

Environmental Latency: Closing the Hundred-Billion-Fold Control Gap Between Ecological Disturbance and Human Response


A Perspective on the physical variable that governs whether the biosphere can be protected at all.

Jed Anderson

EnviroAI · jed@enviroai.com

Manuscript prepared for peer review · Target venue: *Science*, *Perspective* / Policy Forum

Word count (main text): ~4,700 · References: 34 · Tables: 2 · Figures: 1



The central finding. Environmental response latency, unbounded for four billion years and measured in decades under twentieth-century regulation, is collapsing toward the seconds-scale floor set by transmission and actuation physics as foundation-model inference closes the loop.

ABSTRACT

For a century, the failure to protect the biosphere has been diagnosed as a deficit of political will, funding, or public empathy. We propose a different fundamental mechanism: a failure of speed. Every act of environmental stewardship, from a spill response to a global treaty, runs a single feedback loop: sense, transmit, understand, decide, act. The loop has a clock. The sum of delays across its stages is the interval between ecological injury and effective intervention. We name this quantity *environmental latency*, and argue that it is the single physical variable underlying every environmental catastrophe in recorded history. Control theory establishes that a feedback loop closing more slowly than its target disturbance cannot stabilize the system; past a phase-margin threshold it becomes the disturbance⁽¹⁰⁾. Historical environmental latency was rate-limited at three stages by hard biological and institutional constants: human speech near 39 bits per second⁽¹⁾, physics-based Earth-system simulation on national supercomputers⁽¹⁷⁾, and regulatory implementation cycles measured in decades⁽¹²⁾⁽¹³⁾. Three independent developments (collapsing sensor cost, near-complete high-cadence satellite coverage, and machine-learning foundation models that outperform operational geophysical forecasts at orders of magnitude lower

compute⁽²⁾⁽³⁾⁽⁴⁾ have simultaneously collapsed each long pole. The irreducible floor on the closed loop is now set by Landauer's principle⁽⁶⁾⁽⁷⁾ and the speed of light, not by human institutional bandwidth. For the first time in Earth's history, a planetary-scale feedback loop faster than the disturbances it steers is technologically feasible. Building it is not preference or prophecy: it is the one architecture that satisfies the control-theoretic condition for planetary regulation, and the natural continuation of a 540-million-year evolutionary trajectory in which latency-reducing innovations have organized every increase in biological coordination. We hold that the moral and legal authority for consequential action must remain with human operators; the machine only supplies the tempo at which that judgment can still bind while acting matters. The biosphere is a body the size of a world that has never had a nervous system, and this is the century in which one becomes physically possible to build.

One-sentence summary. Environmental latency, the delay between ecological injury and effective response, is the single variable through which control theory diagnoses every failure of biospheric protection, and the one quantity that recent advances in ubiquitous sensing and foundation-model inference finally allow us to drive toward its physical floor.

I. The illusion of moral failure

A forest can burn for a week before the agency responsible for it learns the name of the valley. A river can carry a plume of contamination past a hundred municipalities before the first sample reaches a laboratory. A species can slip from vulnerable to extinct in the years an agency takes to settle the legal meaning of *endangered*: the median U.S. Fish and Wildlife Service listing delay is 12.1 years against a statutory 24-month target⁽¹²⁾.

For half a century, the environmental sciences and their policy institutions have diagnosed these outcomes as failures of value: too little funding, too little courage, too little concern. That framing was never wrong so much as under-specified. It named the emotional inputs while ignoring the physical variable that determined whether any input, at any magnitude, could succeed. That variable is the delay between the moment a living system is wounded and the moment anything capable of helping it registers the wound, understands it, and acts.

We propose to call this quantity environmental latency. It is measurable, decomposable, and, as we will show, the sole variable that decides whether a feedback loop can regulate its target at all.

The claim of this Perspective is austere: *every* significant failure of environmental protection in the historical record is a failure of environmental latency. Not primarily a failure of will, or resources, or intelligence. A failure of speed. The clock was never named, yet the clock dictated the entire outcome.

