

Ontological Re-Entry: What Can We Say Exists?

Ontological Re-Entry: What Can We Say Exists?

Abstract

This paper addresses the "re-entry" problem in ontology: given a theory, formalism, or model built over some primitive vocabulary, what can we legitimately say exists when we project that structure back onto the world, and what does answering the question cost a physically bounded observer? We develop a finite-distinction framework in which existence claims are admissible only when expressible as distinctions at finite resolution. A system of N binary distinctions supports 2^N states and exactly 2^{2^N} extensionally distinct existence predicates, of which only a vanishing fraction is reachable by short formulas; we compute this syntactic bottleneck explicitly. Erasing one distinction at $T = 300$ K costs at least 2.87×10^{-21} J, converting power budgets into distinction budgets and ontological commitments into energy accounting. We define re-entry as the rewriting of a theory's high-level existence claims into predicates over its own base distinctions, modeled on ontological query compilation, and derive the conditions under which re-entry fails: false-premise capture, resolution collapse, and unpayable erasure cost. Worked examples for observer budgets from 2^{40} to 2^{64} states show that under aggressive ultrametric compression the surviving ontology is structure, not individuals. We situate the framework against database query rewriting, primitive-ontology and relational readings of quantum mechanics, false-premise failures in multimodal models, and the definitional stagnation of decentralization, and we state falsification conditions.

1. Introduction

Every theory faces two directions of fit. In one direction, the theory is built: primitives are chosen, axioms are laid down, and a formal structure grows. In the other direction — the direction this paper calls *re-entry* — the finished structure is projected back onto the world, and we ask: what, according to this structure, exists? The re-entry problem is the ontological half of the semantics of theories. A theory may be internally impeccable while its re-entry map — the rule that says which of its internal objects correspond to anything real — is arbitrary, underdetermined, or silently loaded with the very ontology it was supposed to ground.

The question "what can we say exists?" is therefore not a request for a list. It is a request for a *criterion*: a procedure that, given a theory and an observer with finite resources, outputs the set of existence claims the observer is entitled to make, together with the cost of making them. This paper proposes such a criterion. Its core commitment, taken from the finite-distinction program [11], is that a finite-entropy world is a finite-distinction world: uncountable precision is unphysical, while computable depth and p -adic valuation remain physically real. If the world presents only finitely many distinctions to any finite observer, then the raw material of ontology is not substance but *difference*, and the state-space geometry of differences is combinatorial and ultrametric rather than Euclidean.

Three consequences follow. First, existence is graded: an object exists *at distinction depth* d for an observer with budget N if it is invariant under the compression that the budget forces. Second, re-entry is a thermodynamic act with a price: maintaining and erasing distinctions costs energy, connecting ontology to the physics of computation [7]. Third, the classical ontological disputes — particle versus field, substance versus relation, primitive versus structural — can be re-read as disputes about where in the distinction hierarchy re-entry is performed. The relational ontology of contemporary physics [4] and the primitive-ontology programs surveyed in [3] then appear not as rivals but as different re-entry conventions, and our framework supplies the vocabulary in which that choice is made explicit.

The framing generates three questions that structure the paper. The *representational* question: what is the geometry of a space built only from finite distinctions, and how many distinct existence claims can it support? The *physical* question: what does it cost, in thermodynamic terms, to hold or erase a distinction, and when does an ontological commitment become physically unpayable? The *logical* question: when can an ontology's high-level existence claims be rewritten into statements over its own base distinctions without loss — the re-entry condition — and when does rewriting fail?

These questions are not asked in a vacuum. Database research has a mature theory of evaluating and optimizing queries answered against an ontology rather than directly against data [1]; robotics software engineering composes ontological classification systems to make hardware and software components mutually intelligible [8]; and decentralization has been given an explicitly ontological analysis precisely because no cross-domain definition had stabilized [6]. On the physical side, the Primitive Ontology program insists that ontological clarity is a necessary condition on any theoretical framework [3], while a relational reading finds support across classical mechanics, gauge theories, general relativity, quantum field theory, and tentative quantum gravity [4]. On the computational side, erasing a bit carries an intrinsic energy cost, and real-time finite automata pay a complexity price for forgetting [7]. Language-model systems fine-tuned for legal regulation [2] and multimodal models that fail under false premises [5] supply engineering instances of ontological commitment in deployed systems.

