A Novelty–Plausibility Recursive Framework for Creative Reasoning in Epistemic Idea-Spaces

Kalin Stoyanov *

Posted Date: 11 March 2026

doi: 10.20944/preprints202603.0888.v1

Keywords: creativity; analogical reasoning; divergence measures; Jensen–Shannon divergence; energybased evaluation; hierarchical knowledge structures; information theory; information geometry; generative reasoning; computational creativity

 doi:10.20944/preprints202603.0888.v1

Article

A Novelty–Plausibility Recursive Framework for Creative Reasoning in Epistemic Idea-Spaces

Kalin Stoyanov

University of Chemical Technology and Metallurgy (UCTM), Sofia, Bulgaria; kalin.stoyanov@uctm.edu

Highlights

- Introduces a unified information-theoretic geometry of creativity that couples Jensen–Shannon divergence as a diagnostic of structural novelty with energy-based measures as diagnostics of coherence and admissibility.
- Establishes a generative framework for structured creative reasoning based on feature representations and lawful transformation systems, yielding a domain-agnostic foundation applicable across symbolic, scientific, and artistic domains.
- Interprets divergent (within-domain) exploration and analogical (cross-domain) transfer as distinct operating regimes of a single divergence–energy calculus, rather than as separate cognitive or algorithmic mechanisms.
- Derives explicit quantitative criteria governing feasible novelty, peripheral exploration, boundary-sensitive transitions between structured knowledge regions, and admissible cross-domain alignment.
- Proposes a recursive, multiscale creative search strategy operating over stabilized structural references and progressively expanding neighborhoods.
- Demonstrates the framework through a musical case study (*J. S. Bach's* Art of Fugue), while emphasizing applicability to mathematical, scientific, and symbolic reasoning tasks.
- Advances the view that creativity—human or artificial—arises from the construction of task-directed sequences of intermediate representations within a stabilized epistemic structure, evaluated through divergence and plausibility diagnostics.
- Specifies an explicit mechanism for reaching task-realizing epistemic states through structured exploratory processes, without treating intelligence as an intrinsic or static property of agents.

Abstract

We present an information-theoretic framework that models creative reasoning as structured, task-directed motion within hierarchically organized epistemic structures. Creative problem solving is described as a progressive reconfiguration of an initial epistemic state toward a desired target configuration through a sequence of intermediate representations. Each representation is modeled by an empirical structural distribution over extracted features, enabling two complementary quantitative diagnostics: (i) a divergence measure—the Jensen–Shannon (JS) metric—capturing structural departure, novelty, and analogical proximity; and (ii) energy-based plausibility measures expressing conformity to dominant structural regularities and agent-relative constraints. Their interaction induces a geometry in which exploration balances novelty against structural admissibility, and cross-domain transfer is enabled through alignment of compatible representations. We introduce algebraic and probabilistic principles governing the generation, evaluation, and selection of candidate representations, including neighborhood-restricted exploration, history-sensitive evaluation, and non-redundant comparison under progressively refined interpretive conditions. The framework is operationalized at the level of the epistemic structures accessible to an individual reasoning agent, while large language models are interpreted as mechanisms that facilitate access to broader reservoirs of structured knowledge. Although a musical case study (J. S. Bach's *The Art of Fugue*) is used for illustration, the proposed framework is domain-general and applies to any setting involving structured representations and lawful

 doi:10.20944/preprints202603.0888.v1

transformations. The resulting formalism supports principled approaches to task-oriented creative search, analogical reasoning, and autonomous knowledge exploration, with potential implications for machine-assisted discovery and structurally grounded communication across intelligent systems.

Keywords: creativity; analogical reasoning; divergence measures; Jensen–Shannon divergence; energy-based evaluation; hierarchical knowledge structures; information theory; information geometry; generative reasoning; computational creativity

---

1. Introduction and Theoretical Background

What renders an idea creatively "good"? Across philosophy, cognitive science, and artificial intelligence, creative ideas are consistently characterized by a dual requirement: they exhibit *novelty*, departing from established patterns, while maintaining *coherence*, preserving sufficient structure to remain interpretable, viable, and useful. Despite broad agreement on this principle, novelty and coherence are rarely integrated within a unified quantitative framework that operates across multiple epistemic scales and supports explicit mechanisms of search, transition, and termination.

This paper advances the thesis that creativity is best understood as *structured, goal-directed motion within an epistemic space*. Ideas are treated as probabilistic–structural objects, each represented by a distribution over epistemic features. Creative reasoning unfolds as a *Chain of Thought* (CoT): a sequence of intermediate epistemic states that progressively transforms a present configuration into one that is sufficiently aligned with a specified *task or goal*. This process takes place within the agent's internally organized knowledge structure, termed the *nano Ontos*—the finite, structured collection of Idea-Spaces and Idea-Universes accessible to the agent. The nano Ontos is introduced in Section 2 and operationalized in Section 5.

Creativity as task-driven chain construction

A central hypothesis of this work is that creative reasoning is inherently *task-driven*. Cognitive and computational studies of problem solving consistently characterize thinking as navigation from an initial state toward a goal state through intermediate representations [1–3]. Within this view, creativity proceeds through the construction of a path—a Chain of Thought—whose successive elements reduce epistemic distance to the task while remaining admissible within the agent's knowledge structure.

In this respect, the proposed framework exhibits a structural correspondence to classical heuristic accounts of problem solving, most notably those articulated by George Pólya. Identification of a present epistemic state parallels problem understanding; selection of admissible continuations and recursive descent into subdomains parallels plan construction and execution; and task-based termination corresponds to verification. The present framework formalizes this heuristic logic through explicit divergence- and energy-based diagnostics, enabling algorithmic realization while preserving adaptive revision of strategy.

This perspective also aligns with representational change theories of insight, which locate creative breakthroughs at the boundaries of existing representations [4,5]. As exploration within a given Idea-Space saturates, further progress toward the task proceeds through controlled expansion into neighboring Idea-Spaces or Idea-Universes. Such transitions are mediated by analogy and constrained by structural feasibility, rather than by local similarity alone.

Accordingly, creativity is operationally defined in this work as the generation of a task-directed *Chain of Thought*. Individual ideas correspond to intermediate epistemic states along this chain, while creative novelty emerges from their structured sequencing, alignment, and transformation in service of a well-defined task objective.

Divergence and energy as diagnostic coordinates

Two quantitative diagnostics provide the coordinates by which creative motion is evaluated throughout the framework:

 doi:10.20944/preprints202603.0888.v1

- the **Jensen–Shannon (JS) divergence**, which measures structural separation between ideas and canonical reference configurations; and
- **domain and agent energy functionals**, which measure plausibility, grammatical conformity, and adaptation cost relative to canonical structure and agent valuation.

JS divergence quantifies deviation from a reference configuration—whether a prior idea, a target idea representing the task, or a domain canonical. Energy evaluates whether such deviation remains admissible under the structural and evaluative constraints of the active epistemic unit. Together, these quantities function as *diagnostic coordinates* that regulate admissible continuations of a Chain of Thought.

Analogy occupies a central role in cognitive theories of creativity. In particular, Douglas Hofstadter emphasized analogy as a fluid, context-sensitive process through which meaning emerges via structure-preserving reinterpretation across domains [6,7]. Within this view, analogical relevance arises through the trajectory by which representations are accessed and re-framed.

The present framework adopts this process-oriented interpretation. Analogical relevance is treated as a path-dependent property shaped by the active Chain of Thought and by the canonical structures currently in focus. An idea encountered via different epistemic paths may therefore assume distinct functional roles. This dependence is formalized through context-conditioned admissibility and alignment, ensuring that revisiting an idea under refined context constitutes a genuine analogical reinterpretation.

Parallels with heuristic and latent-space search

From an algorithmic standpoint, the framework exhibits structural similarities to heuristic search. As in best-first or A*-like procedures, candidate states are evaluated relative to a goal using diagnostic quantities. Divergence and energy regulate admissibility locally without imposing a global priority ordering over a fixed state space. Search proceeds through canonical-centered expansion of epistemic neighborhoods, with early termination whenever admissible progress toward the task is achieved.

Parallels with world models and latent predictive architectures

Recent advances in self-supervised learning emphasize internal *world models*: latent predictive structures that encode regularities of an environment and support planning through internal dynamics. Architectures such as LeJEPA [8,9] exemplify this approach by operating entirely in structured latent spaces, where prediction and evaluation are carried out through energy-based criteria without explicit generative decoding.

The present framework aligns with this perspective at the level of representation and evaluation. Canonical distributions serve as compact encodings of dominant regularities within epistemic domains, Jensen–Shannon divergence provides a diagnostic coordinate in latent space, and energy functionals regulate the feasibility of transitions within the structured epistemic model. The contribution extends these principles to multiscale, task-driven creative reasoning by introducing canonical-centered search, context-refined admissibility, and analogical transition across hierarchically organized epistemic spaces.

