation of candidate distinctions, and subjects those candidates to successive filters of operationalizability and non-redundancy. The present block details the first four steps; later blocks complete the scoring, selection, and validation phases. The procedure is deliberately modular: each step can be intensified or relaxed according to the resources and aims of a particular inquiry, yet the overall architecture remains fixed so that results remain comparable across applications.

Step 1. Define the universe $\mathcal{U}$.

The algorithm opens by fixing the conceptual universe within which distinctions will be sought. $\mathcal{U}$ is not an ontological inventory of all that exists; it is a pragmatic delimitation of the domain of discourse under investigation. Typical choices include human discourse in general, the specialized vocabulary of a scientific discipline, the decision-making language of economic agents, or the textual output of a particular model or corpus. The choice of $\mathcal{U}$

determines the scope of subsequent generation and the population over which scoring functions will be evaluated. A narrowly drawn universe yields higher internal coherence at the cost of limited transfer; a broadly drawn universe increases the chance of discovering cross-domain regularities while raising the difficulty of reliable operationalization. In either case the definition must be stated with sufficient precision that independent investigators can recognize whether a proposed distinction falls inside or outside the intended scope. Once $\mathcal{U}$ is fixed, every candidate division is understood as an operation performed on some subdomain $X \subseteq \mathcal{U}$, and every coordinate x_i is understood as a location relative to that same universe.

Step 2. Generate candidate binomials $D_1, \dots, D_n$.

With the universe in place, the algorithm proceeds to the generative phase. Candidate distinctions are proposed by a combination of four complementary sources. The first source is systematic enumeration inspired by ideonomic practice: one begins with high-level polar pairs and recursively subdivides, or one cross-applies known distinctions to new subdomains in search of productive refinements. The second source is expert ontology: existing formal taxonomies, controlled vocabularies, and disciplinary dichotomies are harvested and recast as binomial divisions with explicit poles. The third source is corpus analysis: frequency-based, embedding-based, or topic-model methods identify recurrent contrasts latent in the textual material of $\mathcal{U}$. The fourth source is large-language-model-assisted hypothesis generation: a model is prompted to propose novel or under-explored polar pairs that satisfy minimal coherence criteria, after which the proposals are filtered by human review. The output of this step is an unpruned collection $\mathcal{D} = \{D_1, \dots, D_n\}$, each $D_i = (A_i, B_i)$ accompanied by a provisional natural-language gloss of its poles. At this stage no claim is made about independence or operationalizability; the sole requirement is that each candidate be intelligible as a respect of possible variation within $\mathcal{U}$.

Step 3. Test operationalizability $O(D_i)$.

A distinction that cannot be scored reliably cannot function as a computational dimension. Operationalizability is therefore assessed before any further filtering. For each candidate D_i one constructs or retrieves a scoring function f_i that maps entities drawn from $\mathcal{U}$ to values in $[-1, 1]$. The quality of f_i is evaluated along four complementary indicators. Inter-annotator agreement measures the consistency with which human raters assign coordinates when presented with the polar glosses; high agreement supports the claim that the distinction is publicly accessible. Classifier accuracy (or correlation) measures the degree to which an automated model can recover held-out human scores or polar labels; high accuracy supports scalability. Stability across contexts examines whether scores for

the same entity remain consistent under modest changes of framing, genre, or surrounding discourse; low stability signals that the distinction is overly brittle. Semantic coherence assesses whether intermediate values on the scale are interpretable and whether the poles remain opposed rather than orthogonal or synonymous; failure of coherence indicates that the candidate is not a genuine axis. These indicators are aggregated into a composite operationalizability score $O_i \in [0, 1]$. Candidates falling below a pre-specified threshold are discarded or returned to the generative phase for reformulation. The surviving set is smaller and, by construction, equipped with reproducible scoring procedures.

Step 4. Test redundancy R_{ij} and prune.

Even operationalizable distinctions may convey largely overlapping information. The fourth step therefore estimates pairwise redundancy and removes, merges, or transforms highly dependent candidates. For every pair D_i, D_j that survived Step 3 one computes a redundancy statistic R_{ij} on a representative sample of entities scored by f_i and f_j . Suitable statistics include absolute Pearson or Spearman correlation, mutual information (normalized to $[0, 1]$), or the proportion of variance in one coordinate linearly predictable from the other. Conditional measures may be added when hierarchical dependence is suspected. Pairs whose R_{ij} exceeds a chosen threshold are flagged. For each flagged pair the analyst (or an automated rule) chooses among three actions: removal of the lower-utility member, merger into a single composite axis, or transformation (for example, residualization of one coordinate against the other so that the residual becomes the new candidate). The process is iterated until no pair remains above threshold or until a minimum set size is reached. The output is a reduced collection whose members are approximately non-redundant by the chosen metric. This collection is not yet the final selected set $\mathcal{D}^*$; it is the input to the subsequent utility-scoring and ranking stages. The non-redundancy component N_i of each surviving distinction can already be estimated from the residual dependency profile, supplying one of the four terms required by the composite utility function Φ .

Taken together, the first four steps convert an unbounded space of possible distinctions into a finite, operationalizable, and approximately non-redundant candidate pool. The procedure remains fallible (poor choices of $\mathcal{U}$, unimaginative generation, unreliable scorers, or overly aggressive pruning can all degrade the result) yet each source of error is localized to a named step and is therefore open to diagnosis and correction. The algorithm does not claim to recover a unique or privileged set of dimensions; it claims only to implement, in reproducible form, the generative and selective logic that the preceding theoretical sections have defended.

The next section continues the algorithm with the computation of the remaining utility components, the formation of the composite score $\Phi(D_i)$, and the final selection of the retained set $\mathcal{D}^*$.

6.2 Utility Scoring and Tensor Construction

The first four steps of the algorithm produce a finite pool of operationalizable and approximately non-redundant candidate distinctions. Two further steps are required to convert that pool into a selected set of computational dimensions and to assemble those dimensions into an ideonomic tensor. Step 5 computes a composite utility score for each remaining candidate; Step 6 retains the highest-scoring subset and constructs the corresponding multilinear representation. The composite score $\Phi(D_i)$ is presented as one of the paper’s central original contributions: it supplies an explicit, tunable criterion that links the qualitative ambitions of ideonomy to the quantitative demands of dimensional selection.

Step 5. Compute the ideonomic utility score $\Phi(D_i)$.

For each candidate D_i that survived the redundancy prune, define

$$\Phi(D_i) = \alpha M_i + \beta O_i + \gamma N_i + \delta S_i,$$

where the four component scores lie in the unit interval and the non-negative weights sum to one (or are otherwise normalized). Meaningfulness M_i quantifies the semantic clarity and theoretical or practical significance of the polar contrast. It may be obtained by expert rating of the glosses, by measuring the productivity of the distinction in organizing known anomalies or generating novel hypotheses, or by assessing the breadth of subdomains of $\mathcal{U}$ to which the axis usefully applies. Operationalizability O_i is the composite already computed in Step 3, reflecting inter-annotator agreement, classifier recovery, contextual stability, and semantic coherence of intermediate values. Non-redundancy N_i is derived from the residual dependency profile left after Step 4: it rises as the average (or maximum) redundancy R_{ij} between D_i and all other surviving candidates falls. Stability S_i measures robustness under perturbations of sample, scoring method, polar wording, or modest shifts in the boundary of $\mathcal{U}$; high stability indicates that the distinction is not an artifact of a particular operationalization.

