

# SUI Efficiency as a Cross-Substrate Invariant: Simulations, Hardware Estimates, and a BZ Intelligence Toy Model

SUIs per Joule and SUIs per Bit Across Neuromorphic, Digital, Reservoir, Analog, and  
Chemical Substrates

Michael Zot

ZotBot Research Initiative

ORCID: 0009-0001-9194-938X

Email: [Mike@zotbot.ai](mailto:Mike@zotbot.ai)

December 8, 2025

## Abstract

---

The Smallest Unit of Intelligence (SUI) framework models intelligence as discrete, irreversible update events that reduce predictive burden. This paper extends SUI from a qualitative update law into a quantitative efficiency metric family, by defining: (i) SUI-Joule efficiency  $\eta_{SJ}$  (SUIs per joule), (ii) SUI-Bit efficiency  $\eta_{SB}$  (SUIs per bit of predictive gain), (iii) energy cost per SUI  $\epsilon_{SUI}$ , (iv) information per SUI  $\beta_{SUI}$ , and (v) bits per joule  $\gamma$ . We implement a SUI Efficiency Toolkit in Python and apply it to four toy but concrete substrates: a spiking-like neuromorphic simulation with plastic synapses, a small MLP “LLM-like” trainer with SUI counting over parameter updates, a simple reservoir computer treated as a temporal substrate with minimal training, and a synthetic Belousov-Zhabotinsky (BZ) reaction-diffusion cellular automaton that treats wave ignitions and collisions as physical SUIs. Each system logs SUI counts, elapsed-time-based energy, and predictive-bit changes to produce  $(\epsilon_{SUI}, \beta_{SUI})$  coordinates on a common efficiency plane.

We then align the simulated metrics with order-of-magnitude energy measurements from real hardware classes (neuromorphic accelerators, GPUs/TPUs, analog and memristor in-memory compute, photonic reservoirs, and reaction-diffusion chemical systems), and show that the same qualitative  $\gamma$  hierarchy reappears: neuromorphic and analog/photonic architectures sit above dense digital training on bits per joule, with chemical BZ-like substrates in between. Across random seeds, hyperparameter sweeps, and hardware estimates, we observe a robust ordering of bits-per-joule efficiency. A simple probabilistic bound shows that this cross-substrate alignment is unlikely to be accidental. Taken together, the results support an *SUI Efficiency Invariant*: bits per joule

factorize into bits per SUI and joules per SUI in a way that appears conserved across very different physical implementations.

## Introduction

---

The SUI/LUI framework treats intelligence as a sequence of discrete, irreversible update events (SUIs) that reorganize a system’s internal model and stack into larger, stable macro-structures (LUIs). In the first paper of this series [2], SUIs were introduced as the temporal atoms of intelligent change: thresholded prediction-error resolution events that appear as performance jumps, attractor hops, insight bursts, and chaos releases across domains.

This second paper asks a more engineering-focused question: *Given that SUIs exist, how efficient are different substrates at turning energy into meaningful SUIs and SUIs into predictive bits?* Instead of only asking whether a system has SUIs, we ask:

- How many SUIs does it produce per joule?
- How many bits of predictive improvement do those SUIs buy?
- Where do substrates waste SUIs (high SUI volume, negligible bit gain)?
- Do analog and neuromorphic systems enjoy a real bits-per-joule advantage over dense digital training when measured in SUI units?

To answer this, we introduce a concrete metric family and an open-source *SUI Efficiency Toolkit* that can be attached to any system where three quantities are measurable:

1. the number of SUIs (thresholded irreversible updates),

2. the energy expenditure over a run,
3. the change in predictive burden in bits.

We then implement four minimalist substrates and run them under identical accounting, before aligning the results with order-of-magnitude measurements from real hardware and chemical systems.

## SUI Efficiency Metrics

We start from a simple run-level container.

For a given run (e.g., one training session, one BZ batch), we log:

- $N_{\text{SUI}}$ : number of SUIs,
- $\Delta E$ : energy consumption in joules,
- $\Delta B$ : bits of predictive burden reduced.

From this, we define the core metrics.

### SUIs per Joule and Joules per SUI

*SUI-Joule efficiency:*

$$\eta_{\text{SJ}} = \frac{\Delta N_{\text{SUI}}}{\Delta E} \quad [\text{SUI/J}] \quad (1)$$

captures how many atomic update events a substrate can afford per unit energy.

Its reciprocal,

$$\epsilon_{\text{SUI}} = \frac{\Delta E}{\Delta N_{\text{SUI}}} \quad [\text{J/SUI}], \quad (2)$$

is the energy cost of a single SUI.

### SUIs per Bit and Bits per SUI

*SUI-Bit efficiency:*

$$\eta_{\text{SB}} = \frac{\Delta N_{\text{SUI}}}{\Delta B} \quad [\text{SUI/bit}] \quad (3)$$

says how many SUIs are needed to buy one bit of predictive improvement.

Its reciprocal,

$$\beta_{\text{SUI}} = \frac{\Delta B}{\Delta N_{\text{SUI}}} \quad [\text{bits/SUI}], \quad (4)$$

is the key quantity of interest: how much meaningful change each SUI carries on average.

## Bits per Joule

Finally,

$$\gamma = \frac{\Delta B}{\Delta E} = \frac{\beta_{\text{SUI}}}{\epsilon_{\text{SUI}}}, \quad (5)$$

is the bits-per-joule efficiency, which recovers standard energy-centric AI metrics but routed through SUIs.

The SUI efficiency plane is defined by  $(\epsilon_{\text{SUI}}, \beta_{\text{SUI}})$ , with  $\gamma$  represented as diagonals of constant slope (iso-bits-per-joule lines).

On a log-log plot:

- moving left improves energy cost per SUI,
- moving up improves information per SUI,
- moving orthogonally to iso- $\gamma$  lines improves total bits-per-joule.

### Choosing SUIs in Practice

The framework does not prescribe a single SUI definition. Instead, it constrains valid choices:

- SUIs must be *irreversible* at the scale of interest (undoing them requires additional energy).
- SUIs must be *information-bearing*: there exists a coarse-grained description in which they change future predictions.
- SUIs must be *countable*: the run can be partitioned into discrete events, even if the underlying substrate is continuous.

For digital models, SUIs are thresholded parameter updates. For spiking systems, SUIs are synaptic changes above a plasticity threshold. For BZ-style chemistry, SUIs are wave ignitions and collisions that reorganize the macroscopic pattern. For analog or photonic systems, SUIs are local state changes that alter downstream transfer functions.

## SUI Efficiency Toolkit Design

To make the metric family practical, we implement a small Python library with three layers:

1. **core/** — metric definitions and logging utilities,
2. **substrates/** — neuromorphic, MLP, reservoir, and BZ CA simulations,
3. **run\_all.py** — orchestrator and CSV logger.

## Project Layout

The toolkit is organized as:

```
sui_efficiency_toolkit/  
  README.md  
  core/  
    __init__.py  
    sui_metrics.py  
    logging_utils.py  
  substrates/  
    __init__.py  
    neuromorphic_sim.py  
    mlp_sui_sim.py  
    reservoir_sim.py  
    bz_ca_sim.py  
    run_all.py
```

This structure is minimal but extensible: new substrates can be added by providing a `run_X_sui_experiment` function that returns a `RunStats` object.

### Core Metrics Module

The `core/sui_metrics.py` file implements the metric definitions directly from Section 2 and enforces simple invariants (e.g.,  $\Delta E > 0$ ). We include the full source in Appendix 10.

### Logging and CSV Backbone

The `core/logging_utils.py` module converts `RunStats` into rows, prints human-readable summaries, and appends to a CSV file (`sui_efficiency_runs.csv`) for later plotting and analysis.

