

# Complexity & Emergence

No bird is in charge. No bird can see the shape. And yet the flock turns as one.

At dusk over Rome, a few thousand starlings pour into the sky and begin to *breathe*: folding, splitting, tightening into a dark fist and unspooling again into a ribbon, all in eerie silence and without a single collision. It looks choreographed. It looks, frankly, like something is conducting it. For centuries that was the going theory: a leader bird, or some shared avian telepathy. The truth turned out to be stranger and far more profound. **There is no conductor.** Each bird watches a handful of its closest neighbors and obeys a few local rules. The breathing shape exists nowhere in any individual bird's head. It is *emergent*: a pattern that lives only at the level of the whole.

That gap between the simplicity of the parts and the sophistication of the pattern is today's descent. It is one of the deepest and least settled ideas in science, and it hides a question that will stalk us for the rest of the course: when something new appears at a higher level, is it *genuinely* new, or just the old physics seen from far away?

## WHERE WE ARE

This closes Block I's epistemic toolkit and gives a name to a thread that has been flickering since the start. *Emergence* got light first touches on **Day 1**, where knowing was already a system-level property. Today it gets a proper definition. We lean hard on **Day 7** because the day's biggest controversy turns on whether a shiny new complexity measure is more than *compression* in a lab coat. And we reuse **Day 5** because the hardest version of emergence is a claim about *downward causation*: the whole reaching back down to push its parts around.

## THE PHENOMENON

### Six or seven neighbors, and nothing else

Start with what we actually *know*, in the careful Day 1 sense. In the mid-2000s, the STARFLAG project put stereoscopic cameras on a rooftop museum in Rome, photographed real starling flocks of up to **2,700 birds**, and reconstructed the 3-D position of every single one. Then they asked a deceptively simple question: when a bird adjusts to its neighbors, which neighbors?

The intuitive answer is *metric*: a bird reacts to everyone within, say, three meters.

The data said no. Birds track a fixed *number* of neighbors: **six to seven**, regardless of how near or far those neighbors happen to be. Pack the flock tighter and each bird still watches the same six or seven; spread it out and it watches the same six or seven, now farther away. The interaction is *topological*, not metric: it counts neighbors, not meters.

*“Each bird interacts on average with a fixed number of neighbors (six to seven), rather than with all neighbors within a fixed metric distance.”* Ballerini et al., PNAS, 2008.

Why does this matter? Because a topological rule is exactly what you would want if the goal were to hold a group together under attack. When a falcon stabs into the flock and local density explodes, a fixed-radius bird would suddenly track dozens of neighbors and lose the thread. A “watch my nearest seven” bird keeps its cohesion no matter how density warps. The rule is robust. And it is purely **local**. No bird knows the flock’s shape. No bird is solving for the global pattern. The pattern is a side effect of thousands of birds each minding a tiny neighborhood.

## Print model

The live panel varies three local rules. With all rules off, motion is gas-like. With separation, alignment, and cohesion active, order appears as a high alignment score.

RULE	LOCAL INSTRUCTION	SYSTEM EFFECT
Separation	Steer away from neighbors who get too close.	Prevents collisions and keeps individuals distinct.
Alignment	Match the average heading of nearby birds.	Lets the group turn as one.
Cohesion	Drift toward the local center of mass.	Keeps the flock from evaporating.
Nearest-neighbor count	Track about six or seven neighbors, not a radius.	Preserves cohesion when density changes.

### THE MODEL

## Three rules, written in 1986

The simulation is essentially *Boids*, built by Craig Reynolds in 1986 and presented at SIGGRAPH in 1987. Reynolds wanted realistic flocking for animation and discovered he needed three local rules: **separation** (do not crowd), **alignment** (steer with neighbors), and **cohesion** (stay with the group). From those three, the repertoire of flocking falls out: wheeling, splitting around obstacles, fluidly merging again. *Boids* has since flocked movie bats and countless digital armies. It is the

canonical demonstration that **complex global behavior need not have a complex global cause.**

And Boids is one specimen of the same trick: simple parts, surprising wholes.