Creativity as information geometry

Within the proposed framework, creative reasoning corresponds to navigation in a non-Euclidean epistemic landscape. JS divergence governs local exploration within Idea-Spaces, while energy and analogy regulate non-local transitions across Idea-Spaces, Idea-Universes, and ultimately the nano Ontos. This multiscale geometry motivates the recursive zoom-in/zoom-out search strategy developed in later sections. Exploration is interpreted as structured epistemic motion guided by canonical references.

The framework is formulated as an abstract epistemic mechanism, remaining consistent with established patterns of creative reasoning across domains.

 doi:10.20944/preprints202603.0888.v1

A Running Illustrative Example: *The Art of Fugue*

Throughout the paper, J. S. Bach's *The Art of Fugue (BWV 1080)* is employed as a recurring illustrative musical corpus. The work comprises a systematic exploration of contrapuntal technique based on a single principal subject and its derived variants, developed through a controlled sequence of fugal and canonic procedures. Its construction exhibits explicit formal operations, including inversion, augmentation, diminution, stretto, and contrapuntal combination, applied within a rigorously constrained tonal and voice-leading framework. These properties render the work particularly suitable for formalizing musical structure as an Idea-Space governed by lawful transformations and well-defined stylistic constraints.

Counterpoint defines admissible operations—such as inversion, augmentation, diminution, and stretto—that generate closed Idea-Spaces under transformation semigroups. Individual fugues are treated as ideas, while the fugue form constitutes an Idea-Space. The broader musical context functions as an Idea-Universe encompassing neighboring spaces such as the chorale and, at greater structural distance, sonata form, blues, jazz, and other musical practices governed by distinct generative principles.

Within this domain, Jensen–Shannon divergence measures departure from canonical stylistic tendencies, while domain energy quantifies conformity to contrapuntal grammar. The same divergence–energy diagnostics apply, *mutatis mutandis*, to mathematics, logic, engineering design, and other structured creative domains.

Motivation and Contributions

The present work unifies novelty and coherence by treating divergence and plausibility-based energy as complementary diagnostics governing structured creative motion. The main contributions are:

- an information–geometric divergence–energy framework for evaluating creative reasoning;
- a hierarchical organization of epistemic structures linking individual representations to progressively broader knowledge domains;
- a quantitative, path-dependent treatment of analogy based on alignment, divergence, and plausibility diagnostics; and
- a task-driven, recursive creative search procedure operating over the epistemic structures accessible to an individual reasoning agent.

Overall, we advance the view that creativity—human or artificial—arises from the construction of goal-directed sequences of intermediate representations within a stabilized, multiscale epistemic geometry, evaluated through divergence and plausibility rather than driven by them.

Preprint status. This manuscript is presented as a preprint to facilitate early dissemination and discussion of the proposed framework. Several components are intentionally provided in schematic or exploratory form and will be developed further in subsequent work.

2. Epistemic Hierarchy and Accessibility

Creative structure in the proposed framework is organized as a *hierarchical epistemic geometry*. At its core, the framework relies on a compact set of recursively related constructs—*Ideas*, *Idea-Spaces*, *Idea-Universes*, and the *Ontos*[1]—which together define the global architecture within which creative reasoning unfolds.

The same diagnostic quantities—empirical and canonical distributions, divergences, energy functionals, and alignment operators—apply uniformly across all levels of the hierarchy. This uniformity enables recursive, scale-invariant creative dynamics and supports coherent reasoning across representational granularities.

[1] The term *Ontos* is used here as a technical designation for a structured space of epistemic domains together with their analogical relations. It derives from the Greek ὄντος ("of being") and is employed purely as a formal descriptor, without metaphysical commitments

 doi:10.20944/preprints202603.0888.v1

The core structural units further admit finer-grained substructures in the form of *mini Idea-Spaces* and *mini Idea-Universes*, which, at the level of individual instantiation, are encapsulated within distinct *nano Ontoi*.

Core Epistemic Structure

We first introduce the structural backbone of the hierarchy, independent of any particular biological or artificial instantiation. Level numbering is global across the hierarchy; auxiliary representational levels are introduced separately.

Ideas An *idea* constitutes the minimal coherent epistemic object manipulated during creative reasoning (e.g. a musical phrase, a lemma, a visual motif, or a symbolic construction). Formally, an idea S is represented via a feature extractor $\phi(\cdot)$—human- or algorithm-dependent—producing a finite sequence or multiset of features, together with the induced empirical distribution p_S. When internal ordering is relevant, the representation may include an auxiliary ordered descriptor capturing local dependency relations within the idea itself, without imposing any dynamical assumptions on the creative search process.

Ontos The *Ontos* constitutes the global space of epistemic domains. Its nodes correspond to Idea-Universes, and its structure encodes deep cross-domain regularities such as hierarchy, modularity, and recursion. Creative motion at this level consists of cross-domain analogy certified by low divergence after admissible structural alignment and acceptable energy.

Representational Foundations

We next introduce representational levels that supply the atomic units and local coherent groupings from which the core epistemic structure is constructed. These levels provide the scale-relative building blocks underlying the recursive character of the hierarchy.

Epistemic Elements At the foundation lie *epistemic elements*, the atomic units from which higher-level structures are constructed. These elements are domain-dependent (e.g. pitch–duration pairs in music, symbols in logic, geometric primitives in vision, or tokens in machine learning models). Epistemic elements are inherently *scale-relative*: depending on the level of analysis, what functions as an element may itself be an idea, a mini Idea-Space, or an entire Idea-Universe. This scale dependence supports recursive zoom-in/zoom-out exploration.

Mini Idea-Spaces A *mini Idea-Space* is a locally coherent restriction of an Idea-Space, generated by a constrained sub-semigroup of transformations and associated with its own canonical distribution. Mini Idea-Spaces arise naturally as locally stable regions encountered during shell expansion around the canonical center of their parent space.

Example Within the *Art of Fugue* Idea-Space Ω_{AoF}, distinct mini Idea-Spaces arise corresponding to *fugue* and *canon*. Each occupies a different local shell around the shared canonical center q_{AoF} and contributes representative ideas to the BoT during local expansion.

Mini Idea-Universes A *mini Idea-Universe* is a coherent cluster of Idea-Spaces whose canonicals occupy a bounded peripheral shell within a larger Idea-Universe. Mini Idea-Universes represent intermediate-scale specializations that support efficient universe-level exploration and transfer.

Example Jazz constitutes a mini Idea-Universe within the Music Idea-Universe. Its internal Idea-Spaces (e.g. Bebop, Cool Jazz, Modal Jazz [10]) populate neighboring shells around a shared jazz-level canonical.

Operational Epistemic Unit: The Nano Ontos

At the operational level relevant for this paper, the epistemic hierarchy is instantiated as a *nano Ontos*. The nano Ontos constitutes the minimal epistemic unit on which all formal constructions, diagnostics, and algorithms introduced in this work are defined. It represents a finite and operationally accessible instantiation of epistemic structure, corresponding either to the internal organization of a human individual or to that of an appropriately trained artificial agent.

 doi:10.20944/preprints202603.0888.v1

A nano Ontos comprises a localized collection of Idea-Universes, Idea-Spaces, and mini Idea-Spaces, together with their associated canonical distributions, admissible transformations, and alignment constraints. The *Ontos*, by contrast, denotes the unique civilization-scale epistemic structure that subsumes and organizes the totality of such individual instantiations. Although the global Ontos may be regarded as emerging from the aggregation and historical accumulation of multiple nano Ontoi, all creative reasoning processes and algorithmic evaluations considered in this work are carried out exclusively at the level of a single nano Ontos.

Access from a nano Ontos to the global Ontos is mediated through the *representational and inferential capacities* of large language models. As illustrated in Figure 1, an LLM functions as an *access operator*, enabling the selective instantiation, projection, and alignment of Ontos-level structures within the finite epistemic capacity of an operational agent. This mediation preserves the distinction between the Ontos and its operational instantiations, while providing a controlled, task-conditioned interface through which canonical references, structural regularities, and cross-domain relations become locally accessible.

Within this framework, large language models assume a dual and conceptually distinct role. First, they operate as access operators that mediate epistemic interaction between a human nano Ontos and the global Ontos. Second, when considered as autonomous artificial agents, large language models themselves constitute *artificial nano Ontoi*, endowed with internal representational structures, learned canonicals, and admissible transformation rules acquired through training.

Accordingly, the framework admits both human and artificial nano Ontoi as operational epistemic units. In both cases, creative reasoning, analogical transfer, and epistemic exploration proceed internally within a nano Ontos, while remaining structurally grounded in the civilization-scale Ontos through LLM-mediated access.

