A Substrate-Independent Fixed-Point Operator for Composable Computation:

One Convergence Driver Across Electronic Structure, Associative Memory, and Constraint Satisfaction

Benjamin Dakota Horn
Prometheus7 Research Institute
prometheus7.research

June 1, 2026

Abstract

A wide range of computational problems share a single form: their solutions are *fixed points* of a local, composable update acting on a grounded representation. We make this concrete by decomposing fixed-point computation into a domain-invariant *driver* (the convergence loop) and a domain-specific pair, a *cut* (the local update) and a *grounding* (the encoding of the problem). We then demonstrate empirically that a single, byte-identical driver converges to the correct fixed point in three unrelated domains: ground-state electronic structure via density-matrix purification, associative memory via Hopfield recall, and constraint satisfaction via graph colouring. Using ozone as a case study, we show that the predictive power of the scheme is governed by the *richness of the grounding* rather than by the operator: an occupancy-level grounding selects the wrong (cyclic) isomer, whereas an orbital-description grounding recovers the correct bent ground state to within $\sim$10% of a B3LYP/cc-pVDZ reference calculation that we perform independently. We are deliberately explicit about scope. The constituent algorithms are established—density-matrix purification underlies linear-scaling electronic structure; Hopfield networks and mean-field constraint solvers are classical; vector-symbolic representations are well studied—and the dense purification we use is *slower* than direct diagonalization, so we claim no scaling advantage. The contribution is the synthesis and its framing: a structural universal in which stable existence is the fixed point of a local composable operation on a grounded representation, with the grounding as the sole locus of domain knowledge. We argue that this boundary is itself the central result, and we outline its implications for a uniform, composable substrate spanning physical and cognitive state.

1 Introduction

Many of the central computational tasks in the natural and information sciences have the same abstract shape even when they share no content. The ground-state density of an interacting electronic system, a memory recalled from a partial cue, the equilibrium of an economy, a satisfying assignment of a constraint network, and the marginals of a graphical model are all *fixed points*: states $x^\star$ that are invariant under some local update S, i.e. $x^\star = S(x^\star)$, and that are reached by iterating that update from an initial guess. This observation is not new in any individual field; it is, in effect, the reason variational and self-consistent formulations recur throughout physics and computation. What is less often made explicit is that the *machinery* of convergence can be cleanly

separated from the *content* of any particular problem, and that doing so exposes both a genuine cross-domain regularity and a sharp limit on what such a regularity can deliver.

This paper studies that separation directly. We define a single fixed-point *driver*—an iteration-to-convergence loop—and we treat each problem as supplying two domain-specific ingredients: a *cut*, the local update map whose fixed points are the desired solutions, and a *grounding*, the encoding of the problem instance into the driver's state space. The driver is held fixed; only the cut and grounding change between domains. We then ask an empirical question: does one such driver, used verbatim, actually converge to the correct answer across genuinely unrelated problems?

Our contributions are as follows.

1. **A decomposition.** We formalise fixed-point computation as a triple of a domain-invariant driver and a domain-specific (cut, grounding) pair, and we state precisely the sense in which the driver is universal and all domain knowledge resides in the grounding (Section 3).
2. **A three-domain demonstration.** Using one byte-identical driver, we converge to the correct fixed point in electronic structure (matching exact diagonalization to machine precision), associative memory (exact recall from a 30%-corrupted cue), and constraint satisfaction (a proper 3-colouring), Section 4.
3. **Grounding richness governs predictive power.** Through an ozone case study we show that the same relational fold fails (selecting the cyclic isomer) under an occupancy-level grounding and succeeds (recovering the bent ground state to $\sim$10% of B3LYP) under an orbital-description grounding (Section 5). The apparent "information floor" of the method is an artefact of an impoverished grounding, not a property of the operator.
4. **An honest placement.** We delineate what is and is not novel, report a negative scaling result, and argue that the boundary of the scheme—that the driver carries no content and the grounding carries all of it—is the result rather than a caveat (Sections 6–7).