The weights $\alpha, \beta, \gamma, \delta$ encode the relative priority given to each desideratum. They may be set by expert judgment when the inquiry has well-defined downstream aims (for example, maximizing interpretability may elevate α and β ; maximizing statistical independence may elevate γ). Alternatively, they may be tuned by validation performance: one reserves a

held-out collection of tasks or corpora and selects the weight vector that yields a $\mathcal{D}^*$ whose induced coordinates best support those tasks. Intermediate schemes are possible, including hierarchical weighting or context-dependent reweighting. Because $\text{monoi}\Phi$ is monotonic in each component, the score functions as a soft pruning mechanism. Candidates whose utility falls below a threshold can be discarded; the remainder can be ranked. In this way the exponential threat of 2^n terminal categories that accompanies unrestricted recursive division is brought under control: only those distinctions that survive successive filters of operationalizability, non-redundancy, and composite utility are retained for dimensional status.

The formula is original in its explicit aggregation of the four ideonomic desiderata into a single scalar that governs selection. Earlier sections motivated each component separately; the present step unifies them into an operational criterion that can be computed, audited, and revised. The score does not claim to capture every relevant consideration, nor does it eliminate the need for human oversight. It does, however, replace an otherwise informal process of dimensional choice with a transparent, reproducible calculation whose assumptions (the definitions of the components and the values of the weights) are open to inspection and challenge.

Step 6. Select $\mathcal{D}^*$ and construct the ideonomic tensor T .

Given the scored candidates, form the selected set $\mathcal{D}^* = \{D_1, \dots, D_k\}$ by retaining the k highest-scoring distinctions, or by retaining all distinctions whose utility exceeds a pre-specified threshold, or by a hybrid rule that respects both ranking and absolute quality. The integer k may be fixed in advance by considerations of interpretability or computational budget, or it may be determined data-dependently (for example, by locating an elbow in the sorted utility curve). Once $\mathcal{D}^*$ is fixed, each retained distinction D_i is associated with its one-dimensional factor space V_{D_i} . The ideonomic tensor is then the element

$$T \in V_{D_1} \otimes \dots \otimes V_{D_k}$$

whose components are populated by the measured co-variation (or other multilinear statistic) of the corresponding coordinates across a corpus of entities scored by the functions f_i . At this point the generative and selective phases are complete: the algorithm has moved from an open universe $\mathcal{U}$ to a concrete, finite-dimensional multilinear representation whose modes carry explicit polar semantics.

A compact pseudocode summary of the full six-step procedure may be given as follows.

Input: universe $\mathcal{U}$, weight vector $(\alpha, \beta, \gamma, \delta)$, selection rule

1. Define $\mathcal{U}$
2. Generate candidate set $\mathcal{D} = \{D_1, \dots, D_n\}$
3. For each $D_i \in \mathcal{D}$:
compute O_i ; discard if below threshold
4. Estimate R_{ij} for all surviving pairs;
prune/merge/residualize until redundancy criteria met
5. For each remaining D_i :
compute M_i, N_i, S_i ;
 $\Phi(D_i) = \alpha M_i + \beta O_i + \gamma N_i + \delta S_i$
6. Select $\mathcal{D}^*$ according to Φ and selection rule;
construct $T \in \bigotimes_{D_i \in \mathcal{D}^*} V_{D_i}$
Output: $\mathcal{D}^*$, scoring functions $\{f_i\}$, tensor T

The algorithm is deliberately non-deterministic at several choice points—the sources of generation, the precise operationalization of each component score, the redundancy threshold, the weights, and the final selection rule. These degrees of freedom are not defects; they allow the procedure to be adapted to the epistemic aims and resource constraints of different investigations. What remains invariant is the overall flow: generation is followed by successive filters of operationalizability and non-redundancy, the survivors are ranked by a transparent utility score, and the retained dimensions are assembled into a multilinear object whose modes remain semantically interpretable.

With the completion of Steps 5 and 6 the constructive core of the reconstruction is in place. Candidate distinctions have been proposed, tested, scored, and selected; the selected set has been realized both as a coordinate system and as an ideonomic tensor. The remaining sections turn from construction to application and evaluation: first examining how the resulting dimensions and tensors can be used to probe and steer large language models, then assessing the strengths, limitations, and open problems of the overall framework.

The next section develops the experimental protocols by which an ideonomic schema produced by the algorithm can be aligned with, and used to interrogate, the latent geometry of contemporary language models.

VII. Experimental Design (Proof of Concept)

7.1 Corpus and Ideonomic Basis

A proof-of-concept evaluation of the ideonomic framework requires a delimited corpus against which candidate dimensions can be scored, a modest basis of distinctions that satisfy the dimensionality criterion, and a transparent procedure for assigning each textual unit a coordinate vector. The present section addresses these three requirements in turn. The goal is not yet full-scale empirical validation, but the construction of a tractable experimental setting in which the generative algorithm, the projection layer, and the resulting ideonomic tensor can be exercised and inspected.

Four classes of corpus present themselves as natural candidates. Political speeches offer dense concentrations of evaluative and agency-laden language, together with recurrent contrasts between individual leadership and collective fate. Scientific abstracts supply a comparatively constrained register in which ontological and epistemic distinctions (concrete versus abstract entities, certain versus uncertain claims) are foregrounded by disciplinary convention. Business reports emphasize scale, modality, and evaluation, frequently oscillating between descriptive accounts of present conditions and normative projections of future strategy. Literary fiction, finally, spans the full range of the foregoing registers while adding layers of narrative embedding, free indirect discourse, and deliberate conceptual play. Each corpus type could in principle support an ideonomic analysis; the choice among them is governed by the density and variety of conceptual oppositions that the texts reliably exhibit.

Literary fiction is recommended as the primary corpus for an initial proof of concept. The justification is structural rather than aesthetic. Fictional narrative routinely operates at multiple levels of abstraction within a single paragraph or scene: a concrete sensory detail may be immediately taken up into an abstract reflection; a static description of setting may be disrupted by dynamic action; an individual consciousness may be shown in tension with collective social forces. These shifts generate precisely the graded, multi-dimensional variation that continuous ideonomic coordinates are designed to capture. Moreover, literary texts are rich in modal and epistemic shading: characters entertain possibilities, narrators withhold certainty, and evaluative language is often ironized or distributed across competing voices. Such features supply natural test cases for the operationalizability of distinctions that might remain latent or formulaic in more regimented genres. A corpus

assembled from short stories, novel excerpts, or carefully segmented chapters therefore offers a high density of conceptual oppositions per unit of text, while remaining small enough for intensive human annotation and model probing.

Against this corpus a compact ideonomic basis of roughly ten to twenty dimensions is to be constructed. For the proof-of-concept stage a core set of eight high-level distinctions is proposed; the set can be expanded by recursive subdivision or by domain-specific additions once the initial axes have been scored and residualized. The core set is displayed below.

Dimension	Pole A (−1)	Pole B (+1)
Ontological	Concrete	Abstract
Temporal	Static	Dynamic
Scale	Individual	Collective
Modality	Actual	Possible
Epistemic	Certain	Uncertain
Agency	Passive	Active
Structure	Discrete	Continuous
Evaluation	Descriptive	Normative

Each row corresponds to a binomial division $D_i = (A_i, B_i)$ whose poles define the endpoints of a continuous coordinate $x_i \in [-1, 1]$. The ontological axis locates a textual unit according to whether it primarily presents concrete particulars or abstract types and relations. The temporal axis distinguishes descriptions that emphasize enduring or unchanging features from those that emphasize process, change, or event. The scale axis contrasts foci on singular agents or objects with foci on groups, institutions, or aggregates. The modality axis separates claims or depictions of what is from those of what might be, could be, or is imagined. The epistemic axis registers the degree of certainty or tentativeness with which content is advanced. The agency axis measures whether entities are portrayed as undergoing effects or as initiating actions. The structure axis opposes discrete, countable, or sharply bounded presentations to continuous, graded, or fluid ones. Finally, the evaluation axis distinguishes language that primarily describes from language that primarily prescribes, praises, or blames. These eight distinctions are not claimed to be exhaustive or mutually orthogonal a priori; they are offered as a generative starting point whose empirical redundancy and operationalizability will be assessed on the chosen corpus.