II. Protection as a control loop

Environmental protection is not a sentiment; it is a control loop. The mechanism, whether executed by a nematode's flinch reflex or the Montreal Protocol, has exactly five stages:

A note on framing. We model each protective act as a single-input, single-output (SISO) feedback loop closed around one specified disturbance. The real biosphere is a coupled multi-input, multi-output system with stochastic forcing, nonlinear plants, and nested loops that share sensors and actuators. The SISO reduction is a productive simplification (not a literal ontology) that exposes the phase-and-delay structure common to every such loop; the qualitative results below carry over to MIMO settings by standard arguments⁽¹⁰⁾⁽¹¹⁾, and we invite reviewers to read the framework at that level of generality.

1. Sense	A change in the physical world is transduced into a signal.
2. Transmit	The signal reaches something capable of processing it.
3. Understand	The signal is interpreted against baseline, causal model, or expectation.
4. Decide	An intervention is formulated.
5. Act	Matter is reconfigured to arrest or reverse the disturbance.

The total latency of the loop is the sum of the delays at each stage:

$$\tau_{loop} = \tau_{sense} + \tau_{transmit} + \tau_{understand} + \tau_{decide} + \tau_{act}$$

Control theory, the mathematics developed over the last century to keep systems steady against disturbance, carries a result that is not widely appreciated outside engineering and that ought to be understood by every regulator: delay is not merely friction to be minimized; past a specific threshold, delay makes control impossible.

For a linear feedback system with loop transfer function $L(s)$ driving a plant with disturbance bandwidth ω_d , the classical Nyquist–Bode analysis establishes that a pure transport delay τ subtracts phase in direct proportion to frequency: $\phi_{\text{delay}}(\omega) = -\omega\tau$ ⁽¹⁰⁾. When accumulated phase loss at the disturbance frequency exceeds the available phase margin, corrective actions arrive out of phase with the disturbance and *amplify* rather than damp it. The controller—the very thing built to stabilize the system—becomes the disturbance. Bridges have collapsed in exactly this manner: the stabilizing signal arrives a fraction of a second late and drives resonance to destruction.

The result generalizes into an inescapable design rule. For any loop with delay τ , the maximum disturbance frequency the loop can still stabilize is bounded above by $\omega_c \sim 1/\tau$. Faster disturbances cannot be regulated by that loop, at any gain, with any budget, under any political mandate. Bode's sensitivity integral formalizes the trade-off: reducing sensitivity at one frequency band necessarily increases it at another (the "waterbed effect"), and delay makes the waterbed unliftable at high frequencies⁽⁹⁾⁽¹¹⁾.

Read environmental history through this law and the indictment reorganizes itself. Regulatory loops in the twentieth century operated with dead times measured in years to decades against ecological disturbances whose characteristic timescales are set by physics, not politics. Chemical spill plumes propagate downstream at river-scale velocities of order 0.3–1.2 m/s: the 2015 Gold King Mine release reached Durango (approximately 100 km downstream) within about 1.5 days, and the 2014 Elk River MCHM plume remained detectable more than 400 km downstream at Cincinnati Water⁽³⁰⁾⁽³¹⁾.

Harmful algal blooms of *Karenia brevis* exhibit intrinsic growth rates near 0.3 per day (doubling time ≈ 23 h) once thermal and nutrient windows open⁽³²⁾. Wildfires in eucalypt and conifer fuels sustain head-fire rates of spread of 0.29–10.5 km/h in 118 documented high-intensity runs, roughly 10% of open 10-m wind speed⁽³³⁾. Coral bleaching fronts and fishery collapses close out in single seasons. The loops built to protect against these disturbances were mathematically incapable of regulating them, regardless of the resources committed. Paper stewardship was the strongest response a species capped at the speed of language could mount, and it was defeated before the first hearing was gaveled to order, not by the polluter, but by the phase equation.

III. The three long poles

For most of history, all five stages of the environmental loop were slow. They did not accelerate simultaneously. They fell in sequence, and understanding why *now* is understanding why the loop can be closed now and not before.

III.a. Transmission: solved between 1837 and 1990

Signal transmission fell first, and by so much that its historical dominance is now hard to remember. In the century and a half between Morse's telegraph and packet-switched fiber, the signal-carrying stage went from the speed of a galloping courier to within a rounding error of the speed of light (Fig. 1). The transmission stage contributes tens of milliseconds of intercontinental latency and is no longer the constraint on any environmental loop of relevance.