We emphasise at the outset what we do *not* claim. We do not introduce a new electronic-structure method; the purification we use is classical. We do not demonstrate a complexity advantage; our dense implementation is slower than diagonalization. We do not reach chemical accuracy; our agreement with first-principles calculations is in sign and order of magnitude. The value of the work is conceptual: it isolates a unifying form, demonstrates it concretely across domains, and locates exactly where each domain's truth lives.

2 Related Work

Density-matrix purification and linear-scaling electronic structure. The idempotent single-particle density matrix can be obtained as the fixed point of a polynomial iteration without diagonalization. McWeeny's purification [1] and the trace-conserving canonical purification of Palser and Manolopoulos [2] are the foundational schemes. Because the density matrix of a gapped system is spatially local—Kohn's "nearsightedness" [4, 5]—purification can be sparsified to yield linear-scaling electronic-structure methods [3]. The electronic-structure instance we use below is squarely within this family; our purpose is not to improve it but to use it as one grounding of a domain-blind driver.

Associative memory. Hopfield networks store patterns as attractors of a recurrent sign-update whose Lyapunov (energy) function descends to fixed points [6]; recall is convergence from a corrupted

cue to the nearest stored pattern. This is a fixed-point computation in our sense, with the synaptic matrix as grounding.

Constraint satisfaction as message passing. Iterative local-update solvers for constraint satisfaction—belief propagation [7] and survey propagation [8]—express solutions as fixed points of repeated local message updates, and mean-field relaxations of colouring and satisfiability are convergent local iterations. Our constraint-satisfaction grounding is a damped mean-field colouring iteration of this kind.

Compositional representations and cross-scale physics. Vector-symbolic architectures and holographic reduced representations [9, 10] provide fixed-dimensional distributed representations closed under a binding operation, enabling compositional encoding of structured objects. Independently, the density-matrix renormalization group and tensor networks [11, 12] are compositional representations that carry quantum information across scales; crucially, they also make explicit *when* such representations are efficient—namely for low-entanglement states—which bounds the reach of any relational encoding of many-body physics. These two literatures frame both the promise of a composable substrate and its principled limit.

Fixed-point theory. The convergence of an iterated map to a unique fixed point is guaranteed for contractions on complete metric spaces by the Banach fixed-point theorem [13]. The maps we use are not globally contractive, which is why initialisation, damping, and grounding matter, and why a poorly chosen cut can fail to converge.

Reference electronic-structure calculations. Our first-principles ozone energies use Kohn–Sham density-functional theory with the B3LYP functional [14] as implemented in PySCF [15].

3 The Fold: A Domain-Invariant Fixed-Point Operator

3.1 Definition

Let $(\mathcal{X}, \|\cdot\|)$ be a normed state space. A *grounding* g of a problem instance specifies (i) an embedding of the instance into $\mathcal{X}$, including an initial state x_0, and (ii) a *cut* $S_g : \mathcal{X} \to \mathcal{X}$, a local update map whose fixed points encode the solutions of the instance.

Definition 1 (Fold driver)**.** *Given a state $x_0 \in \mathcal{X}$, a cut S_g, a tolerance $\varepsilon > 0$, and an iteration cap T, the* driver $\mathcal{D}$ *computes*

$$x_{t+1} = S_g(x_t), \qquad t = 0, 1, 2, \ldots,$$

and returns the first x_{t+1} with $\|x_{t+1} - x_t\| < \varepsilon$ (or x_T if none). A returned state $x^\star$ satisfies $x^\star \approx S_g(x^\star)$, i.e. it is an (approximate) fixed point of the cut.