Once the basis is fixed, each textual unit—paragraph, scene, or other analytically convenient segment—is assigned a coordinate vector $\mathbf{z}_{\text{ideo}} \in \mathbb{R}^k$, where k is the number

of retained dimensions after pruning. Assignment proceeds by means of the scoring functions f_i developed in the algorithmic phase. In the simplest realization, human annotators read each unit and locate it on every axis, producing a vector of ratings that are then averaged or otherwise aggregated. In a more scalable realization, a supervised probe or a prompted language model is trained or calibrated to approximate the human scores, yielding an automated map from text to $\mathbf{Z}_{\text{ideo}}$. Hybrid protocols are also admissible: a model generates preliminary coordinates that human raters then correct or confirm on a stratified sample. Regardless of the precise instrumentation, the output for each unit is a point in the ideonomic coordinate space whose components are explicitly labeled by the polar semantics of the retained distinctions. The collection of all such points supplies the empirical distribution from which pairwise redundancies R_{ij} , higher-order co-variation, and ultimately the ideonomic tensor T can be estimated.

The resulting experimental object is therefore threefold: a literary corpus segmented into scorable units, a modest basis of ideonomic dimensions equipped with operational scoring procedures, and a matrix of coordinate vectors $\mathbf{Z}_{\text{ideo}}$ that locate every unit within the designed conceptual space. This object is sufficient to support the subsequent proof-of-concept analyses (measurement of inter-annotator reliability, estimation of the redundancy matrix, construction of a low-order tensor, and probing of language-model hidden states for recoverable ideonomic directions) while remaining small enough that every stage remains open to detailed inspection and critique.

The next section specifies the annotation protocol, the reliability metrics, and the procedures by which the coordinate matrix is used to evaluate both the internal coherence of the ideonomic basis and its alignment with latent representations of contemporary language models.

7.2 Comparison Against Embeddings and Tensor Interaction Analysis

The experimental value of an ideonomic representation is measured by what it makes visible that remains opaque in ordinary latent embeddings. Once a literary corpus has been segmented and each unit has been assigned both a high-dimensional hidden state $\mathbf{h}_{\text{LLM}}$ drawn from a contemporary language model and a low-dimensional ideonomic coordinate vector $\mathbf{Z}_{\text{ideo}}$ obtained by scoring or projection onto the retained basis, the two representations can be compared directly. The comparison is organized around five complementary criteria: clustering structure, supervised classification performance, semantic similarity relations, intrinsic interpretability, and human evaluative preference. In each case the ideonomic vectors are treated as an explicit, low-dimensional alternative to the emergent geometry, not as a replacement for it.

Clustering analysis asks whether the ideonomic coordinates induce partitions of the corpus that are more coherent, more stable, or more semantically transparent than those induced by the latent states. Standard metrics include normalized mutual information with respect to independently annotated thematic or stylistic labels, silhouette scores under k -means or hierarchical clustering, and the stability of cluster assignments under bootstrap resampling. Baselines are the clusters obtained from the raw $\mathbf{h}_{\text{LLM}}$ (possibly after dimensionality reduction by PCA or UMAP) and from simple bag-of-words or sentence-embedding baselines. A positive result for the ideonomic side would appear as higher agreement with human-interpretable categories at comparable or lower dimensionality.

Classification probes test whether the ideonomic coordinates support accurate recovery of known polar or categorical distinctions. Linear or kernel classifiers are trained to predict held-out human scores on individual axes or to discriminate pre-defined literary subtypes (for example, passages dominated by interior monologue versus passages dominated by external action). Accuracy, F1, and calibration error serve as primary metrics; the latent states $\mathbf{h}_{\text{LLM}}$ supply the natural high-capacity baseline. Because $\mathbf{z}_{\text{ideo}}$ is already aligned with the target semantics, strong classification performance is expected; the more informative comparison is whether the low-dimensional ideonomic features match or exceed the latent baseline when the latter is restricted to the same number of dimensions.

Semantic similarity evaluation examines whether cosine or Euclidean neighborhoods in the ideonomic space correspond more closely to human judgments of conceptual relatedness than neighborhoods in the latent space. A stratified sample of passage pairs is rated by annotators for similarity along the designed axes and for overall thematic affinity; rank correlation with model-derived similarities then quantifies alignment. Interpretability is assessed both intrinsically (by measuring how reliably human readers can predict an unseen coordinate from the polar glosses alone) and extrinsically, by measuring how quickly and accurately analysts can diagnose model errors or steer generation when given access to $\mathbf{z}_{\text{ideo}}$ rather than to $\mathbf{h}_{\text{LLM}}$. Human evaluation protocols include forced-choice preference tests and timed diagnostic tasks in which participants must identify the conceptual locus of a given passage or of a model failure.

These comparative measures converge on a single empirical question: does an explicit ideonomic representation reveal meaningful structure that ordinary embeddings obscure? The question is deliberately asymmetric. It does not ask whether ideonomic coordinates outperform latent states on every predictive task; the moderate position adopted earlier already concedes that emergent geometry retains superior coverage and capacity. It asks whether the designed axes surface regularities (especially graded, multi-dimensional, and interactional regularities) that remain distributed, entangled, or invisible when only the

high-dimensional latent geometry is examined. Affirmative evidence would consist in higher interpretability scores, tighter alignment with human conceptual judgments, and the discovery of systematic patterns that post-hoc probing of $\mathbf{h}_{\text{LLM}}$ recovers only partially or at greater analytic cost.

The same coordinate matrix that supports the foregoing comparisons also supplies the empirical content of the ideonomic tensor. Pairwise interactions are captured by the second-order slices T_{ij} , which may be estimated as the empirical covariance, mutual information, or excess co-occurrence of coordinates i and j across the corpus. Third-order interactions appear in the components T_{ijk} , which register residual dependence among triples after lower-order effects have been partialled out. Because each mode of T carries an explicit polar semantics, any large interaction term is immediately interpretable: a strong positive T_{ijk} indicates that the joint occurrence of particular polar extremes on axes i , j , and k exceeds what the marginals would predict. In the literary setting such terms may reveal, for example, that passages high on the abstract, collective, and normative poles co-occur more frequently than chance, or that certain combinations of epistemic uncertainty and modal possibility are characteristic of free indirect discourse.

These higher-order statistics permit a direct test of a further hypothesis: that some semantic categories familiar from literary or conceptual analysis are better understood as intersections of continuous dimensions than as discrete, atomic classes. Under the hypothesis, a traditional category label functions as a convenient name for a region in the ideonomic space (an intersection of polar tendencies) rather than as an indivisible primitive. Empirical support would appear if passages bearing the category label cluster tightly in a particular orthant or sub-region of the coordinate space, if the label can be predicted with high accuracy from a sparse linear combination of the ideonomic coordinates, and if the residual variance after such prediction is small. Conversely, if the category remains poorly recoverable from the designed axes, or if it requires a highly non-linear function of the coordinates, the hypothesis is weakened for that category.