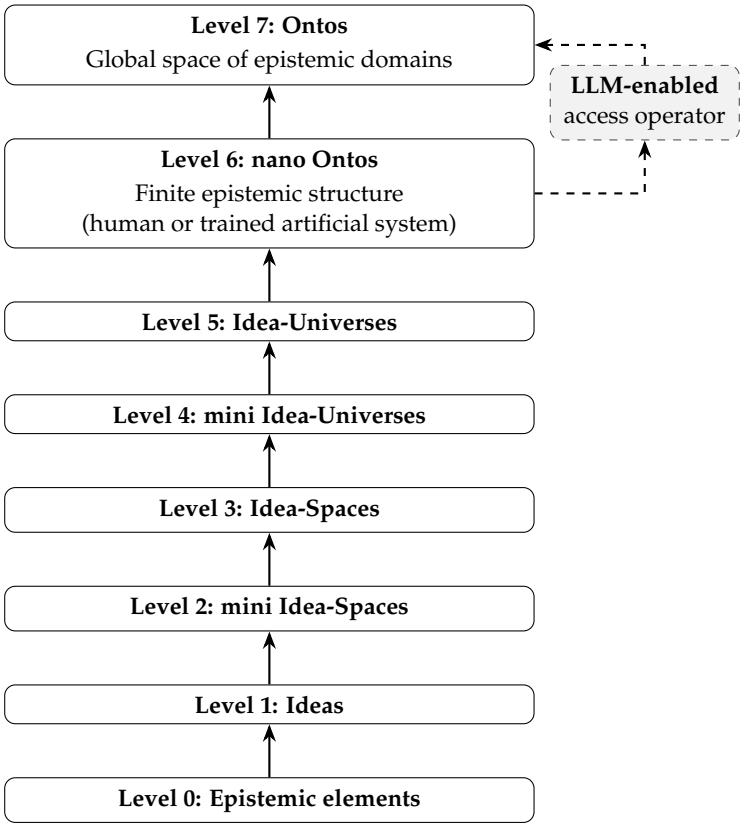

Figure 1. Epistemic hierarchy with explicit LLM-mediated access. Solid arrows encode hierarchical epistemic relations and domain-level structural inclusion, while dashed arrows represent operational access mediated by large language models. The Ontos (Level 7) denotes the global epistemic space, and the nano Ontos (Level 6) denotes a finite, operationally accessible epistemic instantiation corresponding to a human individual or a trained artificial system. Levels 0 through 5 represent the constituent building blocks of the epistemic hierarchy

 doi:10.20944/preprints202603.0888.v1

3. Epistemic Preconditions

We posit the following epistemic preconditions as a single, coherent hypothesis concerning the formation, organization, and operation of creative reasoning. They specify the structural conditions under which the axiomatic and algorithmic framework developed later becomes meaningful and applicable.

The epistemic preconditions operate at a pre-formal level. They motivate the emergence of ideas, canonical references, epistemic boundaries, contextual refinement, and task-directed exploration in creative cognition, while the axioms introduced later specify how these notions behave once assumed.

Formation and stabilization of epistemic clusters Early epistemic experience—whether in human cognitive development or during the training of an artificial system—gives rise to elementary, self-contained epistemic units. These units emerge as localized representations or ideas, shaped by exposure, perceptual regularities, and representational constraints, and may be regarded as autonomous epistemic entities prior to any higher-order organization. Empirical studies of representation learning and language model training corroborate this view, demonstrating that during early phases of training, localized and relatively autonomous representational regularities emerge before the formation of globally coherent semantic structures [11–16].

With continued interaction, task engagement, and repeated structural alignment, these elementary units progressively organize into coherent clusters [16–18]. We hypothesize that such stabilized clusters correspond to *Bags of Thoughts* (BoTs): finite collections of ideas admitting a shared canonical reference that captures their dominant regularities. Canonical references emerge through iterative alignment and reinforcement and stabilize as experience accumulates, yielding persistent epistemic structures that support efficient reasoning and analogical transfer.

With further exposure and specialization, stabilized clusters differentiate into structured epistemic domains [19–21]. At this stage, they acquire the form of Idea-Spaces and Idea-Universes, each organized around a canonical center supporting graded notions of typicality and peripheral deviation.

Epistemic boundaries and adaptive knowledge growth Empirical evidence indicates that functional specialization and modular organization emerge early in cognitive development and become increasingly pronounced with experience [22,23]. Consistent with these findings, we hypothesize that epistemic boundaries arise naturally during the formation of epistemic clusters and persist as stable structural features of the epistemic landscape.

Canonical centers and their associated peripheral regions therefore arise not only within individual clusters (BoTs, Idea-Spaces, and Idea-Universes), but also at the level of the agent's accessible epistemic horizon, the nano Ontos. These boundaries become salient as exploration approaches regions of increased structural deviation and interpretive ambiguity.

Epistemic boundaries function primarily as diagnostic structures. They signal departure from established regularities and indicate conditions under which reinterpretation, analogical transfer, or transition to a different epistemic unit becomes appropriate.

Crossing such boundaries—whether between Idea-Spaces, between Idea-Universes, or at the edge of the nano Ontos—corresponds to epistemic expansion. This expansion is characterized by the stabilization of new canonical references and the enlargement of the agent's active epistemic horizon, providing a structural mechanism for adaptive knowledge growth.

This interpretation aligns with theories of conceptual change and representational restructuring [24,25] and supports the view that lifelong (continual) learning emerges intrinsically from creative search [13,26].

Tasks as epistemic targets Creative reasoning is initiated by a task, understood as a target epistemic configuration to be achieved through exploration. At the level of epistemic preconditions, we assume that tasks can be represented at the relevant epistemic scale—either directly as ideas or through admissible structural reinterpretations—and accordingly treated as target ideas.

The essential requirement is mutual interpretability: the present idea and the target idea admit structure-preserving relations that allow meaningful comparison within an active context. This

 doi:10.20944/preprints202603.0888.v1

assumption enables creative reasoning to be formulated entirely in terms of ideas, without introducing tasks as a separate class of epistemic objects.

Creative search as canonical-centered expansion Given a target epistemic state, represented by target idea S_T and a current epistemic state represented by a present idea S_0, creative reasoning proceeds through the construction of a *Chain of Thought* (CoT),

$$S_0 \rightarrow S_1 \rightarrow \cdots \rightarrow S_T,$$

which incrementally reconfigures the present idea toward the target idea. A Chain of Thought is understood as an epistemic trajectory generated by canonical-guided exploration and context-conditioned evaluation.

Creative search unfolds as a structured expansion centered on canonical references. At each stage, exploration begins in the vicinity of the canonical center of the active epistemic unit and extends outward only when closer regions fail to yield admissible continuations. Candidate ideas are therefore encountered in an order reflecting increasing structural deviation from the dominant regularities of the current domain.

As exploration progresses, peripheral regions are reached in which ideas or subdomains admit multiple admissible interpretations. These regions support analogical transfer and domain transition through local reinterpretation rather than global reorganization of the epistemic landscape. Recursive descent into subdomains corresponds to increased contextual specificity, whereas outward expansion reflects progressively relaxed structural constraints and an expanded analogical reach.

Contextual refinement and non-redundant exploration Creative reasoning is path-dependent: the interpretation of an idea depends on the sequence of epistemic transitions through which it is reached. Accordingly, we hypothesize that creative exploration proceeds through contextual refinement. Revisiting an idea under a strictly more informative or more specific context constitutes a meaningful reinterpretation and may support new admissible continuations.

This principle motivates search procedures based on monotonic contextual refinement, ensuring that identical structural comparisons are not repeatedly applied under equivalent epistemic conditions.

Termination by task realization Creative search terminates when the present idea becomes sufficiently close to the target idea at the active epistemic scale. Exact coincidence is not required; termination occurs when the present idea lies within an admissible neighborhood of the target idea. This notion of termination is formalized later through analogy-based distance measures and acceptance thresholds.

4. Conceptual and Quantitative Foundations

This section formalizes the quantitative machinery used to model *creative search* as structured motion across epistemic units. Creative search is initiated by a specified *task* T, representing an intended epistemic achievement, together with a *present epistemic state* representing existing and already established knowledge. Both the task and the present state are represented as idea-like objects admitting feature-based representations (possibly after alignment), referred to respectively as the *target idea* and the *present idea*. The search process terminates when the present idea enters a sufficiently small canonical neighborhood of the target idea T.

We assume a hierarchy of epistemic units: *Idea-Spaces* (locally coherent spaces generated by lawful transformations), *Idea-Universes* (collections of related Idea-Spaces), and the global *Ontos*. The *nano Ontos* of an agent is the finite subset of this hierarchy that is cognitively or operationally accessible at a given time. Creative search is evaluated and controlled by a designated *agent* (an individual or a collective acting as a single decision authority).

Artificial agents may generate and conditionally approve intermediate ideas within a Chain of Thought, while final admissibility is always determined by the human agent. The ethical scope of artificial agents' authority and the corresponding allocation of human responsibility are beyond the scope of the present work and will be addressed in forthcoming research.

 doi:10.20944/preprints202603.0888.v1

Representations An *idea* S is represented by an empirical distribution p_S over a finite feature alphabet X_D determined by the active epistemic unit D.

A task T is operationally represented by a designated *target idea*. When the present and target ideas reside in different epistemic units, the active context activates an admissible alignment frame used in their comparison.

When internal ordering or sequential structure is relevant, the representation may include an auxiliary ordered descriptor capturing local dependency relations *within the idea itself*. This applies both to intrinsically temporal objects, such as musical works (e.g. fugues), and to non-temporal but order-sensitive constructs, such as mathematical proofs or theorems, where the arrangement and dependency of elements is essential to interpretation.