The contributions are: (i) a formal definition of a finite distinction space and of admissible existence predicates over it, with explicit counting bounds (Section 3, computed in Section 4); (ii) a thermodynamic admissibility criterion derived from the erasure cost of distinctions [7]; (iii) a definition of ontological re-entry modeled on query rewriting against an ontology [1], with a necessary condition for re-entry to be lossless; and (iv) an analysis of failure modes, including false-premise commitments in the sense observed for multimodal models [5] and the definitional fragmentation documented for decentralization [6].

2. Background and Related Work

Ontological queries and rewriting. The database literature has formalized a precise version of re-entry. In [1], ontological queries are evaluated against an ontology rather than directly on a database, and the paper studies two aspects of this problem, query rewriting and query optimization, where rewriting consists of the compilation of an ontological query into an equivalent query against the underlying data. This is exactly a re-entry map in miniature: the ontology is the theory, the database is the world-as-modeled, and the rewriting is the rule for projecting existence claims down to what the substrate can answer. We borrow the rewriting/optimization split — first rewrite, then optimize — as the template for Section 3.4, noting where the analogy is exact and where it breaks: ontologies of existence have no "underlying database" other than the distinction space itself.

Engineering ontologies. At the applied end, [8] presents an ontology-driven software engineering approach for industrial robotics control software, introducing the ReApp architecture that synthesizes model-driven engineering with semantic technologies for the development and reuse of ROS-based components and applications, showing how different ontological classification systems for hardware can be combined. The supplied summary indicates that classification systems are treated as reusable engineering artifacts and gives no further quantitative detail; in our terms, ReApp performs re-entry from a design ontology onto running robot components, and the multiplicity of classification systems it combines illustrates that re-entry maps are not unique.

Primitive ontology and realism. In foundations of physics, [3] observes that Bohm's and Bell's approaches share notable features with the contemporary Primitive Ontology perspective and with Esfeld's and Deckert's minimalist ontology: all these programs consider ontological clarity a necessary condition for every theoretical framework, promote scientific realism also in the quantum domain, and strengthen explanatory power. The supplied summary

states these shared commitments and gives no further technical detail, but the stated position — ontological clarity as a *necessary condition* on theories — is precisely the normative stance this paper operationalizes. Where [3] demands clarity, we supply a quantitative criterion for it.

Relational ontology. Against substance-first readings, [4] argues that quantum theory can be understood as pointing to an ontology of relations, and observes that this reading is supported by the ubiquity of relationality in contemporary fundamental physics, including classical mechanics, gauge theories, general relativity, quantum field theory, and tentative theories of quantum gravity. A distinction-based ontology is naturally relational: a distinction is a difference, and differences are relations. Our framework can thus be read as a resource-bounded refinement of the relational thesis of [4], though Section 6 notes an objection that would invert our direction of analysis.

Non-Archimedean and finite-distinction ontology. The QNFO program supplies the positive ontology. [9], the Unified Theory of Non-Archimedean Ontology, and [10], Super-Universe, are supplied as titles without abstract text in the bibliography; we therefore rely on them only as the umbrella works of the program and make no specific claims about their internal content. [11] states the load-bearing thesis: a finite-entropy world is a finite-distinction world, uncountable precision is unphysical, computable depth and p -adic valuation remain physically real, and the state-space geometry of finite distinctions is combinatorial and ultrametric, with quantum mechanics read as a stochastic theory whose unitary evolution and superposition emerge in the large-distinction limit. [12], a companion to "The Qubit Delusion," argues that if the qubit-gate-circuit model is an epistemic failure — a projection of particle ontology onto relational, field-theoretic reality — then the constructive task is to strip quantum mechanics to its minimal ontological commitments. Our re-entry criterion is the tool that task requires: it says exactly which commitments survive a given observer budget.

Thermodynamic cost of distinction. [7] establishes that each step resulting in a bit of information being "forgotten" by a computing device has an intrinsic energy cost, and that while any Turing machine can be rewritten to be thermodynamically reversible without changing the recognized language, finite automata restricted to scan their input once in real-time fashion can only recognize a proper subset of the regular languages. Two consequences feed directly into our framework: reversibility is always available in principle at the price of keeping distinctions (never erasing them), and forgetting is not free — it is precisely what partitions the expressible from the efficiently expressible. Distinctions are not free, and the devices that maintain them under real-time constraints are strictly weaker than unconstrained ones; ontological budgets are therefore not just epistemic bookkeeping but energy accounting.