Falsification conditions must be stated with equal clarity. The ideonomic proposal would be empirically undermined if, across multiple independent literary samples, the low-dimensional coordinates $\mathbf{z}_{\text{ideo}}$ systematically underperform dimensionality-matched projections of $\mathbf{h}_{\text{LLM}}$ on clustering coherence, classification of human-interpretable distinctions, and semantic-similarity alignment, while simultaneously failing to yield higher human interpretability or diagnostic utility. It would be further undermined if the estimated tensor T exhibited no stable higher-order interactions beyond what is already captured by pairwise correlations, or if the hypothesized intersectional structure of familiar semantic

categories proved irrecoverable from the designed axes. Finally, if the projection matrix W itself proved unstable across model checkpoints or prompt framings, so that the mapping from latent states to ideonomic coordinates could not be reproduced, the practical value of the dual-representation architecture would be called into serious question. Any of these outcomes would count as evidence against the claim that the ideonomic basis supplies a usefully revealing superstructure for the analysis of literary and, by extension, other textual distributions.

The experimental design therefore supplies both positive tests and clear routes to refutation. By placing the designed coordinates and the emergent latent states on equal comparative footing, and by extending the analysis from first-order locations to second- and third-order interactions, the protocol examines whether the generative and selective labor of ideonomy earns its keep in the form of structures that ordinary embeddings leave obscure.

The next section turns from the design of the comparative and tensorial analyses to the interpretation of their anticipated outcomes and to the broader methodological implications for the study of conceptual dimensions in language models.

VIII. Theoretical Implications

8.1 A New Definition of Dimension; Dimensions as Questions

The preceding sections have treated dimensions under a variety of disciplinary descriptions: geometric degrees of freedom, physical coordinates, linear-algebraic basis cardinalities, tensorial index extents, machine-learning features or latent axes, and philosophical distinctions. Each description captures part of the phenomenon; none by itself supplies a definition adequate to the full range of uses that the ideonomic reconstruction requires. The present section therefore advances an explicit working definition that is intended to unify the earlier senses while remaining operational for both conceptual analysis and computational implementation: a dimension is a reproducible distinction along which a system can vary independently enough to support structured representation.

The definition is deliberately composite. “Reproducible” requires that the distinction can be applied by different observers or instruments with acceptable inter-subjective or inter-method agreement; without reproducibility the distinction remains a private intuition rather than a public coordinate. “Distinction” preserves the ideonomic origin in binomial (or multi-polar) opposition: the dimension is generated by the isolation of poles that orient an axis of possible variation. “Along which a system can vary” insists that intermediate positions are at least conceivable and, ideally, measurable; a merely classificatory cut that admits no

gradation falls short of dimensional status. “Independently enough” encodes the non-redundancy demand formalized earlier as N_i and estimated via R_{ij} : perfect linear independence is unnecessary, yet the residual variation along the axis must not be wholly predictable from the other retained dimensions. Finally, “to support structured representation” ties the definition to its functional purpose. A distinction earns dimensional status when the coordinates it induces can participate in larger geometric, statistical, or multilinear constructions (vectors $\mathbf{x}$ or $\mathbf{z}$, tensors T , projections $W\mathbf{h}$) that make relational structure among entities visible and manipulable.

This definition bridges four intellectual domains that rarely share a common vocabulary. In philosophy and classification theory it recovers the classical concern with respects of sameness and difference, while adding the modern requirements of reproducibility and graded variation. In mathematics it aligns with the linear-algebraic and tensorial notions of independent directions, yet relaxes the demand for strict basis independence in favor of a pragmatic threshold of residual variation. In information theory it resonates with the idea of a coordinate that reduces uncertainty about a system’s state; a dimension that carries no mutual information with any observable is vacuous, while a dimension that is wholly redundant with others adds no new information. In machine learning it supplies an explicit criterion for deciding which directions in a latent space, or which engineered features, deserve to be treated as interpretable axes rather than as opaque components of a distributed representation. The same definition therefore serves as a translation manual: a philosophical distinction that meets the reproducibility and independence thresholds can be admitted into a mathematical coordinate system; a statistically recovered latent direction that can be given a stable polar semantics can be admitted into an ideonomic schema.

The ideonomic contribution to this definitional synthesis lies in the further insight that a dimension is the answer space to a structured question. Consider the interrogative form “Is this phenomenon more individual or more collective?” The question is structured: it isolates two poles, licenses intermediate answers, and thereby defines a continuum of possible responses. The dimension is precisely the space of those possible answers, together with the reproducible procedure that maps any given entity onto a location within the space. In prose form the relation may be stated: dimension equals structured question plus space of possible answers. The binomial division $D_i = (A_i, B_i)$ is the question; the interval $[-1, 1]$ (or any order-isomorphic continuum) is the answer space; the scoring function f_i is the answering procedure. Once several such question–answer pairs have been assembled and residualized, the resulting coordinate vector $\mathbf{x}$ or $\mathbf{z}$ becomes a

simultaneous answer to a battery of structured questions, and the ideonomic tensor T becomes a record of how answers to different questions co-vary.

This interrogative framing is useful as a bridge between Gunkel's ideonomy and contemporary machine learning for three interlocking reasons. First, it renders the generative phase of ideonomy computationally natural. Large language models are already optimized to answer questions; prompting a model to locate a passage on a polar axis is a direct exploitation of that capacity. The same models can be used to propose new structured questions, thereby closing the loop between generation and scoring. Second, the framing supplies a clear semantics for the projection layer $\mathbf{z} = W\mathbf{h}$.

Each row of W may be interpreted as extracting the latent direction that best answers a particular structured question; the quality of the answer is measurable by how well z_i recovers human or gold-standard responses to that question. Third, the framing clarifies the moderate position adopted earlier concerning emergent and designed dimensions. Emergent latent directions are answers to questions that the training objective implicitly posed: questions about next-token predictability, syntactic well-formedness, or distributional similarity. Ideonomic dimensions are answers to questions that human investigators have chosen to pose explicitly. The two sets of answers need not coincide; the value of the ideonomic superstructure lies in making a selected subset of questions and their answer spaces publicly inspectable and controllable, while leaving the remainder of the emergent geometry free to capture whatever further regularities the data contain.

The definition and the interrogative insight together reorient the status of dimensionality itself. Dimensionality is no longer an intrinsic property that a system either possesses or lacks; it is a relational achievement that obtains when a reproducible distinction has been isolated, operationalized, and shown to support structured representation. Systems do not "have" dimensions in the way they have mass or charge; they are assigned dimensions by acts of conceptual and computational labor. That labor can be performed well or badly, and the criteria of success (meaningfulness, operationalizability, non-redundancy, stability) are precisely those formalized in the utility score Φ . The ideonomic algorithm is therefore not merely a method for discovering pre-existing dimensions; it is a method for constructing dimensional representations that meet explicit standards of quality.

Several consequences follow for the theory of representation. Because dimensions are answer spaces to structured questions, the choice of which questions to ask becomes an inseparable part of the representational act. Different batteries of questions yield different coordinate systems, and the comparative evaluation of those systems is itself an empirical and normative enterprise. Because the answer spaces are continuous, the theory avoids the ontological rigidity that would follow from treating every distinction as a binary

classification. Because independence is required only “enough” to support structured representation, the theory accommodates the residual hierarchical and interactive dependencies that recursive division and tensorial analysis both bring into view. And because the entire construction is required to be reproducible, the theory remains accountable to inter-subjective and inter-method checks rather than to private intuition.

