The Computational Transition Model

When Machine Interpreter Capacity Outruns Organizational Complexity

Eli Mikel
Independent Researcher; Founder and CEO, Clockwise Capital, LLC
Miami, FL, USA
Corresponding author: eli@clockwisecapital.com
ORCID: 0009-0003-1481-6027
Article type: Theoretical Article
Word count: 4,561 (excluding references)

Abstract

Technological forecasts usually ask how rapidly artificial intelligence improves. For organizations, however, absolute machine capability is not the relevant transition variable. A more capable model can remain external, advisory, opaque, nonportable, or disconnected from the machinery that reconstructs an organization after turnover. This article develops the Computational Transition Model (ACTM), a dynamic theory of when machine capability becomes sufficient relative to the complexity of the system it must interpret. ACTM separates raw capability from effective machine interpreter capacity, defines reconstructive organizational complexity, and combines them in a dimensionless relative-sufficiency variable, χ. The model yields a simple directional result: relative interpreter sufficiency rises when the growth rate of effective interpreter capacity exceeds complexity-adjusted organizational growth, remains stationary when the rates match, and falls when complexity grows faster. When executability and developmental heredity are monotone functions of relative sufficiency, the same inequality determines the predicted direction of change in organizational reconstruction fidelity. The model also derives system-specific threshold conditions and an analytically convenient exponential benchmark without assuming a universal law of computational growth. ACTM therefore predicts a transition surface rather than a universal technology threshold or calendar date. The theory is falsifiable through prospective measurement, architectural interventions, complexity-reduction interventions, and out-of-sample threshold tests.

Keywords: artificial intelligence; organizational complexity; executability; developmental heredity; technological change; evolutionary individuality

Introduction

Discussion of artificial intelligence often treats technological capability as the independent variable that matters most. Model performance rises, inference becomes cheaper, context windows expand, and systems acquire broader tool use. It is therefore tempting to infer that sufficiently rapid capability growth must eventually transform the organizations that deploy those systems. That inference is incomplete. A powerful model that cannot access canonical state, carry authorized memory, execute organizational rules, preserve permissions, or survive component turnover may contribute little to the machinery by which an organization reconstructs itself. Conversely, a narrower machine layer with persistent state, authority, and reproducible control logic may have a larger organizational effect than a more capable but weakly integrated model.

The same problem appears from the opposite direction. Organizations do not remain structurally constant while artificial intelligence improves. They accumulate dependencies, data, roles, policies, software, counterparties, interfaces, and environmental couplings. They can become easier to reconstruct through modularization and standardization, or harder to reconstruct through entanglement and path dependence. A theory of computational transition therefore requires two moving quantities rather than one: the effective interpreter capacity available to the system and the reconstructive complexity of the system being interpreted.

The Computational Transition Model (ACTM) formalizes this race. Its central claim is not that computation grows exponentially, that artificial intelligence necessarily replaces human interpretation, or that a sufficiently capable model creates a new evolutionary individual. The core claim is relational: a system moves toward greater machine-mediated reconstructive sufficiency only when effective interpreter capacity grows faster than the complexity-adjusted burden placed on that capacity. Absolute technological progress can therefore coexist with declining relative sufficiency.

ACTM makes four contributions. First, it separates raw machine capability $C(t)$ from effective machine interpreter capacity $I_A(t)$ through an explicit integration architecture. Second, it defines relative interpreter sufficiency $\chi(t)$ as interpreter capacity normalized by reconstructive complexity. Third, it derives the growth-rate condition that determines whether relative sufficiency rises or falls and, under stated monotonicity assumptions, links that condition to the direction of change in developmental heredity. Fourth, it converts a registered developmental-heredity criterion into a prospective, system-specific threshold without implying a universal date of technological or evolutionary transition.

The theory is motivated by the interpreter problem developed in work on organizational heredity: copying a specification does not necessarily reconstruct the organization that specification describes. Evolutionary economics has long emphasized that routines contain tacit and developmental components that are not reducible to documentation (Nelson and Winter 1982), while generative-replication accounts distinguish transmitted information from the machinery required to rebuild organized form (Hodgson and Knudsen 2008, 2010). ACTM adds a dynamic question: when should machine-mediated interpretation become increasingly sufficient relative to the burden of reconstruction?

The model is compatible with, but does not depend on, stronger claims about evolutionary individuality. Within the AI-Enabled Individuation Hypothesis (AEIH), developmental heredity is one condition relevant to higher-order individuation. ACTM can be evaluated earlier. It can fail as a model of organizational reconstruction even if artificial intelligence remains economically important, and it can succeed as a model of reconstruction without establishing that the organization is a biological or Darwinian individual.