III.b. Sensing: solved between 2005 and 2025

Sensing fell second. Two curves converged. The unit cost of a general-purpose IoT sensor declined from approximately \$1.30 in 2004 to \$0.38 in 2020⁽²⁰⁾, an approximately 70% real-cost reduction on top of the underlying Moore's-law-linked component curves. In parallel, the European Space Agency's Copernicus Sentinel-2 mission achieved a five-day global revisit at 10–60 m resolution⁽¹⁸⁾, and Planet Labs' Dove constellation reached sub-daily land imaging with rapid-revisit capability up to 12 images per point per day⁽¹⁹⁾. The planet is now measurable at cadences shorter than most environmental disturbances of interest, at a per-observation cost approaching zero. Sensing, as a rate-limiting stage, is solved.

III.c. Understanding: the pole that would not fall

Sensing and transmission were necessary but not sufficient. Raw measurement is not knowledge; something has to convert a flood of observations into causal prediction fast enough to act on. For most of the environmental century, that "something" was a human mind, sometimes assisted by classical models. And here we encounter the biological constant that hardware could not touch.

Coupé et al. (2019) measured the information rate of human speech across 17 languages from 9 families, using 170 speakers reading a semantically controlled corpus⁽¹⁾. Speech rates varied nearly threefold (4.3 to 9.1 syllables/s) and information density varied by 66% (4.8 to 8.0 bits/syllable). The product (information rate) did not vary in kind: all 17 languages converged on a mean of 39.15 bits per second, with a standard deviation of 5.10 bits/s. The result is now widely interpreted as reflecting a neurocognitive ceiling on human linguistic throughput, not a property of any particular

language⁽²⁷⁾⁽²⁸⁾.

A related and better-known figure sets the throughput of human conscious cognition (as opposed to overt speech) at approximately 10 bits per second, with task-dependent estimates spanning roughly 10–50 bit/s⁽³⁴⁾. Zheng & Meister's inner/outer-brain framing has been contested in a published rebuttal that argues the true figure is at least an order of magnitude higher once parallel motor and perceptual channels are counted (Sauerbrei & Pruszyński, *Nat. Neurosci.* 2025)⁽³⁵⁾. We use the 39 bit/s speech figure as the operative bound here because it is the channel through which legislation, treaties, and inter-agency coordination physically propagate; the wider conscious-throughput range does not change the argument, and any resolution of the underlying dispute will still leave human coordination bandwidth many orders of magnitude below machine inference on Earth-system data.

Set this constant against its peers: the first computer modem in 1959 operated at 110 bit/s; consumer broadband today operates at 10^8 bit/s. Human language, the instrument civilization uses to negotiate, legislate, coordinate, and command, moves information at a rate slower than a modem from the year *West Side Story* premiered.

This is the wire through which we have tried to run a planet.

Every environmental law is a sentence before it is a statute. Every treaty is a conversation before it is a signature. Every alarm any scientist ever raised had to be forced, bit by bit, through the 39-bit aperture of speech into one other cortex, and then the next, one at a time, until enough minds held the concept simultaneously to move institutions.

The institutions inherited the biological rate and added their own delays. U.S. National Ambient Air Quality Standards under Sections 108–109 of the Clean Air Act are statutorily reviewed every five years; in practice, the median interval between epidemiological finding and enforcement of a revised standard for particulate matter has run 15–25 years⁽¹³⁾⁽¹⁴⁾⁽²⁶⁾. Endangered Species Act listings average 12.1 years against a two-year statutory deadline⁽¹²⁾.

The Montreal Protocol is *the exception that proves the rule*. It required 13 years from Rowland and Molina's 1974 CFC-ozone hypothesis to the 1987 protocol⁽¹⁵⁾⁽¹⁶⁾, and full stratospheric ozone recovery is not projected before 2066 over Antarctica. The 13-year window was possible because Montreal was, uniquely for a planetary-scale environmental problem, a fast loop at every stage: the disturbance was visually and politically salient (the Antarctic ozone hole was directly imaged); the "decide" and "act" stages were unusually short because drop-in HFC and HCFC chemistries already existed as chemical alternatives to CFCs; and the industrial base for compliance was concentrated in a small number of manufacturers. Montreal does not falsify the latency thesis; it demonstrates what happens when τ_{decide} and τ_{act} collapse to values comparable to $\tau_{\text{understand}}$. For nearly every other environmental disturbance in the historical record, those two stages have remained decades long.