The proposed definition does not erase the differences among the six senses of dimension surveyed at the outset. It supplies a common functional core that each specialized sense can be seen to elaborate. Geometric and physical dimensions emphasize measurable variation in the external world; linear-algebraic and tensorial dimensions emphasize formal independence and multilinear interaction; machine-learning dimensions emphasize predictive utility within an optimized model; philosophical dimensions emphasize conceptual intelligibility.

The ideonomic reconstruction shows that these emphases are compatible once dimension is understood as a reproducible, sufficiently independent distinction that supports structured representation and, more pointedly, as the answer space to a structured question that investigators have chosen to pose.

The next section examines the consequences of this definitional stance for the ontology of conceptual structure and for the ethics of dimensional design in large-scale learning systems.

8.2 The Generative Nature of Division

The elementary act of ideonomy is the binomial division $X \rightarrow (A, B)$. What this act accomplishes is not the discovery of a pre-existing cut in the fabric of the world, but the creation of a new possibility for representation. Prior to the division, entities belonging to the domain X may be compared along whatever axes happen already to be available; after the division, a previously unavailable respect of variation has been isolated, oriented by a pair of poles, and thereby rendered eligible for coordinate assignment. The distinction does not merely label; it enlarges the space of admissible descriptions. Each subsequent division, whether applied to a polar region of a prior division or introduced as an independent axis, further expands that space. The iterative process therefore generates a sequence $D_1, D_2, \dots, D_n$ whose cumulative effect is the progressive construction of a dimensional representation.

This generative character distinguishes ideonomy from purely taxonomic or purely statistical approaches to dimensionality. Taxonomy typically begins with a fixed repertoire of characters and seeks optimal partitions; it does not treat the invention of new characters as its primary task. Statistical dimension-reduction techniques begin with an observed

data matrix and extract directions of maximal variance or mutual information; they do not treat the proposal of semantically interpretable axes as their primary task. Ideonomy, by contrast, treats the proposal of distinctions as the inaugural operation. The subsequent apparatus (scoring functions f_i , redundancy estimates R_{ij} , utility scores $\Phi(D_i)$, and the selection of $\mathcal{D}^*$) exists in order to discipline and validate what generation has proposed. Generation comes first; validation and interaction modeling follow.

Once this priority is recognized, ideonomy can be redescribed as a form of conceptual dimensional calculus. Ordinary differential and integral calculi supply rules for generating new functions from old ones and for extracting quantitative information about rates and accumulations. A conceptual dimensional calculus supplies rules for generating new distinctions from old domains and for extracting quantitative information about location, independence, and interaction.

The elementary operation is division; the derived objects are coordinates $x_i \in [-1, 1]$, vectors $\mathbf{x}$ or $\mathbf{z}$, and tensors T . The differential aspect appears in the capacity to refine a coarse distinction into finer residual axes; the integral aspect appears in the capacity to assemble a collection of axes into a joint representation that supports multilinear analysis. The calculus is conceptual rather than numerical in its generative phase, yet it becomes numerical as soon as scoring functions and empirical co-variation are introduced. The result is a hybrid instrument: qualitative in the questions it permits investigators to pose, quantitative in the representations it constructs and the interactions it records.

The same generative orientation clarifies the division of labor that has structured the entire reconstruction. Ideonomy generates. It supplies the structured questions whose answer spaces become the candidate dimensions. Statistics and linear algebra validate. They furnish the measures of operationalizability, redundancy, and residual independence that determine which candidates survive into $\mathcal{D}^*$ and that underwrite the composite utility score Φ . Tensor calculus models interaction.

Once the selected dimensions have been realized as factor spaces, the ideonomic tensor T records how answers to different structured questions co-vary, thereby capturing higher-order conceptual structure that no single axis can express. Large language models provide the experimental testbed. Their hidden states $\mathbf{h}$ supply a high-dimensional geometry against which designed dimensions can be projected, probed, and evaluated; their generative capacities can themselves be recruited into the proposal of new distinctions. Each component performs a task that the others cannot: generation without validation risks unconstrained proliferation; validation without generation risks sterility; interaction modeling without prior selection risks uninterpretable multilinear noise; and an

experimental testbed without an explicit dimensional schema risks leaving the latent geometry permanently opaque.

This division of labor is not a contingent methodological preference; it follows from the generative nature of division itself. Because each new distinction creates a representational possibility that did not previously exist, the space of possible dimensional schemes is open-ended. No finite set of statistical extractions or architectural inductive biases can exhaust it. Systematic generation is therefore indispensable.

At the same time, because unconstrained generation rapidly produces redundant, unstable, or non-operationalizable candidates, systematic validation is equally indispensable. Because validated axes still interact, multilinear modeling is required to recover the relational architecture they jointly determine. And because the ultimate proving ground for any dimensional scheme that aspires to relevance for contemporary representation learning is the latent geometry of large models, those models must serve as the empirical arena in which generation, validation, and interaction modeling are tested.

The generative view also supplies a final perspective on the moderate stance toward emergent and designed dimensions. Emergent dimensions are the answers that a training objective has implicitly generated and validated for its own purposes. Designed ideonomic dimensions are the answers that an explicit conceptual calculus has generated and validated for human purposes of interpretation and control. The two generative processes need not compete. The designed calculus can be used to interrogate the emergent geometry, to surface directions that the training objective left entangled, and to impose a sparse, legible coordinate system on selected subspaces without claiming to replace the high-capacity substrate. In this sense the generative nature of division is what makes the dual-representation architecture possible: each new ideonomic question creates a new readout direction, and the accumulation of such directions yields a low-dimensional interface that remains accountable to the same standards of reproducibility, independence, and structured representability that the definition of dimension demands.

By treating division as a generative operation that enlarges representational possibility, ideonomy becomes more than a historical curiosity or a taxonomic heuristic. It becomes a calculus for the controlled construction of conceptual dimensions—an instrument whose elementary steps are qualitative and whose derived structures are quantitative, and whose proper evaluation requires the joint resources of statistics, multilinear algebra, and large-scale learning systems.

The next section considers the limitations of this calculus and the open problems that remain for its further development and empirical assessment.

IX. Limitations and Critiques

9.1 Five Objections and Responses

Any proposal that elevates designed conceptual dimensions to a formal role alongside emergent latent representations invites systematic criticism. Five recurrent objections deserve explicit statement and response. The first concerns the risk of false dichotomies. Critics observe that binomial divisions appear to force continuous or multifaceted phenomena into binary molds, thereby distorting the very domains they purport to clarify. The response is that the poles of a binomial division function as anchors rather than as exhaustive categories. The continuous coordinate $x_i \in [-1, 1]$ explicitly admits intermediate, mixed, and context-dependent positions; an entity may occupy any location on the axis without being compelled to inhabit one pole or the other. Complementary divisions that demand mutual exclusivity and joint exhaustiveness are only one of three recognized kinds, and even they are treated as idealizations that may be relaxed once empirical distributions are examined. The framework therefore does not impose a binary ontology; it supplies reference extremes that orient a graded space of possible answers to a structured question. When a distinction proves too coarse, recursive subdivision or residualization remains available. The dichotomous appearance of the generative act is an artifact of the initial polar contrast, not a commitment to dichotomous reality.