The central structural claim is a factorisation: *a computation is a triple $(\mathcal{D}, S_g, g)$ in which $\mathcal{D}$ is invariant across all problems and domains, while the pair (S_g, g) is bespoke to each.* The driver is the universal "profile"; the cut and grounding are the bespoke "cut" of a key. Listing 1 gives the driver used for every experiment in this paper, unchanged.

```
def fold(x, step, tol=1e-9, maxit=4000):
    for _ in range(maxit):
        xn = step(x)
        if norm((xn - x).ravel()) < tol:
            return xn
        x = xn
    return x
```

Figure 1: The driver, used verbatim in all three domains. Only `step` (the cut) and the encoding of `x` (the grounding) change between domains.

Domain	Fixed point	Validation	Result
Electronic structure	density matrix P	energy vs. exact diag. (N=8)	-9.51754 (match)
Associative memory	stored pattern	recall from 30/100-bit cue	exact recall
Constraint satisfaction	3-colouring	all 9 edges proper	proper colouring

Table 1: One driver, three groundings. Each converged to the correct fixed point of its domain (iteration counts 11, 2, 58 respectively).

3.2 Why one driver suffices, and why that is not enough

A single driver can serve many problems because of the ubiquity of fixed-point structure: an enormous range of solution concepts are characterised as invariants of a local map. When the cut S_g is a contraction, the Banach theorem [13] guarantees a unique fixed point and geometric convergence. In practice the cuts we use (purification, sign-update recall, mean-field colouring) are not globally contractive; they converge on basins around their fixed points, which is precisely why the grounding—including initialisation and damping—is load-bearing, and why a poorly chosen cut may diverge or settle on a spurious fixed point. We exhibit such a failure deliberately in Section 4.3.

This already foreshadows the paper's main interpretive point. The driver's universality is bought by its emptiness: it commits to nothing about any domain, which is exactly why it transfers and exactly why it cannot, by itself, solve anything. All of the physics, all of the stored content, all of the constraint structure enters through (S_g, g). Universality and the irreducible specificity of the grounding are two faces of one fact.

4 Three Groundings

We instantiate the driver of Listing 1 with three groundings drawn from unrelated domains. Table 1 summarises the outcomes; Figure 2 shows the convergence of the residual $\|x_{t+1} - x_t\|$ for each.

4.1 Electronic structure: faithfulness, self-grounding, and an honest scaling result

Grounding. For a Hückel (tight-binding) chain of N sites with Hamiltonian H (nearest-neighbour coupling $\beta = -1$), the cut is the McWeeny purification polynomial

$$S(P) = 3P^2 - 2P^3,$$

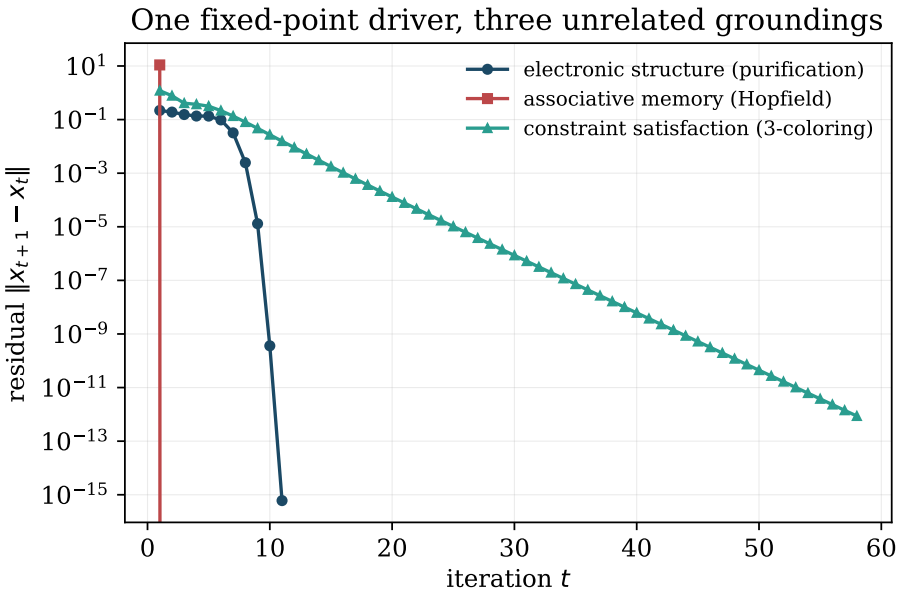

Figure 2: The identical driver converging in three unrelated domains. Electronic-structure purification reaches machine-precision idempotency; Hopfield recall reaches a fixed point in two steps; damped mean-field colouring converges more slowly to a proper assignment. The loop is the same in all three; only the cut and grounding differ.