- **Ant colonies** find short paths to food with no surveyor and no map. Each ant lays a chemical trail and follows others' trails; shorter routes get reinforced faster. Biologists call this *stigmergy*: coordination through traces left in the environment rather than direct command.
- **A pot of water** can spontaneously organize itself. Heat it from below and, past a threshold, smooth liquid breaks into ordered convection cells: structure sustained by energy flow, a preview of the physics-of-life arc on Days 83-85.
- **Your brain** is reading this sentence with no homunculus inside doing the reading. Meaning emerges from neurons that individually understand nothing, the steep version of the puzzle we return to on Days 123-126.

WHEN THE PARTS COMPUTE

## The maddening case of the slime mold

If the murmuration is emergence as beauty, slime mold is emergence as intelligence arising in something with no brain. *Physarum polycephalum* is a single giant cell, a sprawling yellow blob with protoplasm and an appetite for oats.

In 2010, a team led by Toshiyuki Nakagaki and Atsushi Tero laid out oat flakes in the geographic pattern of **Tokyo and 35 surrounding towns**, then dropped a slime mold on Tokyo. The blob spread out to engulf the food, then pruned itself. Redundant tubes withered; efficient ones thickened. Within about a day it had sculpted a network similar in total length, efficiency, and fault tolerance to the actual **Tokyo rail system**.

*The model “quantitatively mimics phenomena that can be neither captured nor quantified by verbal description alone.”* Wolfgang

Marwan, *Science Perspective*, 2010.

The slime mold is not clever. It is running a simple local rule: reinforce tubes where protoplasm flows well, abandon the rest. Applied everywhere at once, that rule *computes* a near-optimal network. The colony solves a problem no ant understands. The blob designs a railway it cannot conceive of. **The computation lives in the interactions, not in the components.**

## A useful warning

It is tempting to narrate all this as the parts being secretly smart. Resist it. The payoff of emergence is that the parts can be as simple as they look. The sophistication is manufactured by organization: by the pattern of who affects whom. That is also why emergence is seductive to overclaim. Once you accept that simple rules can yield astonishing wholes, it is a short step to declaring that your favorite simple rule explains everything.

### THE MODEL, SHARPENED

## The phrase that launched a field

For a long time, emergence was a word physicists said with suspicion: vague, faintly mystical, the sort of thing you invoked when you lacked equations. Then in 1972, the condensed-matter physicist **Philip W. Anderson** published a short essay in *Science* with a title that became a slogan: “**More Is Different.**”

Anderson’s target was *reductionism*: the assumption that once you know the fundamental laws, everything else is “just” applied physics. He did not deny reductionism. He granted it, then drove a wedge through its smug corollary.

*“The ability to reduce everything to simple fundamental laws does not imply the ability to start from those laws and reconstruct the universe... at each level of complexity entirely new properties appear.”* P. W. Anderson, *Science*, 1972.

Knowing how two water molecules interact does not hand you wetness, fluidity, convection cells, or the exact way a tsunami breaks. Those are real, lawful properties of collectives. New levels of organization have their own rules, concepts, and sciences. Chemistry is not merely physics in inconvenient notation; biology is not merely chemistry with extra vocabulary. Each level is emergent from the one below and yet, in a meaningful sense, autonomous from it.

### THE DEBATE

## Weak vs. strong: the fault line

Everyone agrees murmurations and slime molds are emergent. The fight is over what that word is allowed to mean.

## Weak emergence: the scientist's version

The philosopher **Mark Bedau** gave the careful account most working scientists can use. A property is *weakly emergent* if you cannot predict it from the rules by a clever shortcut: the only way to find out what the system does is to **run it** or simulate it. The murmuration is weakly emergent. There is no simple formula that takes the three Boids rules and spits out tomorrow's flock shape; you have to let it fly. Crucially, weak emergence is not spooky. It is fully grounded in the parts and their interactions.

## Strong emergence: the radical version

A property is *strongly emergent* if facts about the whole are *in principle* not deducible from complete knowledge of the parts, not even by a godlike simulator with unlimited compute. The high level would possess genuinely new causal powers, including downward causation. **David Chalmers** argues that there may be exactly one candidate for strong emergence in nature: **consciousness**. Everything else looks weakly emergent: surprising, irreducible in practice, but not irreducible in principle.

