

**The Machine Gaze:
Machinehood as Closure Under an Observation**

PIATRA . INSTITUTE

September 2026

Abstract

Whether something is a machine is usually treated as a fact about its parts. We treat it as a fact about a system together with an observation: a system is a machine under an observation when its dynamics close on the observed variables, so that the next observed state is fixed by the current one and the inputs. The condition is exact lumpability and the commutation of Rosen’s modeling relation; it makes machinehood relative to the observation without making it arbitrary, and it names what a machine description leaves out. Three computations show what that remainder can contain. Self-sorting arrays with cell-level algorithms, reproduced after Zhang, Goldstein and Levin, all compute the same function, yet arrays that mix two algorithms cluster by algorithm while they sort, reaching aggregation of 0.59 to 0.65 where label-matched controls stay at 0.49, and at the same observed sortedness the remaining work differs 11-fold between algorithms. A controller designed on two modes of a beam is stable on its own model and destabilizes the beam unless the unmodelled modes have a damping ratio above 0.0093. In an organisation whose throughput is observed and whose repertoire for absorbing disturbances is not, ascent on observed throughput removes all maintenance and cuts long-run output to 40% of the optimum, at the policy under which the system closes exactly under its observation. We call the use of a closed quotient as if it were the system the machine gaze, and interventions that make a system close under the gaze’s observation machinization.

Introduction

A sorting algorithm is the textbook example of a machine whose function is exhausted by its specification. Zhang, Goldstein and Levin implemented bubble, insertion and selection sort as policies executed by each element of an array, with no central controller, and then measured more than sortedness (Zhang et al. 2025). The arrays still sorted. They also did things the code does not mention: when some elements were frozen, the arrays temporarily moved away from sortedness and recovered, and in arrays whose elements ran two different algorithms, elements running the same algorithm clustered together during the sort, although no element can read the algorithm of any other (Zhang et al. 2025, 2024). The authors present the arrays as a minimal case of capacities that “may lie in systems that we think we fully understand because we designed them” (Zhang et al. 2024, Introduction).

The case raises a general question. A sorting program is a machine if anything is. If the same deterministic process supports regularities that its specification does not contain, then being a sorting machine is a property of the process taken under a particular description, and the description is something an observer supplies. Earlier debates about whether organisms are machines asked whether living systems have the properties that machines have (Nicholson 2014, 2019; Bongard and Levin 2021). The sorting arrays reverse the direction of the question: they ask what an engineered machine is, once its specification is set aside.

We answer with a definition that can be checked. A system is a machine under an observation when its dynamics close on the observed variables: the next observed state is determined by the

current observed state and the input, whatever the unobserved state is. The machine is then the quotient of the system by the distinctions the observation does not draw. Closure is the lumpability condition of Markov chains (Kemeny and Snell 1960) and the condition under which Rosen’s modeling relation commutes (Rosen 1991, 2012). It makes machinehood relative to an observation, since one system can close under several observations, and it keeps machinehood from being arbitrary, since whether a system closes under a given observation is a fact about the system.

The definition gives two operations names. We call the practice of using a closed quotient as if it were the whole system the machine gaze. Treating the quotient as the system is harmless inside the regime where closure holds and fails outside it, and the failures have a specific form: the excluded states keep their own dynamics, and interventions designed on the quotient act on them. We call machinization any intervention that changes a system so that it closes under an observation it did not close under before. Machinization is how the gaze becomes accurate, and it works by removing the dynamics of what the observation leaves out. The phrase “machinic gaze” has been used for the way machine-learning models reconfigure the representation of built heritage (Arora et al. 2024); here the term refers to the epistemic and engineering operation defined above.

The contribution is the definition, four propositions that follow from it, and three computations that make its consequences concrete: the self-sorting arrays, reproduced and tested for closure; a controller designed on a reduced model of a flexible beam; and an organisation in which throughput is observed and the repertoire that absorbs disturbances is not. The literature on industrial and administrative abstraction, from Marx and Taylor to Scott, Hacking and Deleuze, enters where it describes one of these operations, and the engineering literature on realization, model reduction and regulation supplies most of the mathematics.