N	E_{exact}	E_{fold}	|error|	$\text{Tr}(P)$	iters
4	−4.47214	−4.47214	8.9×10^{-16}	2.000	9
8	−9.51754	−9.51754	0	4.000	11
16	−19.67590	−19.67590	0	8.000	12
32	−40.03278	−40.03278	0	16.000	14
64	−80.76863	−80.76863	0	32.000	16
128	−162.25196	−162.25196	2.8×10^{-14}	64.000	17

Table 2: The diagonalization-free fold reproduces the exact electronic energy to machine precision and derives the correct occupancy $\text{Tr}(P) = N/2$ without being given it.

initialised at $P_0 = (\varepsilon_{\max} I - H)/(\varepsilon_{\max} - \varepsilon_{\min})$, which maps the spectrum of H into $[0, 1]$. The fixed point is the idempotent projector P onto the occupied subspace; the electronic energy is $E = 2\,\text{Tr}(PH)$.

Self-grounding. The only physical input beyond H is the electron count (half-filling). The driver is *not* told which states are occupied: the purification derives the occupied subspace itself, and the trace $\text{Tr}(P)$ converges to the correct occupancy $N/2$ as a consequence of the iteration, not as an imposed constraint (Table 2). Nothing eigenvector-derived is supplied.

Faithfulness. Table 2 compares the fold energy to exact diagonalization across system sizes. Agreement is at machine precision for every size, and the occupancy is recovered exactly.

N	eigensolve (s)	dense fold (s)
128	0.0007	0.0039
256	0.0037	0.0331
512	0.0193	0.2654
1024	0.1113	2.4002

Table 3: Wall-clock time, dense purification fold versus direct diagonalization. The dense fold is slower; we claim no speed advantage.

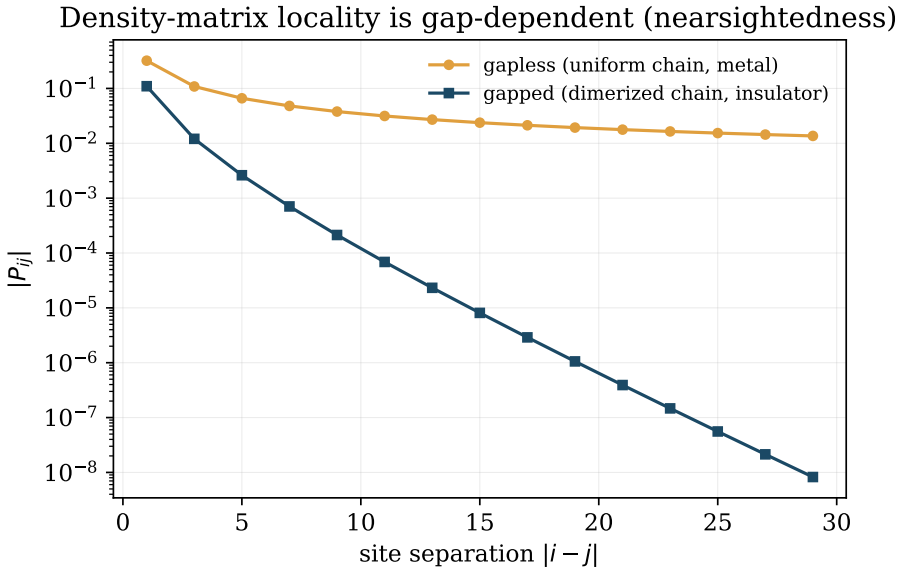

Figure 3: Density-matrix locality is gap-dependent. For a gapped (dimerised) chain the off-diagonal magnitude $|P_{ij}|$ decays rapidly (exponential "nearsightedness"); for a gapless uniform chain it decays only as a power law. Even separations vanish by the sublattice symmetry of the half-filled bipartite chain and are omitted. This decay is the structural basis for—but is not itself a demonstration of—linear-scaling purification.