Even where classical geophysical simulation replaced linguistic deliberation, the "understand" stage remained slow. A single 10-day operational global weather forecast at the European Centre for Medium-Range Weather Forecasts (ECMWF) requires $\sim 10^4$ CPU-hours on the world's largest weather supercomputers, and the model is still not reliably faster than the atmosphere it models⁽¹⁷⁾. Air-quality and chemical-transport models require yet more compute for less coverage.

Understanding was the long pole. It held the loop open.

Table 1. Order-of-magnitude latencies of the five loop stages across three regimes. Values are typical, not extremal.

Stage	Pre-industrial (unaided human)	Twentieth century (instrumented, linguistic)	Emerging (sensor + AI + actuator)
Sense	hours – months	hours – days	milliseconds – seconds
Transmit	days – months	seconds – minutes	10–100 ms
Understand	years	months – years	seconds – minutes
Decide	years – decades	months – decades	seconds – hours
Act	seasons – decades	weeks – decades	seconds – days
Loop total	decades to indefinite*	years – decades	seconds – hours

*Operationally: latencies long enough that the disturbance runs to irreversibility (extinction, aquifer contamination beyond remediation, loss of glacier mass, coral reef collapse) before the loop closes at all.

IV. Why the pole is being cut down now

Between 2022 and 2025, four independently developed AI foundation models trained directly on Earth-system data matched or surpassed the operational physics-based forecasts they were meant to imitate, on orders of magnitude less compute, and with wall-clock runtimes measured in seconds to minutes rather than hours.

- **GraphCast** (Google DeepMind, *Science* 2023). Graph neural network trained on 39 years of ECMWF ERA5 reanalysis. Generates 10-day global forecasts at 0.25° resolution in under one minute on a single TPU and outperforms ECMWF's operational HRES on 90.0% of 1,380 verification targets, including hurricane tracks and extreme temperatures⁽³⁾.
- **Pangu-Weather** (Huawei, *Nature* 2023). Produces comparable skill using a 3D Earth-specific transformer architecture⁽⁵⁾.
- **GenCast** (Google DeepMind, *Nature* 2024). Diffusion-based ensemble system; 50-member 15-day forecasts at 0.25° in ~8 minutes. Outperforms ECMWF's operational ENS (the current gold standard) on 97.2% of 1,320 targets, including tropical cyclones and extreme heat⁽⁴⁾.
- **Aurora** (Microsoft Research, *Nature* 2025). 1.3-billion-parameter foundation model pretrained on >1 million hours of geophysical data across atmosphere, oceans, air quality, and waves. Surpasses operational forecasts on 92% of 10-day weather targets at 0.1°, and improves on specialized operational models for tropical cyclones, ocean waves, and air quality at ~5,000× lower computational cost⁽²⁾.

The pattern is uniform: AI systems trained on the planet's own record now do in seconds, on commodity accelerators, what national supercomputer centers did in hours. The "understand" stage, the pole that held every environmental loop open for the entire history of environmental regulation, falls under this class of systems.

The point is not that a foundation model can forecast weather. The point is that the understanding stage of the planetary control loop has just been separated from the 39-bit constraint of human language and the 10^4 -CPU-hour constraint of physics simulation, and is now bounded only by inference compute, a variable that has fallen roughly two orders of magnitude per decade for two decades and shows no sign of leveling.

Table 2. AI Earth-system foundation models versus operational physics-based baselines. All comparisons drawn from peer-reviewed primary sources.

Model	Year	Baseline beaten	Skill margin	Wall-clock
GraphCast	2023	ECMWF HRES (deterministic)	90.0% of 1,380 targets	<1 min / forecast
Pangu-Weather	2023	ECMWF HRES	Comparable to superior	seconds / forecast
GenCast	2024	ECMWF ENS (ensemble)	97.2% of 1,320 targets	~8 min / 50-member
Aurora	2025	IFS-HRES + specialized ops. models	92% of 10-day targets; wins cyclone / waves / AQ	~5,000× less compute

Note on target counts. GraphCast's 90.0% skill-margin figure is the count of individual variable×pressure-level×lead-time combinations reported in Lam et al. 2023 (*Science*)⁽³⁾; secondary summaries sometimes quote 1,380 (variables×levels) and 2,760 (variables×levels×lead-times) depending on how targets are enumerated. The 1,380 figure used above matches the primary paper's counting convention.