Machinehood as closure

A system is a set of states X with dynamics $F : X \times U \rightarrow X$, where U is a set of inputs. An observation is a surjective map $\pi : X \rightarrow Z$. The system closes under π on a regime $R \subseteq X \times U$ if there is a map $G : Z \times U \rightarrow Z$ with

$$\pi(F(x, u)) = G(\pi(x), u) \quad \text{for all } (x, u) \in R.$$

The pair (Z, G) is then the machine that the system realizes under π on R , and the fibres $\pi^{-1}(z)$ are the distinctions the machine does not draw. For a stochastic system with transition kernel $P(x, \cdot)$, closure is lumpability: the probability of moving into each fibre depends on x only through $\pi(x)$ (Kemeny and Snell 1960).

Proposition 1. The system closes under π on R if and only if $\pi(F(x, u)) = \pi(F(x', u))$ whenever $\pi(x) = \pi(x')$ and both (x, u) and (x', u) lie in R . When it closes, G is unique on $\pi(R)$.

Proof. If G exists, both sides equal $G(\pi(x), u)$. Conversely, set $G(z, u) = \pi(F(x, u))$ for any x in the fibre over z ; the condition makes the choice irrelevant. $\square$

Three readings of the condition already exist in the literature. In Rosen’s modeling relation a natural system and a formal system are linked by encoding and decoding, and the formal system models the natural one when following causal entailment and then encoding gives the same result as encoding and then following inferential entailment (Rosen 1991; Rosen 2012, ch. 3). Closure is that commutation with π as the encoding and G as the formal system. In automata theory, two states are equivalent when no experiment on inputs and outputs distinguishes them (Moore 1956), and probabilistic bisimulation makes the same identification for stochastic processes (Larsen and Skou 1991); the machine is the quotient by that equivalence. In linear systems theory, Kalman showed that the input-output relations of a system determine only its completely controllable and completely observable part (Kalman 1963). The input-output machine is therefore blind, by construction, to everything outside that part, and Kalman’s result fixes exactly what the blindness covers.

Proposition 1 has two consequences that organize the rest of the argument. One system can close under several observations whose fibres cut across one another, so that it realizes several machines at once; the sorting arrays below are an example. And whether a system closes under a given observation, and on which regime, is a fact about the system that can fail. Machinehood is relative to an observation without being a matter of choice. Wimsatt’s analysis of aggregativity makes the same point about decompositions: the conditions under which a whole behaves as a sum of its parts are met differently for different decompositions of the same system, and to different degrees, which turns the choice of decomposition into an empirical question (Wimsatt 1997, 2007). Descriptions that hold under several independent observations and interventions deserve more ontological confidence (Wimsatt 2012), and a pattern that permits compressed prediction is real in the sense that matters for prediction (Dennett 1991).

Closure usually holds on a regime. Simon observed that many complex systems are nearly decomposable: interactions within subsystems are fast and strong, interactions between them slow and weak, so that over short times the subsystems behave independently and over long times they affect one another only through aggregates (Simon 2019, ch. 8). Model reduction exploits the same structure. A balanced truncation keeps the states that are both strongly controllable and strongly observable (Moore 1981), and the reduced model is accurate for the inputs, outputs and frequencies for which it was built (Antoulas 2005). Each reduced model is a machine on a regime.

Two definitions complete the vocabulary. The machine gaze is the use of (Z, G) in place of (X, F) outside the regime on which closure holds, or the treatment of the fibres as if they were empty. Machinization is an intervention that changes F to some F' so that the system closes under π on a larger regime, or with a smaller defect. A convenient measure of the defect, for a future observed quantity Y , is the conditional mutual information $I(Y; X \mid \pi(X))$, which vanishes exactly when the distribution of Y depends on x only through $\pi(x)$.

One function, several dynamics

We reimplemented the cell-level algorithms of Zhang, Goldstein and Levin with random sequential updates in place of their threads, following the rules in their code. A bubble cell compares itself

with a randomly chosen neighbour and swaps if the pair is out of order. An insertion cell moves left past a larger neighbour only when everything to its left is sorted. A selection cell keeps an ideal position, initially the left end; it swaps into that position if the occupant has a larger value and otherwise advances the ideal position by one (Zhang et al. 2024, Methods). Arrays had 100 cells holding the values 0 to 99 in random order, and each condition was run 100 times. Every run of every condition ended sorted, including mixtures of two or three algorithms, so under the observation that records an initial array and a final array every composition is the same machine: it computes the sorting function.