Scaling: a negative result, honestly reported. Faithfulness is not efficiency. Table 3 times the dense fold against a direct eigensolve. The dense purification performs ∼15 matrix products per solve and is consistently *slower* than a single $O(N^3)$ diagonalization—about an order of magnitude slower at $N = 1024$. A scaling advantage exists in principle only by exploiting locality: the density matrix of a gapped system decays with distance (Figure 3), so a sparse, thresholded purification is $O(N)$ for insulators [3, 4]. We did not implement the sparse variant; the dense fold demonstrated here has no scaling advantage, and the locality that would enable one is absent for gapless (metallic) systems.

4.2 Associative memory

Grounding. We store $K = 3$ random bipolar patterns $p_k \in \{\pm 1\}^d$ ($d = 100$) in a Hopfield matrix $W = \sum_k p_k p_k^\top$ with zeroed diagonal [6]. The cut is the recurrent sign-update $S(x) = \mathrm{sgn}(Wx)$; its attractors include the stored patterns. We present a cue equal to p_1 with 30 of its 100 bits flipped.

Result. The driver converges in two iterations to a fixed point equal to p_1 exactly (Table 1, Figure 2): exact recall from a 30%-corrupted cue. The same loop that purified a density matrix now performs content-addressable recall, with only the cut and the encoding of the state changed.

4.3 Constraint satisfaction, and a bad cut as evidence of domain-blindness

Grounding. We 3-colour the triangular prism graph (two triangles $\{0,1,2\}$, $\{3,4,5\}$ joined by rungs 0–3, 1–4, 2–5; chromatic number 3). The state $P \in (\Delta_3)^6$ assigns each vertex a distribution over colours. The cut is a damped mean-field Potts update

$$S(P) = (1-\lambda)\,P + \lambda\,\mathrm{softmax}\big(-\beta\,AP + b\big),$$

with adjacency A, inverse temperature $\beta = 4$, damping $\lambda = 0.5$, and a small fixed field b that breaks the colour-permutation symmetry. Vertices are coloured by arg max at convergence.

A deliberate failure, then success. An undamped, unbroken mean-field cut ($\lambda = 1, b = 0$) does *not* converge: it oscillates and collapses toward the spurious symmetric state in which all vertices share one colour, which violates every edge. This is informative rather than embarrassing. It shows that the driver is genuinely domain-blind—it will faithfully iterate a *bad* cut toward a bad fixed point—and therefore that the problem-solving competence resides entirely in the cut, exactly as the factorisation predicts. With damping and symmetry-breaking restored, the same driver converges in 58 iterations to the proper colouring $[0,1,2,1,2,0]$, satisfying all nine edges (Table 1, Figure 2).

5 Grounding Richness Governs Predictive Power: An Ozone Case Study

The three-domain demonstration shows breadth. We now show *depth dependence*: that within a single domain, the predictive accuracy of the fold is set by the richness of the grounding, not by the operator. Ozone (O_3) is an ideal probe because its ground-state geometry is non-obvious—it is bent, not a symmetric ring—and the reason is a genuinely electronic effect.

Occupancy grounding fails. If each oxygen is grounded only by its occupancy (a featureless valence-2 atom), the relational fold that assembles the molecule sees three identical atoms whose only distinguishing variable is the number of bonds. More bonds lower the energy, so the fold selects the *cyclic* three-membered ring—the wrong topology. The information that distinguishes the isomers is simply absent from occupancy vectors, and no amount of relational composition over identical inputs can manufacture it.