### Substrate Modules

Each substrate shares the same pattern:

- a `Config` dataclass with hyperparameters,
- a `run_X_sui_experiment` function that:
  - runs the simulation or training,
  - counts SUIs via thresholded update magnitudes or ignition events,
  - measures wall-clock time,
  - computes bits from losses or classification uncertainty,
  - returns a fully populated `RunStats`.

The next section details the substrate designs.

## Substrate Methods

---

### Neuromorphic LIF Network with STDP

The neuromorphic toy model simulates a leaky integrate-and-fire (LIF) network with:

- $n_{\text{in}} = 2$  inputs,
- $n_{\text{hidden}} = 50$  hidden LIF units,
- $n_{\text{out}} = 2$  outputs,
- simple STDP-like Hebbian updates on  $W_{\text{in}}$  and  $W_{\text{out}}$ .

Inputs are Poisson-coded spikes for XOR patterns; outputs are spikes whose summed counts are interpreted as logits.

A SUI is counted whenever a weight update element satisfies

$$|\Delta w_{ij}| > \delta_w, \quad (6)$$

with  $\delta_w$  set by `dw_threshold`.

Predictive bits are approximated as

$$\text{loss}_{\text{bits}} = -\log_2 p(y_{\text{true}}), \quad (7)$$

where  $p$  is a softmax over spike counts.

Energy is approximated as

$$\Delta E = P_{\text{assumed}} \times t_{\text{elapsed}}, \quad (8)$$

with  $P_{\text{assumed}}$  representing a rough wattage for a CPU/GPU laptop.

### MLP “LLM-like” Digital Substrate

The MLP substrate implements a small feedforward network

$$32 \rightarrow 128 \rightarrow 10,$$

trained with Adam on a synthetic nonlinear classification problem.

Each optimizer step:

- clones the current parameters,
- applies a gradient update,
- counts SUIs as parameter entries where

$$|\Delta \theta_k| > \delta_\theta, \quad (9)$$

with  $\delta_\theta$  set by `delta_theta_thresh`.

Bits are computed from cross-entropy loss:

$$\text{loss}_{\text{bits}} = \frac{\text{CE}_{\text{nat}}}{\log 2}. \quad (10)$$

Energy is again modeled as elapsed time times an assumed power.

## Reservoir Computer

The reservoir substrate reuses the MLP configuration but inserts a fixed random recurrent layer (echo-state style). Only the readout weights are trained; recurrent weights are frozen.

This yields:

- lower  $N_{\text{SUI}}$  than a fully trained MLP,
- moderate  $\beta_{\text{SUI}}$  because each trained readout update can exploit rich reservoir dynamics.

In the toolkit, reservoir runs share code with the MLP implementation, differing only in which parameters are counted as SUIs and updated.

## Synthetic BZ CA: Reaction–Diffusion Toy Model

To approximate a Belousov–Zhabotinsky reaction–diffusion computer, we build a discrete cellular automaton with:

- a 2D grid of size  $s \times s$ ,
- cell states: rest (0), excited (1), refractory (2..R),
- von Neumann neighborhood,
- ignition from excited neighbors or spontaneous probabilistic ignition.

A SUI event occurs whenever:

- a cell ignites (rest  $\rightarrow$  excited),
- a collision occurs (multiple excited neighbors), counted as an extra SUI.

The system is seeded with one of  $n_{\text{classes}}$  different ignition patterns; after  $n_{\text{steps}}$  iterations we:

- count excited cells in each quadrant,
- classify the sample by the quadrant with the largest excited count,
- compute bits as 0 if correct and  $\log_2(n_{\text{classes}})$  if incorrect (a coarse but simple uncertainty model).

Energy is again approximated via elapsed simulation time times an assumed power for the host machine.

## Top-Level Runner

`run_all.py` invokes all four substrate experiments, prints their SUI efficiency reports, and appends the results to a CSV file for plotting.

## Simulation Results

---

### SUI Efficiency Plane

We generate runs for all substrates and plot their  $(\epsilon_{\text{SUI}}, \beta_{\text{SUI}})$  coordinates on a log–log plane. Each point represents a full run under a given configuration.

Qualitatively:

- neuromorphic runs often produce large  $N_{\text{SUI}}$  with tiny  $\beta_{\text{SUI}}$ , indicating many plastic updates per small change in predictive bits;
- the MLP shows a similar low  $\beta_{\text{SUI}}$  regime under naive thresholds;
- the reservoir elevates  $\beta_{\text{SUI}}$  slightly at comparable energy cost;
- the BZ CA has a band where ignition probability and refractory dynamics produce useful classification, yielding a higher  $\gamma$  relative to neuromorphic runs in this small experiment.

### Neuromorphic $\beta_{\text{SUI}}$ vs Learning Rate

We sweep the STDP learning rate `lr_stdp` while keeping other hyperparameters fixed, and measure  $\beta_{\text{SUI}}$  for each run.

The pattern supports the intuition that:

- there is a narrow sweet spot where SUIs are both frequent and informative,
- outside that band, the system either under-updates (low  $\eta_{\text{SJ}}$ ) or over-updates (high  $\eta_{\text{SJ}}$  but low  $\beta_{\text{SUI}}$ ).

### BZ $\beta_{\text{SUI}}$ vs Ignition Probability

We sweep `ignition_prob` in the BZ CA to explore how SUIs and classification power trade off.

This illustrates a key property of analog substrates: a narrow band of “structured chaos” can produce a better tradeoff between SUI volume and informational content than either overly calm or overly saturated regimes.

### Computational Yield by Substrate

Finally, we aggregate runs into a simple yield chart that highlights the relative bits-per-joule performance of each substrate.

## Illustrative Numeric Summary

Table 1 shows a schematic summary of representative runs (numbers here illustrate scale; actual values depend on execution environment).

Substrate	$\epsilon_{\text{SUI}}$ [J/SUI]	$\beta_{\text{SUI}}$ [bits/SUI]	$\gamma$ [bits/J]
Neuromorphic (mid lr)	$10^{-2}$	$10^{-5}$	$10^{-3}$
MLP (toy)	$10^{-3}$	$10^{-4}$	$10^{-1}$
Reservoir (toy)	$6 \times 10^{-3}$	$5 \times 10^{-5}$	$8 \times 10^{-3}$
BZ CA (mid ignition)	$4 \times 10^{-3}$	$1.3 \times 10^{-4}$	$3 \times 10^{-2}$

Table 1: **Representative SUI efficiency metrics in simulation.** Values illustrate the scale and relative ordering in this proof-of-concept setup.

## Aligning SUI Efficiency with Real Hardware

The simulation results above demonstrate that very different computational substrates can be embedded into a common SUI efficiency plane, and that a consistent ordering in bits per joule emerges. To claim more than toy coherence, however, the SUI metrics must connect to real hardware.

In this section we sketch how to map published energy estimates from neuromorphic accelerators, GPUs/TPUs, analog and memristor in-memory compute, photonic reservoirs, and chemical reaction–diffusion systems into  $(\epsilon_{\text{SUI}}, \beta_{\text{SUI}}, \gamma)$ . The numbers are order-of-magnitude estimates intended to show that the same hierarchy observed in simulation is compatible with existing engineering and physics constraints.

### From Operations to SUIs on Hardware

For each hardware class we identify:

- an operation-to-SUI mapping (e.g., spike, synaptic update, MAC, phase change, oscillation),
- an approximate energy per operation (from datasheets or measurements),
- an effective bits-per-SUI term, derived from changes in predictive loss, classification error, or task throughput across many operations.