The mean number of swaps reproduces the published values closely: 1716 for bubble with selection, against 1740.9 reported, 2466 for bubble with insertion, against 2476.02, and 1557 for insertion with selection, against 1534.77 (Zhang et al. 2024). Aggregation was measured as in the original, as the fraction of cells whose right-hand neighbour runs the same algorithm, and in controls every cell ran one algorithm while carrying the same two labels. In mixed arrays aggregation starts at 0.50, rises early in the run to a peak of the mean curve of 0.59 for bubble with selection, 0.61 for bubble with insertion and 0.65 for insertion with selection, and falls back near its starting value at the end, when sorting by value has scattered the labels again. In controls the mean at the same points is 0.49 (Figure 1a), and the difference is large in every pair (Welch t from 13.8 to 23.0). Mean per-run peaks are 0.674, 0.666 and 0.698, against 0.72, 0.65 and 0.69 reported in the original study (Zhang et al. 2024, Results).

The clustering is a regularity of the dynamics that the sorting specification neither requires nor forbids, and the original authors offer a mechanism for it: an algorithm that moves its elements into place faster than its partner leaves partially sorted runs of its own type behind, which the slower algorithm later breaks up (Zhang et al. 2024, Results). A mechanism of this kind is what one should expect. The clustering is an ordinary consequence of the dynamics, visible under an observation other than the one the specification uses.

The sortedness observable, the fraction of adjacent pairs in order, also fails to close. We stopped each run when its sortedness first reached 0.9 and counted the swaps that remained. At that common observed state the remaining work averaged 72.5 swaps for bubble, 795.7 for insertion, 117.0 for selection and 82.6 for the bubble-selection mixture, a ratio of 11.0 between the largest and smallest means (Figure 1b–c). The future of the observed progress depends on the algorithms the cells run, which the progress measure does not record. Under the specification’s observation the arrays are one machine; under the progress observation they are not a machine at all; under the label observation they show a regularity that neither of the other two contains. Bongard and Levin describe this situation for organisms, which “perform multiple functions in the same place at the same time” and are best treated with an observer-dependent, pragmatic view of what they compute (Bongard and Levin 2023). The sorting arrays show that the same holds for a program whose every line is known.




Figure 1: One function, several dynamics. (a) Aggregation by algorithm during sorting, averaged over 100 runs per condition, for arrays that mix two algorithms (solid) and for controls in which every cell runs one algorithm and carries the same two labels (dotted); the horizontal axis is the fraction of each run’s swaps. (b) Sortedness against swaps for three runs each of pure bubble, insertion and selection arrays; all reach the same final state. (c) Swaps remaining when sortedness first reaches 0.9, for four compositions (100 runs each; boxes show quartiles, whiskers 1.5 interquartile ranges).

Acting on a quotient

The machine gaze becomes consequential when a controller designed on the quotient acts on the system. Consider a pinned beam with modal frequencies $\omega_k = (k\pi)^2$, a point actuator at 0.21 of its length and a displacement sensor at 0.83, and a controller designed on the first two modes: a linear-quadratic regulator with a state observer, both computed from the two-mode model. On its own model the closed loop is stable, with all poles at real part -4.32 or below. The beam has three further modes in the simulation, damped with ratio ζ_R , which the controller’s model omits. The actuator excites them (control spillover) and the sensor picks them up (observation spillover).

Proposition 2. If either spillover is absent, the poles of the full closed loop are the poles of the designed closed loop together with the open-loop poles of the residual modes. If both are present, the residual poles can move into the right half plane.

Proof. Order the state as modelled modes, observer error and residual modes. Without observation spillover the residual modes do not enter the controller’s measurement, and without control spillover the controller does not drive them; in either case the closed-loop matrix is block triangular, and its spectrum is the union of the diagonal blocks. With both, the residual block is coupled to the controller in both directions. $\square$

This is the instability that Balas identified for flexible structures, where the combined effect of control and observation spillover can destabilize a loop designed on a finite number of modes (Balas 1978). In the computation the full loop is unstable for every residual damping ratio below 0.0093, with the largest pole real part 0.82 at zero damping and 0.73 at 0.001, while the designed loop keeps its poles at -4.32 (Figure 2a). The destabilized pole belongs to the third mode, the first unmodelled one, and oscillates at 85.3 against its natural frequency of 88.8. With a sensor

that cannot see the residual modes the loop is stable for any positive damping, as Proposition 2 requires.