Why not embrace strong emergence everywhere? Because it collides with **Jaegwon Kim's causal exclusion problem**. Suppose a high-level state causes a physical effect. But that effect already has a complete physical cause at the level of parts. Either the high-level cause is redundant, the effect is overdetermined, or the claim violates the physics. Defenders have answers, but no knockout. This is why the strong/weak line remains live under biology, mind, and social systems.

● WORKABLE SCIENCE

● CONTESTED

## Print model

A two-rule cellular automaton can support gliders, oscillators, and universal computation. The important lesson is not the animation but the absence of a shortcut: many starting patterns must be run to reveal their fate.

CELL STATE	NEIGHBOR COUNT	NEXT STATE
Alive	2 or 3 live neighbors	Survives
Alive	0, 1, 4, 5, 6, 7, or 8 live neighbors	Dies
Dead	Exactly 3 live neighbors	Becomes alive
Dead	Any other live-neighbor count	Stays dead

Sit with what the Game of Life proves. Two rules, applied to a grid, are enough to build logic gates, memory, clocks, and a universal computer. The “computer” is

nowhere in the rules; it is a pattern *in* the cells. This is the cleanest illustration of weak emergence: you can know the rules perfectly and still have no way, short of running them, to know whether a starting pattern will fizzle, freeze, repeat, or compute. **Simple law, unboundedly complex behavior, no shortcut.**

A NECESSARY DETOUR

## Nobody can agree how to measure “complex”

Here is an embarrassing secret of complexity science: there is **no single agreed definition of complexity**. Seth Lloyd once compiled a famous list of dozens of proposed measures, and they do not agree with each other. This is not mere sloppiness. “Complex” is doing several jobs at once: how hard is it to describe, how hard is it to build, and how organized is it?

That last family hides the key idea. **The most complex things are not the most random ones.** Consider a crystal, a living cell, and a canister of gas. By a pure “hard to describe” measure, the gas wins because perfect randomness needs the longest description. But a gas is not more complex than a cell. A crystal is too ordered; a gas is too disordered. Complexity lives in the structured, surprising, neither-frozen-nor-random middle.

### Diagram · the complexity-entropy hump

Disorder runs left to right: perfect order on the left, perfect randomness on the right. Effective complexity is the curve. It collapses at both extremes and peaks in the organized middle where life, language, and minds live.

Murray Gell-Mann’s effective complexity tries to measure the regularities of a thing while throwing away the purely random part. Keep this hump in mind: the Assembly Theory fight is partly about whether a new metric has captured this middle, or smuggled “hard to describe” back in under a new name.

THE FRONTIER · 2026

## Three live edges, and the hype filter at full stretch

Emergence is where some of the boldest and most overhyped claims in current science live. Today’s frontier contains a genuine, bruising fight in the peer-reviewed literature.

EDGE 01 · THE BIG ONE ● ESTABLISHED ● PROMISING ● CONTESTED

## Assembly Theory: a complexity ruler, or a theory of everything?

In 2023, a group led by **Leroy Cronin** and **Sara Walker** published a *Nature* paper titled “*Assembly theory explains and quantifies selection and evolution.*” The core idea is clever. Take an object, such as a molecule, and ask: what is the minimum number of steps to build it from basic parts, if you can reuse any piece already made? That number is the *assembly index*. Pair it with copy number, and Assembly Theory claims to detect the fingerprint of a process that selects and reproduces structures. The killer application is astrobiology: a molecule with assembly index above roughly **15** is argued to be extremely unlikely to form abiotically.

What is solid here is genuinely solid. The assembly index is a real, computable, experimentally measurable quantity. A 2024 *ACS Central Science* paper showed it can be read from mass spectrometry, NMR, and infrared spectroscopy, and the numbers agree. It separates many biological from abiotic samples in the lab.

