

**Subthreshold Worlds:
Sorites, Spikes, and Accumulation under a Coarse-Grained
Readout**

PIATRA . INSTITUTE

July 2026

Abstract

The oldest puzzle about heaps asks for a number: at how many grains does sand become a heap? This paper argues the number was never the point, and that the heap belongs to a wide family of systems in which a latent state accumulates small changes while a coarse readout classifying it stays fixed, until at some point the readout flips. Write the state update as $x_{t+1} = F(x_t, u_t)$ and the readout as $y_t = R(x_t)$. The sorites condition is that for many steps $R(x_{t+1}) = R(x_t)$ although $x_{t+1} \neq x_t$, and the paradox is the confusion of readout-invariance with state-invariance, of a tolerance relation that is not transitive with an equivalence relation that is. The paper develops this as one architecture across five domains and grounds four of its claims in a reproducible simulation. A population of leaky integrators driven by a stream of small inputs holds the accumulated count in full, so that a linear decoder recovers it exactly, at an R-squared of 1.0 that a deterministic linear state guarantees by construction rather than reports as a finding, while the thresholded readout stays silent for the first 37 steps and carries no information at all about how much has arrived, which is the exact sense in which a subthreshold input is nothing only to the spike that has not fired. The indistinguishability relation, in which adjacent states differ by less than a just-noticeable amount, holds for every neighboring pair yet its transitive closure merges a thousand plainly distinct states into a single class, and three quarters of its chained triples already fail transitivity, so the paradox is a relation error, not a mystery about grains. A soft threshold makes local tolerance and global transition coexist without contradiction: the largest change in category probability produced by a single grain is 0.0062 while the total change across the range is 1.0. And a bistable latent state driven slowly across a fold reproduces the emergence of a new category, a spike, a heap, a multicellular individual, with the coarse readout giving no warning at all until the jump, while the variance of the latent state rises roughly 15-fold beforehand, its autocorrelation climbs from 0.9 toward 0.99, and a reservoir trained only on pre-transition data raises an alarm 198 steps before the readout moves. The synthesis is deliberately bounded. It does not solve the sorites by naming a grain number, it does not claim that neurons are heaps or that cells are reservoirs, and it takes the sheaf-theoretic reading of contextuality only as far as the honest slogan that locally adequate readouts need not glue into one global predicate. What it offers is a single formal spine, latent accumulation under coarse-grained readout, on which vagueness, psychophysics, synaptic integration, reservoir dynamics, and the major transitions of individuality can be read as one design and not five metaphors. ## 1. The Wrong Question

How many grains make a heap? The question has the form of a request for a boundary, a number N such that $N-1$ grains are not a heap and N grains are, and the whole trouble is that no such number survives inspection. Add one grain to a non-heap and surely you still have a non-heap, since a single grain cannot be the difference between a heap and its absence; iterate that impeccable step and you prove that a million grains are not a heap either, which is absurd. Twenty-four centuries of work have not produced the number, and this paper takes that failure as information rather than as a debt still owed. The number was never there to be found, because heap is not a predicate of grain count.

The move this paper makes is to relocate the puzzle. What looks like a question about a threshold

is better read as a question about two levels that have been run together: a latent state that changes with every grain, and a coarse classification of that state that does not. A grain is not nothing. It changes the pile, the configuration, the mass, the latent variable, at every step. It is nothing only to a particular readout, a coarse device, perceptual or linguistic or practical, that maps many distinct states to the same label. The heap paradox is what happens when the invariance of the label is mistaken for the invariance of the state.

Stated that way the heap stops being unique. The same shape appears wherever small increments accumulate in a hidden variable that a coarse channel only occasionally reclassifies: a neuron summing subthreshold inputs until it spikes, a material accumulating microdamage until it fails, a population of cells accumulating couplings until it becomes an individual, a learner accumulating small updates until it is skilled. The rest of this paper treats these as one architecture, says what is shared and what is not, and puts numbers on the shared part.