The second objection addresses cultural and historical contingency. Distinctions that appear natural within one linguistic, disciplinary, or historical setting may be unintelligible or differently valenced in another. An ideonomic schema constructed from contemporary English-language literary fiction, for example, cannot be assumed to transfer unaltered to classical Chinese historiography or to indigenous oral traditions. The response distinguishes three strata of distinctions. Some axes—such as those tracking concrete versus abstract reference or static versus dynamic presentation—approximate cross-cultural and cross-temporal recurrence and may be treated as near-universal for practical purposes. Others are domain-specific, arising reliably within particular fields of inquiry yet lacking broader scope. Still others are culturally or historically contingent, reflecting local conceptual priorities that may shift under translation or social change. The algorithm itself does not decide the stratum; it requires only that each candidate be operationalizable and non-redundant within the chosen universe $\mathcal{U}$. Contingency is therefore not a refutation but a parameter of application: the stability component S_i of the utility score can be evaluated across cultural or historical partitions of the data, and low stability correctly signals that a distinction should be treated as local rather than portable. The generative apparatus remains usable; only the claims made on behalf of any particular $\mathcal{D}^*$ must be scoped accordingly.

The third objection notes that genuine statistical independence among conceptual axes is difficult to achieve and perhaps impossible in complex domains. Natural language, social life, and literary representation are rife with entanglement; insisting on orthogonal dimensions may therefore produce an idealized coordinate system that systematically misrepresents the phenomena. The response adopts approximate rather than absolute orthogonality. The non-redundancy term N_i and the pairwise measures R_{ij} are thresholded, not zeroed. Residual dependence is tolerated provided it falls below a level that would collapse the effective dimensionality or render coordinates mutually predictable. Hierarchical residues inherited from recursive division are managed by residualization or by hierarchical tensor formats rather than by denial. Approximate orthogonality is already the operative standard in much of multivariate statistics and representation learning; the ideonomic framework simply makes the same pragmatic concession while retaining an explicit semantic interpretation for each retained axis. Perfect independence remains a regulative ideal that guides pruning, not an empirical prerequisite that must be satisfied before any dimensional work can begin.

The fourth objection highlights the threat of dimensional proliferation. Unconstrained recursive division yields 2^n terminal categories after n levels, rapidly exceeding human interpretive capacity and computational convenience. Even after initial filtering, a large candidate pool may survive into the scoring stage. The response points to the composite utility score $\Phi(D_i) = \alpha M_i + \beta O_i + \gamma N_i + \delta S_i$ as the principal pruning mechanism. Because Φ aggregates meaningfulness, operationalizability, non-redundancy, and stability, candidates that are obscure, unscorable, redundant, or unstable receive low scores and are eliminated by ranking or thresholding. The weights themselves can be tuned to favor parsimony when interpretive clarity is paramount. In practice the algorithm therefore converts an open-ended generative process into a finite, ranked list whose length is controlled by explicit numerical criteria rather than by ad-hoc truncation. Proliferation is real at the generative stage; it is disciplined at the selective stage. The same score that elevates high-quality distinctions simultaneously suppresses the combinatorial explosion that would otherwise render the representation unusable.

The fifth objection contrasts interpretability with expressive power. Low-dimensional, human-legible coordinate systems necessarily discard information that high-dimensional latent representations retain; any gain in transparency may therefore be purchased at the price of predictive or generative capacity. The response endorses a hybrid architecture as the likely optimum. The ideonomic projection $\mathbf{z} = W\mathbf{h}$ leaves the underlying latent state $\mathbf{h}$ intact while extracting a sparse, labeled interface. Downstream tasks that require maximal coverage continue to operate on $\mathbf{h}$ or on task-specific fine-tunings thereof; tasks

that require diagnosis, steering, or human-readable explanation operate on $\mathbf{Z}$ or on the associated tensor T . The two representations are complementary rather than competitive. Empirical comparison protocols already built into the experimental design test precisely where the interpretable layer adds diagnostic or control value without imposing it as a universal substitute. Where the hybrid fails to improve interpretability or control relative to latent baselines, the ideonomic schema can be revised or abandoned for that domain; where it succeeds, the expressive power of the emergent geometry is preserved while selected directions of variation become publicly accountable.

These five responses do not eliminate the corresponding difficulties; they relocate them inside a framework that makes the difficulties measurable and, in principle, manageable. False dichotomies are mitigated by continuous coordinates and by the typology of division kinds. Cultural contingency is accommodated by stratified stability assessment and scoped claims. Approximate independence replaces an unattainable ideal with a tunable threshold. Proliferation is controlled by an explicit utility calculus. The tension between interpretability and power is addressed by architectural hybridization rather than by forced choice. Each response is itself open to empirical pressure: if continuous scoring proves infeasible, if stability across cultural partitions collapses, if residual dependence remains catastrophic, if Φ fails to prune effectively, or if the hybrid layer yields no interpretive gain, the corresponding element of the framework must be revised or withdrawn. The objections therefore function not as external barriers but as internal design constraints that the ideonomic reconstruction is obliged to satisfy.

The next section turns from the enumeration of objections to a broader assessment of the framework's residual limitations and to the open research problems that the present reconstruction leaves unresolved.

X. Conclusion

10.1 Conclusion and Future Work

The central argument of this paper may now be restated with the precision that the intervening constructions have made available. Patrick Gunkel's ideonomic system of binomial divisions, originally advanced as a general science of ideas, can be reconstructed as an unfinished theory of conceptual dimensionality. The reconstruction does not claim that Gunkel anticipated tensor calculus or neural networks. It claims only that the generative operator he isolated, the division $X \rightarrow (A, B)$, supplies a principled starting point for a methodology that identifies distinctions, organizes them recursively, tests their independence, converts them into continuous coordinates, and models their interactions

by means of tensors. What begins as a qualitative act of polar differentiation ends as a quantitative representational system whose axes remain semantically transparent.

The methodological chain that realizes this reconstruction runs as follows. Ideonomy generates candidate conceptual distinctions through systematic, expert, corpus-based, and model-assisted proposal. Those distinctions are filtered by the ideonomic dimensionality criterion and ranked by the composite utility score

$\Phi(D_i) = \alpha M_i + \beta O_i + \gamma N_i + \delta S_i$, yielding a selected set $\mathcal{D}^*$. Each retained distinction is realized as a continuous coordinate $x_i \in [-1, 1]$, so that entities become points $\mathbf{x}$ or projected points $\mathbf{z} = W\mathbf{h}$ in a vector space whose dimensions carry explicit polar semantics. The joint structure of the selected dimensions is then encoded in an ideonomic tensor $T \in V_{D_1} \otimes \cdots \otimes V_{D_k}$, whose components record measured co-variation and higher-order interaction. In compressed form the chain is: Ideonomy $\rightarrow$ Conceptual Distinctions $\rightarrow$ Dimensions $\rightarrow$ Vector Spaces $\rightarrow$ Tensor Relations. Each arrow corresponds to a concrete procedural step defended in the body of the paper; none is merely metaphorical.

For large language models the practical upshot is an interpretable coordinate system layered onto latent semantic space. The hidden state $\mathbf{h} \in \mathbb{R}^d$ continues to support the fluent prediction and broad generalization for which the model was trained. The ideonomic projection $\mathbf{z} = W\mathbf{h}$ extracts a low-dimensional, human-legible shadow of that state, enabling diagnosis, controlled steering, and the construction of semantically labeled tensors. The architecture is hybrid by design: designed dimensions do not replace learned embeddings; they instrument them. The moderate position adopted throughout remains in force. Where the ideonomic layer demonstrably improves interpretability or controllability without catastrophic loss of capacity, it earns its place; where it fails to do so, the schema can be revised or set aside for that domain.