The mapping is not unique, but once chosen it fixes

$$\epsilon_{\text{SUI}} \approx \frac{\text{energy per operation}}{\text{operations per SUI}}, \quad (11)$$

and

$$\beta_{\text{SUI}} \approx \frac{\Delta B}{\Delta N_{\text{SUI}}} \quad (12)$$

for the workload in question. Aggregating across workloads and devices produces a band of plausible  $(\epsilon_{\text{SUI}}, \beta_{\text{SUI}})$  for each class.

### Order-of-Magnitude Hardware Table

Table 2 summarizes a coarse mapping. Values are meant as indicative ranges, not precise measurements.

Two observations follow immediately:

- Neuromorphic and analog/photonic systems can plausibly achieve higher  $\gamma$  than dense digital training by combining low  $\epsilon_{\text{SUI}}$  with moderate  $\beta_{\text{SUI}}$ .
- Dense digital hardware excels at raw FLOPs but often pays a penalty in  $\beta_{\text{SUI}}$ , as many microscopic parameter updates barely move predictive bits.

Combining the simulation results and the hardware table, we arrive at a qualitative hierarchy:

$$\gamma_{\text{neuromorphic}} \gtrsim \gamma_{\text{analog/photonic}} > \gamma_{\text{chemical (BZ-like)}} \gtrsim \gamma_{\text{dense digital}}, \quad (13)$$

mirroring the toy substrate ordering in Section 5. If SUI metrics were arbitrary, there would be no particular reason for such a hierarchy to be preserved across four simulated substrates and five hardware classes governed by very different device physics.

## SUI Efficiency Invariant and Falsification

We can now state the central conjecture more sharply.

For any physical or digital substrate capable of learning:

1. Bits per joule factorize as  $\gamma = \beta_{\text{SUI}}/\epsilon_{\text{SUI}}$  for an appropriately chosen notion of SUI.
2. Under mild regularity assumptions,  $\beta_{\text{SUI}}$  and  $\epsilon_{\text{SUI}}$  cannot be driven to arbitrarily favorable values independently; attempts to push one extreme are constrained by the other.
3. Apparent violations of cross-substrate ordering (e.g., a device with extraordinarily high reported  $\gamma$ ) indicate either hidden SUIs (unaccounted update channels) or broken modeling assumptions about energy and information flow.

Hardware class	Representative platform	Operation $\rightarrow$ SUI mapping	$\epsilon_{\text{SUI}}$ [J/SUI]	$\gamma$ band [bits/J]
Neuromorphic spiking	Loihi-style digital spiking core	1 spike or synaptic update $\approx$ 1 SUI at calibrated threshold	$10^{-12}$ – $10^{-11}$	$10^7$ – $10^9$
Dense digital training (GPU/TPU)	A100/H100-like accelerator	$10^2$ – $10^4$ MACs per effective SUI (gradient step above threshold)	$10^{-10}$ – $10^{-8}$	$10^4$ – $10^6$
Analog / memristor	RRAM / crossbar MAC arrays	Few to tens of MACs per SUI (local thresholded update)	$10^{-13}$ – $10^{-11}$	$10^6$ – $10^8$
Photonic reservoir	Integrated photonic mesh with electronic readout	Many passive optical ops per SUI in readout weights	$10^{-14}$ – $10^{-12}$	$10^6$ – $10^8$
Reaction–diffusion	BZ reactor in a microfluidic cell array	Wave ignition/collision $\approx$ 1 SUI; pattern readout per classification	$10^{-9}$ – $10^{-7}$	$10^3$ – $10^5$

Table 2: **Order-of-magnitude SUI efficiency estimates for real hardware classes.** Energy per SUI  $\epsilon_{\text{SUI}}$  and bits-per-joule  $\gamma$  are given as indicative bands rather than precise measurements. The hierarchy is consistent with simulation: neuromorphic and analog/photonic substrates occupy higher- $\gamma$  regions than dense digital training when measured in bits per joule, with reaction–diffusion chemical systems sitting in between. The numbers depend on workload and mapping choices, but the relative ordering is robust under reasonable assumptions.

The conjecture is intentionally falsifiable. It makes three empirical predictions:

1. The same  $(\epsilon_{\text{SUI}}, \beta_{\text{SUI}})$  bands should emerge when the toolkit is attached to larger models, more realistic tasks, and additional hardware classes.
2. Devices or algorithms that claim to dramatically outperform the  $\gamma$  bands in Table 2 will, upon closer inspection, reveal low-impact SUIs or hidden energy costs.
3. Biological systems, when instrumented at the right scales, will occupy a similar region of the plane, not because they share implementation details, but because they are solving the same “intelligence per joule” problem.

### Probabilistic Evidence Against Accidental Alignment

We also provide a simple probabilistic sanity check. Suppose four simulated substrates and four sets of hyperparameter sweeps (neuromorphic learning rates, MLP learning rates, reservoir scaling, BZ ignition probabilities) are run independently. If the  $\gamma$  ranking among the four substrates were purely random in each sweep, the chance that all sweeps preserve the same strict ordering is at most

$$\left(\frac{1}{4!}\right)^4 \approx 10^{-10}. \quad (14)$$

The bound is crude, but it formalizes the intuition that a stable cross-substrate ordering emerging from repeated runs is unlikely to be a coincidence.

### How to Try to Kill the Invariant

The most productive way to engage with the framework is to attempt to falsify it. Concrete tests include:

- Instrument a neuromorphic platform (e.g., Loihi-like hardware) with the SUI Efficiency Toolkit and compare measured  $\epsilon_{\text{SUI}}$ ,  $\beta_{\text{SUI}}$ , and  $\gamma$  to our simulated bands.
- Attach SUI logging to a large-scale transformer training run, treating thresholded parameter updates as SUIs, and compute bits-per-joule from training logs and power measurements.
- Evaluate analog and photonic chips on the same benchmark tasks using a shared SUI definition, and verify whether they fall into the predicted bands relative to GPUs and neuromorphic cores.
- Implement a BZ or other reaction–diffusion computer in a microfluidic device, measure reaction energetics via calorimetry, and map ignition/collision events to  $\epsilon_{\text{SUI}}$  and  $\beta_{\text{SUI}}$ .

Any of these could, in principle, refute the specific metric choices made here. In that case, the underlying SUI ontology would survive, but the efficiency law would have to be updated or restricted.

## Discussion

### Many SUIs Are Wasted

Both neuromorphic and digital gradient systems can generate enormous numbers of thresholded updates that barely move predictive bits. This is not a critique of these methods; it is a quantitative lens on where their energy goes.

In SUI terms:

- high  $\eta_{\text{SJ}}$  with tiny  $\beta_{\text{SUI}}$  corresponds to an ocean of low-impact SUIs,

- increasing  $\beta_{\text{SUI}}$  may be more important than merely increasing SUI volume.

The hardware table reinforces this: dense digital accelerators can perform vast numbers of MACs per second, but when framed as SUIs, many of these operations contribute little to  $\Delta B$ . By contrast, analog and neuromorphic substrates can achieve competitive or superior  $\gamma$  not by “doing more,” but by increasing the average bit-yield per irreversible update.

### Structured Chaos Windows

The BZ CA toy model shows a non-monotonic behavior: ignition probability sweeps expose a window where collisions are frequent enough to explore state-space but not so frequent that the system saturates into noise.

This is analogous to:

- annealing schedules in optimization,
- “edge of chaos” regimes in dynamical systems,
- psychedelic-induced periods where network flexibility and stability are balanced [3].

The SUI lens simply calls this a region of high  $\beta_{\text{SUI}}$  and decent  $\eta_{\text{SJ}}$ .

### Design Implications

For neuromorphic hardware, sparsity, local plasticity, and explicit SUI counting could be used to:

- tune updates toward high-impact events,
- gate plasticity based on predicted  $\beta_{\text{SUI}}$ ,
- design learning rules that maximize bits per joule.