False premises as re-entry failures. In machine perception, [5] reports that open-source Large Multimodal Models excel at open-vocabulary grounding and segmentation but suffer under false premises when queries imply the existence of something not actually present in the image, and that fine-tuning an LMM to segment images significantly degrades its ability to reliably determine ("see") whether an object is present. This is an empirical instance of re-entry failure: the model's internal structure licenses an existence claim the world does not support. The tension the summary reports — segmentation ability bought at the price of presence-detection — is a resource trade-off of exactly the kind our budget formalism predicts.

Definitional stagnation. Finally, [6] documents that decentralization, despite its fundamental role across security, distributed computing, artificial intelligence, cloud infrastructures, and IoT architectures and a history of over half a century, has no universally accepted definition applicable across computer communication systems, a situation the summary indicates has become problematic. We take this as a cautionary case study: a concept in wide cross-domain use without an explicit re-entry criterion accumulates incompatible ontological commitments; the term persists while its extension drifts. The same risk attaches to "existence" itself, which is our target concept. The supplied summary gives no further detail on the proposed remedy, and we do not attribute one to it.

Language-model ontologies. [2] fine-tunes Large Language Models with a curated supervised legal dataset and integrates Retrieval-Augmented Generation to support policymakers in understanding, analyzing, and crafting legal regulations. The supplied summary describes the engineering and cuts off before reporting results; its relevance here is structural: a legal regulation is an ontology enforced by re-entry onto human conduct, and RAG is a retrieval-bounded re-entry mechanism whose false-premise costs are institutional rather than perceptual. We cite it only for the architectural pattern and draw no findings from it.

3. Methods

3.1 Distinction spaces

Definition 3.1 (Distinction space). A *distinction space* is a pair $D_N = (S_N, d_u)$ where $S_N = \{0, 1\}^N$ is the set of N -bit distinction records and $d_u : S_N \times S_N \rightarrow \mathbb{Q}_{\geq 0}$ is an ultrametric, i.e., $d_u(x_i, x_j) \leq \max\{d_u(x_i, x_k), d_u(x_k, x_j)\}$ for all $x_i, x_j, x_k \in S_N$. A *distinction* d_i is a binary partition of the observer's state space; maintaining it means the observer can reliably re-identify which side of the partition the world is on. The ultrametric requirement implements the finite-distinction geometry of [11]: balls in D_N are hierarchically nested, mirroring the combinatorial, p -adic structure in which depth of distinction, not metric distance, carries the physical content. We realize the geometry by a b -ary tree of depth L with branching factor b , so that

$$N = L \log_2 b, \quad |S_N| = b^L = 2^N.$$

3.2 Admissible existence predicates and the counting bound

Definition 3.2 (Admissible existence predicate). An *existence predicate* over D_N is a subset $E_k \subseteq S_N$; the claim " E_k is instantiated" is *admissible* iff E_k is definable by a finite Boolean formula over the N base coordinates. Two predicates are *extensionally equivalent* iff $E_i = E_j$ as sets.

Proposition 3.1. The number of extensionally distinct admissible existence predicates over D_N is exactly 2^{2^N} .

Proof. Each Boolean formula over N coordinates defines a subset of S_N , and every subset of S_N is defined by its characteristic formula (a disjunction over its members). Hence admissible predicates correspond bijectively to subsets of S_N , of which there are $2^{|S_N|} = 2^{2^N}$. $\square$

The bound is tight but syntactically unreachable: the number of *short* formulas is far smaller, and Section 4.1 quantifies the gap for a concrete budget.