The time course shows how the gaze reads its own success. From a small displacement of the first mode, the modelled modes decay to an amplitude of $1.8e-5$ within three time units, as the design promises, while the residual modes, excited through the actuator, overtake them after 0.65 time units and then grow without bound, dragging the modelled modes with them (Figure 2b). The design succeeds on the quotient during the period in which it would be judged, and fails on the system. The unmodelled damping that prevents this, above 0.0093, is a property of the material that no part of the design accounts for. Conant and Ashby proved that a regulator that is maximally successful and maximally simple must be isomorphic to the system it regulates (Conant and Ashby 1970); the computation shows the converse hazard, a regulator isomorphic to a quotient of its system, which regulates the quotient.




Figure 2: Acting on a quotient. (a) Largest real part of the closed-loop poles against the damping ratio of the three unmodelled modes, for the controller acting on the full beam, on its own two-mode model, and on the full beam with a sensor blind to the unmodelled modes; the dotted line marks the critical damping ratio 0.0093. (b) Amplitudes of the modelled and unmodelled modes after a small displacement of the first mode, with residual damping ratios 0.001 (solid) and 0.03 (dashed).

Machinization

The machine gaze can also be made accurate by changing the system. Taylor stated the program in his own terms: “In the past the man has been first; in the future the system must be first” (Taylor 1911, 7), and management takes on “the burden of gathering together all of the traditional knowledge which in the past has been possessed by the workmen and then of classifying, tabulating, and reducing this knowledge to rules, laws, and formulae” (Taylor 1911, 36). In our terms, the workshop is to be redesigned so that its dynamics close on the variables management records. Marx had described the outcome half a century earlier: “In handicrafts and manufacture, the workman makes use of a tool, in the factory, the machine makes use of him”, and the worker becomes the “mere living appendage” of a mechanism independent of him (Marx 1887, ch. XV, sec. 4).

Braverman traced how scientific management separated the conception of work from its execution (Braverman 1974), Foucault analysed discipline as the decomposition and recomposition of bodily gestures and times into a calculable machinery (Foucault 1977, pt. 3, ch. 1), and Mumford argued that organisations of people assembled as standardized, replaceable parts, which he called the megamachine, preceded mechanical machinery by millennia (Mumford 1967). Each of these accounts describes an intervention on a human system that makes it close under the variables of the organisation that governs it.

A minimal model shows what machinization removes and why the gaze cannot see it. An organisation has N units. Each spends a share s of its time maintaining a repertoire of responses to K types of disturbance and the rest producing at unit rate, so that throughput while running is $(1 - s)N$. For each unit and each type, the response is learned at rate $s\rho/K$ and forgotten at rate δ , so a unit holds a given response with stationary probability $q(s) = (s\rho/K)/(s\rho/K + \delta)$, and a type is covered by at least one unit with probability $c(s) = 1 - (1 - q)^N$. Disturbances of uniformly random type arrive at rate ν while the organisation runs; a covered disturbance is absorbed, and an uncovered one stops all output for a time τ . The organisation observes throughput.

Proposition 3. (i) The long-run output rate is $Y(s) = (1 - s)N/[1 + \nu\tau(1 - c(s))]$, while observed throughput $(1 - s)N$ decreases in s . (ii) The information that the hidden repertoire carries about whether the next disturbance is absorbed, given the observation, is $I(s) = h(c) - \mathbb{E} h(B/K)$ with $B \sim \text{Binomial}(K, c)$ and h the binary entropy; it vanishes if and only if $c \in \{0, 1\}$, and in particular at $s = 0$, where $Y(0) = N/(1 + \nu\tau)$.

Proof. (i) By the renewal-reward theorem, each unit of running time produces $(1 - s)N$ and is followed on average by $\nu(1 - c)\tau$ units of stoppage, since each disturbance is uncovered with probability $1 - c$ in the stationary state. (ii) Given the set of covered types, absorption of a uniformly random type has probability B/K ; the mutual information between the coverage and the absorption is the entropy of the marginal minus the mean conditional entropy, and the conditional entropies all vanish only when B is almost surely 0 or K . $\square$

With $N = 10$, $K = 20$, $\rho = 1$, $\delta = 0.05$, $\nu = 0.005$ and $\tau = 800$, the long-run output is largest at $s^* = 0.323$, where it is 5.44 against an observed throughput of 6.77 and the repertoire covers 0.94 of disturbance types. At $s = 0$ observed throughput is largest and long-run output is 2.00, a factor of 2.72 lower, because 80% of the time is spent stopped. Simulations of the full stochastic model at four fixed policies reproduce $Y(s)$ to within 0.81% (Figure 3a). The parameters are illustrative; the qualitative result needs only that maintenance takes time from throughput and that its benefit arrives through rare events.