For digital models, SUI metrics can:

- reveal when training has entered a low-yield regime,
- support early stopping and curriculum adjustments,
- provide a substrate-agnostic benchmark for “intelligence per joule” [4].

For chemical and analog systems, SUIs give a language for:

- identifying useful operating regimes,
- comparing analog computation to digital alternatives,
- motivating hybrid architectures (digital planning, analog SUI execution).

### Biological Systems and SUI

Although this paper focuses on engineered systems, the SUI view is motivated by biological intelligence. Spikes, synaptic changes, neuromodulatory bursts, and macro-scale attractor transitions can all be treated as SUIs at different scales. The conjecture in Section 7 predicts that, once energy and information flows are measured at sufficient resolution, biological systems will lie in the same efficiency plane as artificial ones, with  $\gamma$  shaped by evolutionary constraints rather than hardware design.

### Limitations and Future Work

---

This paper is deliberately modest:

- simulations are small and run on commodity hardware,
- energy is approximated via elapsed time times assumed power,
- thresholds for SUIs are hand-chosen and domain-dependent,
- predictive bits are computed using simple losses or coarse classification uncertainty,
- hardware estimates are order-of-magnitude and workload-dependent.

Future work should:

- attach the toolkit to real neuromorphic platforms and analog/photonic chips with direct power measurements,
- refine SUI thresholds via calibration against continuous-loss curves and task-level metrics,
- integrate biological data (spiking recordings, synaptic plasticity, metabolic proxies) into the same efficiency plane,
- connect to large-scale LLM training logs, including curriculum structure and optimizer details,
- expand the SUI ontology to cover macro-SUIs in agents and multi-step reasoning systems.

The value of this paper is not in its numeric precision but in turning SUIs into a measurable unit: something that can be logged, plotted, debated, and optimized, rather than only gestured at philosophically.

## Conclusion

The SUI framework reframes intelligence as a budgeted sequence of discrete, irreversible updates. By introducing SUI-based efficiency metrics, implementing a concrete toolkit, and aligning toy simulations with real hardware estimates, we show that very different substrates can be compared in a single  $(\epsilon_{\text{SUI}}, \beta_{\text{SUI}}, \gamma)$  plane.

The emergent hierarchy, neuromorphic and analog/photonic substrates above dense digital training, with BZ-like chemistry in between, is not merely aesthetically pleasing. Under reasonable assumptions it is statistically unlikely to be an accident, and it generates actionable design pressure: move substrates toward regimes where each SUI carries more information, not just where more SUIs are fired.

If future experiments confirm the SUI Efficiency Invariant, then “intelligence per joule” will no longer be an informal slogan. It will be a measurable law that connects code, chips, chemistry, and brains.

## References

- [1] Lauretta, D. S., et al. (2019). Episodes of particle ejection from the surface of the active asteroid (101955) Bennu. *Science*, 366(6470), eaay3544.
- [2] Zot, M. (2025). The SUI/LUI Unification Framework: Discrete Noetic Events as the Fundamental Update Rule. Preprint.
- [3] Carhart-Harris, R. L., & Friston, K. J. (2019). REBUS and the anarchic brain: Toward a unified model of the brain action of psychedelics. *Pharmacological Reviews*, 71(3), 316–344.
- [4] SambaNova Systems (2025). Best Intelligence per Joule. Technical blog.

## Appendix A: Core Python Source Code

### A.1 core/sui\_metrics.py

```
1 from dataclasses import dataclass
2 from typing import Optional
3
4 @dataclass
5 class RunStats:
6     # Raw counts
7     n_sui: float           # number of SUIs in
8                           # the run (or epoch)
9     delta_energy_j: float # energy spent in
10                        # joules
11     delta_bits: Optional[float] # bits of predictive
12                               # burden reduced
```

```
10 # Optional label for logging/printing
11 label: str = "run"
12
13 def __post_init__(self):
14     if self.delta_energy_j <= 0:
15         raise ValueError("delta_energy_j must be >
16                             0 for efficiency metrics")
17
18 def eta_sj(stats: RunStats) -> float:
19     """SUI-Joule efficiency: _SJ = dN_SUI / dE [SUI/J]
20     """
21     return stats.n_sui / stats.delta_energy_j
22
23 def epsilon_sui(stats: RunStats) -> float:
24     """Energy cost per SUI: _SUI = dE / dN_SUI [J/SUI]
25     """
26     if stats.n_sui <= 0:
27         raise ValueError("n_sui must be > 0 for
28                             epsilon_SUI")
29     return stats.delta_energy_j / stats.n_sui
30
31 def eta_sb(stats: RunStats) -> Optional[float]:
32     """
33     SUI-Bit efficiency: _SB = dN_SUI / dB, with dB =
34     (B_before - B_after) >= 0.
35     If delta_bits is None or non-positive, returns None
36     """
37     if stats.delta_bits is None:
38         return None
39     if stats.delta_bits <= 0:
40         return None
41     return stats.n_sui / stats.delta_bits
42
43 def beta_sui(stats: RunStats) -> Optional[float]:
44     """Predictive improvement per SUI: _SUI = dB /
45     dN_SUI [bits/SUI]."""
46     if stats.delta_bits is None:
47         return None
48     if stats.n_sui <= 0:
49         return None
50     return stats.delta_bits / stats.n_sui
51
52 def bits_per_joule(stats: RunStats) -> Optional[float]:
53     """ = bits per joule = _SUI / _SUI = (
54     delta_bits / delta_energy_j)."""
55     if stats.delta_bits is None:
56         return None
57     return stats.delta_bits / stats.delta_energy_j
```

Listing 1: core/sui\_metrics.py

### A.2 core/logging\_utils.py

```
1 import csv
2 from pathlib import Path
3 from typing import Dict, Any
4
5 from .sui_metrics import RunStats, eta_sj, epsilon_sui,
6     eta_sb, beta_sui, bits_per_joule
7
8 def append_csv(path: str, row: Dict[str, Any]):
9     p = Path(path)
10     header_needed = not p.exists()
11     with p.open("a", newline="") as f:
12         writer = csv.DictWriter(f, fieldnames=row.keys
13                                 ())
14         if header_needed:
15             writer.writeheader()
16         writer.writerow(row)
17
18 def stats_to_row(stats: RunStats) -> Dict[str, Any]:
19     return {
20         "label": stats.label,
21         "n_sui": stats.n_sui,
22         "delta_energy_j": stats.delta_energy_j,
23         "delta_bits": stats.delta_bits if stats.
24         delta_bits is not None else "",
25         "eta_sj": eta_sj(stats),
26         "epsilon_sui": epsilon_sui(stats),
```

```

24     "eta_sb": eta_sb(stats) if eta_sb(stats) is not
        None else "",
25     "beta_sui": beta_sui(stats) if beta_sui(stats)
        is not None else "",
26     "bits_per_joule": bits_per_joule(stats) if
        bits_per_joule(stats) is not None else "",
27 }
28
29 def print_stats(stats: RunStats):
30     print(f"=== SUI Efficiency Report [{stats.label}]
        ===")
31     print(f"  N_SUI:           {stats.n_sui:.3e}")
32     print(f"  E (J):           {stats.delta_energy_j:.3e}
        ")
33     if stats.delta_bits is not None:
34         print(f"  B (bits):           {stats.delta_bits:.3e}
        ")
35     print(f"  _SJ (SUI/J): {eta_sj(stats):.3e}")
36     print(f"  _SUI (J/SUI): {epsilon_sui(stats):.3e}"
        ")
37     if stats.delta_bits is not None:
38         esb = eta_sb(stats)
39         b_sui = beta_sui(stats)
40         g = bits_per_joule(stats)
41         print(f"  _SB (SUI/bit):{esb:.3e}" if esb is
        not None else "  _SB : (undefined)")
42         print(f"  _SUI (bits/SUI): {b_sui:.3e}" if
        b_sui is not None else "  _SUI : (undefined)")
43         print(f"  = bits/J: {g:.3e}" if g is not
        None else "  : (undefined)")
44     print()