3.3 Thermodynamic admissibility

Following [7], a step that forgets a bit has an intrinsic energy cost. We adopt the standard Landauer form as a modeling assumption (not derived here): erasing one bit at temperature T_{env} costs at least

$$E_{\text{bit}} = k_B T_{\text{env}} \ln 2,$$

where k_B is Boltzmann's constant. A commitment is *thermodynamically admissible* for an agent with erasure budget E_{budget} if the distinctions it requires can be maintained without erasure exceeding budget. If the agent must cycle its memory at rate f_{cycle} , the erasure power is

$$P_{\text{erase}} = f_{\text{cycle}} N_{\text{dist}} k_B T_{\text{env}} \ln 2,$$

with N_{dist} the number of distinctions cycled per cycle. Symmetrically, maintaining N distinctions against noise over a coherence time τ requires refreshing them, so an observer's power budget P_{obs} bounds

$$N \leq \frac{P_{\text{obs}} \tau}{k_B T_{\text{env}} \ln 2}.$$

This inequality is the bridge between physics and ontology: it converts a power budget into a distinction budget, and hence into an existence criterion. Reversibility, per [7], removes the erasure cost but forces the agent to retain all distinctions — trading power for unbounded memory growth.

3.4 Ontological re-entry

Definition 3.3 (Re-entry). Let $\mathcal{O}$ be an ontology whose commitments are expressed in a high-level vocabulary V_{high} , and let V_{base} be the base coordinates of D_N . *Re-entry* is a map $r : V_{\text{high}} \rightarrow \mathcal{P}(S_N)$ assigning each high-level predicate an admissible existence predicate. Re-entry is *lossless* iff for every pair of commitments $c_i, c_j \in \mathcal{O}$, all entailments among them are preserved under r .

This mirrors the compilation of an ontological query into an equivalent query against the underlying data [1], with D_N playing the role of the underlying data. The re-entry condition fails when V_{high} contains predicates whose extension depends on precision beyond the N -bit budget — the "uncountable precision is unphysical" clause of [11].

4. Analysis

Every input number below is stated with its source; every arithmetic step is shown. No empirical data are used. Two illustrative budgets appear: $N = 64$ (one machine word) for the counting analysis and $N = 40$ ($b = 2$, $L = 40$) for the ultrametric tree scenario; both are modeling choices of this paper, and the divergence between drafts on this parameter is documented in Appendix A.

4.1 Counting admissible existence predicates

Inputs. $N = 64$ binary distinctions (modeling choice). Formula: Proposition 3.1.

Step 1. $|S_N| = 2^N = 2^{64} = 18,446,744,073,709,551,616 \approx 1.8447 \times 10^{19}$ states.

Step 2. Number of admissible predicates: $2^{2^N} = 2^{2^{64}}$. In base-10 logarithms:

$$\log_{10} \left(2^{2^{64}} \right) = 2^{64} \log_{10} 2 \approx 1.8447 \times 10^{19} \times 0.30103 \approx 5.553 \times 10^{18}.$$

So there are $\approx 10^{5.553 \times 10^{18}}$ extensionally distinct existence predicates — combinatorially vast but finite.

Step 3 (syntactic reachability). Assume formulas are conjunctions of at most $k_{\text{lit}} = 8$ literals chosen from $2N = 128$ literals (modeling choice). The dominant term is $\binom{128}{8} = \frac{128 \times 127 \times 126 \times 125 \times 124 \times 123 \times 122 \times 121}{40320}$. Numerator, stepwise: $128 \times 127 = 16,256$; $\times 126 = 2,048,256$; $\times 125 = 256,032,000$; $\times 124 = 31,747,968,000$; $\times 123 = 3,905,000,064,000$; $\times 122 = 476,410,007,808,000$; $\times 121 = 57,645,610,944,768,000$. Dividing: $57,645,610,944,768,000 / 40,320 \approx 1.4298 \times 10^{12}$. The remaining terms $\sum_{j=0}^7 \binom{128}{j}$ are smaller by factors of at least $121/8 \approx 15$ each, so the total is $\approx 1.43 \times 10^{12} \times (1 + 1/15 + \dots) \approx 1.53 \times 10^{12}$ short conjunctions.

Conclusion 4.1. With $N = 64$, the fraction of admissible predicates reachable by short formulas is $\approx 1.53 \times 10^{12} / 10^{5.553 \times 10^{18}}$, i.e., effectively zero relative to the full space. Ontological expressiveness is therefore bottlenecked by *syntax*, not by the state space: re-entry must route through structured (hierarchical,

ultrametric) definitions, consistent with the nested-ball geometry of [11].

4.2 Erasure cost of a distinction

Inputs. $k_B = 1.380649 \times 10^{-23} \text{ J K}^{-1}$ (SI defined value); $T_{\text{env}} = 300 \text{ K}$ (room-temperature modeling choice); $\ln 2 = 0.693147$ (mathematical constant). Formula: $E_{\text{bit}} = k_B T_{\text{env}} \ln 2$ [7].