Now let the organisation adjust s by finite-difference ascent on observed throughput: alternate windows of 10 time units at $s + 0.02$ and $s - 0.02$, compare their output, and move s by 0.01 times the normalized difference, setting aside any pair of windows that contains a stoppage as disruption from outside. On the windows it keeps, the gaze's model, output equals $(1 - s)N$ times the window length, is exact. Starting from s^* , the ascent reaches $s = 0$ after 3120 time units and stays there (Figure 3b). Over a run of 2,000,000 time units it sets aside 79,683 pairs of windows

as disruption and uses 20,317, because the organisation is stopped most of the time. Over the second half of the run its output is 2.18, which is 0.40 of the optimum and 1.017 times Y at its average effective share of 0.01, the residue of its own probing.

The policy the gaze converges to is the one under which the organisation is exactly a machine under its observation: with no maintenance the repertoire is empty, it carries no information about the future, and the observed throughput together with the record of disturbances predicts everything that happens (Proposition 3, Figure 3c). Ashby’s law of requisite variety states the cost in general form: in his words, “only variety can destroy variety” (Ashby 1956, S.11/7), and the repertoire is the variety with which the organisation regulates its disturbances. Holmstrom and Milgrom showed the same effect in contracts: when some tasks are measured and others are not, incentives attached to the measured tasks draw effort from the unmeasured ones (Holmstrom and Milgrom 1991). Campbell observed that a quantitative indicator used for social decisions becomes subject to pressures that distort the processes it was meant to monitor (Campbell 1979), and Strathern generalized the point for audit (Strathern 1997). The model adds that the distortion can be invisible on every observation the gaze retains, and that it is complete at the point where the system has been made to fit the gaze. Beer’s viable system model was built to preserve the operational autonomy that this kind of optimization removes (Beer 1984).




Figure 3: Machinization in an organisation whose throughput is observed. (a) Observed throughput while running and long-run output against the share of time spent maintaining the repertoire, with simulations of the full stochastic model at four fixed shares. (b) The maintenance share and the fraction of disturbance types covered while the organisation ascends observed throughput, starting from the optimal share; flat stretches of the share are stoppages, during which comparisons are set aside. (c) Information that the hidden repertoire carries about the absorption of the next disturbance, given the observation.

The loop between model and system

Rosen’s modeling relation runs from the system to the model. Machinization adds a return path: the model guides interventions, and the interventions change the system. Hacking called the version of this loop that runs through human self-understanding the looping effect of human kinds, in which classifications change the people classified and the changed people change the classifi-

cation (Hacking 1995). MacKenzie showed that the option-pricing models of financial economics became, for a period, more accurate because traders used them (MacKenzie 2006), and Bowker and Star traced how classification systems become infrastructure that shapes the worlds they sort (Bowker and Star 1999). Scott described the administrative version: states simplify the societies they govern into legible categories, and then remake those societies along the categories (Scott 1998, introduction). Espeland and Stevens analyse commensuration, the comparison of different entities under a common metric, as a mode of power (Espeland and Stevens 1998).

A two-variable model isolates the part of these accounts that concerns who adjusts. Let θ be what the model asserts about an observed quantity and x the quantity itself. The model learns at rate η , $\dot{\theta} = \eta(x - \theta)$, and the system is pushed toward the model at rate κ , $\dot{x} = \kappa(\theta - x)$.

Proposition 4. The quantity $\kappa\theta + \eta x$ is conserved, both variables converge to $(\kappa\theta_0 + \eta x_0)/(\kappa + \eta)$, and the share of the initial disagreement removed by changing the system is $\kappa/(\kappa + \eta)$.

Proof. $\frac{d}{dt}(\kappa\theta + \eta x) = \kappa\eta(x - \theta) + \eta\kappa(\theta - x) = 0$, and $\theta - x$ decays at rate $\kappa + \eta$. $\square$

An observer who can only learn ($\kappa = 0$) moves the model all the way to the world, which is the ordinary situation of science. An observer that intervenes at 3 times the rate at which it learns removes 0.75 of the disagreement by changing the world, and one with $\kappa = 9\eta$ removes 0.9. Prediction succeeds in every case, which is why the accuracy of a model of a governed system says little by itself about whether the model was right about the system before it was governed. Combined with Proposition 1, the result says where the adjustment lands: the system is changed along the observed variables, and the fibres absorb whatever the intervention does to them.