Such ordered descriptors serve as representational devices that encode the internal dependency structure of an idea. They characterize how constituent elements are arranged and related within the idea, independently of the dynamics governing creative search or transitions between ideas.

Each epistemic unit D is summarized at its scale by a canonical distribution q_D, which serves as a reference for structural comparison and plausibility evaluation.

Diagnostics Creative evaluation relies on three coupled diagnostics: (i) Jensen–Shannon (JS) divergence, quantifying structural separation between ideas; (ii) a *domain energy* functional, quantifying structural admissibility and adaptation cost relative to the canonical structure of the active epistemic unit; and (iii) an *agent energy* functional, quantifying subjective cost or resistance relative to the agent's internal valuation and goals. These components are integrated through an enriched analogy functional that combines divergence, domain energy, agent energy, and alignment regularity to assess whether a candidate idea can serve as the next step in the Chain of Thought toward the task.

The term energy is used in a diagnostic and evaluative sense and does not imply an energy-based generative model or a learned probabilistic dynamics.

Termination criterion Exploration terminates when the present idea enters a task neighborhood

$$d_T(S,T) = \sqrt{\mathcal{J}_\sigma(S;T)} \le \varepsilon_T, \tag{1}$$

or when the peripheral boundary of the active unit is exhausted, triggering transition to adjacent units through overlapping peripheral regions.

4.1. Algebraic, Probabilistic, and Analogical Axioms

The axioms below define the admissible objects, relations, and evaluation principles governing creative search within structured epistemic units. They specify how ideas are generated through lawful transformations, how canonical structures and divergence constrain exploration, and how context-dependent diagnostics regulate recursive traversal across Idea-Spaces and Idea-Universes.

(A1) Closure under lawful transformation Each Idea-Space is generated as the closure of a seed idea under a lawful transformation system:

$$\Omega_\ell = \{\, g(S_0) \mid g \in \mathcal{G}_\ell \,\}. \tag{2}$$

This closure defines the total set of ideas that may become accessible through exploration, though only subsets are examined at any given stage.

(A2) Associativity of transformations Within any Idea-Space, admissible transformations compose associatively:

$$(f \circ g) \circ h = f \circ (g \circ h), \qquad f,g,h \in \mathcal{G}_\ell. \tag{3}$$

This ensures path-consistent generation of ideas under recursive exploration.

(A3) Feature representation Each idea admits a finite structural feature representation over a domain-specific alphabet X_D, inducing an empirical distribution p_S on X_D.

(A4) Canonical structure Each epistemic unit D (Idea-Space, Idea-Universe, or Ontos-level unit) possesses a canonical structural reference q_D, defined as the barycentric distribution of the unit. The

 doi:10.20944/preprints202603.0888.v1

canonical structure serves as the geometric reference point for exploration, typicality, and neighborhood construction.

(A5) Jensen–Shannon divergence and induced metric Structural separation between distributions is quantified by Jensen–Shannon divergence. Given distributions p and q on a finite alphabet X, define

$$m(x) = \tfrac{1}{2}\big(p(x) + q(x)\big). \tag{4}$$

Then

$$D_{\mathrm{JS}}(p,q) = \tfrac{1}{2}D_{\mathrm{KL}}(p\|m) + \tfrac{1}{2}D_{\mathrm{KL}}(q\|m), \tag{5}$$

with induced metric

$$d_{\mathrm{JS}}(p,q) = \sqrt{D_{\mathrm{JS}}(p,q)}. \tag{6}$$

When feature alphabets differ, divergence is evaluated within an admissible alignment frame associated with the active comparison.

(A6) Domain energy and agent energy Each epistemic unit D induces a *domain energy* $E_D(S)$ measuring structural strain relative to q_D. In parallel, creative evaluation involves an *agent energy* $E_{\mathrm{agent}}(S,T)$ encoding task- and agent-dependent valuation.

These energies are mapped to plausibility scores via thresholded acceptance functions:

$$\begin{aligned} P_{\mathrm{dom}}(S) &= \sigma\big(-\beta_{\mathrm{dom}}(E_D(S) - \theta_{\mathrm{dom}})\big), \\ P_{\mathrm{agent}}(S) &= \sigma\big(\beta_{\mathrm{agent}}(E_{\mathrm{agent}}(S,T) - \theta_{\mathrm{agent}})\big). \end{aligned} \tag{7}$$

(A7) Conditional creative cost Given a task representation p_T, candidate ideas are evaluated through a context-dependent creative cost

$$\mathcal{J}_\sigma(S;T) = D_{\mathrm{JS}}(p_S, p_T) + \lambda_2\big(1 - P_{\mathrm{dom}}(S)\big) - \lambda_3 P_{\mathrm{agent}}(S). \tag{8}$$

Lower values correspond to ideas that are task-aligned, structurally admissible, and agent-endorsed.

Importantly, $\mathcal{J}_\sigma$ is evaluated only over candidates admissible under the active context and neighborhood restriction. It therefore defines a context-conditioned diagnostic applied within a history-dependent search process, rather than a memoryless transition rule.

(A8) Neighborhood-restricted evaluation At any stage of exploration, admissible candidates are drawn exclusively from the currently explored region of the active epistemic unit, defined as a canonical neighborhood centered at q_D:

$$B_D(r) = \{\, S \in D : d_{\mathrm{JS}}(p_S, q_D) \le r \,\}. \tag{9}$$

Evaluation proceeds hierarchically, beginning with the representative core and extending to successively expanded peripheral layers. **(A9) Context** We introduce a *context* π as a structured epistemic descriptor that specifies the conditions under which candidate ideas are interpreted and evaluated at a given stage of creative exploration. When a Chain of Thought (CoT) is viewed as a directed epistemic graph whose nodes are ideas and whose edges represent admissible epistemic transitions, a context corresponds to a parent subgraph that induces the admissible interpretation and evaluation conditions for the present idea.

Formally, a context is defined as a tuple

$$\pi = (G_\pi, D, r, \mathcal{C}, \mathcal{A}), \tag{10}$$

where G_π is a directed subgraph of the CoT encoding the relevant epistemic history, D denotes the active epistemic unit (Idea-Space or Idea-Universe), r specifies the admissible divergence neighborhood, $\mathcal{C}$ encodes the admissibility constraints currently in force, and $\mathcal{A}$ denotes the active analogical alignment frame.

 doi:10.20944/preprints202603.0888.v1

The subgraph G_π represents a structured projection of prior exploration and need not coincide with the full history of the Chain of Thought. Contexts are not associated with scalar utility values; rather, they define structured conditions that determine which comparisons and transitions are admissible.

(A10) Contextual non-redundancy A candidate idea S is evaluated under a context π only if there exists no prior evaluation of S under a context π' such that

$$\pi \preceq \pi'. \tag{11}$$

Re-evaluation is permitted only under strict context refinement, ensuring that repeated encounters with the same idea correspond to increasingly specific epistemic interpretations rather than redundant comparisons.

(A11) Canonical-level analogy and unit viability Let S_0 denote the present idea with representation p_{S_0}. For a candidate epistemic unit $\bar{D}$ with canonical structure $q_{\bar{D}}$, analogy is assessed through alignment operators

$$\Pi \in \mathbb{R}_{\geq 0}^{|X_D| \times |X_{\bar{D}}|}, \qquad \sum_y \Pi_{xy} = 1.$$

The canonical-level analogy cost is defined as

$$\mathcal{A}(S_0, q_{\bar{D}}; \Pi) = D_{\mathrm{JS}}(\Pi^\top p_{S_0}, q_{\bar{D}}) + \lambda_2 \big(1 - P_{\mathrm{dom}}(q_{\bar{D}})\big) + \lambda_R \mathcal{R}(\Pi). \tag{12}$$

An epistemic unit is viable if this cost is sufficiently small for at least one alignment.

Alignment operators are epistemic reinterpretation maps and are not learned parameters, attention mechanisms, or components of a data-driven training procedure.

(A12) Recursive descent and idea selection If a unit $\bar{D}$ is viable, idea-level exploration proceeds recursively within $\bar{D}$ via sphere expansion around $q_{\bar{D}}$. The next idea is selected as

$$S_j = \arg \min_{S \in B_{\bar{D}}(r)} \big[D_{\mathrm{JS}}(p_S, q_{\bar{D}}) + \lambda E_{\bar{D}}(S) \big], \tag{13}$$

with early termination upon meeting acceptance criteria.

The minimization is local to the current admissible region and does not define a global transition function, as candidate sets are themselves shaped by contextual refinement and prior neighborhood expansion.

Axioms (A1)–(A12) characterize creative search as a canonical-centered, context-refined, and recursively expanding process in which analogy arises through structured traversal of epistemic neighborhoods. Selection operators function as local decision rules applied to context-restricted candidate sets, supporting adaptive progression through the search space. The task neighborhood serves as an epistemic stopping condition that signals sufficient alignment with the target objective.