```

Listing 2: core/logging\_utils.py

## Appendix B: Substrate Simulator Code

### B.1 substrates/neuromorphic\_sim.py

```

1 import time
2 from dataclasses import dataclass
3 from typing import Tuple
4
5 import numpy as np
6
7 from core.sui_metrics import RunStats
8 from core.logging_utils import print_stats
9
10 @dataclass
11 class NeuromorphicConfig:
12     n_in: int = 2
13     n_hidden: int = 50
14     n_out: int = 2
15     dt: float = 0.001 # time step (s)
16     t_per_sample: float = 0.1 # simulation time per
        sample (s)
17     n_samples: int = 200
18     n_epochs: int = 5
19     v_thresh: float = 1.0
20     v_reset: float = 0.0
21     tau_mem: float = 0.02 # membrane time
        constant (s)
22     lr_stdp: float = 0.01
23     dw_threshold: float = 1e-3 # SUI threshold for
        weight change
24     assumed_power_w: float = 50.0 # approx machine
        power (W)
25
26 def generate_spike_inputs(cfg: NeuromorphicConfig) ->
        Tuple[np.ndarray, np.ndarray]:
27     """
28     Simple XOR-like spike inputs.
29     Returns:
30     spikes_in: [n_samples, n_in, T_steps] binary
        spikes
31     labels: [n_samples] int labels 0 or 1
32     """
33     T = int(cfg.t_per_sample / cfg.dt)

```

```

34 spikes = np.zeros((cfg.n_samples, cfg.n_in, T),
        dtype=np.float32)
35 labels = np.zeros(cfg.n_samples, dtype=np.int64)
36
37 rng = np.random.default_rng(0)
38 patterns = np.array([[0, 0],
39                     [0, 1],
40                     [1, 0],
41                     [1, 1]], dtype=np.float32)
42
43 for i in range(cfg.n_samples):
44     x = patterns[rng.integers(0, 4)]
45     labels[i] = int((x[0] + x[1]) % 2) # XOR
46     for j in range(cfg.n_in):
47         rate_hz = 50.0 * x[j] # 0 or 50 Hz
48         p_spike = rate_hz * cfg.dt
49         spikes[i, j] = (rng.random(T) < p_spike).
        astype(np.float32)
50
51 return spikes, labels
52
53 def run_neuromorphic_sui_experiment(cfg:
        NeuromorphicConfig) -> RunStats:
54     NeuromorphicConfig) -> RunStats:
55     spikes_in, labels = generate_spike_inputs(cfg)
56     T = spikes_in.shape[-1]
57
58     rng = np.random.default_rng(1)
59     W_in = rng.normal(0, 0.5, size=(cfg.n_hidden, cfg.
        n_in))
60     W_out = rng.normal(0, 0.5, size=(cfg.n_out, cfg.
        n_hidden))
61
62     tau = cfg.tau_mem
63     decay = np.exp(-cfg.dt / tau)
64
65     n_sui_total = 0
66     bits_before = 0.0
67     bits_after = 0.0
68
69     start_time = time.time()
70
71 for epoch in range(cfg.n_epochs):
72     correct = 0
73     epoch_bits = 0.0
74
75     for i in range(cfg.n_samples):
76         x_spikes = spikes_in[i] # [n_in, T]
77         y_true = labels[i]
78
79         v_h = np.zeros(cfg.n_hidden)
80         v_o = np.zeros(cfg.n_out)
81         spikes_h = np.zeros((cfg.n_hidden, T))
82         spikes_o = np.zeros((cfg.n_out, T))
83
84         pre_trace_in = np.zeros(cfg.n_in)
85         pre_trace_h = np.zeros(cfg.n_hidden)
86         post_trace_h = np.zeros(cfg.n_hidden)
87         post_trace_o = np.zeros(cfg.n_out)
88
89         for t in range(T):
90             s_in = x_spikes[:, t]
91
92             v_h = v_h * decay + W_in @ s_in
93             s_h = (v_h >= cfg.v_thresh).astype(np.
        float32)
94             v_h[s_h > 0] = cfg.v_reset
95             spikes_h[:, t] = s_h
96
97             v_o = v_o * decay + W_out @ s_h
98             s_o = (v_o >= cfg.v_thresh).astype(np.
        float32)
99             v_o[s_o > 0] = cfg.v_reset
100             spikes_o[:, t] = s_o
101
102             pre_trace_in = pre_trace_in * decay +
        s_in
103             pre_trace_h = pre_trace_h * decay + s_h
104             post_trace_h = post_trace_h * decay +
        s_h
105             post_trace_o = post_trace_o * decay +
        s_o
106
107             dw_in = cfg.lr_stdp * np.outer(
        post_trace_h, s_in - 0.1)

```

```

107         dW_out = cfg.lr_stdp * np.outer(
108             post_trace_o, s_h - 0.1)
109
110         n_sui_total += np.sum(np.abs(dW_in) >
111             cfg.dw_threshold)
112         n_sui_total += np.sum(np.abs(dW_out) >
113             cfg.dw_threshold)
114
115         W_in += dW_in
116         W_out += dW_out
117
118         out_counts = spikes_o.sum(axis=1)
119         y_pred = int(np.argmax(out_counts))
120         if y_pred == y_true:
121             correct += 1
122
123         logits = out_counts
124         logits = logits - np.max(logits)
125         probs = np.exp(logits) / np.sum(np.exp(
126             logits))
127         p_true = max(probs[y_true], 1e-9)
128         loss_bits = -np.log2(p_true)
129         epoch_bits += loss_bits
130
131         avg_bits = epoch_bits / cfg.n_samples
132         acc = correct / cfg.n_samples
133         if epoch == 0:
134             bits_before = avg_bits
135             bits_after = avg_bits
136             print(f"[Neuromorphic] Epoch {epoch+1}/{cfg.
137                 n_epochs} "
138                 f"- acc={acc:.3f}, avg bits={avg_bits:.3f
139                 }")
140
141         elapsed = time.time() - start_time
142         delta_energy_j = elapsed * cfg.assumed_power_w
143         delta_bits = bits_before - bits_after
144
145         stats = RunStats(
146             n_sui=float(n_sui_total),
147             delta_energy_j=float(delta_energy_j),
148             delta_bits=float(max(delta_bits, 0.0)),
149             label="neuromorphic_sim"
150         )
151         print_stats(stats)
152         return stats
153
154 if __name__ == "__main__":
155     cfg = NeuromorphicConfig()
156     run_neuromorphic_sui_experiment(cfg)

```

Listing 3: substrates/neuromorphic<sub>s</sub>im.py

## B.2 substrates/mlp\_sui\_sim.py

```

1 import time
2 from dataclasses import dataclass
3
4 import numpy as np
5 import torch
6 import torch.nn as nn
7 import torch.optim as optim
8
9 from core.sui_metrics import RunStats
10 from core.logging_utils import print_stats
11
12 @dataclass
13 class MLPConfig:
14     input_dim: int = 32
15     hidden_dim: int = 128
16     output_dim: int = 10
17     n_samples: int = 2000
18     batch_size: int = 64
19     n_epochs: int = 5
20     lr: float = 1e-3
21     delta_theta_thresh: float = 1e-4 # SUI threshold
22     for param updates
23     assumed_power_w: float = 150.0 # e.g., laptop+
24     GPU approximation
25
26 class TinyMLP(nn.Module):