The broader question that the reconstruction leaves in view is not new, yet it receives a sharpened formulation. Given an effectively unlimited universe of possible ideas, how can we systematically determine the fundamental distinctions by which those ideas may be organized? Taxonomy alone cannot answer, because it presupposes the very axes it would need to justify. Purely statistical extraction cannot answer, because the directions it recovers are optimized for predictive objectives that need not coincide with human conceptual priorities. Ideonomy, reconstructed as a generative dimensional calculus, offers one coherent response: treat the proposal of structured questions as the primary act, convert their answer spaces into coordinates, validate those coordinates for reproducibility and residual independence, and assemble the survivors into representations that make both location and interaction visible. The response is

incomplete (cultural contingency, approximate orthogonality, and the residual tension between hierarchy and free axial combination all remain live constraints) yet it is precise enough to be tested, revised, and extended.

Future work must therefore move from reconstruction to sustained empirical and formal development. First, the proof-of-concept design outlined for literary fiction requires full implementation: annotation protocols must be executed at scale, reliability metrics reported, redundancy matrices and low-order tensors estimated, and comparative evaluations against latent baselines completed under pre-registered falsification conditions. Second, the tensorial side of the framework should be enriched by systematic exploration of hierarchical and tensor-train decompositions whose topologies are informed by the original division trees, together with quantitative assessment of how well such networks capture residual hierarchical dependence. Third, cross-cultural and cross-lingual dimension sets must be constructed and compared so that the stability component S_i can be evaluated across genuine variation in conceptual repertoire rather than within a single linguistic tradition. Additional directions include the integration of ideonomic regularizers into model training, the development of interactive tools for real-time inspection of $\mathbf{Z}$ during generation, and the extension of the generative algorithm to multimodal settings in which visual or auditory polar contrasts supplement textual ones.

The unfinished character of Gunkel's original project is therefore not a defect to be lamented but a resource to be worked. By supplying an explicit generative operator, a disciplined selection calculus, and a clear interface to contemporary representation learning, the reconstruction converts that unfinished project into a research program whose next steps are concrete. Whether the program ultimately yields durable dimensional schemata or merely sharper tools for probing latent geometry, it will have clarified the conditions under which conceptual distinctions can be made to function as computational dimensions. In an era when the most powerful representational systems remain largely opaque, that clarification is worth pursuing.

References

- Alain, G., & Bengio, Y. (2016). Understanding intermediate layers using linear classifier probes. arXiv:1610.01644.
- Altshuller, G. (1984). Creativity as an Exact Science: The Theory of the Solution of Inventive Problems. Gordon & Breach.
- An, J., Kwak, H., & Ahn, Y.-Y. (2018). SemAxis: A lightweight framework to characterize domain-specific word semantics beyond sentiment. ACL.
- Aristotle. Categories; Porphyry. Isagoge (the "Tree of Porphyry").
- Ashby, W. R. (1956). An Introduction to Cybernetics. Chapman & Hall.
- Bateson, G. (1972). Steps to an Ecology of Mind. Chandler.
- Belinkov, Y. (2022). Probing classifiers: Promises, shortcomings, and advances. Computational Linguistics, 48(1), 207–219.
- Bengio, Y., Courville, A., & Vincent, P. (2013). Representation learning: A review and new perspectives. IEEE TPAMI, 35(8), 1798–1828.
- Björck, Å., & Golub, G. H. (1973). Numerical methods for computing angles between linear subspaces. Mathematics of Computation, 27(123), 579–594.
- Bolukbasi, T., Chang, K.-W., Zou, J., Saligrama, V., & Kalai, A. (2016). Man is to computer programmer as woman is to homemaker? Debiasing word embeddings. NeurIPS 29.
- Borges, J. L. (1942). El idioma analítico de John Wilkins.
- Bowker, G. C., & Star, S. L. (1999). Sorting Things Out: Classification and Its Consequences. MIT Press.
- Boyd, R. L., Blackburn, K. G., & Pennebaker, J. W. (2020). The narrative arc: Revealing core narrative structures through text analysis. Science Advances, 6(32), eaba2196.
- Bricken, T., et al. (2023). Towards monosemanticity: Decomposing language models with dictionary learning. Transformer Circuits Thread (Anthropic).
- Brown, T. B., et al. (2020). Language models are few-shot learners. NeurIPS 33.
- Chen, X., et al. (2016). InfoGAN: Interpretable representation learning by information maximizing generative adversarial nets. NeurIPS 29.
- Coecke, B., Sadrzadeh, M., & Clark, S. (2010). Mathematical foundations for a compositional distributional model of meaning. Linguistic Analysis, 36, 345–384.

- Cohen, J. (1960). A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20(1), 37–46.
- Cover, T. M., & Thomas, J. A. (2006). *Elements of Information Theory* (2nd ed.). Wiley.
- Cruse, D. A. (1986). *Lexical Semantics*. Cambridge University Press.
- Cunningham, H., Ewart, A., Riggs, L., Huben, R., & Sharkey, L. (2023). Sparse autoencoders find highly interpretable features in language models. arXiv:2309.08600.
- Deerwester, S., et al. (1990). Indexing by latent semantic analysis. *JASIS*, 41(6), 391–407.
- Destutt de Tracy, A. L. C. (1801–1815). *Éléments d'idéologie*. Paris.
- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. *NAACL-HLT*.
- Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. arXiv:1702.08608.
- Eco, U. (1995). *The Search for the Perfect Language*. Blackwell.
- Elhage, N., et al. (2022). Toy models of superposition. *Transformer Circuits Thread* (Anthropic).
- Engels, J., et al. (2024). Not all language model features are linear. arXiv:2405.14860.
- Enthoven, G. (2018). Towards a science of ideas [Keynote]. IFKAD 2018. <https://www.imi.nu/userfiles/imi.nu/files/Towards%20a%20science%20of%20ideas%20august%202018.pdf>
- Foucault, M. (1966/1970). *The Order of Things*. Pantheon.
- Ganter, B., & Wille, R. (1999). *Formal Concept Analysis: Mathematical Foundations*. Springer.
- Gärdenfors, P. (2000). *Conceptual Spaces: The Geometry of Thought*. MIT Press.
- Gärdenfors, P. (2014). *The Geometry of Meaning: Semantics Based on Conceptual Spaces*. MIT Press.
- Garg, N., Schiebinger, L., Jurafsky, D., & Zou, J. (2018). Word embeddings quantify 100 years of gender and ethnic stereotypes. *PNAS*, 115(16), E3635–E3644.
- Goldberg, L. R. (1990). An alternative "description of personality": The Big-Five factor structure. *Journal of Personality and Social Psychology*, 59(6), 1216–1229.

Grand, G., Blank, I. A., Pereira, F., & Fedorenko, E. (2022). Semantic projection recovers rich human knowledge of multiple object features from word embeddings. *Nature Human Behaviour*, 6, 975–987.

Grasedyck, L. (2010). Hierarchical singular value decomposition of tensors. *SIAM Journal on Matrix Analysis and Applications*, 31(4), 2029–2054.

Gruber, T. R. (1993). A translation approach to portable ontology specifications. *Knowledge Acquisition*, 5(2), 199–220.

Gunkel, P. (1978). A bibliography on the search for extraterrestrial intelligence (NASA). <https://ntrs.nasa.gov/api/citations/19780013076/downloads/19780013076.pdf>

Gunkel, P. M. (1987–1994). *Ideonomy: The Science of Ideas: Introduction, Foundations, and Applications* (multi-volume monographs; self-published). Digitized at MIT: <https://ideonomy.mit.edu>

Gunkel, P. M. (n.d.). The ideonomic division: Ignorances and agnology. Patrick M. Gunkel personal archives (MC-0772), Series 3, Box 2, Folder 8. MIT Libraries, Department of Distinctive Collections.