5. Canonical-Centered, Context-Dependent Exploratory Search

This section formalizes a context-dependent exploratory algorithm induced by the canonical geometry of domains. The algorithm describes a structured process of idea generation guided by canonical centers, divergence-based neighborhoods, and recursive domain transitions. The resulting exploration reflects an analogical mechanism in which meaning emerges progressively as increasingly distant structures become relevant within a given contextual trajectory.

 doi:10.20944/preprints202603.0888.v1

5.1. Domains, Ideas, and Canonical Centers

Let D denote a domain, understood as a set of ideas. Each idea $x \in D$ is represented by a probability distribution P_x in a suitable statistical space. The canonical center of a domain D is defined as the averaged (barycentric) distribution

$$c(D) = \mathbb{E}_{x \sim D}[P_x]. \tag{14}$$

The canonical center serves as a prototypical reference for the domain. There exists a sufficiently small radius $\varepsilon_D > 0$ such that all ideas satisfying

$$\mathcal{D}(P_x \,\|\, c(D)) \leq \varepsilon_D \tag{15}$$

form a representative core of the domain. Ideas in this core are maximally consistent with the dominant structure of the domain and provide the initial reference frame for analogical exploration.

5.2. Present Idea, Present Domain, and Context

At the beginning of each exploratory phase, a *present idea* x_p is specified. Represented by its associated distribution P_{x_p}, the present idea constitutes the most recent state in the evolving chain of thought. Exploration proceeds within a *present domain* D_p, which comprises admissible ideas as well as structured subdomains defined by the domain grammar.

Associated with the present idea is a *context* π, encoding the trajectory through which x_p has been reached. The context determines which structural features of candidate ideas are considered salient and how analogical similarity is assessed. Consequently, analogical relevance is a path-dependent property shaped by the preceding sequence of transformations.

5.3. Divergence Neighborhood and Initial Exploration

Exploration begins by constructing a divergence-based neighborhood of radius r around the canonical center of the present domain,

$$B_{D_p}(r) = \{\, y \in D_p \;:\; \mathcal{D}(P_y \,\|\, c(D_p)) \leq r \,\}. \tag{16}$$

All candidate ideas $y \in B_{D_p}(r)$ are then compared with the present idea x_p using a context-dependent admissibility criterion

$$\mathcal{A}(y \mid x_p, \pi). \tag{17}$$

This stage corresponds to the context-driven retrieval of structurally closest, domain-typical analogues relative to the present idea. If an admissible idea is identified at this level, it is selected as the next element in the chain of thought, and the current exploratory phase terminates.

5.4. Recursive Exploration of Subdomains

The divergence neighborhood $B_{D_p}(r)$ may contain subdomains $D_{\text{sub}} \subset D_p$. Each subdomain is characterized by its own canonical center and induced periphery, which together define the structural organization of the subdomain. When such a subdomain enters the explored region, the exploratory procedure is applied recursively, starting from the canonical center of the subdomain and evaluated relative to the same present idea and context.

This recursive descent reflects a refinement of the analogical frame: at the level of the parent domain, structural regularities are specified implicitly through proximity to the canonical center and its periphery, whereas focusing on a more specialized subdomain renders these regularities explicit by introducing more restrictive divergence thresholds and admissibility conditions. If an admissible idea is identified during recursive exploration, the result is propagated upward and the search terminates.

 doi:10.20944/preprints202603.0888.v1

5.5. Conditional Neighborhood Expansion and Peripheral Exploration

If neither the ideas contained in $B_{D_p}(r)$ nor the recursive exploration of subdomains yields an admissible idea, the divergence neighborhood is expanded to a larger radius $r' > r$. The newly explored region is defined as the peripheral band

$$\text{Periphery}_{D_p}(r, r') = B_{D_p}(r') \setminus B_{D_p}(r). \tag{18}$$

Exploration within the periphery is conducted independently of previously examined regions. Candidate ideas and subdomains located in this peripheral band are processed using the same priority order: ideas are evaluated first, followed by recursive exploration of subdomains. As expansion proceeds, increasingly distant structures become available for analogical consideration, enabling the emergence of non-obvious correspondences.

On divergence neighborhoods and peripheral completeness

The notion of a divergence neighborhood is used in an operational sense and does not imply a geometric ball in either Euclidean or non-Euclidean space. The epistemic space induced by divergence from a canonical center is generally anisotropic and may exhibit highly irregular peripheral structure. Consequently, the boundary of a divergence neighborhood should not be interpreted as smooth, convex, or geometrically uniform.

Exploration therefore does not rely on geometric coverage but on a scalar ordering induced by divergence from the canonical center. The expanding neighborhood $B_{D_p}(r)$ is defined as a level set of the divergence function $\mathcal{D}(P_x \| c(D_p))$, and expansion proceeds monotonically in the scalar parameter r. Every idea $x \in D_p$ is associated with a unique value

$$r_x = \mathcal{D}(P_x \| c(D_p)), \tag{19}$$

at which it becomes included in the explored region.

As a result, no structurally admissible region of the domain can be bypassed due to geometric irregularities. Peripheral regions that may appear disconnected, narrow, or recessed when viewed geometrically are nevertheless encountered when the expansion parameter reaches the corresponding divergence level. Exploration is therefore exhaustive with respect to the divergence ordering, even though the induced peripheral structure may be highly non-uniform.

This interpretation reflects the role of the canonical center as a prototype rather than as a coordinate origin: expansion corresponds to increasing deviation from dominant structural regularities rather than movement along spatial directions. Completeness of exploration is thus guaranteed by the monotonic traversal of divergence levels, not by assumptions of geometric regularity.

5.6. Contextual Non-Repetition and Analogical Refinement

Creative exploration proceeds under progressively refined epistemic contexts. Each context encodes the structural conditions under which candidate ideas are interpreted and evaluated. Contexts are partially ordered by informational refinement: we write $\pi_1 \prec \pi_2$ when context π_2 extends π_1 by introducing increased domain specificity, deeper recursive descent, or stronger structural constraints.

An idea is evaluated only when encountered under a context that has not been previously subsumed. Formally, a candidate idea y is evaluated under a context π only if there exists no prior evaluation under a context π' such that

$$\pi \preceq \pi'. \tag{20}$$

Accordingly, repeated encounters with the same idea occur exclusively under strictly refined contexts. Each such encounter constitutes a distinct epistemic act, corresponding to reinterpretation under enhanced structural resolution. This mechanism reflects the analogical character of creative reasoning, whereby an unchanged structure may assume different functional or conceptual roles as the surrounding epistemic frame becomes more specific.

 doi:10.20944/preprints202603.0888.v1

5.7. Termination and Domain Transition

Neighborhood expansion and contextual refinement proceed iteratively until an admissible idea is identified or the divergence neighborhood reaches the boundary of the present domain. Termination within the present domain is triggered whenever

$$\exists\, y \in B_{D_p}(r) \quad \text{such that} \quad \mathcal{A}(y \mid x_p, \pi) \leq \tau(\pi). \tag{21}$$

If no admissible idea is identified for any r satisfying

$$B_{D_p}(r) = D_p, \tag{22}$$

the present domain is considered exhausted.

Exploration then proceeds in adjacent domains whose peripheral regions overlap with that of the present domain. A domain transition from D_p to a neighboring domain D_q occurs whenever

$$\text{Periphery}_{D_p}(r, r') \cap \text{Periphery}_{D_q}(r_q, r'_q) \neq \varnothing, \tag{23}$$

indicating the presence of ideas or subdomains that admit coherent interpretations under both domain-level structures.

These transitions arise naturally through peripheral overlap, enabling the continuation of analogical exploration across domains without introducing external switching criteria.

6. Numerical Examples

Rationale The numerical examples in this section are intentionally low-dimensional and serve a diagnostic purpose. Their role is to make the operational objects of the framework explicit and to demonstrate how they interact within canonical-centered search. Specifically, the examples (i) show how feature counts induce empirical distributions p_S, (ii) illustrate how a step-level Bag of Thoughts (BoT) yields a canonical reference q for an Idea-Space (IS), (iii) instantiate a domain-level energy functional that penalizes structurally strained deviations relative to the IS canonical, and (iv) expose how the next-idea mechanism summarized in Algorithm 1 operates as controlled exploration of divergence neighborhoods around q. The emphasis is on interpretability of the construction rather than on musical exhaustiveness, making the procedure suitable as a preprint-stage illustration.

The *Art of Fugue* (AoF) is selected as an illustrative domain because it constitutes a stylistically homogeneous corpus with rich internal variation and well-defined transformational devices (e.g. stretto, inversion, augmentation). The chosen feature alphabet captures essential contrapuntal properties while remaining compact, allowing all steps of the construction to be followed numerically without loss of transparency. The counts in Table 1 are extracted under a simplified annotation protocol and are used exclusively to expose the mechanics of distribution construction, canonical estimation, and energy-based admissibility.