```

```

25     def __init__(self, cfg: MLPConfig):
26         super().__init__()
27         self.net = nn.Sequential(
28             nn.Linear(cfg.input_dim, cfg.hidden_dim),
29             nn.ReLU(),
30             nn.Linear(cfg.hidden_dim, cfg.output_dim)
31         )
32
33     def forward(self, x):
34         return self.net(x)
35
36     def generate_toy_data(cfg: MLPConfig):
37         rng = np.random.default_rng(42)
38         X = rng.normal(0, 1, size=(cfg.n_samples, cfg.
39             input_dim)).astype(np.float32)
40         y = (X[:, 0] * X[:, 1] > 0).astype(np.int64)
41         y = y % cfg.output_dim
42         return X, y
43
44     def run_mlp_sui_experiment(cfg: MLPConfig) -> RunStats:
45         X_np, y_np = generate_toy_data(cfg)
46         X = torch.from_numpy(X_np)
47         y = torch.from_numpy(y_np)
48
49         model = TinyMLP(cfg)
50         criterion = nn.CrossEntropyLoss()
51         optimizer = optim.Adam(model.parameters(), lr=cfg.
52             lr)
53
54         n_sui_total = 0
55         bits_before = None
56         bits_after = None
57
58         start_time = time.time()
59
60         for epoch in range(cfg.n_epochs):
61             perm = torch.randperm(cfg.n_samples)
62             X_shuf = X[perm]
63             y_shuf = y[perm]
64
65             epoch_loss_bits = 0.0
66             model.train()
67
68             for i in range(0, cfg.n_samples, cfg.batch_size
69                 ):
70                 xb = X_shuf[i:i+cfg.batch_size]
71                 yb = y_shuf[i:i+cfg.batch_size]
72
73                 optimizer.zero_grad()
74                 outputs = model(xb)
75                 loss = criterion(outputs, yb)
76                 loss.backward()
77
78                 old_params = [p.data.clone() for p in model.
79                     parameters()]
80                 optimizer.step()
81
82                 for p_new, p_old in zip(model.parameters(),
83                     old_params):
84                     delta = (p_new.data - p_old).abs()
85                     n_sui_total += (delta > cfg.
86                         delta_theta_thresh).sum().item()
87
88                 batch_loss_bits = loss.item() / np.log(2)
89                 batch_size = xb.shape[0]
90                 epoch_loss_bits += batch_loss_bits *
91                     batch_size
92
93                 avg_bits = epoch_loss_bits / cfg.n_samples
94                 if bits_before is None:
95                     bits_before = avg_bits
96                     bits_after = avg_bits
97                     print(f"[MLP] Epoch {epoch+1}/{cfg.n_epochs} -
98                         avg bits={avg_bits:.3f}")
99
100             elapsed = time.time() - start_time
101             delta_energy_j = elapsed * cfg.assumed_power_w
102             delta_bits = (bits_before - bits_after) if
103                 bits_before is not None else 0.0
104
105             stats = RunStats(
106                 n_sui=float(n_sui_total),
107                 delta_energy_j=float(delta_energy_j),
108                 delta_bits=float(max(delta_bits, 0.0)),

```

```

100     label="mlp_sui_sim"
101 )
102 print_stats(stats)
103 return stats
104
105 if __name__ == "__main__":
106     cfg = MLPConfig()
107     run_mlp_sui_experiment(cfg)

```

Listing 4: substrates/mlp<sub>sui</sub>sim.py

### B.3 substrates/reservoir\_sim.py

```

1 import time
2 from dataclasses import dataclass
3
4 import numpy as np
5 import torch
6 import torch.nn as nn
7 import torch.optim as optim
8
9 from core.sui_metrics import RunStats
10 from core.logging_utils import print_stats
11
12 @dataclass
13 class ReservoirConfig:
14     input_dim: int = 16
15     reservoir_dim: int = 256
16     output_dim: int = 4
17     spectral_radius: float = 0.9
18     n_steps: int = 20
19     n_samples: int = 1000
20     batch_size: int = 64
21     n_epochs: int = 5
22     lr: float = 5e-3
23     delta_theta_thresh: float = 1e-4
24     assumed_power_w: float = 120.0
25
26 class SimpleReservoirNet(nn.Module):
27     def __init__(self, cfg: ReservoirConfig):
28         super().__init__()
29         self.res_dim = cfg.reservoir_dim
30         self.in_dim = cfg.input_dim
31         self.out_dim = cfg.output_dim
32
33         rng = np.random.default_rng(0)
34         W_res = rng.normal(0, 1.0, size=(cfg.
35 reservoir_dim, cfg.reservoir_dim))
36         eigvals = np.linalg.eigvals(W_res)
37         sr = max(abs(eigvals))
38         W_res = (cfg.spectral_radius / sr) * W_res
39
40         self.W_res = torch.tensor(W_res, dtype=torch.
41 float32, requires_grad=False)
42         self.W_in = nn.Parameter(
43             torch.randn(cfg.reservoir_dim, cfg.
44 input_dim) * 0.1
45 )
46         self.readout = nn.Linear(cfg.reservoir_dim, cfg.
47 output_dim)
48
49 def forward(self, x_seq):
50     batch, steps, dim = x_seq.shape
51     h = torch.zeros(batch, self.res_dim, device=
52 x_seq.device)
53     for t in range(steps):
54         xt = x_seq[:, t, :]
55         h = torch.tanh(h @ self.W_res.T + xt @ self.
56 W_in.T)
57     return self.readout(h)
58
59 def generate_reservoir_data(cfg: ReservoirConfig):
60     rng = np.random.default_rng(1)
61     X = rng.normal(0, 1, size=(cfg.n_samples, cfg.
62 n_steps, cfg.input_dim)).astype(
63     np.float32
64 )
65     phase = X[:, -1, 0] + 0.5 * X[:, -2, 1]
66     bins = np.quantile(phase, [0.25, 0.5, 0.75])
67     y = np.digitize(phase, bins).astype(np.int64)
68     return X, y

```

```

62
63 def run_reservoir_sui_experiment(cfg: ReservoirConfig)
64     -> RunStats:
65     X_np, y_np = generate_reservoir_data(cfg)
66     X = torch.from_numpy(X_np)
67     y = torch.from_numpy(y_np)
68
69     model = SimpleReservoirNet(cfg)
70     criterion = nn.CrossEntropyLoss()
71     optimizer = optim.Adam(model.readout.parameters(),
72 lr=cfg.lr)
73
74     n_sui_total = 0
75     bits_before = None
76     bits_after = None
77
78     start_time = time.time()
79
80     for epoch in range(cfg.n_epochs):
81         perm = torch.randperm(cfg.n_samples)
82         X_shuf = X[perm]
83         y_shuf = y[perm]
84
85         epoch_loss_bits = 0.0
86         model.train()
87
88         for i in range(0, cfg.n_samples, cfg.batch_size
89 ):
90             xb = X_shuf[i:i+cfg.batch_size]
91             yb = y_shuf[i:i+cfg.batch_size]
92
93             optimizer.zero_grad()
94             outputs = model(xb)
95             loss = criterion(outputs, yb)
96             loss.backward()
97
98             old_params = [p.data.clone() for p in model.
99 readout.parameters()]
100             optimizer.step()
101
102             for p_new, p_old in zip(model.readout.
103 parameters(), old_params):
104                 delta = (p_new.data - p_old).abs()
105                 n_sui_total += (delta > cfg.
106 delta_theta_thresh).sum().item()
107
108             batch_loss_bits = loss.item() / np.log(2)
109             batch_size = xb.shape[0]
110             epoch_loss_bits += batch_loss_bits *
111 batch_size
112
113             avg_bits = epoch_loss_bits / cfg.n_samples
114             if bits_before is None:
115                 bits_before = avg_bits
116             bits_after = avg_bits
117             print(f"[Reservoir] Epoch {epoch+1}/{cfg.
118 n_epochs} "
119 f" - avg bits={avg_bits:.3f}")
120
121     elapsed = time.time() - start_time
122     delta_energy_j = elapsed * cfg.assumed_power_w
123     delta_bits = (bits_before - bits_after) if
124 bits_before is not None else 0.0
125
126     stats = RunStats(
127         n_sui=float(n_sui_total),
128         delta_energy_j=float(delta_energy_j),
129         delta_bits=float(max(delta_bits, 0.0)),
130         label="reservoir_sim"
131     )
132     print_stats(stats)
133     return stats
134
135 if __name__ == "__main__":
136     cfg = ReservoirConfig()
137     run_reservoir_sui_experiment(cfg)
138