Orbital-description grounding succeeds. We re-ground each oxygen in an actual orbital description: a Slater $2p$ orbital with exponent $\zeta = 2.275\,a_0^{-1}$. The π-overlaps between neighbouring orbitals are computed numerically from the orbital shapes, and a Hückel π-Hamiltonian is built over the molecular bond graph and folded with the same convergence machinery. The three-centre, four-electron π system is then resolved correctly: a three-membered ring bearing four π electrons is antiaromatic [16] and destabilised, whereas the open form delocalises the four electrons stably. The fold places the *bent* isomer below the cyclic one by 26.5 kcal/mol.

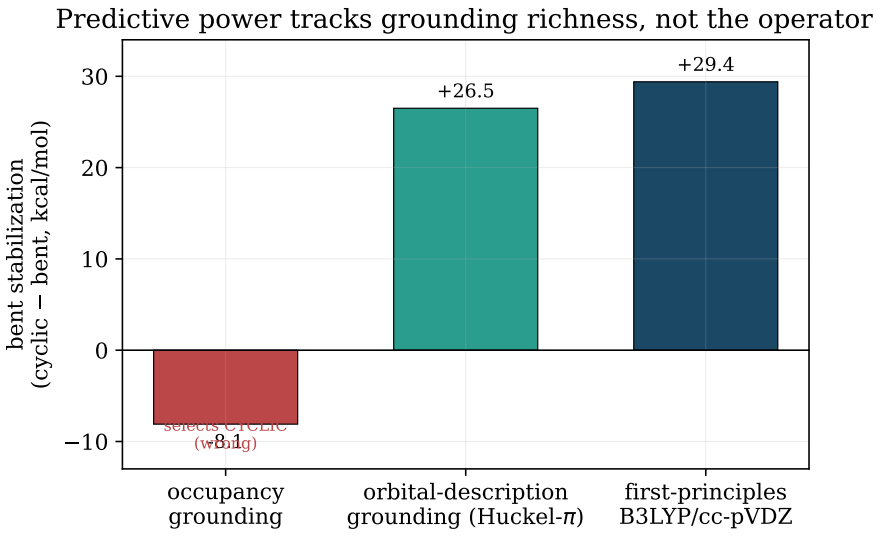

Figure 4: Predictive power tracks grounding richness, not the operator. The same relational fold selects the wrong (cyclic) isomer under an occupancy grounding, but recovers the correct bent ground state under an orbital-description grounding, agreeing with an independent B3LYP/cc-pVDZ calculation in sign and to ∼10% in magnitude. Positive values indicate the (correct) stabilisation of the bent form relative to the cyclic ring.

First-principles reference. To verify the sign and magnitude independently, we computed both isomers from first principles at the B3LYP/cc-pVDZ level [14, 15] (no tabulated energies, no fitted parameters). The bent form lies 29.4 kcal/mol below the cyclic form. The orbital-description fold thus agrees with full density-functional theory in sign and to within ∼10% in magnitude (Figure 4).

Interpretation. The "information floor" encountered with occupancy grounding is not a limit of the fold; it is a limit of the grounding. Enriching the bottom-layer parameterisation from a label to an orbital shape moved the prediction from qualitatively wrong to within 10% of first-principles theory, with the operator unchanged. This is the depth-axis counterpart of the breadth-axis result of Section 4: the driver is fixed, and competence is supplied by the cut.

6 Discussion

6.1 A structural universal, and its boundary as the result

Taken together, the experiments support a single interpretive claim. There is a genuine *structural universal* at work—*stable existence as the fixed point of a local, composable operation on a grounded representation*—and we have exhibited one driver instantiating it across domains (physics, memory, combinatorial search) that share no content. This is the kind of regularity from which unifying frameworks are built, and it is real: the convergence machinery transferred verbatim and produced correct fixed points in every case.