The trouble is everything stacked on top. The grander claims that Assembly Theory explains selection, evolution, time, and the physics-biology bridge drew fierce backlash. A 2024 *PLOS Complex Systems* analysis by Abrahao, Zenil, and colleagues argued that the assembly index is mathematically an approximation to standard data compression and cannot outperform Shannon entropy for the job it claims. Other critiques attacked the biosignature classification and the evolutionary claims.

Cronin and Walker have pushed back. Their defense is the crux: compression measures statistical redundancy in strings, they argue, while assembly index is about physical causal history: the minimal path by which a real object could be built in the world, tied to how many copies exist. Whether that distinction is deep or cosmetic is exactly what remains unsettled. **Verdict: the ruler is real; the theory-of-everything claims are not established.**

### ASSEMBLY THEORY CLAIM

### HYPE-FILTER STATUS

The assembly index is a measurable molecular-complexity metric.

● ESTABLISHED

Index above about 15 can serve as a biosignature in unknown samples.

● PROMISING

It explains and quantifies selection and evolution in a way biology recognizes.

● CONTESTED

It is fundamentally novel versus compression, Kolmogorov complexity, and Shannon entropy.

● CONTESTED

## Self-organized criticality: the sandpile that tunes itself

In 1987, Per Bak, Chao Tang, and Kurt Wiesenfeld proposed a mesmerizing idea: some systems drive themselves to the knife-edge of a phase transition with no external tuning. Their model was a **sandpile**. Drop grains one at a time and the pile steepens until avalanches of all sizes keep it near a critical slope. This *self-organized criticality* became a candidate explanation for earthquakes, forest fires, extinctions, market crashes, and  $1/f$  noise.

SOC is real: sandpile-type models do self-organize to criticality. But the grand claim that SOC is the universal engine of complexity has not held up. The most hyped application, the “critical brain” hypothesis, remains contested. A beautiful mechanism was oversold as a theory of everything.

EDGE 03 ● PROMISING ● CONTESTED

## Can you put a number on emergence itself?

The newest frontier tries to move weak versus strong emergence from armchair debate into mathematics. **Erik Hoel** and colleagues developed *causal emergence*, using effective information to show that a macro-scale model can sometimes carry more causal information than the micro-scale model it is built from. **Fernando Rosas**, **Pedro Mediano**, and colleagues use information decomposition to test whether a macro feature predicts a system’s future better than any individual part can.

This is careful, fertile work, so it earns “promising.” The caveat is interpretation. Critics argue these measures may capture epistemic emergence, about our best models and descriptions, rather than the ontological strong emergence philosophers mean. The math is landing; what it measures is still under debate.

OPEN QUESTIONS

## What’s genuinely unsettled

- **Is strong emergence real anywhere?** Or is everything weakly emergent, surprising but reducible, with Kim’s exclusion argument quietly winning?
- **Is downward causation a real force or a figure of speech?** Can a whole genuinely push its parts around in a way the micro-laws do not already determine?
- **Is there one principled measure of complexity?** Or will we always have a drawer full of measures, each right for a different job?
- **Does Assembly Theory carve nature at a new joint?** Or is it useful compression and information theory wrapped in bigger claims than it can cash?
- **What about AI emergence?** When a large language model seems to gain a new skill at scale, is that genuine emergence or an artifact of the metric? Hold that

thought for Day 139.

## The day in three sentences

### Big idea

Simple parts following local rules can generate global order with no leader and no blueprint. “More is different” names the fact that higher levels can have real patterns and laws of their own.

### Best analogy

The murmuration that breathes though no bird sees its shape, and the slime mold that computes a railway it cannot understand.

### Live controversy

Assembly Theory is a real molecular-complexity ruler, but its claims to explain selection, evolution, and time remain contested.

emergence gets its proper definition · information returns as compression and assembly index · computation appears in Game of Life · energy flow hints toward life.