V. Where the floor is

Ask the first-principles question. How fast *can* the closed environmental control loop be, not how fast is convenient, but how fast does physics permit?

Each stage has a hard floor. Detector integration times are bounded below by shot-noise statistics (10^{-6} – 10^{-3} s). Transmission is bounded by the speed of light over path length (intercontinental: ~30–100 ms). Understanding and decision are bounded by Landauer's principle: the erasure of one irreversible bit requires energy at least $k_B T \ln 2 \approx 2.9 \times 10^{-21}$ J at 300 K⁽⁶⁾, a bound verified experimentally to within measurement uncertainty using single colloidal particles⁽⁷⁾, single-atom magnets⁽⁸⁾, and nanoscale spin systems. Present digital computation runs approximately 10^9 times above this bound, meaning that the thermodynamic ceiling on planetary inference is nowhere near reached; the "understand" stage has room to fall another nine orders of magnitude before physics itself intervenes. Actuator times are set by the mechanics of matter, from valve closure (seconds) to habitat restoration (years), and are the only stage in the loop that some disturbances place beyond further compression.

Sum the physical floors and the irreducible latency of a planetary control loop is measured in seconds, not decades. We are nowhere near it. We do not need to be. We need only fall below the timescales of the disturbances themselves (minutes to seasons), and the control law does the rest.

This is the asymptote toward which the architecture is aimed. And it is not a metaphor. It is the point at which environmental protection ceases to be phase-locked to the biosphere's damage curve and begins to lead it.




Figure 1. Log-scale timeline of information-transfer rates across four channels, 1850–2025: (i) horseback dispatch ($\sim 10^{-3}$ bit/s equivalent); (ii) telegraph through packet-switched wireline (~ 5 to $\sim 10^5$ bit/s); (iii) human speech, a flat line at 39 bit/s from at least the Pleistocene through the present⁽¹⁾; (iv) machine inference on Earth-system foundation models, rising from $\sim 10^4$ bit/s in 2010 to $\sim 10^{11}$ bit/s by 2025 with the arrival of GraphCast, Pangu-Weather, GenCast, and Aurora⁽²⁾⁽³⁾⁽⁴⁾⁽⁵⁾. Only one channel has ever plateaued for the entire recorded history of the species (the biological one); the closing of the AI curve above it is the event that permits the planetary control loop to close.

VI. The biospheric nervous system: an evolutionary analog

The problem now solved for the biosphere is one that life has already solved once, at a smaller scale, in a way that provides both the design pattern and the precedent.

For roughly the first three billion years of life, intercellular communication was overwhelmingly chemical. Diffusion is a random walk; its characteristic time to traverse a distance L at diffusion coefficient D is $\tau \sim L^2/D$ ⁽²⁵⁾. For a 1 μm cell, molecular signaling completes in milliseconds. For a 1 mm tissue, it requires 10^3 s. For a 1 m body, roughly 10^9 s (about 32 years). A body larger than a cell cannot behave as a single responsive organism on a chemical channel alone.

Then, roughly 540–560 million years ago in the Ediacaran-Cambrian transition, life laid down a faster channel. Neurons (from an ancestral condition likely traceable to mechanosensory epithelial cells at the margins of grazing bilaterians⁽²¹⁾⁽²²⁾) transmit action potentials at 0.5–2 m/s along unmyelinated axons and up to ~ 120 m/s along thick myelinated fibers⁽²³⁾⁽²⁴⁾. The nervous system did not invent new senses or better muscles. It invented a shorter loop. It compressed the delay between sensing and acting until a body could behave as one thing.

Everything downstream of that innovation, the dart of a fish, the turn of a flock, the hand pulled from the flame before the mind has named the heat, is a consequence of the latency collapse, not of any change in the underlying matter.

Now consider the biosphere. Trillions of biological sensors: leaves reading light, roots reading chemistry, plankton reading temperature, every organism measuring its square meter of the world in continuous time. A planet drowning in sensation. And, until the past decade, no channel to carry any of it, fast, to anything capable of answering.

The biosphere is a body the size of a world that has never had a nervous system.