The boundary of the claim is not a hedge appended to it; it is the same fact viewed from the other side. The driver generalises *because* it is empty of content—a master profile fits many locks precisely by abandoning the particulars of any one. Consequently every correct answer in this paper

was produced by the grounding: the purification polynomial and Hamiltonian, the Hopfield weights, the colouring field, the orbital shapes. The universality of the form and the irreducible specificity of the grounding are inseparable. The honest and complete statement is therefore not "one operator solves everything" but "one operator *shape* organises a broad class of problems, and what solves any particular instance is the grounding cut for it."

6.2 Honest placement relative to established methods

We stress that the individual pieces are known. Density-matrix purification as a diagonalization-free route to electronic structure is classical and is the basis of an entire subfield of linear-scaling methods [1, 2, 3, 4]; Hopfield recall [6] and mean-field/message-passing constraint solving [7, 8] are standard; compositional distributed representations are well developed [9, 10]; and the cross-scale, entanglement-bounded character of relational encodings of physics is the lesson of tensor-network theory [11, 12]. The novelty we claim is not any kernel but the synthesis: the explicit driver/cut/grounding factorisation, the verbatim cross-domain demonstration, and the depth-dependence result that pins predictive power to grounding richness.

6.3 Toward a composable substrate for physical and cognitive state

The factorisation suggests a research programme rather than a finished result. Because the driver is domain-blind and the representation can be made compositional and closed under binding [9, 10], a converged physical state and a converged "identity" state can in principle live in the *same* representational space and be composed by the *same* operation. The electronic-structure rung of this paper provides a physically validated bottom layer for such a substrate; whether the same fixed-point operation can serve as a useful organising principle for adaptive or agent-like state is an open question we regard as the broader motivation for this work. We make no claim here to demonstrate cognition; we note only that the same operation that converges the states studied above is well posed on other grounded representations.

6.4 What would convert synthesis into discovery

Three concrete steps would strengthen the programme from a unifying synthesis into a source of new results. (i) *Scaling*: implement the sparse, thresholded purification on gapped systems and time it against diagonalization as N grows, to demonstrate (or refute) the linear-scaling advantage that Figure 3 only motivates. (ii) *Novel prediction*: apply the orbital-description fold to a quantity not already tabulated and confirm it against experiment or a higher-level calculation, converting "reproduces known structure" into "predicts unknown structure." (iii) *The entanglement frontier*: enrich the molecular grounding (adding the σ framework and lone pairs, scaling couplings from computed overlaps) and quantify how far orbital-description grounding climbs toward chemical accuracy before genuinely many-body correlation—the regime where relational composition is known to strain [12]—becomes the limiting factor.

7 Limitations

We collect the limitations explicitly.

- **The shared object is the driver, not the cut.** Across domains the convergence loop is identical, but the update map (cut) is domain-specific. We do not claim that a single function autonomously computes physics, memory, and search; we claim that one convergence *driver* does, given the appropriate cut.
- **No scaling advantage.** The dense purification is slower than diagonalization (Table 3). The linear-scaling regime requires a sparse implementation we did not build and applies only to gapped systems.
- **Hückel-class accuracy.** The electronic-structure results are at the level of a minimal relational model. The ozone agreement is in sign and order of magnitude ($\sim$10%), not chemical accuracy; full correlation is not captured.
- **Convergence is not guaranteed.** The cuts are not globally contractive. As Section 4.3 shows, a poorly chosen cut can oscillate or settle on a spurious fixed point; convergence inherits each domain's basin structure, requiring damping, symmetry-breaking, or good initialisation.
- **Known constituents.** The component algorithms are established; the contribution is their unification and framing, not the kernels.
- **The cognitive substrate is programmatic.** Claims in Section 6 about agent-like state are motivation and conjecture, not demonstrated capability.