## Sources & further reading

1. Ballerini, M., Cabibbo, N., Candelier, R., et al. (2008). “Interaction ruling animal collective behavior depends on topological rather than metric distance.” *PNAS* 105(4): 1232-1237. doi:10.1073/pnas.0711437105.
2. Cavagna, A., Cimarelli, A., Giardina, I., et al. (2010). “Scale-free correlations in starling flocks.” *PNAS* 107(26): 11865-11870.
3. Reynolds, C. W. (1987). “Flocks, Herds, and Schools: A Distributed Behavioral Model.” *Computer Graphics* 21(4): 25-34.
4. Anderson, P. W. (1972). “More Is Different.” *Science* 177(4047): 393-396. doi:10.1126/science.177.4047.393.
5. Tero, A., Takagi, S., Saigusa, T., et al. (2010). “Rules for Biologically Inspired Adaptive Network Design.” *Science* 327(5964): 439-442. doi:10.1126/science.1177894.
6. Gardner, M. (1970). “Mathematical Games: The fantastic combinations of John Conway’s new solitaire game ‘life.’” *Scientific American* 223(4): 120-123.
7. Bedau, M. A. (1997). “Weak Emergence.” *Philosophical Perspectives* 11: 375-399.
8. Chalmers, D. J. (2006). “Strong and Weak Emergence.” In *The Re-Emergence of Emergence*, Oxford University Press.
9. Kim, J. (1999). “Making Sense of Emergence.” *Philosophical Studies* 95: 3-36.
10. Bak, P., Tang, C., and Wiesenfeld, K. (1987). “Self-organized criticality.” *Physical Review Letters* 59(4): 381-384. doi:10.1103/PhysRevLett.59.381.
11. Bonachela, J. A., et al. (2010). “Self-organization without conservation: are neuronal avalanches generically critical?” *J. Stat. Mech.* P02015. arXiv:1001.3256.

12. Sharma, A., Czege, D., Lachmann, M., Kempes, C. P., Walker, S. I., and Cronin, L. (2023). "Assembly theory explains and quantifies selection and evolution." *Nature* 622: 321-328. doi:10.1038/s41586-023-06600-9.
13. Marshall, S. M., Mathis, C., Carrick, E., et al. (2021). "Identifying molecules as biosignatures with assembly theory and mass spectrometry." *Nature Communications* 12: 3033.
14. Jirasek, M., Sharma, A., Bame, J. R., et al. (2024). "Investigating and Quantifying Molecular Complexity Using Assembly Theory and Spectroscopy." *ACS Central Science* 10(5): 1054-1064.
15. Abrahao, F. S., Hernandez-Orozco, S., Kiani, N. A., Tegner, J., and Zenil, H. (2024). "Assembly Theory is an approximation to algorithmic complexity based on LZ compression that does not explain selection or evolution." *PLOS Complex Systems* 1(1): e0000014.
16. Uthamacumar, A., Abrahao, F. S., Kiani, N. A., and Zenil, H. (2024). "On the salient limitations of the methods of assembly theory and their classification of molecular biosignatures." *npj Systems Biology and Applications* 10: 82.
17. Jaeger, J. (2024). "Assembly Theory: What It Does and What It Does Not Do." *Journal of Molecular Evolution* 92: 87-92.
18. Walker, S. I., Cronin, L., et al. (2024). "Reply to 'Experimental measurement of assembly indices are required to determine the threshold for life.'" *Journal of the Royal Society Interface* 21: 20240367.
19. Gell-Mann, M., and Lloyd, S. (1996). "Information measures, effective complexity, and total information." *Complexity* 2(1): 44-52.
20. Hoel, E. P., Albantakis, L., and Tononi, G. (2013). "Quantifying causal emergence shows that macro can beat micro." *PNAS* 110(49): 19790-19795.
21. Rosas, F. E., Mediano, P. A. M., Jensen, H. J., et al. (2020). "Reconciling emergences: An information-theoretic approach to identify causal emergence in multivariate data." *PLOS Computational Biology* 16(12): e1008289.
22. Dewhurst, J. (2021). "Causal emergence from effective information: Neither causal nor emergent?" *Thought* 10(3): 158-168.
23. Schaeffer, R., Miranda, B., and Koyejo, S. (2023). "Are Emergent Abilities of Large Language Models a Mirage?" *NeurIPS 2023*.