2. Two Containers

The architecture is a pair of equations. A latent state evolves under small inputs,

$$x_{t+1} = F(x_t, u_t),$$

and a readout classifies it,

$$y_t = R(x_t),$$

where R is many-to-one, a projection that discards most of what distinguishes states. The sorites situation is the regime in which the state moves but the readout does not,

$$x_{t+1} \neq x_t \quad \text{while} \quad R(x_{t+1}) = R(x_t),$$

sustained across many steps, until enough has accumulated that $R(x_{t+k}) \neq R(x_t)$ and the category finally flips. Two containers, one that fills continuously and one that reports in coarse steps, and the confusion between them is the whole paradox.

This is why the tempting premise, that one grain makes no difference, is both true and false depending on the container it is read in. In the readout container it is true: $R(S+1) = R(S)$, adding a grain does not change the label. In the state container it is false: $S+1 \neq S$, the pile is different. The paradox smuggles a truth about the projection, $R(S+1) = R(S)$, into a claim about the state, $S+1 = S$, and once that substitution is made the march to a million non-heaps is unstoppable. Keep the containers apart and the march never starts. The sections that follow show that this is not a verbal trick but a structure with a measurable signature, first in a relation, then in a neuron, a reservoir, and a tipping point.

3. Indistinguishability Is Not Identity

The sharpest form of the error lives in a relation. Let two states count as indistinguishable when their heap-relevant magnitudes differ by less than a just-noticeable amount, a tolerance. This relation is reflexive and symmetric, and locally it is nearly total: in the simulation, on a chain of a thousand states one unit apart under a tolerance of three units, every adjacent pair is indistinguishable, a rate of 1.0. The trap is transitivity. Indistinguishability does not have it. State 1 is indistinguishable from state 2 and 2 from 3, but 1 and 1000 are not indistinguishable at all, differing by 333 times the tolerance.

The paradox is now locatable. It is the treatment of a non-transitive tolerance relation as though it were a transitive equivalence relation. If you take the transitive closure of indistinguishability, which is what the sorites premise silently does when it chains one grain to the next, the whole chain collapses into a single class: all thousand states, from the plainly-not-a-heap to the plainly-heap, become one, because they are connected by a path of locally tolerant steps. The simulation reports that closure as 1 class against 1000 genuinely distinct states, and it finds that three quarters of the chained triples, the triples where the first is tolerant of the second and the second of the third, already fail to be transitive. The absurd conclusion is not a discovery about sand. It is what a non-transitive relation does when you force transitivity on it.

This diagnosis is close to the strict-tolerant treatment of vague predicates, which keeps the tolerant step, if x is a heap and y is similar enough then y is tolerantly a heap, while denying that the tolerant steps chain into a classical proof (Cobreros et al., 2012). It is also what careful psychophysics finds when it tests the premise directly rather than assuming it: perceptual matching is probabilistic, so a sequence of pairwise-matched stimuli can have endpoints that are clearly discriminable, exactly the soritical structure realized in the laboratory rather than the armchair (Dzhafarov and Perry, 2014). The older families of theory each capture a piece of this, epistemicism a sharp unknowable line, supervaluationism many admissible lines, fuzzy logic a graded truth, but they underweight the mechanism, which is that a coarse channel imposes a tolerance relation whose closure is not to be trusted (Williamson, 1994; Fine, 1975; Raffman, 2014).

4. Why No Grain Is Noticed

The readout is coarse for a reason, and the reason is in the perceiver. Human number sense is not uniform across magnitude. Very small collections, up to about four, are apprehended at once, each item individuated, through the system called subitization or parallel individuation; larger collections are handed to an approximate number system that represents them as one noisy magnitude, with discriminability governed by ratio rather than difference (Hyde, 2011; Feigenson, Dehaene, and Spelke, 2004). A parieto-frontal population code for numerosity underlies the approximate side, so that beyond the subitizing range number is registered as an estimate, not a count (Nieder and Miller, 2004). For a pile of any size worth calling a heap, no one is counting. The relevant channel is not numerosity at all but a coarse read of mass, spread, and height.