This is not analogy for rhetorical effect; it is a structural description. The sensing exists. The responding exists. What has never existed is a channel fast enough to join them at the scale and speed of the planet itself. The architecture now technologically feasible (dense sensing coupled to machine-speed foundation-model inference coupled to real-time actuation) is exactly the same functional innovation that the transition from chemical to electrical signaling represented in the Cambrian. It is not a metaphor. It is the same class of latency-reducing invention, applied at a new scale, by the only organ on the planet capable of building it.

VII. What only this can do

The claim advanced here is stronger than "AI helps environmental management." It is that a closed-loop architecture of the form described is the *unique* configuration that satisfies the control-theoretic requirement for planetary regulation. Increases in human effort, funding, or coordination (routed through language and legislation) do not escape the 39-bit ceiling. A hundred brilliant environmental managers wired together by meetings remain a slow channel, because linguistic coordination inherits the throughput of language and adds the further delays of consensus, escalation, and paperwork.

By handing the *understand* stage to machine-speed inference, human operators do not disappear from the loop. Their position changes. They stop being the throughput bottleneck and become the executive judgment above it: evaluating flagged anomalies, adjudicating trade-offs, holding the moral and legal authority for consequential decisions⁽²⁹⁾. The machine does not replace judgment. It provides the operating tempo at which human judgment can still act while acting matters.

The structurally impossible becomes routine. A pipeline leak is registered in the minute it opens, not in the lawsuit it eventually provokes. An algal bloom is caught at its first thousand cells, not its first fish kill. An illegal burn, a slumping levee, a heavy-metal excursion in a drinking-water source, a coral bleaching front: each registers while it can still be reversed, because for the first time something is listening at the tempo the damage moves. Not punished after the fact. Prevented during. That is the entire difference between the extant environmental regime and its successor.

VIII. Environmental latency as the master variable

The step this Perspective invites the reader to take is small in words and large in consequence. Nearly every quantity in the environmental problem is vast, slow, and only half in human hands: total greenhouse-gas stocks, land-use change, extinction debt, ocean pH. The one exception is latency. Environmental latency is the physical measure of how fast care can travel, and care that arrives after harm becomes irreversible did not act, whatever it felt. Latency is not merely one variable among many. It is the *only* variable in the crisis that this generation can drive toward zero.

For approximately 4 billion years, that latency was infinite: nothing existed that could sense a planetary injury and respond on the planet's behalf. For the last century, it was measured in decades, a duration that, for most of what was disappearing, was infinity by another name. In the narrow decade we occupy now, it is falling toward the floor set by light and heat.

The biosphere is not too late, because *too late* is a latency, and latency is the variable this civilization has just learned how to compress. Every hour of loop time removed is prevention that no longer has to be paid for in remediation, litigation, replacement, or grief. Every ton of contamination avoided because the loop closed in time is a ton that never has to be scrubbed, trucked, treated, or mourned.

We are the channel a planet has grown after four billion years of slow chemistry, the way an organism grows the cell that becomes its first neuron: blind to what it is reaching for, driven only by the need to feel itself faster. The assignment, stated as plainly as the physics permits: finish the nervous system, and keep its eyes open.

Materials and Methods

Concept formalization. Environmental latency τ_{loop} is defined as the sum of transduction, transmission, inference, decision, and actuation delays in the closed protective loop associated with a specified ecological disturbance. Stability requires $\omega_c \cdot \tau_{\text{loop}} < \phi_m$, where ω_c is disturbance bandwidth and ϕ_m is available phase margin⁽¹⁰⁾.

Historical latency estimation. Regulatory implementation timelines are sourced from EPA and CRS chronologies of NAAQS revisions⁽¹³⁾⁽¹⁴⁾⁽²⁶⁾, Endangered Species Act listing analyses⁽¹²⁾, and Montreal Protocol phase-out schedules⁽¹⁵⁾⁽¹⁶⁾.

AI forecast benchmarking. Skill comparisons for GraphCast, Pangu-Weather, GenCast, and Aurora against operational ECMWF systems (HRES, ENS) are as reported in the referenced primary publications⁽²⁾⁽³⁾⁽⁴⁾⁽⁵⁾.

Data and materials availability. No new data were generated. All quantitative claims are drawn from cited primary literature. The framework is developed further in the author's public corpus at <https://enviroai.com>.