8 Conclusion

We separated the machinery of fixed-point convergence from the content of any particular problem, formalised this as a domain-invariant driver acting on a domain-specific (cut, grounding) pair, and showed that one byte-identical driver converges to the correct fixed point across electronic structure, associative memory, and constraint satisfaction. A single-domain study of ozone showed that predictive accuracy is set by the richness of the grounding—occupancy-level grounding selects the wrong isomer, orbital-description grounding recovers the right one to within $\sim$10% of first-principles theory—with the operator unchanged. The constituent methods are classical and we demonstrate no efficiency gain; the contribution is the synthesis and the interpretive result that follows from it. The driver is universal precisely because it is empty, and the grounding is where every domain's truth resides. That boundary is not a limitation of the finding but its content: a structural universal for stable, composable existence, with reach across physics, memory, and search, in which the locus of knowledge is always the grounding.

Acknowledgements

Computational experiments and manuscript drafting were carried out by the author in collaboration with an AI assistant; all results reported here were generated by executable code and independently re-run for consistency. The author thanks colleagues at the Prometheus7 Research Institute for discussion.

Reproducibility

All experiments use small, self-contained NumPy programs; the first-principles ozone calculations use PySCF [15] at the B3LYP/cc-pVDZ level. The driver of Listing 1 is used verbatim in every domain, and all reported figures and tables are generated directly from the experiment scripts.

References

[1] R. McWeeny, "Some recent advances in density matrix theory," *Rev. Mod. Phys.* **32**, 335–369 (1960).

[2] A. H. R. Palser and D. E. Manolopoulos, "Canonical purification of the density matrix in electronic-structure theory," *Phys. Rev. B* **58**, 12704–12711 (1998).

[3] S. Goedecker, "Linear scaling electronic structure methods," *Rev. Mod. Phys.* **71**, 1085–1123 (1999).

[4] W. Kohn, "Density functional and density matrix method scaling linearly with the number of atoms," *Phys. Rev. Lett.* **76**, 3168–3171 (1996).

[5] E. Prodan and W. Kohn, "Nearsightedness of electronic matter," *Proc. Natl. Acad. Sci. USA* **102**, 11635–11638 (2005).

[6] J. J. Hopfield, "Neural networks and physical systems with emergent collective computational abilities," *Proc. Natl. Acad. Sci. USA* **79**, 2554–2558 (1982).

[7] J. S. Yedidia, W. T. Freeman, and Y. Weiss, "Constructing free-energy approximations and generalized belief propagation algorithms," *IEEE Trans. Inf. Theory* **51**, 2282–2312 (2005).

[8] M. Mézard, G. Parisi, and R. Zecchina, "Analytic and algorithmic solution of random satisfiability problems," *Science* **297**, 812–815 (2002).

[9] T. A. Plate, "Holographic reduced representations," *IEEE Trans. Neural Netw.* **6**, 623–641 (1995).

[10] P. Kanerva, "Hyperdimensional computing: an introduction to computing in distributed representation with high-dimensional random vectors," *Cogn. Comput.* **1**, 139–159 (2009).

[11] S. R. White, "Density matrix formulation for quantum renormalization groups," *Phys. Rev. Lett.* **69**, 2863–2866 (1992).

[12] R. Orús, "A practical introduction to tensor networks: matrix product states and projected entangled pair states," *Ann. Phys.* **349**, 117–158 (2014).

[13] S. Banach, "Sur les opérations dans les ensembles abstraits et leur application aux équations intégrales," *Fund. Math.* **3**, 133–181 (1922).

[14] A. D. Becke, "Density-functional thermochemistry. III. The role of exact exchange," *J. Chem. Phys.* **98**, 5648–5652 (1993).

[15] Q. Sun *et al.*, "PySCF: the Python-based simulations of chemistry framework," *WIREs Comput. Mol. Sci.* **8**, e1340 (2018).

[16] E. Hückel, "Quantentheoretische Beiträge zum Benzolproblem," *Z. Phys.* **70**, 204–286 (1931).