Perception then sharpens what magnitude leaves smooth. Categorical perception warps a con-

tinuous stimulus dimension into discrete regions, compressing differences within a category and expanding them across a boundary, so that two stimuli one step apart look the same inside a category and different across it, which is the readout doing its coarsening in the perceptual system itself (Harnad, 1987; Bonnasse-Gahot and Nadal, 2022). The heap inherits this. It begins where the perceptual task stops being enumeration and becomes material classification, and the boundary it seems to have is a feature of the category geometry the perceiver imposes, not of the sand. That the boundary is soft rather than sharp is the subject of the next mechanism.

5. The Subthreshold Neuron

The cleanest physical realization of the two containers is a neuron, and it is worth dwelling on because it turns the philosophical picture into biophysics. A neuron sums the small postsynaptic potentials arriving on its membrane; each is typically far below the threshold for firing, and individually none produces a spike. But they accumulate, through temporal and spatial summation, in the membrane potential, a genuine latent state with its own continuous dynamics, until the summed depolarization reaches threshold and an action potential fires (Hodgkin and Huxley, 1952; Burkitt, 2006). The spike is the readout. It is binary, and for a long stretch it reads zero while the membrane potential underneath it climbs.

The simulation makes the information structure explicit with a population of leaky integrators driven by a stream of small inputs. Two facts come out together. The thresholded readout is silent for the first 37 steps, its value never changing, its variance zero over the window, so that across a stretch in which the accumulated input grows 37-fold the readout says nothing has happened and carries no information about how much has arrived. And yet the latent population state holds that same accumulated count in full: a linear decoder reconstructs it at an R-squared of 1.0, exactly, as it must, since the integrator's state is a deterministic linear function of the accumulated input and a linear decoder inverts it without residual. The precision is guaranteed by construction and is not the point; the point is qualitative, that the information the readout discards is not lost. It is sitting in the latent state, fully recoverable, waiting for either a finer decoder or the moment the coarse one finally flips. A subthreshold input is not nothing. It is nothing only to the spike that has not yet fired.

The softness of the boundary is the last piece here, and it dissolves the last of the paradox. Model the category probability as a soft threshold on a heap-relevant magnitude, a sigmoid. Then two things that seemed to conflict hold at once. The largest change in heap-probability produced by any single grain is tiny, 0.0062 in the simulation, so locally no grain is decisive, just as the premise insists. And the total change in heap-probability across the range is 1.0, the difference between definitely-not and definitely-yes, so globally the accumulation is complete. A small per-step change and a large total change are not in tension; the total is the sum of many small changes, and the ratio of the two, about 160 here, is just the number of steps over which the transition is spread. Widen the tolerance and the borderline zone widens with it, from 175 grains to 703 in the two cases run, which is contextualism given a knob rather than a slogan.

6. Reservoirs: Where History Hides

The neuron shows a latent state holding what a readout discards; reservoir computing gives the general model class for it. A reservoir is a high-dimensional recurrent system with fixed internal connections, driven by an input stream, from which a trained linear readout extracts whatever is wanted; its virtue is that the recurrent state is a rich, fading memory of input history, so that a simple readout on top of it can classify temporal patterns a direct readout of the input never could (Jaeger and Haas, 2004; Lukoševičius and Jaeger, 2009). The spiking version, the liquid state machine, makes the tie to the neuron explicit, since it is precisely a recurrent network of integrate-and-fire units serving as the reservoir (Maass, Natschläger, and Markram, 2002). The population of integrators in the previous section was a small reservoir, and its decodability was the reservoir property in miniature: history accumulates in the latent state before it is legible at any coarse readout.

The point to hold onto, and the one the paper is careful not to overstate, is what reservoir computing is being used for. It is not a solution to the sorites and it is not a claim that heaps or brains or cells are literally reservoir computers. It is a model class for one thing, the accumulation of history in a hidden high-dimensional state that a coarse channel reads only in part, and it earns its place here by making that accumulation both precise and, as the next section shows, predictable in advance of the readout it will eventually move.