Competing interests. The author is the founder of EnviroAI, an environmental intelligence company developing closed-loop environmental monitoring and permitting systems.

References and Notes

1. Coupé, C., Oh, Y. M., Dediu, D. & Pellegrino, F. Different languages, similar encoding efficiency: comparable information rates across the human communicative niche. *Sci. Adv.* 5, eaaw2594 (2019). <https://www.science.org/doi/10.1126/sciadv.aaw2594>
2. Bodnar, C. *et al.* A foundation model for the Earth system. *Nature* 641, 1180–1187 (2025). <https://www.nature.com/articles/s41586-025-09005-y>
3. Lam, R. *et al.* Learning skillful medium-range global weather forecasting. *Science* 382, 1416–1421 (2023). <https://www.science.org/doi/10.1126/science.adi2336>
4. Price, I. *et al.* Probabilistic weather forecasting with machine learning. *Nature* 637, 84–90 (2024). <https://www.nature.com/articles/s41586-024-08252-9>
5. Bi, K. *et al.* Accurate medium-range global weather forecasting with 3D neural networks. *Nature* 619, 533–538 (2023). <https://www.nature.com/articles/s41586-023-06185-3>
6. Landauer, R. Irreversibility and heat generation in the computing process. *IBM J. Res. Dev.* 5, 183–191 (1961). <https://ieeexplore.ieee.org/document/5392446>
7. Bérut, A. *et al.* Experimental verification of Landauer's principle linking information and thermodynamics. *Nature* 483, 187–189 (2012). <https://www.nature.com/articles/nature10872>
8. Hong, J., Lambson, B., Dhuey, S. & Bokor, J. Experimental test of Landauer's principle in single-bit operations on nanomagnetic memory bits. *Sci. Adv.* 2, e1501492 (2016). <https://pmc.ncbi.nlm.nih.gov/articles/PMC4795654/>
9. Bode, H. W. *Network Analysis and Feedback Amplifier Design* (Van Nostrand, 1945). <https://archive.org/details/networkanalysisf00bode>
10. Åström, K. J. & Murray, R. M. *Feedback Systems: An Introduction for Scientists and Engineers*, 2nd edn (Princeton Univ. Press, 2021). <https://fbswiki.org/wiki/index.php?title=FBS>
11. Karimi, A. & Freudenberg, J. S. Bode's sensitivity integral constraints. *arXiv* 1902.11302 (2019). <https://arxiv.org/pdf/1902.11302.pdf>
12. Puckett, E. E., Kesler, D. C. & Greenwald, D. N. Taxa, petitioning agency, and lawsuits affect time spent awaiting listing under the U.S. Endangered Species Act. *Biol. Conserv.* 201, 220–229 (2016). https://biologicaldiversity.org/campaigns/esa/pdfs/Puckett_et_al_2016.pdf
13. Bachmann, J. Will the circle be unbroken: A history of the U.S. National Ambient Air Quality Standards. *J. Air Waste Manag. Assoc.* 57, 652–697 (2007). https://datafedwiki.wustl.edu/images/6/69/Bachmann_AWMA.pdf
14. U.S. Environmental Protection Agency. Timeline of Particulate Matter (PM) National Ambient Air Quality Standards (2024). <https://www.epa.gov/pm-pollution/timeline-particulate-matter-pm-national-ambient-air-quality-standards-naaqs>
15. UNEP Ozone Secretariat. Ozone timeline (2023). <https://ozone.unep.org/ozone-timeline>
16. UN News. Ozone layer recovery is on track, due to success of Montreal Protocol (Jan 2023). <https://news.un.org/en/story/2023/01/1132277>
17. Bauer, P., Thorpe, A. & Brunet, G. The quiet revolution of numerical weather prediction. *Nature* 525, 47–55 (2015). <https://www.nature.com/articles/nature14956>
18. European Space Agency. Sentinel-2 images the globe every 5 days (2024). <https://sentinels.copernicus.eu/web/success-stories/-/sentinel-2-images-the-globe-every-5-days>
19. Planet Labs. Planet's rapid revisit platform to capture up to 12 images per day (2020). <https://www.planet.com/pulse/12x-rapid-revisit-announcement/>
20. World Economic Forum. IoT can help small and medium businesses implement sustainability strategies (2022). <https://www.weforum.org/stories/2022/07/iot-small-medium-businesses-profitable-sustainable/>
21. Jékely, G., Paps, J. & Nielsen, C. The phylogenetic position of ctenophores and the origin(s) of nervous systems. Introduction to theme issue, *Phil. Trans. R. Soc. B* (2015). <https://royalsocietypublishing.org/doi/pdf/10.1098/rstb.2015.0033>
22. Cattell, M. V. *et al.* Events in early nervous system evolution. *Top. Cogn. Sci.* (2019). <https://onlinelibrary.wiley.com/doi/10.1111/tops.12461>
23. Debanne, D., Campanac, E., Bialowas, A., Carlier, E. & Alcaraz, G. Axon physiology. *Physiol. Rev.* 91, 555–602 (2011). <https://pmc.ncbi.nlm.nih.gov/articles/PMC3849146/>
24. Nave, K.-A. Myelination and the trophic support of long axons. *Nat. Rev. Neurosci.* 11, 275–283 (2010). <https://www.ncbi.nlm.nih.gov/books/NBK10921/>