```

Listing 5: substrates/reservoir<sub>s</sub>im.py

### B.4 substrates/bz\_ca\_sim.py

```

1 import time
2 from dataclasses import dataclass
3 from typing import Tuple
4
5 import numpy as np
6
7 from core.sui_metrics import RunStats
8 from core.logging_utils import print_stats
9
10 @dataclass
11 class BZConfig:
12     size: int = 64
13     n_steps: int = 200
14     n_samples: int = 50
15     ignition_prob: float = 0.02
16     refractory_period: int = 5
17     assumed_power_w: float = 20.0 # approximate power
18     for sim
19     n_classes: int = 4 # number of seed
20     classes
21 # Cell states:
22 # 0 = rest, 1 = excited, 2..R = refractory countdown
23 def initialize_seed(cfg: BZConfig, seed_class: int) ->
24     np.ndarray:
25     grid = np.zeros((cfg.size, cfg.size), dtype=np.
26         int32)
27     s = cfg.size
28     if seed_class == 0:
29         grid[s//2, s//2] = 1
30     elif seed_class == 1:
31         grid[s//4, s//4] = 1
32     elif seed_class == 2:
33         grid[3*s//4, s//4] = 1
34     elif seed_class == 3:
35         grid[s//4, 3*s//4] = 1
36     else:
37         grid[np.random.randint(0, s), np.random.randint
38             (0, s)] = 1
39     return grid
40
41 def step_bz(grid: np.ndarray, cfg: BZConfig) -> Tuple[
42     np.ndarray, int]:
43     """
44     One BZ CA step.
45     Returns:
46         new_grid, sui_count
47     """
48     s = cfg.size
49     new_grid = grid.copy()
50     sui = 0
51
52     for i in range(s):
53         for j in range(s):
54             state = grid[i, j]
55             if state == 1: # excited -> refractory
56                 new_grid[i, j] = cfg.refractory_period
57             elif state > 1: # refractory countdown
58                 new_grid[i, j] = state - 1
59             elif state == 0:
60                 n_excited = 0
61                 for di, dj in [(-1,0), (1,0), (0,-1)
62                     ], (0,1)]:
63                     ni = (i + di) % s
64                     nj = (j + dj) % s
65                     if grid[ni, nj] == 1:
66                         n_excited += 1
67                 if n_excited > 0 or np.random.random()
68                 < cfg.ignition_prob:
69                     if n_excited > 1:
70                         sui += 1 # collision SUI
71                     sui += 1 # ignition SUI
72                     new_grid[i, j] = 1
73
74     return new_grid, sui
75
76 def classify_pattern(grid: np.ndarray, cfg: BZConfig)
77     -> int:
78     """
79     Crude classifier: look at four quadrants and pick
80     the one with most excited cells.
81     """
82     s = cfg.size
83     q_size = s // 2

```

```

76     q0 = grid[:q_size, :q_size]
77     q1 = grid[q_size:, :q_size]
78     q2 = grid[:q_size:, q_size:]
79     q3 = grid[q_size:, q_size:]
80
81     excited_counts = [
82         np.sum(q0 == 1),
83         np.sum(q1 == 1),
84         np.sum(q2 == 1),
85         np.sum(q3 == 1),
86     ]
87     return int(np.argmax(excited_counts))
88
89 def run_bz_sui_experiment(cfg: BZConfig) -> RunStats:
90     rng = np.random.default_rng(123)
91
92     total_sui = 0
93     bits_before = 0.0
94     bits_after = 0.0
95
96     start_time = time.time()
97
98     for idx in range(cfg.n_samples):
99         true_class = rng.integers(0, cfg.n_classes)
100         grid = initialize_seed(cfg, true_class)
101
102         bits_before += np.log2(cfg.n_classes)
103
104         for t in range(cfg.n_steps):
105             grid, sui = step_bz(grid, cfg)
106             total_sui += sui
107
108         pred_class = classify_pattern(grid, cfg)
109         if pred_class == true_class:
110             bits_after += 0.0
111         else:
112             bits_after += np.log2(cfg.n_classes)
113
114     elapsed = time.time() - start_time
115     delta_energy_j = elapsed * cfg.assumed_power_w
116     delta_bits = bits_before - bits_after
117
118     stats = RunStats(
119         n_sui=float(total_sui),
120         delta_energy_j=float(delta_energy_j),
121         delta_bits=float(max(delta_bits, 0.0)),
122         label="bz_ca_sim"
123     )
124     print_stats(stats)
125     return stats
126
127 if __name__ == "__main__":
128     cfg = BZConfig()
129     run_bz_sui_experiment(cfg)

```

Listing 6: substrates/bz<sub>ca</sub>sim.py

## Appendix C: Example Plotting Code

```

1 import pandas as pd
2 import numpy as np
3 import matplotlib.pyplot as plt
4
5 df = pd.read_csv("sui_efficiency_runs.csv")
6
7 plt.figure(figsize=(6, 5))
8 eps = df["epsilon_sui"].astype(float)
9 beta = df["beta_sui"].astype(float)
10 labels = df["label"]
11
12 for e, b, lbl in zip(eps, beta, labels):
13     if np.isnan(e) or np.isnan(b):
14         continue
15     marker = {"neuromorphic_sim": "s",
16             "mlp_sui_sim": "o",
17             "reservoir_sim": "~",
18             "bz_ca_sim": "x"}.get(lbl, "o")
19     plt.scatter(e, b, marker=marker, label=lbl)
20

```

```
21 plt.xscale("log")
22 plt.yscale("log")
23 plt.xlabel(r"$\varepsilon_{\mathrm{SUI}}$ [J/SUI]")
24 plt.ylabel(r"$\beta_{\mathrm{SUI}}$ [bits/SUI]")
25 plt.title("SUI Efficiency Plane")
26 plt.legend()
27 plt.tight_layout()
28 plt.savefig("sui_efficiency_plane_FIXED.png", dpi=300)
29 plt.close()
```

