

The Universal Constraint Engine: Emergent Neuromorphic Architectures from Declarative Constraint Rules

Stephen C. Kinney

Bee Tree Holdings LLC, Freeport, NY

sckinney73@gmail.com

April 2026

Abstract

We introduce the Universal Constraint Engine (UCE), a system for generating emergent multi-state architectures from declarative constraint rules over conserved quantities. Unlike conventional neural network architectures that rely on learned weights, gradient descent, and massive training corpora, UCE derives computational behaviors—including memory, logic, hysteresis, and oscillation—directly from symbolic constraints without any training phase. The system comprises four layers: a Rule Definition Layer for specifying conserved quantities and transition constraints; a Constraint Solver Layer for state-space generation and constraint propagation; an Emergent Behavior Engine for identifying attractors, metastability, and higher-order dynamics; and an Embodiment Mapper for translating symbolic architectures into hardware implementations spanning FPGA, neuromorphic, spintronic, and quantum substrates. We present worked examples demonstrating that minimal rule sets (2–4 constraints over 3–4 conserved quantities) produce non-trivial emergent behaviors analogous to SR latches, biological oscillators, and write-gated memory cells. These results suggest that constraint-based architectures represent a fundamentally different approach to neuromorphic computing—one that discovers structure rather than learning it, and that operates independently of the matrix-arithmetic pipelines that dominate current AI hardware.

Keywords: neuromorphic computing, constraint satisfaction, emergent behavior, state-space generation, declarative rules, conserved quantities, hardware-agnostic architecture

1. Introduction

The dominant paradigm in artificial intelligence is built on a single computational primitive: the weighted sum followed by a nonlinear activation, organized into layers and trained by gradient descent over large datasets. This paradigm has produced remarkable results in language modeling, image recognition, and reinforcement learning. However, it carries fundamental limitations that become increasingly apparent as the field matures.

First, weight-based architectures are opaque. The knowledge encoded in billions of floating-point parameters cannot be inspected, verified, or formally reasoned about. Interpretability remains an open research problem. Second, they are expensive. Training frontier models requires thousands of GPUs running for months, consuming megawatts of power and hundreds of millions of dollars. Third, they are substrate-locked. The entire hardware ecosystem—from NVIDIA's CUDA to Google's TPUs to custom ASICs—is optimized for one operation: dense matrix multiplication. Any architecture that does not reduce to tensor operations cannot leverage this infrastructure.

We propose a fundamentally different approach. The Universal Constraint Engine (UCE) generates computational architectures not by learning weights from data, but by discovering emergent behaviors from declarative constraint rules over conserved quantities. The user specifies what is allowed and what is forbidden; the system derives everything else—including attractors, memory, logic, hysteresis, and oscillation—automatically.

This paper presents the UCE architecture, demonstrates its capabilities through worked examples, and discusses its implications for the future of neuromorphic computing. A provisional patent application covering the system has been filed with the United States Patent and Trademark Office (Application No. 64/036,854; filed April 12, 2026).

2. Background and Related Work

2.1 Weight-Based vs. Constraint-Based Computation

In weight-based systems (neural networks, transformers, diffusion models), behavior is encoded implicitly in parameter matrices. The system has no explicit model of what it knows or why it produces a given output. In constraint-based systems, behavior is encoded explicitly in rules. The system does exactly what the constraints allow—nothing more and nothing less. This distinction has profound implications for verifiability, energy efficiency, and hardware design.

2.2 Constraint Satisfaction Problems

Constraint satisfaction problems (CSPs) have a long history in computer science, from arc-consistency algorithms [1] to modern SAT solvers [2]. However, classical CSP work focuses on finding solutions that satisfy a given set of constraints. UCE inverts this: it uses constraints to generate an entire state-space and transition network, then analyzes the emergent structure of that network for higher-order behaviors.

2.3 Neuromorphic Computing

Neuromorphic computing seeks to build hardware that mimics biological neural systems [3]. Current approaches include Intel's Loihi [4], IBM's TrueNorth [5], and various memristive crossbar arrays [6]. These systems are designed to efficiently execute spiking neural networks or weight-based inference. UCE proposes a different foundation for neuromorphic hardware: constraint-driven state machines that exhibit emergent physical-like behavior without requiring neural-network-style computation.

2.4 Cellular Automata and Emergence

Cellular automata (CA) demonstrate that complex behavior can emerge from simple local rules [7]. UCE shares this philosophy but differs in three key ways: (a) UCE operates over conserved quantities rather than neighborhood-based cell states; (b) UCE explicitly constructs and analyzes the full state-space and transition network; and (c) UCE includes an embodiment mapper that translates symbolic architectures into physical hardware specifications.