Empirical distributions from counts Fix the AoF feature alphabet

$$X_{\text{AoF}} = \{x_1, x_2, x_3, x_4, x_5\}, \tag{24}$$

where x_1 denotes stepwise motion, x_2 leaps, x_3 suspensions, x_4 stretto entries, and x_5 augmentation/diminution cues. For each piece S, let $c_S(x_k)$ be its count (Table 1) and $N_S = \sum_{k=1}^{5} c_S(x_k)$. The associated empirical distribution is

$$p_S(x_k) = \frac{c_S(x_k)}{N_S}, \qquad k = 1, \ldots, 5. \tag{25}$$

 doi:10.20944/preprints202603.0888.v1

Algorithm 1 Canonical-Guided Search for the Next Idea

Require: Present idea S_0, initial epistemic unit D_0, target idea T, initial context π_0, initial radius r_0
Ensure: Next idea S_{next} or FAILURE

Initialize a priority queue $\mathcal{Q}$ of search states
Insert $(D_0, q_{D_0}, r_0, \pi_0)$ into $\mathcal{Q}$
Initialize a contextual registry $\mathcal{H}$ for non-redundant evaluation
while $\mathcal{Q}$ is not empty **do**
  Pop (D, q_D, r, π) from $\mathcal{Q}$
  $r_{\text{prev}} \leftarrow$ PREVIOUSRADIUS(D, π)
  ▷ (1) Neighborhood-band exploration around q_D
  $\mathcal{C} \leftarrow \{ S \in D \ : \ r_{\text{prev}} < d_{\text{JS}}(p_S, q_D) \leq r \}$
  Order $\mathcal{C}$ increasingly by $\mathcal{J}_\sigma(S; T, \pi)$
  for all $S \in \mathcal{C}$ in this order **do**
    if NONREDUNDANT$(\mathcal{H}, S, \pi)$ **then**
      if DOMAINGATE(S, D, π) **then**
        if INTERMEDIATEADMISSIBLE(S, S_0, T, D, π) **then**
          return S
        end if
      end if
      REGISTER$(\mathcal{H}, S, \pi)$
    end if
  end for
  ▷ (2) Recursive descent into viable sub-units entering the same band
  $\mathcal{S} \leftarrow \{ D' \subset D \ : \ r_{\text{prev}} < d_{\text{JS}}(q_{D'}, q_D) \leq r \}$
  for all $D' \in \mathcal{S}$ **do**
    $\pi' \leftarrow$ REFINECONTEXT(π, D', r)
    if NONREDUNDANT$(\mathcal{H}, D', \pi')$ **then**
      if UNITVIABLE(S_0, D', π') **then**
        $r_{\text{init}} \leftarrow$ INITRADIUS(π', r_0)
        Insert $(D', q_{D'}, r_{\text{init}}, \pi')$ into $\mathcal{Q}$
      end if
      REGISTER$(\mathcal{H}, D', \pi')$
    end if
  end for
  ▷ (3) Neighborhood expansion and transition via peripheral overlap
  $r \leftarrow$ EXPANDRADIUS(r, π)
  if WITHINUNIT(D, r) **then**
    Insert (D, q_D, r, π) into $\mathcal{Q}$
  else
    $\mathcal{N}(D) \leftarrow$ NEIGHBORUNITS(D, π)
    for all $\bar{D} \in \mathcal{N}(D)$ **do**
      if PERIPHERALOVERLAP$(D, \bar{D})$ **then**
        if UNITVIABLE$(S_0, \bar{D}, \pi)$ **then**
          $r_{\text{init}} \leftarrow$ INITRADIUS(π, r_0)
          Insert $(\bar{D}, q_{\bar{D}}, r_{\text{init}}, \pi)$ into $\mathcal{Q}$
        end if
      end if
    end for
  end if
end while
return FAILURE

Step-level BoT and IS canonical At a given step t of the Chain of Thought, the step-level BoT for the AoF Idea-Space is a finite set of currently available epistemic units,

$$\mathcal{B}_t^{\text{AoF}} = \{S^{(1)}, \ldots, S^{(m)}\}. \tag{26}$$

For the present numerical instantiation, the step-level BoT is fixed as a diagnostic proxy: it consists of the $m = 18$ pieces constituting *The Art of Fugue*. This choice yields a stable and fully transparent canonical estimate while keeping the construction low-dimensional and reproducible. In the full

 doi:10.20944/preprints202603.0888.v1

search procedure, the BoT is step-dependent and typically smaller, as it is formed from candidates encountered within the active divergence neighborhood during exploration.

Table 1. Extracted feature counts for *Die Kunst der Fuge*. Here Cp. denotes Contrapunctus, inv. inversus, rec. recursus, and unf. unfinished. Counts are illustrative and follow a simplified annotation protocol intended to demonstrate the computational pipeline rather than to provide an exhaustive musicological analysis.

Piece	Stepwise	Leap	Susp.	Stretto	Aug./Dim.	Measure count
Cp. I	36	12	18	0	0	78
Cp. II	36	12	22	0	0	84
Cp. III	32	16	25	0	0	72
Cp. IV	32	16	31	0	0	138
Cp. V	28	18	42	14	0	90
Cp. VI	44	22	38	11	12	79
Cp. VII	48	24	45	13	18	76
Cp. VIII	92	54	56	6	0	188
Cp. IX	52	30	28	4	0	130
Cp. X	64	38	48	8	0	120
Cp. XI	115	68	72	10	0	184
Cp. XII(rec.)	34	14	21	0	0	112
Cp. XII(inv.)	34	14	19	0	0	112
Cp. XIII(rec.)	26	12	14	0	0	90
Cp. XIII(inv.)	26	12	16	0	0	90
Cp. XIV(unf.)	148	82	95	12	0	239
Canon 1	18	10	12	2	0	103
Canon 2	22	14	28	4	0	79
Canon 3	20	16	30	2	0	80
Canon 4	148	82	95	12	0	82

The AoF canonical is defined, for simplicity, as the barycenter of the empirical distributions associated with the BoT,

$$q_{\mathrm{AoF}}(x) = \frac{1}{m}\sum_{j=1}^{m} p_{S^{(j)}}(x). \tag{27}$$

The resulting canonical distribution is

$$q_{\mathrm{AoF}} = (0.3978,\ 0.2212,\ 0.3241,\ 0.0388,\ 0.0181). \tag{28}$$

Energy relative to the IS canonical Plausibility of an idea relative to the AoF Idea-Space is quantified by a cross-entropy energy (bits),

$$E_{\mathrm{AoF}}(S) = -\sum_{k=1}^{5} p_S(x_k)\log_2 q_{\mathrm{AoF}}(x_k). \tag{29}$$

Lower values of $E_{\mathrm{AoF}}(S)$ correspond to greater typicality under the AoF canonical grammar. In canonical-centered search, this energy functions as a domain-level admissibility diagnostic applied prior to task- and context-conditioned evaluation.

Specialization of the divergence geometry For the numerical examples in this section, the abstract divergence $\mathcal{D}$ is instantiated as the Jensen–Shannon divergence d_{JS}, the canonical center $c(D)$ is identified with the empirical canonical q_{AoF}, and idea representations are identified with the empirical distributions p_S. Divergence neighborhoods and peripheral bands are therefore defined directly in terms of $d_{\mathrm{JS}}(p_S, q_{\mathrm{AoF}})$.

6.1. Procedural Realization of Creative Search

The following procedure specifies how a single admissible continuation is selected within the proposed framework. The search is organized around canonical references, exploring divergence-bounded neighborhoods in increasing bands and prioritizing candidates by a combined divergence–

 doi:10.20944/preprints202603.0888.v1

energy cost $\mathcal{J}_\sigma$ under the current context. Contextual refinement and non-redundant evaluation ensure path dependence, while recursive descent and peripheral overlap enable transitions across epistemic units when local exploration is exhausted.

Recomputed distributions and energies For transparency, the empirical distributions p_S and the corresponding AoF energies $E_{\mathrm{AoF}}(S)$ used by the procedure are listed in Table 2. These quantities are recomputed directly from the feature counts in Table 1 and the canonical q_{AoF} above.