The interpreter problem

An organization can preserve its written specification while failing to preserve its operational form. Code can be copied while governance collapses. Procedures can survive while the tacit practices that make them executable disappear. A franchise manual can be transferred without recreating the same service quality, conflict-resolution norms, or decision architecture. The distinction is between information that describes organization and machinery that interprets that information into organization.

Let G denote a portable organizational specification: code, governance rules, procedures, schemas, templates, models, or other transferable state. Let D denote developmental or interpreter machinery: trained roles, tacit knowledge, routines, control processes, and machine-mediated interpreters. Let Z denote a validated organizational phenotype, and let ε collect environmental and stochastic influences. The general reconstruction problem is

$$Z = \Phi(G, D, \varepsilon).$$

The form is deliberately general. It does not assume that human and machine interpreters are additive or that organizational phenotype can be reduced to a single scalar. It only states that reconstructing organized form depends on a specification, interpreter machinery, and context.

Let $E \in [0, 1)$ denote **executability**: the fraction of parent-relevant interpreter function effectively carried by or reconstructable from inherited artifacts. Executability is not the fraction of tasks automated, not benchmark model performance, and not available FLOPs. It concerns whether organizational information participates directly in reconstruction rather than waiting for transient members to supply all interpretation.

A local first-order model makes the distinction measurable. Around a parent state, let syntactic perturbation, interpreter mismatch, and environmental perturbation have per-dimension variances σ_s^2, σ_d^2, and σ_ε^2. Let a, b, and c summarize the local sensitivity of organizational phenotype to those perturbations. The expected per-dimension developmental reconstruction error is

$$V_d(E) = a\,\sigma_s^2 + b(1 - E)^2\,\sigma_d^2 + c\,\sigma_\varepsilon^2.$$

For a monotone exponential fidelity mapping,

$$H_d(E) = \exp\left[-\lambda_d V_d(E)\right), \lambda_d > 0,$$

where λ_d sets how rapidly measured developmental fidelity decays with reconstruction error. For $E < 1$, $b > 0$, and $\sigma_d^2 > 0$,

$$\frac{dH_d}{dE} = 2\lambda_d b(1-E)\sigma_d^2 H_d(E) > 0.$$

The result is limited but important. Increasing executability raises developmental fidelity by reducing sensitivity to interpreter mismatch. It does not remove specification error or the environmental floor. The result also does not imply that more artificial intelligence automatically means greater executability. ACTM begins precisely where that static relation stops: what determines whether E rises through time?

From capability to effective interpreter capacity

Let $C(t)$ denote machine capability available to a candidate organization at time t. Depending on the empirical domain, C could be represented by a preregistered composite of model performance, inference economics, memory, tool-use reliability, planning horizon, or other capabilities relevant to the organizational task. ACTM does not require one universal capability metric.

Raw capability is filtered through an **integration architecture**, denoted $A_r(t)$. The architecture includes the organizational properties that determine whether capability enters reconstructive machinery: access to canonical state, authority to execute or block actions, persistent memory, machine-readable rules, reproducibility of configuration, portability of state and permissions, auditability, identity controls, and dependence on inaccessible vendor-side components.

Define effective machine interpreter capacity as

$$I_A(t) = \Omega\left[C(t), A_r(t)\right].$$

The function Ω is empirical rather than definitional. A model upgrade can increase C while leaving I_A nearly unchanged if the model remains advisory or lacks durable access to organizational state. A smaller model can generate a larger increase in I_A if architectural changes move it into authoritative and persistent workflows. Likewise, opacity or vendor dependence can cause a nominally more capable system to become less reconstructable.

The minimum directional assumption is local: within a preregistered architectural regime, increases in capability should not reduce effective interpreter capacity unless a measurable architectural countereffect offsets them. The theory therefore predicts mediation rather than a direct capability effect:

$$C(t) \rightarrow I_A(t) \rightarrow E(t),$$

not

$$C(t) \rightarrow E(t)$$

by definition.

This distinction separates ACTM from a simple technological-scaling argument. Historical work on technological progress has found regularities in cost and performance, but no single universal law applies to every technology or every relevant measure (Nagy et al. 2013; Farmer and Lafond 2016). Machine-learning compute has also grown rapidly across historical eras (Sevilla et al. 2022). ACTM treats such trajectories as candidate inputs to $C(t)$. They do not determine the organizational transition unless architecture converts capability into interpreter capacity.