25. Francis, K. & Palsson, B. O. Effective intercellular communication distances are determined by the relative time constants for cyto/chemokine secretion and diffusion. *PNAS* 94, 12258–12262 (1997). <https://www.pnas.org/doi/pdf/10.1073/pnas.94.23.12258>
26. Congressional Research Service. 2013 National Ambient Air Quality Standard for Fine Particulate Matter. Report R43953 (2015). https://www.congress.gov/crs_external_products/R/PDF/R43953/R43953.8.pdf
27. Reed, T., Levy, R. & Sanborn, S. From "um" to "yeah": Producing, predicting, and regulating information flow in human conversation. *arXiv* 2403.08890 (2024). <https://arxiv.org/pdf/2403.08890v1.pdf>
28. Cambridge University Press & Assessment. The neurocognitive bandwidth of language. Preprint (2024). <https://www.cambridge.org/engage/api-gateway/coe/assets/orp/resource/item/68a7dc29a94eede1540325dd/original/the-neurocognitive-bandwidth-of-language-tempo-information-density-balance-linguistic-structure-and-philosophical-thinking.pdf>
29. Anderson, J. Author's corpus on environmental latency and superintelligence, EnviroAI (2024–2026). <https://enviroai.com/>
30. U.S. Environmental Protection Agency Office of Research and Development. Gold King Mine release: analytical modeling of transport in the Animas and San Juan Rivers (2016). https://www.epa.gov/sites/default/files/2016-06/documents/epa_ord_gkm_project_review_session_2_final_june222016.pdf
31. Weaver, J. W. *et al.* Simulation of the 2014 Freedom Industries chemical spill in the Elk River, West Virginia. *J. Environ. Eng.* 141, 05014006 (2015). [https://ascelibrary.org/doi/abs/10.1061/\(ASCE\)EE.1943-7870.0000930](https://ascelibrary.org/doi/abs/10.1061/(ASCE)EE.1943-7870.0000930)
32. Van Dolah, F. M., Lidie, K. B., Monroe, E. A., Bhattacharya, D., Campbell, L., Doucette, G. J. & Kamykowski, D. The Florida red tide dinoflagellate *Karenia brevis*: New insights into cellular and molecular processes underlying bloom dynamics. *Harmful Algae* 8, 562–572 (2009). <https://ui.adsabs.harvard.edu/abs/2009HAlga...8..573V/abstract>
33. Cruz, M. G. & Alexander, M. E. The 10% wind speed rule of thumb for estimating a wildfire's forward rate of spread in forests and shrublands. *Ann. For. Sci.* 76, 44 (2019). <https://hal.science/hal-02554937/document>
34. Zheng, J. & Meister, M. The unbearable slowness of being: Why do we live at 10 bits/s? *Neuron* 113, 192–204 (2025). [https://www.cell.com/neuron/fulltext/S0896-6273\(24\)00808-0](https://www.cell.com/neuron/fulltext/S0896-6273(24)00808-0)
35. Sauerbrei, B. A. & Pruszynski, J. A. The brain works at more than 10 bits per second. *Nat. Neurosci.* 28, 1365–1366 (2025). <https://pmc.ncbi.nlm.nih.gov/articles/PMC12320479/>