Table 2. Empirical distributions p_S and AoF energy $E_{\mathrm{AoF}}(S)$ (bits)

Piece	$p_S = (p_1, p_2, p_3, p_4, p_5)$	$E_{\mathrm{AoF}}(S)$
Cp. I	(0.545, 0.182, 0.273, 0.000, 0.000)	1.531
Cp. II	(0.514, 0.171, 0.314, 0.000, 0.000)	1.539
Cp. III	(0.438, 0.219, 0.342, 0.000, 0.000)	1.597
Cp. IV	(0.405, 0.203, 0.392, 0.000, 0.000)	1.603
Cp. V	(0.275, 0.176, 0.412, 0.137, 0.000)	2.088
Cp. VI	(0.345, 0.173, 0.298, 0.086, 0.094)	1.930
Cp. VII	(0.324, 0.162, 0.304, 0.088, 0.122)	2.004
Cp. VIII	(0.442, 0.260, 0.269, 0.029, 0.000)	1.553
Cp. IX	(0.456, 0.263, 0.246, 0.035, 0.000)	1.548
Cp. X	(0.405, 0.241, 0.304, 0.051, 0.000)	1.659
Cp. XI	(0.434, 0.257, 0.272, 0.038, 0.000)	1.569
Cp. XII(rec.)	(0.492, 0.203, 0.304, 0.000, 0.000)	1.540
Cp. XII(inv.)	(0.507, 0.209, 0.284, 0.000, 0.000)	1.535
Cp. XIII(rec.)	(0.500, 0.231, 0.269, 0.000, 0.000)	1.530
Cp. XIII(inv.)	(0.481, 0.222, 0.296, 0.000, 0.000)	1.536
Cp. XIV(unf.)	(0.439, 0.243, 0.281, 0.036, 0.000)	1.592
Canon 1	(0.429, 0.238, 0.286, 0.048, 0.000)	1.637
Canon 2	(0.324, 0.206, 0.412, 0.059, 0.000)	1.808
Canon 3	(0.294, 0.235, 0.441, 0.029, 0.000)	1.837
Canon 4	(0.439, 0.243, 0.281, 0.036, 0.000)	1.592

Interpretative remarks The present numerical instantiation exposes how canonical references, divergence neighborhoods, energy-based screening, and contextual refinement interact in a concrete setting. Its purpose is diagnostic: to demonstrate how the abstract objects of the framework admit explicit numerical realization and support canonical-centered exploration, without claiming empirical completeness or musicological exhaustiveness.

7. Discussion

This work outlines a unified information–geometric framework in which creativity is conceptualized as *structured motion* through a hierarchy of epistemic spaces (ideas → Idea-Spaces → Idea-Universes → Ontos). At all levels, creative motion is characterized through two coupled diagnostics: (i) **divergence**, capturing structural deviation via the Jensen–Shannon metric, and (ii) **energy**, capturing plausibility relative to the canonical grammars active at the corresponding scale.

7.1. Two Modes Governed by a Common Calculus

Within this framework, creative behavior manifests in two complementary regimes:

- **Divergent creativity**: increased structural divergence under maintained plausibility, corresponding to boundary-reaching yet grammatically admissible variation within an Idea-Space or Idea-Universe.
- **Analogical creativity**: structurally aligned reinterpretation across epistemic units, in which admissible alignments preserve deeper relational invariants while enabling task-relevant transfer across forms or domains.

These regimes correspond to different operating points of the same divergence–energy calculus rather than to distinct mechanisms. Creative motion may thus be viewed as controlled displacement in a joint diagnostic space, supporting both exploratory variation and structurally coherent transfer within a single formal calculus.

 doi:10.20944/preprints202603.0888.v1

7.2. Canonical References Across Scales

At each epistemic level, structural comparison is performed relative to a canonical distribution. Canonicals are defined as barycentric summaries of available representatives at the corresponding scale, providing stable references for typicality and for the ordering of exploration through divergence neighborhoods.

Divergence (via $\sqrt{D_{\mathrm{JS}}}$) and plausibility (via energy-to-acceptance mappings) extend consistently across Idea-Spaces, Idea-Universes, and the Ontos. The induced geometry is non-Euclidean in an operational sense: exploration is indexed by divergence levels around canonical references, while admissibility is regulated by nonlinear acceptance functions. Canonicals thus function as anchors for exploration and comparison, supporting shell-based expansion and alignment-based screening across epistemic units.

7.3. Creative Trajectories as Context-Refined Paths

Within the proposed formulation, creative reasoning is modeled as the selection of admissible trajectories in a curved information–geometric landscape, where lawful transformations define candidate moves and plausibility is regulated through energy-based acceptance.

A central consequence of the context mechanism is that creative trajectories are path-dependent: the admissibility and functional role of a candidate depend on the Chain of Thought through which it is reached. Redundant evaluation is regulated through monotonic contextual refinement rather than through explicit state memory. Within this view, a creative trajectory is governed by the coupled balance

$$\text{divergence to references} \;+\; \text{alignment feasibility} \;+\; \text{plausibility gates},$$

with termination occurring as soon as an admissible continuation satisfies the task constraints.

7.4. Connections to Classical and Contemporary Accounts of Reasoning

At an interpretative level, the proposed procedure can be related to classical heuristic accounts of problem solving, particularly those articulated by Polya [27]. In Pólya's formulation, progress emerges through cycles of local exploration, structural insight, and revision of perspective. The present framework provides a canonical and divergence-ordered realization of this heuristic logic, in which exploration, refinement, and termination are governed by explicit information-theoretic diagnostics rather than informal judgment.

From the perspective of contemporary artificial intelligence, the framework can be related to recent non-generative world-model approaches, such as LeCun's JEPA framework [8]. At a conceptual level, both approaches emphasize structured latent representations, predictive consistency, and the avoidance of explicit reconstruction. In the present work, creative search unfolds through divergence- and energy-regulated motion in canonical latent spaces, supporting predictive selection of admissible continuations rather than combinatorial generation.

Finally, the framework admits a natural cognitive interpretation in the spirit of Hofstadter's account of analogy as a fluid, context-sensitive process [28]. Canonical references play the role of prototypes, divergence ordering implements graded structural slippage, and recursive descent into subdomains corresponds to shifts in conceptual framing. The contribution here differs in that these mechanisms are expressed through explicit probabilistic and information–geometric operators, enabling algorithmic realization while preserving the path-dependent and interpretive character emphasized in Hofstadter's analysis.

7.5. Geometry, Scale Dependence, and Epistemic Fractality

The geometry induced by the creative cost

$$\mathcal{J} = \mathcal{J}(D_{\mathrm{JS}}, E)$$

 doi:10.20944/preprints202603.0888.v1

is scale-consistent yet scale-dependent. At the level of Idea-Spaces, divergence often dominates, and local neighborhoods exhibit approximately metric behavior, supporting fine-grained variation around a stable canonical grammar. As scale increases, multiple canonicals, alignment constraints, and energy-gated admissibility introduce curvature and non-monotonic accessibility. At the Ontos level, reachability is increasingly determined by alignment feasibility and reinterpretive plausibility rather than by proximity to a single reference.

This scale dependence yields a fractal epistemic organization: the same divergence–energy–analogy calculus governs creative motion at all scales, while the relative influence of its components varies continuously with the active epistemic unit and context. Canonical-centered expansion and recursive descent thus provide a uniform operational mechanism for exploration across the entire hierarchy.

7.6. Scope, Limitations, and Future Directions

The present formulation instantiates ideas through symbolic feature representations, canonicals through corpus-based estimation, and lawful transformations through explicitly defined systems. Natural extensions include richer feature geometries, expanded families of transformations, adaptive canonical drift under sustained exploration, and empirical evaluation across creative domains beyond music. These directions preserve the central structural property of the framework—canonical-centered, context-refined exploration under coupled divergence and plausibility diagnostics—while extending its expressive scope and empirical reach.

Within this framework, context π is operationalized as a structured epistemic descriptor that constrains interpretation, admissibility, and evaluation during creative search. Its specification is sufficient to regulate non-redundant exploration, enforce history dependence, and support termination through progressive refinement, without introducing scalar utilities or explicit update rules. Future work may explore explicit context-refinement operators or parametric formulations that support concrete algorithmic realizations while preserving the interpretive role of context.

In the present numerical instantiation, the feature extraction operator ϕ underlying the empirical distributions is instantiated through expert human judgment, operating at the level of symbolic and perceptual interpretation of the musical material, as reflected in Table 1. Within the proposed framework, this may be interpreted as the involvement of an external nano-level Ontos, providing epistemic input at the level of feature realization rather than at the level of structural inference. The framework itself remains agnostic to the nature of this source: analogous inputs could be supplied by automated symbolic analyzers, learned perceptual models, or other epistemic agents. This example illustrates how canonical-centered search may operate across interacting Ontoi, each contributing specialized structure at different stages of the creative process.

Beyond the present numerical illustration, the framework naturally supports applications in machine-assisted and autonomous discovery. By formulating creative reasoning as task-directed traversal of structured epistemic representations under explicit divergence and plausibility diagnostics, it provides a principled basis for systems that can autonomously explore, evaluate, and refine hypotheses within well-defined knowledge domains, while remaining grounded in interpretable structural criteria rather than opaque generative behavior.

At a more speculative but conceptually grounded level, the framework also admits potential relevance to the problem of inter-civilizational communication, as encountered for example in the context of SETI research. The hypothesis that knowledge within a civilization organizes into hierarchical, canonical-centered epistemic structures governed by admissible transformations and diagnostic regularities suggests the possibility of representation-independent alignment across intelligent systems. While preliminary, this perspective motivates future exploration of epistemic-structural approaches to communication that do not presuppose shared semantics or biological commonality.

Several open problems and development directions naturally arise from the present formulation. These include the explicit characterization of feasible novelty regions, the development of quantitative energy–divergence trade-offs, the extension to additional cognitive and symbolic domains, and the

 doi:10.20944/preprints202603.0888.v1

construction of empirical evaluation protocols. Together, these directions define a clear roadmap for subsequent theoretical and applied investigations.