Hackbusch, W. (2012). *Tensor Spaces and Numerical Tensor Calculus*. Springer.

Hackbusch, W., & Kühn, S. (2009). A new scheme for the tensor representation. *Journal of Fourier Analysis and Applications*, 15, 706–722.

Halmos, P. R. (1958). *Finite-Dimensional Vector Spaces* (2nd ed.). Van Nostrand.

Hewitt, J., & Manning, C. D. (2019). A structural probe for finding syntax in word representations. *NAACL-HLT*.

Higgins, I., et al. (2017). β -VAE: Learning basic visual concepts with a constrained variational framework. *ICLR*.

Ideonomy: The Science of Ideas (n.d.). MIT legacy site, compiled by Whitman Richards and others. <https://ideonomy.mit.edu>

Jockers, M. L. (2013). *Macroanalysis: Digital Methods and Literary History*. University of Illinois Press.

Kanerva, P. (2009). Hyperdimensional computing. *Cognitive Computation*, 1, 139–159.

Keeney, R. L., & Raiffa, H. (1976). *Decisions with Multiple Objectives*. Wiley.

Kelly, G. A. (1955). *The Psychology of Personal Constructs*. Norton.

- Kim, B., et al. (2018). Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (TCAV). ICML.
- Koh, P. W., et al. (2020). Concept bottleneck models. ICML.
- Kolda, T. G., & Bader, B. W. (2009). Tensor decompositions and applications. *SIAM Review*, 51(3), 455–500.
- Kornblith, S., Norouzi, M., Lee, H., & Hinton, G. (2019). Similarity of neural network representations revisited. ICML.
- Kozłowski, A. C., Taddy, M., & Evans, J. A. (2019). The geometry of culture: Analyzing the meanings of class through word embeddings. *American Sociological Review*, 84(5), 905–949.
- Krippendorff, K. (2018). *Content Analysis: An Introduction to Its Methodology* (4th ed.). Sage.
- Kruskal, J. B., & Wish, M. (1978). *Multidimensional Scaling*. Sage.
- Lakoff, G. (1987). *Women, Fire, and Dangerous Things*. University of Chicago Press.
- Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, 104(2), 211–240.
- Leibniz, G. W. (1666). *Dissertatio de Arte Combinatoria*.
- Lenci, A. (2018). Distributional models of word meaning. *Annual Review of Linguistics*, 4, 151–171.
- Li, K., et al. (2023). Inference-time intervention: Eliciting truthful answers from a language model. *NeurIPS* 36.
- Lipton, Z. C. (2018). The mythos of model interpretability. *Communications of the ACM*, 61(10), 36–43.
- Llull, R. (c. 1305). *Ars Magna / Ars Generalis Ultima*.
- Locatello, F., et al. (2019). Challenging common assumptions in the unsupervised learning of disentangled representations. ICML.
- Marks, S., & Tegmark, M. (2023). The geometry of truth: Emergent linear structure in large language model representations of true/false datasets. *arXiv:2310.06824*.

- Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. arXiv:1301.3781.
- Mikolov, T., Yih, W., & Zweig, G. (2013). Linguistic regularities in continuous space word representations. NAACL-HLT.
- Murphy, M. L. (2003). *Semantic Relations and the Lexicon*. Cambridge University Press.
- Nickel, M., Tresp, V., & Kriegel, H.-P. (2011). A three-way model for collective learning on multi-relational data (RESCAL). ICML.
- Novikov, A., Podoprikin, D., Osokin, A., & Vetrov, D. (2015). Tensorizing neural networks. NeurIPS 28.
- Nunnally, J. C., & Bernstein, I. H. (1994). *Psychometric Theory* (3rd ed.). McGraw-Hill.
- Orús, R. (2014). A practical introduction to tensor networks: Matrix product states and projected entangled pair states. *Annals of Physics*, 349, 117–158.
- Oseledets, I. V. (2011). Tensor-train decomposition. *SIAM Journal on Scientific Computing*, 33(5), 2295–2317.
- Osgood, C. E., Suci, G. J., & Tannenbaum, P. H. (1957). *The Measurement of Meaning*. University of Illinois Press.
- Park, K., Choe, Y. J., & Veitch, V. (2024). The linear representation hypothesis and the geometry of large language models. ICML.
- Patrick M. Gunkel personal archives, 1984–2017 (MC-0772), including "Ideonomic divisions drafts." MIT Libraries, Department of Distinctive Collections.
<https://archivesspace.mit.edu/repositories/2/resources/1353>
- Pennebaker, J. W., Boyd, R. L., Jordan, K., & Blackburn, K. (2015). The development and psychometric properties of LIWC2015. University of Texas at Austin.
- Pennington, J., Socher, R., & Manning, C. D. (2014). GloVe: Global vectors for word representation. EMNLP.
- Plate, T. A. (1995). Holographic reduced representations. *IEEE Transactions on Neural Networks*, 6(3), 623–641.
- Raghu, M., Gilmer, J., Yosinski, J., & Sohl-Dickstein, J. (2017). SVCCA: Singular vector canonical correlation analysis. NeurIPS 30.
- Ritchey, T. (2011). *Wicked Problems–Social Messes: Decision Support Modelling with Morphological Analysis*. Springer

- Rosch, E. (1975). Cognitive representations of semantic categories. *Journal of Experimental Psychology: General*, 104(3), 192–233.
- Shepard, R. N. (1962). The analysis of proximities: Multidimensional scaling with an unknown distance function. *Psychometrika*, 27, 125–140, 219–246.
- Shi, Y.-Y., Duan, L.-M., & Vidal, G. (2006). Classical simulation of quantum many-body systems with a tree tensor network. *Physical Review A*, 74, 022320.
- Smolensky, P. (1990). Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artificial Intelligence*, 46(1–2), 159–216.
- Smolensky, P., & Legendre, G. (2006). *The Harmonic Mind*. MIT Press.
- Socher, R., et al. (2013). Recursive deep models for semantic compositionality over a sentiment treebank. *EMNLP*.
- Stipp, D. (1987, June 1). Patrick Gunkel is an idea man who thinks in lists. *The Wall Street Journal*, p. 1.
- Stoudenmire, E. M., & Schwab, D. J. (2016). Supervised learning with tensor networks. *NeurIPS* 29.
- Templeton, A., et al. (2024). Scaling monosemanticity: Extracting interpretable features from Claude 3 Sonnet. *Transformer Circuits Thread (Anthropic)*
- Turner, A., et al. (2023). Activation addition: Steering language models without optimization. *arXiv:2308.10248*.
- Turney, P. D., & Pantel, P. (2010). From frequency to meaning: Vector space models of semantics. *JAIR*, 37, 141–188.
- Underwood, T. (2019). *Distant Horizons: Digital Evidence and Literary Change*. University of Chicago Press.
- Vaswani, A., et al. (2017). Attention is all you need. *NeurIPS* 30.
- Wierzbicka, A. (1996). *Semantics: Primes and Universals*. Oxford University Press.
- Zadeh, L. A. (1965). Fuzzy sets. *Information and Control*, 8(3), 338–353.
- Zou, A., et al. (2023). Representation engineering: A top-down approach to AI transparency. *arXiv:2310.01405*.
- Zwicky, F. (1969). *Discovery, Invention, Research Through the Morphological Approach*. Macmillan.