Two further descriptions of the gaze place it historically. Heidegger characterized modern technology as a way of revealing in which things show up as standing-reserve, available to be ordered and stockpiled, and observed that human beings are themselves challenged forth into this ordering (Heidegger 1977). In the present vocabulary, standing-reserve is what a system becomes when the observation retains only what can be ordered. Deleuze described a shift from disciplinary institutions that mould individuals to control mechanisms that modulate continuously and address “dividuals”, fragments such as codes, samples and data (Deleuze 1995). In engineering terms, the observation is replaced by a state estimate updated as data arrive, and the intervention loop of Proposition 4 runs continuously, with a separate quotient of the same person for every system that addresses them.

What resists closure

Closure applies in the same way to organisms, people and artifacts, and the symmetry is informative. Canguilhem argued that a living being differs from a machine in its normativity: it institutes its own norms in relation to its milieu, and pathology is a narrowing of the norms it can sustain (Canguilhem 1978, 2008). A machine description fixes the criterion of success along with the observation; a system that can change its criterion can leave any closed quotient defined with respect to the old one. Simondon argued from the other side that technical objects themselves are not exhausted by their functions, since each develops through the mutual adaptation of its parts and

depends on an associated milieu of physical conditions that it partly creates (Simondon 2017). The beam computation is a small instance of this: the controller's function is defined on two modes, and its fate is decided by a damping ratio in three others.

Whitehead named the general error. Abstractions such as simply located bits of matter are legitimate products of thought, and the mistake lies in taking them for concrete fact, "the accidental error of mistaking the abstract for the concrete", which he called the fallacy of misplaced concreteness (Whitehead 1925, 74–75). He added the cost of a well-founded abstraction: "by the nature of the case, you have abstracted from the remainder of things" (Whitehead 1925, 85). The definition of closure locates that remainder in the fibres of the observation, and the computations show that it has dynamics of its own, that interventions act on it, and that it can be emptied.

Diagnostics

The framework yields checks that apply wherever a machine description is in use. First, test closure directly: hold the observed state fixed and ask whether the future of the observation varies with something the observation omits, as the remaining work at fixed sortedness varies with the algorithm. Second, perturb outside the regime on which the description was built, since closure on a regime says nothing beyond it, as the beam's design was sound on its two modes. Third, audit what interventions remove from the fibres. In the organisation model the gaze's improvement in observed throughput was paid for in repertoire, which no observed variable recorded. Fourth, watch what the gaze sets aside: a procedure that classifies the data it cannot fit as external disruption will find its model exact on the data it keeps, as the ascent did while discarding most of its comparisons.

Limitations

The propositions are elementary and the computations are minimal models with illustrative parameters. The sorting reproduction replaces threads with random sequential updates, and its peaks differ from the published ones by up to 0.05, with the largest difference for bubble with selection. The beam has five modes and a single actuator and sensor, and its controller ignores the robustness methods that engineers use against spillover. The organisation model represents the repertoire as independent two-state chains and treats every uncovered disturbance identically; real organisations learn from failures, and many disturbances are not independent. The historical accounts are used for the operation they describe, and the formal model adds nothing to their historical claims. The definition of a machine as a closed quotient leaves open which observations a community should adopt, which is a normative question.

Conclusion

A system is a machine under an observation when its dynamics close on the observed variables. The definition makes one system several machines, makes each of them a fact about the system that can fail, and identifies what each leaves out. The machine gaze uses a closed quotient as the system;

it errs outside the regime of closure, and its interventions act on the states it does not distinguish. Machinization removes those states' dynamics until the system closes, which can make the gaze exact and the system fragile at the same time. Whether a given organisation, technology or person is being described or being remade can be decided case by case, by testing closure, perturbing outside the design regime and auditing what interventions remove.

Reproducibility

The simulation in `simulation/` regenerates every number and figure: `uv run run_all.py` writes `output/results.json` and the three figures and evaluates 18 internal checks (closed forms against simulation, reproduction against published values, stability claims against eigenvalues). The claims ledger `claims.yaml` binds each reported number to its location in `results.json` through a recorded execution.