Step 1. $k_B T_{\text{env}} = 1.380649 \times 10^{-23} \times 300 = 4.141947 \times 10^{-21} \text{ J}$.

Step 2. $E_{\text{bit}} = 4.141947 \times 10^{-21} \times 0.693147 = 2.8709 \times 10^{-21} \text{ J}$.

Step 3 (cycling cost). Assume an agent maintains $N_{\text{dist}} = 10^9$ distinctions and rewrites (erases-and-replaces) each once per second, $f_{\text{cycle}} = 1 \text{ s}^{-1}$ (modeling choices). Then

$$P_{\text{erase}} = 1 \times 10^9 \times 2.8709 \times 10^{-21} = 2.8709 \times 10^{-12} \text{ W}.$$

Step 4 (scaling projection). Under the stated linear-scaling assumption and constant T_{env} , an agent cycling 10^{20} distinctions at 10^9 s^{-1} would require $P_{\text{erase}} = 10^{20} \times 10^9 \times 2.8709 \times 10^{-21} = 2.87 \times 10^8 \text{ W}$, which is physically prohibitive. This is a labeled projection, not a measurement.

Conclusion 4.2. At 10^9 distinctions cycled at 1 s^{-1} , the minimal erasure power is $\approx 2.87 \times 10^{-12} \text{ W}$ — negligible for silicon-scale hardware but a hard floor scaling linearly in both N_{dist} and f_{cycle} . Thermodynamic admissibility bites at the large-distinction limit that [11] identifies with quantum-mechanical behavior.

4.3 Distinction budget from a power budget

Inputs. Power $P_{\text{obs}} = 10^{-12} \text{ W}$ (illustrative nanowatt-scale bound, modeling choice), coherence time $\tau = 1 \text{ s}$, $T = 300 \text{ K}$. From Step 2 of Section 4.2, $k_B T \ln 2 = 2.8709 \times 10^{-21} \text{ J/bit}$.

$$N_{\text{max}} = \frac{P_{\text{obs}} \tau}{k_B T \ln 2} = \frac{10^{-12}}{2.8709 \times 10^{-21}} = 3.4832 \times 10^8.$$

So a nanowatt observer with 1 s coherence can in principle sustain at most $N_{\text{max}} \approx 3.48 \times 10^8$ bit-distinctions per second against the erasure floor of [7]. This is an upper bound on *refresh rate*, not on stored distinctions; we use it as the budget ceiling in the re-entry criterion, noting that real devices operate far above the Landauer bound, so the true budget is smaller by an unquantified efficiency factor.

4.4 Ultrametric ball counting and compression

Inputs. $b = 2$, $L = 40$ (modeling choices; any b, L work). Then $N = L \log_2 b = 40 \times 1 = 40$ and

$$|S_N| = 2^{40} = 1,099,511,627,776 \approx 1.10 \times 10^{12}.$$

The entropy of a uniform distribution over S_N is $H = N$ bits, so the thermodynamic entropy is

$$S = N k_B \ln 2 = 40 \times 1.380649 \times 10^{-23} \times 0.693147.$$

Step 1: $1.380649 \times 10^{-23} \times 0.693147 = 9.56963 \times 10^{-24} \text{ J/K}$ (since $1.380649 \times 0.693147 = 0.956963$). Step 2: $S = 40 \times 9.56963 \times 10^{-24} = 3.82785 \times 10^{-22} \text{ J/K}$. The full erasure cost at $T = 300 \text{ K}$ is

$$E_{\text{erase}} = S \cdot T = 3.82785 \times 10^{-22} \times 300 = 1.14836 \times 10^{-19} \text{ J}.$$