7. When Accumulation Tips

Some transitions are not soft. When the latent state is bistable, slow accumulation does not raise a probability gently; it drifts a control parameter until a stable state collides with an unstable one and vanishes, and the system jumps to a distant attractor. This is the form the emergence of a genuinely new category takes, a spike, a phase, an individual, and it is where the coarse readout is most misleading, because right up to the jump the readout is flat and afterward it is simply elsewhere.

The simulation drives a bistable latent state slowly across such a fold. The binary readout, one above the tipping point and zero below, gives no warning whatever: it holds at zero through the entire approach and flips only at the jump, at step 4483, with a warning count of zero. But the latent state is not silent, and its approach to the tipping point has a signature that the readout cannot show. As the fold nears, the state recovers ever more slowly from fluctuations, so that its variance rises, roughly 15-fold before the jump, and its lag-one autocorrelation climbs from 0.9 toward 0.99, the generic early-warning signature of a critical transition (Scheffer et al., 2009). A reservoir trained only on pre-transition data raises an alarm 198 steps before the readout moves, and the simple variance measure crosses its own threshold far earlier still. The transition is invisible to the coarse readout and legible in the latent state long before it happens, which is the whole thesis stated dynamically: everything changes in the latent container while nothing changes in the readout, until it does, and the change was readable all along.

8. Emergent Individuals

The most consequential version of this architecture is the appearance of a new individual. No single cell-cell adhesion, no one signaling event, makes an organism, just as no grain makes a heap; but repeated local coupling, adhesion, communication, division of labor, the management of internal conflict, can carry a collection of cells across a transition into a new unit of organization that reproduces and is selected as a whole (Szathmáry and Maynard Smith, 1995; Brunet and King, 2017). Complex multicellularity arose independently several times, by these routes, which is what one expects if it is a transition that accumulation can reach rather than a singular accident (Knoll, 2011). The readout, individual or aggregate, flips once; the latent coupling that brings the group there accumulates for a very long time before it does.

There is a formal frame for the local-to-global step this involves, and it should be introduced only now, once the pattern is visible, so that it reads as diagnosis rather than decoration. What counts as a heap in a laboratory, on a beach, in a legal dispute, and at ant scale need not agree, and the question of whether locally adequate classifications determine one global classification is the question a sheaf answers. The sheaf-theoretic treatment of contextuality makes this precise in physics, where locally consistent data can provably fail to admit any global section, and that is the honest shape of vagueness too: local readouts, each fine in its own context, need not glue into a single context-free predicate (Abramsky and Brandenburger, 2011; Mac Lane and Moerdijk, 1992). The paper takes the sheaf idea just this far and no further. It names the obstruction; it does not pretend to compute the cohomology of heaps.

9. One Design, Not Five Metaphors

A synthesis across this many fields earns trust only by being explicit about the join, and the honest statement is that these systems share an architecture, not an ontology. A neuron is not a heap, a heap is not an organism, and the transitions differ in kind: a spike is a physical event with a biophysical threshold, a heap is a category judgment with a perceptual and pragmatic one, a multicellular individual is a unit of selection. What they share is the two-container structure, latent accumulation under coarse-grained readout, and the specific pathologies it produces, the non-transitive tolerance relation, the silent readout over a moving state, the soft or sharp transition that the readout announces late. To claim more than shared architecture would be to melt real differences into a metaphor, which is the failure mode this framing is most exposed to and must most guard against.

Three further limits bound the claim. It is not fuzzy logic wearing new clothes: fuzzy logic assigns graded truth to a static predicate, whereas the object here is a dynamics, a latent state and a readout and the accumulation between them, of which a degree of truth is at most a snapshot. It is not mere thresholding: some cases do have a fixed physical threshold, but the sorites cases have context-sensitive readouts and a tolerance relation whose non-transitivity, not any single cutoff, is the source of the paradox. And its novelty is of a modest and stateable kind, not a new physics or a new neuroscience but a cross-domain formalization, the claim that logical tolerance, perceptual

indistinguishability, synaptic summation, reservoir dynamics, and the transitions of individuality instantiate one abstract design, latent accumulation under coarse-grained readout, and that seeing them as one clarifies each.