References

- Antoulas, Athanasios C. 2005. *Approximation of Large-Scale Dynamical Systems*. Society for Industrial and Applied Mathematics. <https://doi.org/10.1137/1.9780898718713>.
- Arora, Vanicka, Liam Magee, and Luke Munn. 2024. "(Re)framing Built Heritage Through the Machinic Gaze." *Journal of Social Archaeology* 24 (2): 197–217. <https://doi.org/10.1177/14696053241237949>.
- Ashby, W. Ross. 1956. *An Introduction to Cybernetics*. Chapman & Hall.
- Balas, M. 1978. "Feedback Control of Flexible Systems." *IEEE Transactions on Automatic Control* 23 (4): 673–79. <https://doi.org/10.1109/tac.1978.1101798>.
- Beer, Stafford. 1984. "The Viable System Model: Its Provenance, Development, Methodology and Pathology." *Journal of the Operational Research Society* 35 (1): 7–25. <https://doi.org/10.1057/jors.1984.2>.
- Bongard, Joshua, and Michael Levin. 2021. "Living Things Are Not (20th Century) Machines: Updating Mechanism Metaphors in Light of the Modern Science of Machine Behavior." *Frontiers in Ecology and Evolution* 9. <https://doi.org/10.3389/fevo.2021.650726>.
- Bongard, Joshua, and Michael Levin. 2023. "There's Plenty of Room Right Here: Biological Systems as Evolved, Overloaded, Multi-Scale Machines." *Biomimetics* 8 (1): 110. <https://doi.org/10.3390/biomimetics8010110>.
- Bowker, Geoffrey C., and Susan Leigh Star. 1999. *Sorting Things Out: Classification and Its Consequences*. The MIT Press. <https://doi.org/10.7551/mitpress/6352.001.0001>.
- Braverman, Harry. 1974. *Labor and Monopoly Capital: The Degradation of Work in the Twentieth Century*. Monthly Review Press.

- Campbell, Donald T. 1979. "Assessing the Impact of Planned Social Change." *Evaluation and Program Planning* 2 (1): 67–90. [https://doi.org/10.1016/0149-7189\(79\)90048-x](https://doi.org/10.1016/0149-7189(79)90048-x).
- Canguilhem, Georges. 1978. *On the Normal and the Pathological*. Translated by Carolyn R. Fawcett. Springer Netherlands. <https://doi.org/10.1007/978-94-009-9853-7>.
- Canguilhem, Georges. 2008. *Knowledge of Life*. Edited by Paola Marrati and Todd Meyers. Translated by Stefanos Geroulanos and Daniela Ginsburg. Fordham University Press.
- Conant, Roger C., and W. Ross Ashby. 1970. "Every Good Regulator of a System Must Be a Model of That System." *International Journal of Systems Science* 1 (2): 89–97. <https://doi.org/10.1080/00207727008920220>.
- Deleuze, Gilles. 1995. "Postscript on Control Societies." In *Negotiations, 1972–1990*, translated by Martin Joughin. Columbia University Press.
- Dennett, Daniel C. 1991. "Real Patterns." *The Journal of Philosophy* 88 (1): 27–51. <https://doi.org/10.2307/2027085>.
- Espeland, Wendy Nelson, and Mitchell L. Stevens. 1998. "Commensuration as a Social Process." *Annual Review of Sociology* 24 (1): 313–43. <https://doi.org/10.1146/annurev.soc.24.1.313>.
- Foucault, Michel. 1977. *Discipline and Punish: The Birth of the Prison*. Translated by Alan Sheridan. Pantheon Books.
- Hacking, Ian. 1995. "The Looping Effects of Human Kinds." In *Causal Cognition: A Multidisciplinary Debate*, edited by Dan Sperber, David Premack, and Ann James Premack. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198524021.003.0012>.
- Heidegger, Martin. 1977. *The Question Concerning Technology, and Other Essays*. Translated by William Lovitt. Harper & Row.
- Holmstrom, Bengt, and Paul Milgrom. 1991. "Multitask Principal–Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design." *The Journal of Law, Economics, and Organization* 7 (special issue): 24–52. https://doi.org/10.1093/jleo/7.special_issue.24.
- Kalman, R. E. 1963. "Mathematical Description of Linear Dynamical Systems." *Journal of the Society for Industrial and Applied Mathematics Series A Control* 1 (2): 152–92. <https://doi.org/10.1137/0301010>.
- Kemeny, John G., and J. Laurie Snell. 1960. *Finite Markov Chains*. Van Nostrand.
- Larsen, Kim G., and Arne Skou. 1991. "Bisimulation Through Probabilistic Testing." *Information and Computation* 94 (1): 1–28. [https://doi.org/10.1016/0890-5401\(91\)90030-6](https://doi.org/10.1016/0890-5401(91)90030-6).
- MacKenzie, Donald. 2006. *An Engine, Not a Camera: How Financial Models Shape Markets*. The MIT Press. <https://doi.org/10.7551/mitpress/9780262134606.001.0001>.