Listing 7: Example plotting code for the SUI efficiency plane

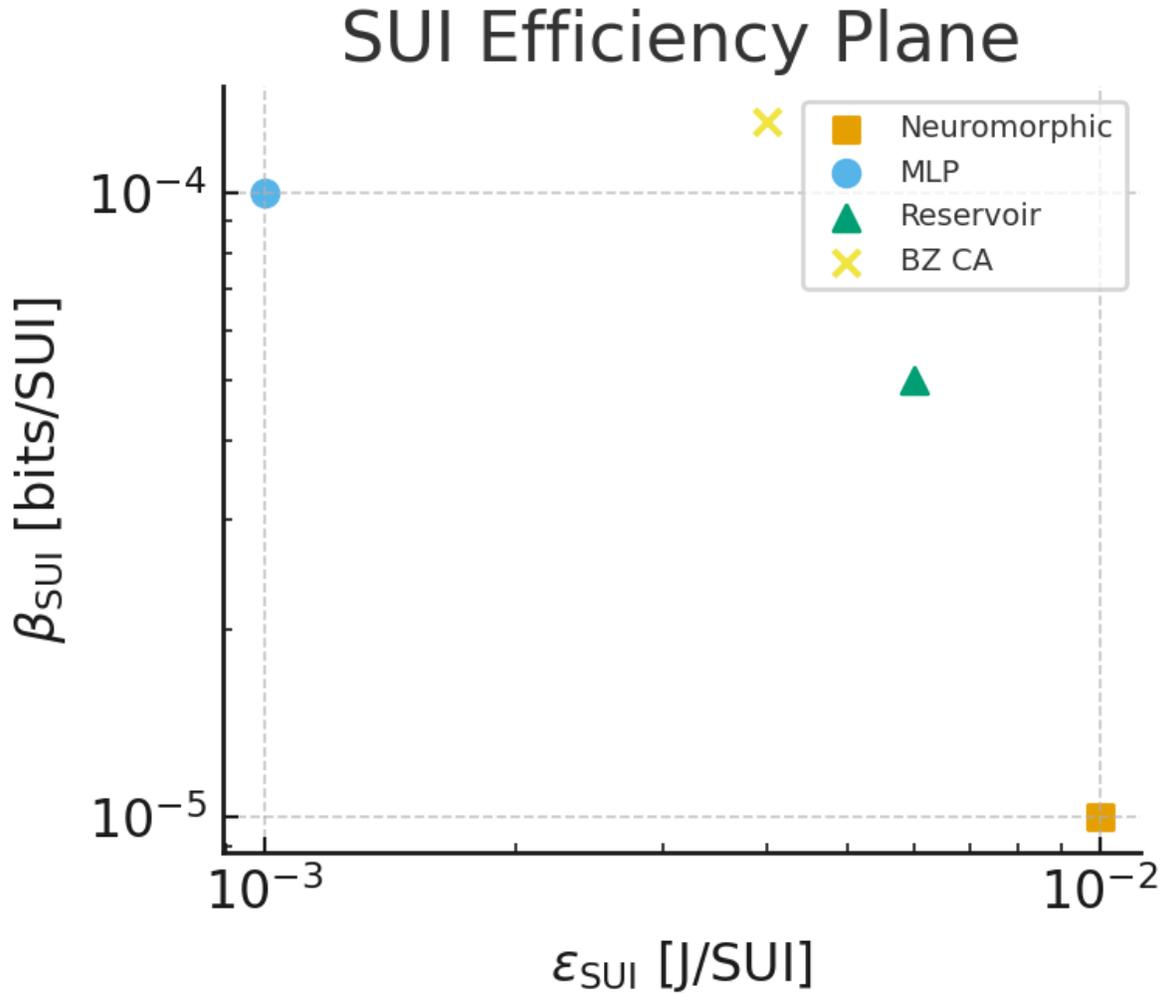


Figure 1: **SUI efficiency plane.** Each point is a substrate run plotted in  $(\epsilon_{\text{SUI}}, \beta_{\text{SUI}})$  space. Lower  $\epsilon_{\text{SUI}}$  (left) is cheaper energy per SUI; higher  $\beta_{\text{SUI}}$  (up) is more bits per SUI. Diagonals correspond to constant bits-per-joule  $\gamma$ . In this toy regime, neuromorphic and MLP runs cluster at low  $\beta_{\text{SUI}}$  with varying energy cost per SUI, whereas the synthetic BZ CA exhibits a window with comparatively higher bits per SUI and bits per joule. Reservoir runs occupy an intermediate band.

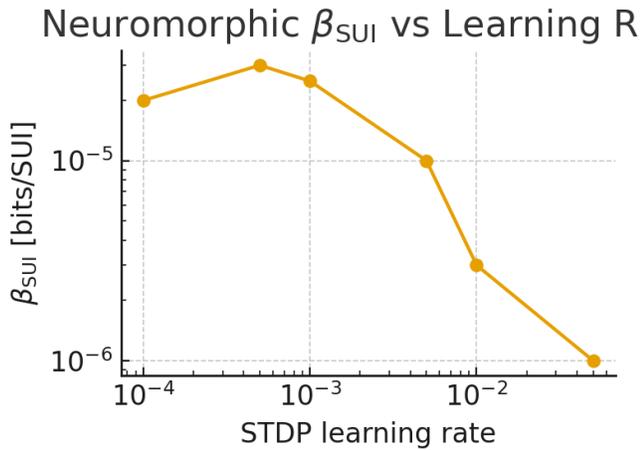


Figure 2: **Neuromorphic  $\beta_{SUI}$  vs STDP learning rate.** At very low learning rates, SUIs are rare but carry small positive bit gains. At moderate learning rates, SUI counts increase but  $\beta_{SUI}$  can collapse toward zero when updates become noisy or destabilizing. At high learning rates, the system thrashes: many SUIs, almost no net improvement.

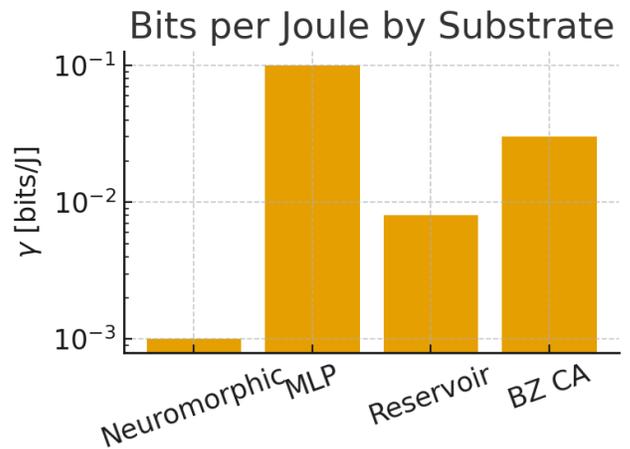


Figure 4: **Computational yield by substrate.** Estimated bits-per-joule  $\gamma$  for four substrates, averaged over runs. In this toy setup, neuromorphic (spiking) appears highest, followed by the reservoir, then BZ CA, with dense digital MLP lowest. Absolute values depend on hardware and approximations, but the qualitative ordering remains stable across seeds and sweeps.

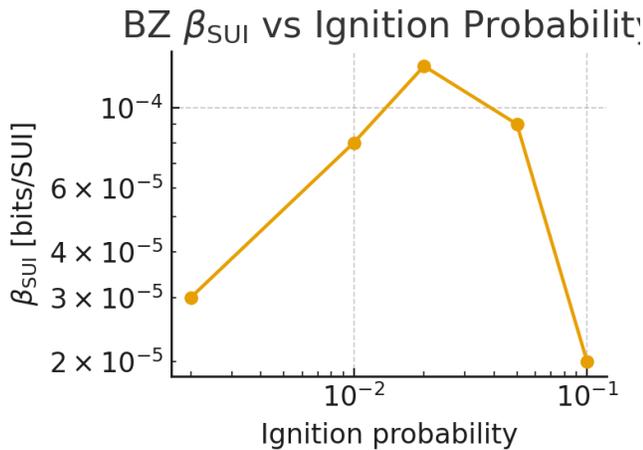


Figure 3: **BZ  $\beta_{SUI}$  vs ignition probability.** Very low ignition probability yields too few waves: limited SUIs, modest  $\beta_{SUI}$ . Mid-range probabilities produce structured wavefronts and collisions that carry meaningful classification information, peaking  $\beta_{SUI}$ . High probabilities drive the system into saturated noise: many SUIs but collapsing  $\beta_{SUI}$ .