What is left when the limits are subtracted is a clean reversal of the original question. The heap was never hidden at grain N , and asking for N mistook a coarse readout for the world. Nothing changes locally because the readout is coarse, and everything changes globally because the accumulation is real; and the accumulation, the simulation shows, is not merely real but legible, decodable from the latent state at an R-squared of 1.0, and forewarned by that state hundreds of steps before the readout admits that anything has happened at all. A grain is not nothing. It is nothing only to the reader who is not counting.

References

- Abramsky, S., and Brandenburger, A. (2011). The sheaf-theoretic structure of non-locality and contextuality. *New Journal of Physics*, 13, 113036.
- Bonnasse-Gahot, L., and Nadal, J.-P. (2022). Categorical perception: a groundwork for deep learning. *Neural Computation*, 34(2), 437–475.
- Brunet, T., and King, N. (2017). The origin of animal multicellularity and cell differentiation. *Developmental Cell*, 43(2), 124–140.
- Burkitt, A. N. (2006). A review of the integrate-and-fire neuron model: I. Homogeneous synaptic input. *Biological Cybernetics*, 95(1), 1–19.
- Cobrerros, P., Égré, P., Ripley, D., and van Rooij, R. (2012). Tolerant, classical, strict. *Journal of Philosophical Logic*, 41(2), 347–385.
- Dzhafarov, E. N., and Perry, L. (2014). Perceptual matching and sorites: experimental study of an ancient Greek paradox. *Attention, Perception, and Psychophysics*, 76(8), 2441–2464.
- Feigenson, L., Dehaene, S., and Spelke, E. (2004). Core systems of number. *Trends in Cognitive Sciences*, 8(7), 307–314.
- Fine, K. (1975). Vagueness, truth and logic. *Synthese*, 30(3–4), 265–300.
- Harnad, S. (Ed.). (1987). *Categorical Perception: The Groundwork of Cognition*. Cambridge University Press.
- Hodgkin, A. L., and Huxley, A. F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *Journal of Physiology*, 117(4), 500–544.
- Hyde, D. C. (2011). Two systems of non-symbolic numerical cognition. *Frontiers in Human Neuroscience*, 5, 150.
- Jaeger, H., and Haas, H. (2004). Harnessing nonlinearity: predicting chaotic systems and saving energy in wireless communication. *Science*, 304(5667), 78–80.

- Knoll, A. H. (2011). The multiple origins of complex multicellularity. *Annual Review of Earth and Planetary Sciences*, 39, 217–239.
- Lukoševičius, M., and Jaeger, H. (2009). Reservoir computing approaches to recurrent neural network training. *Computer Science Review*, 3(3), 127–149.
- Maass, W., Natschläger, T., and Markram, H. (2002). Real-time computing without stable states: a new framework for neural computation based on perturbations. *Neural Computation*, 14(11), 2531–2560.
- Mac Lane, S., and Moerdijk, I. (1992). *Sheaves in Geometry and Logic: A First Introduction to Topos Theory*. Springer.
- Nieder, A., and Miller, E. K. (2004). A parieto-frontal network for visual numerical information in the monkey. *Proceedings of the National Academy of Sciences*, 101(19), 7457–7462.
- Raffman, D. (2014). *Unruly Words: A Study of Vague Language*. Oxford University Press.
- Scheffer, M., Bascompte, J., Brock, W. A., Brovkin, V., Carpenter, S. R., Dakos, V., Held, H., van Nes, E. H., Rietkerk, M., and Sugihara, G. (2009). Early-warning signals for critical transitions. *Nature*, 461(7260), 53–59.
- Szathmáry, E., and Maynard Smith, J. (1995). The major evolutionary transitions. *Nature*, 374(6519), 227–232.
- Williamson, T. (1994). *Vagueness*. Routledge.