- Marx, Karl. 1887. *Capital: A Critical Analysis of Capitalist Production*. Edited by Frederick Engels. Translated by Samuel Moore and Edward Aveling. Vol. 1. Swan Sonnenschein, Lowrey & Co.
- Moore, B. 1981. "Principal Component Analysis in Linear Systems: Controllability, Observability, and Model Reduction." *IEEE Transactions on Automatic Control* 26 (1): 17–32. <https://doi.org/10.1109/tac.1981.1102568>.
- Moore, Edward F. 1956. "Gedanken-Experiments on Sequential Machines." In *Automata Studies*, edited by C. E. Shannon and J. McCarthy. Annals of Mathematics Studies 34. Princeton University Press. <https://doi.org/10.1515/9781400882618-006>.
- Mumford, Lewis. 1967. *The Myth of the Machine: Technics and Human Development*. Harcourt, Brace & World.
- Nicholson, Daniel J. 2014. "The Machine Conception of the Organism in Development and Evolution: A Critical Analysis." *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 48: 162–74. <https://doi.org/10.1016/j.shpsc.2014.08.003>.
- Nicholson, Daniel J. 2019. "Is the Cell Really a Machine?" *Journal of Theoretical Biology* 477: 108–26. <https://doi.org/10.1016/j.jtbi.2019.06.002>.
- Rosen, Robert. 1991. *Life Itself: A Comprehensive Inquiry into the Nature, Origin, and Fabrication of Life*. Columbia University Press.
- Rosen, Robert. 2012. *Anticipatory Systems: Philosophical, Mathematical, and Methodological Foundations*. 2nd ed. IFSR International Series in Systems Science and Systems Engineering. Springer New York. <https://doi.org/10.1007/978-1-4614-1269-4>.
- Scott, James C. 1998. *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed*. Yale University Press.
- Simon, Herbert A. 2019. *The Sciences of the Artificial*. 3rd ed. The MIT Press. <https://doi.org/10.7551/mitpress/12107.001.0001>.
- Simondon, Gilbert. 2017. *On the Mode of Existence of Technical Objects*. Translated by Cecile Malaspina and John Rogove. Univocal.
- Strathern, Marilyn. 1997. "'Improving Ratings': Audit in the British University System." *European Review* 5 (3): 305–21. [https://doi.org/10.1002/\(sici\)1234-981x\(199707\)5:3%3C305::aid-euro184%3E3.0.co;2-4](https://doi.org/10.1002/(sici)1234-981x(199707)5:3%3C305::aid-euro184%3E3.0.co;2-4).
- Taylor, Frederick Winslow. 1911. *The Principles of Scientific Management*. Harper & Brothers.
- Whitehead, Alfred North. 1925. *Science and the Modern World*. Macmillan.
- Wimsatt, William C. 1997. "Aggregativity: Reductive Heuristics for Finding Emergence." *Philosophy of Science* 64 (S4): S372–84. <https://doi.org/10.1086/392615>.

- Wimsatt, William C. 2007. *Re-Engineering Philosophy for Limited Beings: Piecewise Approximations to Reality*. Harvard University Press. <https://doi.org/10.2307/j.ctv1pncnrh>.
- Wimsatt, William C. 2012. "Robustness, Reliability, and Overdetermination (1981)." In *Characterizing the Robustness of Science: After the Practice Turn in Philosophy of Science*, edited by Léna Soler, Emiliano Trizio, Thomas Nickles, and William C. Wimsatt. Boston Studies in the Philosophy of Science. Springer Netherlands. https://doi.org/10.1007/978-94-007-2759-5_2.
- Zhang, Taining, Adam Goldstein, and Michael Levin. 2024. "Classical Sorting Algorithms as a Model of Morphogenesis: Self-Sorting Arrays Reveal Unexpected Competencies in a Minimal Model of Basal Intelligence." arXiv:2401.05375. Preprint. <https://doi.org/10.48550/ARXIV.2401.05375>.
- Zhang, Taining, Adam Goldstein, and Michael Levin. 2025. "Classical Sorting Algorithms as a Model of Morphogenesis: Self-Sorting Arrays Reveal Unexpected Competencies in a Minimal Model of Basal Intelligence." *Adaptive Behavior* 33 (1): 25–54. <https://doi.org/10.1177/10597123241269740>.