In a b -ary ultrametric tree of depth L , the number of balls at radius level r (subtrees of height r) is b^{L-r} . At $r = 10$: $2^{30} = 1,073,741,824 \approx 1.07 \times 10^9$ fine balls. At $r = 30$: $2^{10} = 1,024$ coarse balls. Compression factor:

$$\frac{2^{30}}{2^{10}} = 2^{20} = 1,048,576 \approx 1.05 \times 10^6.$$

4.5 Re-entry classification under compression

Inputs. A theory T posits $m = 10^5$ objects (modeling choice), spread uniformly over the 2^{30} fine balls of Section 4.4. Each coarse ball at level 30 contains on average

$$\lambda = \frac{10^5}{2^{10}} = \frac{100,000}{1,024} = 97.65625$$

fine states occupied by objects. With threshold $\theta = 1$ (an object is "real" if some observer state isolates it), an object remains resolvable only if it is the sole object in its coarse ball. Under a Poisson model with mean λ , the probability a given ball contains exactly one object is

$$p_1 = \lambda e^{-\lambda} = 97.65625 \times e^{-97.65625}.$$

Since $e^{-97.65625} = 10^{-97.65625 / \ln 10} = 10^{-97.65625 / 2.302585} = 10^{-42.410} \approx 3.87 \times 10^{-43}$, we get

$$p_1 \approx 97.65625 \times 3.87 \times 10^{-43} \approx 3.78 \times 10^{-41}.$$

The expected number of resolvable objects among 10^5 is

$$m_{\text{real}} \approx 10^5 \times 3.78 \times 10^{-41} = 3.78 \times 10^{-36} \approx 0.$$

Interpretation: under uniform spreading and this compression, essentially no individual object survives re-entry as "real"; the ontology that survives is the coarse partition itself — the 1,024 level-30 balls — not the 10^5 objects. This is a precise, if idealized, statement of the finite-distinction thesis of [11]: what exists at re-entry is what survives the observer's compression, and under aggressive compression that is structure, not individuals.

Scaling of the survival threshold. For objects to survive as individuals, we need $\lambda \lesssim 1$, i.e., $m \lesssim b^{L-r}$. With $b = 2$, $L = 40$, $r = 30$: $m \lesssim 1,024$. Solving for the compression depth at which $m = 10^5$ objects just survive:

$$b^{L-r} = 10^5 \implies L - r = \log_2 10^5 = \frac{\ln 10^5}{\ln 2} = \frac{5 \times 2.302585}{0.693147} = \frac{11.512925}{0.693147} = 16.6096.$$

So $r = L - 16.6096 = 40 - 16.6096 = 23.39$: the observer may compress to level ≈ 23 but not to level 30 if 10^5 individuals are to remain resolvable. Each additional level of compression divides resolvable capacity by $b = 2$.

4.6 Resolution collapse

Inputs. Budget $N = 64$ (Section 4.1). A high-level predicate c_i whose extension depends on $n_{\text{prec}} = 100$ bits of precision (modeling choice: the predicate separates states that agree on the first 64 coordinates). Then $n_{\text{prec}} > N$ and no admissible predicate over D_{64} separates the states: re-entry for c_i fails by Definition 3.3, since $r(c_i)$ cannot be defined. The excess

References

- [1] Ontological Queries: Rewriting and Optimization (Extended Version). arXiv:1112.0343v1. <https://arxiv.org/abs/1112.0343v1> [2] ASVRI-Legal: Fine-Tuning LLMs with Retrieval Augmented Generation for Enhanced Legal Regulation. arXiv:2511.03563v1. <https://arxiv.org/abs/2511.03563v1> [3] Beables, Primitive Ontology and Beyond: How Theories Meet the World. arXiv:2104.13859v1. <https://arxiv.org/abs/2104.13859v1> [4] The relational ontology of contemporary physics. arXiv:2201.00907v2. <https://arxiv.org/abs/2201.00907v2> [5] See, Say, and Segment: Teaching LMMs to Overcome False Premises. arXiv:2312.08366v1. <https://arxiv.org/abs/2312.08366v1> [6] Defining Decentralization: An Ontological Perspective. arXiv:2608.09748v2. <https://arxiv.org/abs/2608.09748v2> [7] Energy Complexity of Regular Languages. arXiv:2204.06025v2. <https://arxiv.org/abs/2204.06025v2> [8] A Model-Driven Engineering Approach for ROS using Ontological Semantics. arXiv:1601.03998v1. <https://arxiv.org/abs/1601.03998v1> [9] QNFO: Unified Theory of Non-Archimedean Ontology [10] QNFO: Super-Universe [11] QNFO: Finite-Distinction Quantum Mechanics: Unitary Evolution and Superposition as the Large-Distinction Limit of Stochastic Thermodynamics [12] QNFO: Beyond the Qubit: Constructive Paradigms for Post-Particle Computation