8. Conclusion

In this work, intellectual and creative processes are modeled as the construction, stabilization, and systematic exploitation of epistemic structures. Central to this perspective is the emergence of canonical references and their associated peripheral regions, which together organize the space of admissible ideas within a nano Ontos. Creative reasoning proceeds through the selective reorganization and transformation of existing ideas relative to these canonical structures, rather than through isolated acts of generation.

By operating on canonicals and their peripheries, the proposed framework supports both local refinement and structural reconfiguration of the epistemic content contained within a nano Ontos. This mechanism enables the principled extension of the nano Ontos itself, allowing new Idea-Spaces and Idea-Universes to be incorporated while maintaining coherence with existing epistemic structure. In this sense, creativity is formalized as a process that simultaneously exploits established organization and expands the operational epistemic horizon of the agent.

The paper introduced a unified information–geometric framework in which creative reasoning—human or artificial—is modeled as motion through a hierarchy of epistemic structures. Across all levels of this hierarchy, from atomic epistemic elements to Idea-Spaces, Idea-Universes, and the global Ontos, creative dynamics are governed by two fundamental quantities: *divergence*, capturing structural departure via the Jensen–Shannon metric, and *energy*, capturing plausibility relative to canonical grammars. Their interaction defines a landscape in which novelty and coherence are not competing objectives but jointly determine feasible creative trajectories.

Within an Idea-Space, divergence enables exploratory variation, while energy anchors candidate ideas to domain-appropriate constraints. Across Idea-Spaces and Idea-Universes, the same calculus identifies low-divergence, moderate-energy pathways that support analogical transfer and stylistic reinterpretation. At the highest level, the Ontos organizes entire epistemic domains, providing a principled basis for cross-domain analogy and integrative conceptual innovation.

The resulting epistemic hierarchy is inherently recursive: epistemic elements combine into ideas; ideas generate spaces; spaces assemble into universes; and universes inhabit the Ontos. Creative progress therefore consists not only of movement *within* epistemic structures but also of transformations *across* them, enabling the expansion, reweighting, and recombination of the structures themselves. This perspective supports the design of computational systems capable of controlled creative search, principled analogy-making, and knowledge-guided innovation.

Future work will extend the framework toward richer feature representations, learned transformation operators, adaptive canonical distributions, multimodal domains, and empirical evaluation of energy–divergence trade-offs. These directions aim to bridge theoretical insight with practical implementations of machine-assisted creativity, grounding creative reasoning in a unified information–geometric foundation.

A further long-term objective is the formulation of a global epistemic structure of the Ontos itself, capable of quantitatively reflecting the developmental state of a civilization through the organization, diversity, and interconnectivity of its epistemic domains.

Such a global epistemic structure may also offer a principled foundation for future investigations into abstract communication protocols with other civilizations, grounded in canonical structure and information-theoretic regularities rather than culture-specific semantics.

Acknowledgments: The author expresses deep gratitude to Dr. Eng. Vasil Metodiev (UCTM, Sofia), a trained musician, lecturer, and engineer, whose combined musical expertise and algorithmic insight effectively assumed the role of a *human feature extractor* $\phi(S)$ in the preparation and formalization of the *Art of Fugue* examples. His contribution was essential in translating intricate contrapuntal structures into explicit representations compatible with the proposed probabilistic and algorithmic framework. Artificial intelligence tools were used to assist with

 doi:10.20944/preprints202603.0888.v1

linguistic polishing and formatting of the manuscript. The intellectual content and scientific responsibility remain entirely with the author.

References

1. Newell, A.; Simon, H.A. *Human Problem Solving*; Prentice-Hall: Englewood Cliffs, NJ, 1972.
2. Johnson-Laird, P.N. *Mental Models: Towards a Cognitive Science of Language, Inference, and Consciousness*; Harvard University Press: Cambridge, MA, 1983.
3. Anderson, J.R. *Rules of the Mind*; Lawrence Erlbaum Associates: Hillsdale, NJ, 1993.
4. Ohlsson, S. Information-Processing Explanations of Insight and Related Phenomena. *Advances in the Psychology of Thinking* **1992**. Please replace with your existing BibTeX entry; this key is cited in the section.
5. Ohlsson, S. *Deep Learning: How the Mind Overrides Experience*; Cambridge University Press: Cambridge, 2011.
6. Hofstadter, D.R. *Gödel, Escher, Bach: An Eternal Golden Braid*; Basic Books: New York, 1979.
7. Hofstadter, D.R. *Fluid Concepts and Creative Analogies: Computer Models of the Fundamental Mechanisms of Thought*; Basic Books: New York, 1995.
8. LeCun, Y. A Path Towards Autonomous Machine Intelligence. *arXiv preprint arXiv:2301.08241* **2023**. Introduces the Joint Embedding Predictive Architecture (JEPA) and its learning-based variant LeJEPA.
9. Balestriero, R.; LeCun, Y. LeJEPA: Provable and Scalable Self-Supervised Learning Without the Heuristics. *arXiv preprint* **2025**, *arXiv:2511.08544*. Provable and scalable formulation of Joint-Embedding Predictive Architectures (LeJEPA), https://doi.org/10.48550/arXiv.2511.08544.
10. Gridley, M.C. *Jazz Styles: History and Analysis*, 10 ed.; Pearson: Upper Saddle River, NJ, 2009.
11. Piaget, J. *The Origins of Intelligence in Children*; International Universities Press: New York, 1952.
12. Carey, S. *Conceptual Change in Childhood*; MIT Press: Cambridge, MA, 1985.
13. Bengio, Y.; Courville, A.; Vincent, P. Representation Learning: A Review and New Perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **2013**, *35*, 1798–1828. https://doi.org/10.1109/TPAMI.2013.50.
14. Mikolov, T.; Chen, K.; Corrado, G.; Dean, J. Efficient Estimation of Word Representations in Vector Space. In Proceedings of the Proceedings of the International Conference on Learning Representations (ICLR), 2013.
15. Saxe, A.M.; McClelland, J.L.; Ganguli, S. A mathematical theory of semantic development in deep neural networks. *Proceedings of the National Academy of Sciences* **2019**, *116*, 11537–11546. https://doi.org/10.1073/pnas.1820226116.
16. Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I. Language Models are Unsupervised Multitask Learners. *OpenAI Technical Report* **2019**.
17. Murphy, G.L. *The Big Book of Concepts*; MIT Press: Cambridge, MA, 2002.
18. Kaplan, J.; McCandlish, S.; Henighan, T.; Brown, T.; Chess, B.; Child, R.; Gray, S.; Radford, A.; Wu, J.; Amodei, D. Scaling Laws for Neural Language Models. *arXiv preprint arXiv:2001.08361* **2020**.
19. Nosofsky, R.M. Attention, Similarity, and the Identification–Categorization Relationship. *Journal of Experimental Psychology: General* **1986**, *115*, 39–57. https://doi.org/10.1037/0096-3445.115.1.39.
20. Chi, M.T.H. Laboratory Methods for Assessing Experts' and Novices' Knowledge. *The Cambridge Handbook of Expertise and Expert Performance* **2006**, pp. 167–184. Please set editor/publisher fields if you maintain this as an incollection entry.
21. Gobet, F.; Simon, H.A. Templates in chess memory: A mechanism for recalling several boards. *Cognitive Psychology* **1996**, *31*, 1–40. https://doi.org/10.1006/cogp.1996.0011.
22. Sporns, O.; Chialvo, D.R.; Kaiser, M.; Hilgetag, C.C. Organization, Development and Function of Complex Brain Networks. *Trends in Cognitive Sciences* **2004**, *8*, 418–425. https://doi.org/10.1016/j.tics.2004.07.008.
23. Kanwisher, N. Functional Specificity in the Human Brain: A Window into the Functional Architecture of the Mind. *Proceedings of the National Academy of Sciences* **2010**, *107*, 11163–11170. https://doi.org/10.1073/pnas.1005062107.
24. Elman, J.L.; Bates, E.A.; Johnson, M.H.; Karmiloff-Smith, A.; Parisi, D.; Plunkett, K. Rethinking Innateness: A Connectionist Perspective on Development. *Trends in Cognitive Sciences* **1996**. If you intended the book, use @book with publisher = MIT Press and year = 1996.
25. Lake, B.M.; Ullman, T.D.; Tenenbaum, J.B.; Gershman, S.J. Building Machines That Learn and Think Like People. *Behavioral and Brain Sciences* **2017**, *40*, e253. https://doi.org/10.1017/S0140525X16001837.
26. Parisi, G.I.; Kemker, R.; Part, J.L.; Kanan, C.; Wermter, S. Continual Lifelong Learning with Neural Networks: A Review. *Neural Networks* **2019**, *113*, 54–71. https://doi.org/10.1016/j.neunet.2019.01.012.

 doi:10.20944/preprints202603.0888.v1

27. Polya, G. *How to Solve It: A New Aspect of Mathematical Method*; Princeton University Press: Princeton, NJ, 1945.
28. Hofstadter, D.R. The Copycat Project: A Model of Mental Fluidity and Analogy-Making. In *Fluid Concepts and Creative Analogies*; Basic Books, 1995; pp. 205–267.