3. System Architecture

The UCE comprises four principal layers, each receiving input from the preceding layer and producing structured output for the next. The pipeline is deterministic: given the same rule set, it always produces the same architecture.

3.1 Rule Definition Layer (RDL)

The RDL provides the interface through which a user defines the symbolic structure of a system. It accepts: (a) a set of conserved quantities $Q = \{q_1, q_2, \dots, q_n\}$, where each quantity represents a symbolic attribute governed by rules; (b) allowed transitions specifying which state changes are permitted; (c) forbidden transitions specifying which state changes are prohibited; (d) symmetry rules specifying invariances; and (e) control gates that enable or disable specific transition axes.

3.2 Constraint Solver Layer (CSL)

The CSL generates the complete state-space and transition network implied by the rule set. This involves: (a) enumerating all possible states as combinations of conserved quantity values; (b) constructing a directed graph $G = (S, T)$ where nodes are states and edges are allowed transitions; (c) applying constraint propagation to discover emergent constraints not explicitly defined by the user, including asymmetric reversibility, multi-step dependencies, and conservation-driven pruning; and (d) detecting emergent forbidden transitions.

3.3 Emergent Behavior Engine (EBE)

The EBE analyzes the refined transition network to identify higher-order behaviors. These include: **Attractors**—states or state-sets toward which the system naturally evolves; **Metastability**—states that are stable under certain conditions but transition under others; **Hysteresis**—path-dependent cycles producing history-dependent behavior; **Multi-state memory**—configurations where final states depend on transition sequence rather than instantaneous inputs; and **Logic behavior**—constraint-based Boolean-equivalent operations.

3.4 Embodiment Mapper

The Embodiment Mapper translates symbolic architectures into physical implementations. Supported targets include: software finite-state machines; synthesizable digital hardware (FPGA/ASIC via Verilog export); neuromorphic systems (memristive arrays with asymmetric pulse rules); spintronic devices (spin-polarization states with charge-conditioned switching); quantum systems (QUBO/Ising-model representations); and correlated-material embodiments (phase-change or coupled-order-parameter systems).

4. Worked Examples

We present three examples demonstrating that minimal constraint rule sets produce non-trivial emergent behaviors spanning digital logic, biological circuits, and memory elements. In each case, the emergent behaviors arise without explicit definition in the rule set.

4.1 Charge-Spin-Orbital System

We define three binary conserved quantities— C (charge-like), S (spin-like), and O (occupancy-like)—each taking values in $\{0, 1\}$, producing an 8-state initial state-space.

Rules:

```
allow_toggle(C) only if O = 1
allow_toggle(S) only if C = 1
allow_toggle(O) freely
forbid_simultaneous(C, S)
forbid_toggle(S) when C = 0
```

Emergent behaviors discovered: The EBE identifies (C=1, S=1, O=1) as an attractor; detects metastability at (C=1, S=0, O=1); discovers a hysteresis loop (1,0,1) → (1,1,1) → (1,1,0) → (1,0,0) → (1,0,1) that cannot be traversed in reverse; and identifies multi-state memory behavior where final states depend on transition sequence. These behaviors are physical analogues of phenomena found in spintronic and correlated-electron materials—derived entirely from five symbolic rules.

4.2 Bistable Latch (SR-Latch Analogue)

We define three binary conserved quantities: S (set), R (reset), Q (output).

Rules:

```
allow_toggle(Q) only if S=1 or R=1
forbid_simultaneous(S, R)
forbid_toggle(Q) if S=0 and R=0
```

Emergent behaviors discovered: The system produces 6 valid states (2 eliminated by the simultaneous-forbid rule). The EBE identifies dual attractors (Q=0 and Q=1 when S=0, R=0); logic-like behavior (S forces Q=1, R forces Q=0); and state retention (Q holds its value when no input is active). This is a functional SR latch—a fundamental digital memory element—arising from three constraint rules without any explicit logic gate definition.

4.3 Biological Oscillator (Repressilator Analogue)

We define four binary conserved quantities: A, B, C (protein-like quantities) and E (energy/activation gate), with cyclic dependency A→B→C→A.

Rules:

```
allow_toggle(A) only if C=1 and E=1
allow_toggle(B) only if A=1 and E=1
allow_toggle(C) only if B=1 and E=1
forbid_simultaneous(A, B)
```

Emergent behaviors discovered: The system produces 12 valid states. The EBE identifies an oscillatory attractor (A→B→C→A cycle); energy-gated behavior (oscillation occurs only when E=1); and metastability (system freezes when E drops to 0, persisting indefinitely until energy is restored). This mirrors a synthetic genetic repressilator—demonstrating that UCE can model biological circuit dynamics from purely symbolic constraints.

5. Discussion

5.1 Constraint-Based vs. Weight-Based: A Paradigm Comparison

Property	Weight-Based (Neural Nets)	Constraint-Based (UCE)
Knowledge encoding	Implicit in parameters	Explicit in rules
Training required	Yes (gradient descent)	No
Interpretability	Low (black box)	High (inspectable rules)
Hardware dependency	Matrix multiply units	Substrate-agnostic
Energy scaling	O(parameters)	O(states x transitions)
Correctness guarantee	Statistical	Formal (by construction)

Behavior source	Learned from data	Emergent from constraints
-----------------	-------------------	---------------------------

Table 1: Comparison of weight-based and constraint-based computing paradigms.

5.2 Implications for Hardware Design

Current AI hardware is optimized for a single computational primitive: the multiply-accumulate operation arranged in systolic arrays. This design choice reflects the dominance of weight-based architectures. Constraint-based systems like UCE do not require dense matrix multiplication. Their computational core—state enumeration, graph construction, and constraint propagation—maps naturally to graph processors, cellular architectures, and reconfigurable fabrics. This opens the possibility of AI-capable hardware that operates outside the CUDA/tensor-core ecosystem entirely.

The Embodiment Mapper's ability to generate synthesizable Verilog, QUBO representations, and neuromorphic specifications from a single rule set means that a single UCE design can target FPGA prototyping, quantum optimization, and memristive deployment without re-engineering the architecture.

5.3 Scalability Considerations

The state-space of a UCE system grows exponentially with the number of conserved quantities (2^n for binary quantities). For small systems ($n < 25$), classical enumeration is tractable. For larger systems, the QUBO/Ising-model embodiment provides a path to quantum-accelerated constraint solving, where quantum annealers or gate-based quantum computers can explore the state-space more efficiently than classical machines. The crossover point at which quantum acceleration becomes advantageous is estimated at approximately 25–30 conserved quantities.

5.4 Formal Verifiability

Because UCE architectures are derived from explicit rules and their state-spaces are fully enumerable, the resulting systems are formally verifiable by construction. Every possible state and transition is known. This property is significant for safety-critical applications where neural-network opacity is a liability—autonomous vehicles, medical devices, aerospace systems, and infrastructure control.

6. Conclusion

We have presented the Universal Constraint Engine, a system that generates emergent computational architectures from declarative constraint rules over conserved quantities. Through worked examples, we have demonstrated that minimal rule sets produce non-trivial behaviors including memory, logic, hysteresis, and oscillation—without training, without weights, and without dependence on tensor arithmetic hardware.

UCE represents a different trajectory for neuromorphic computing: one that prioritizes interpretability, formal verifiability, and substrate independence over parameter scale. As the AI industry confronts the rising costs and diminishing returns of brute-force scaling, constraint-based architectures offer an alternative foundation—one where intelligence emerges from structure rather than statistics.

The UCE prototype, documentation, and further technical details are available at universalconstraintengine.net. The underlying technology is protected by U.S. Provisional Patent Application No. 64/036,854 (filed April 12, 2026), assigned to Bee Tree Holdings LLC.

References

- [1] Mackworth, A. K. (1977). Consistency in networks of relations. *Artificial Intelligence*, 8(1), 99–118.
- [2] Biere, A., Heule, M., van Maaren, H., & Walsh, T. (Eds.). (2009). *Handbook of Satisfiability*. IOS Press.

- [3] Mead, C. (1990). Neuromorphic electronic systems. *Proceedings of the IEEE*, 78(10), 1629–1636.
- [4] Davies, M., et al. (2018). Loihi: A neuromorphic manycore processor with on-chip learning. *IEEE Micro*, 38(1), 82–99.
- [5] Merolla, P. A., et al. (2014). A million spiking-neuron integrated circuit with a scalable communication network and interface. *Science*, 345(6197), 668–673.
- [6] Yang, J. J., Strukov, D. B., & Stewart, D. R. (2013). Memristive devices for computing. *Nature Nanotechnology*, 8, 13–24.
- [7] Wolfram, S. (2002). *A New Kind of Science*. Wolfram Media.
- [8] D-Wave Systems. (2022). U.S. Patent No. 11,263,547 B2. Emergent behavior detection in quantum computing systems.
- [9] Angstadt, K., et al. (2022). AUTOMATASYNTH: Synthesizing automata from behavioral traces for FPGA deployment.
- [10] Marrows, C. H., et al. (2024). Neuromorphic spintronic computing architectures. *Nature Reviews Physics*.

This work is licensed under CC BY-NC-ND 4.0. Patent pending: U.S. Provisional Application No. 64/036,854. © 2026 Bee Tree Holdings LLC. All rights reserved.