Reconstructive organizational complexity

The second moving quantity is reconstructive organizational complexity, $K(t)$. K is not headcount, revenue, codebase size, or the number of organizational roles by itself. It is the effective burden placed on whatever machinery must reconstruct the organizational phenotype.

Candidate contributors include state dimensionality, interdependence among roles or modules, number and depth of control dependencies, turnover, environmental coupling, exception frequency, path dependence, heterogeneity of local contexts, and the amount of tacit coordination required to make formal rules operational. A large modular organization can therefore have lower reconstructive complexity than a smaller but densely entangled one.

This distinction matters because organizations can respond to rising complexity in two fundamentally different ways. They can increase interpreter capacity, or they can reduce the reconstructive burden through modularity, standardization, decomposition, markets, hierarchy, stable interfaces, or other architectures. ACTM treats these as partially substitutable routes because both change the same relative-sufficiency state.

For empirical work, I_A and K should be represented by preregistered positive indices with stable interpretation over the observation window. To avoid arbitrary units determining the ratio, define normalized quantities relative to a baseline or reference system,

$$\tilde{I}_A(t)=\frac{I_A(t)}{I_{A,0}}, \tilde{K}(t)=\frac{K(t)}{K_0}.$$

ACTM then defines **relative interpreter sufficiency** as

$$\chi(t)=\chi_0\frac{\tilde{I}_A(t)}{\tilde{K}(t)^{\gamma}}, \gamma>0,$$

where χ_0 is baseline relative sufficiency and γ measures how strongly increases in reconstructive complexity load the interpreter. If I_A and K are already dimensionless normalized indices, this reduces to the simpler expression

$$\chi(t)=\frac{I_A(t)}{K(t)^{\gamma}}.$$

The numerical value of χ is not assumed to be universal across domains. Its scientific role is relational: it locates a system relative to its own calibrated reconstructive burden and enables directional and threshold predictions.

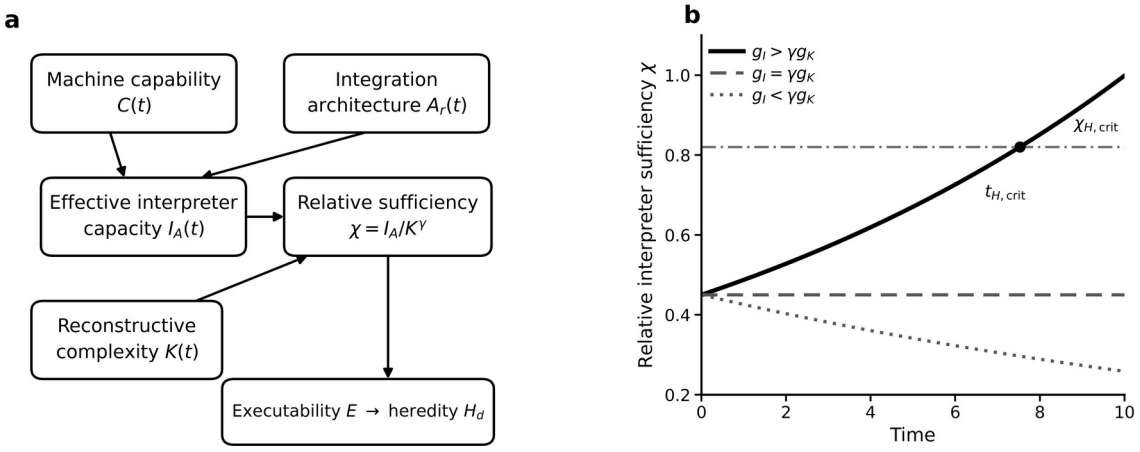


Fig. 1 The Computational Transition Model. **a** Machine capability affects developmental heredity only when integration architecture converts it into effective interpreter capacity. Relative interpreter sufficiency compares that capacity with reconstructive organizational complexity. **b** Relative sufficiency rises, remains stationary, or falls depending on the difference between interpreter-capacity growth and complexity-adjusted organizational growth. The plotted curves are analytic illustrations, not empirical data

The capability-complexity race

Let

$$g_I(t)=\frac{d\ln I_A}{dt}, g_K(t)=\frac{d\ln K}{dt}.$$

Because baseline normalization adds only constants in log space,

$$\frac{d \ln \chi}{d t} = g_I(t) - \gamma\, g_K(t).$$

This identity gives ACTM its three dynamical regimes.

Convergent regime. If

$$g_I(t) > \gamma\, g_K(t),$$

then $d \ln \chi / d t > 0$ and relative interpreter sufficiency rises. Machine-mediated reconstruction becomes progressively more capable relative to the burden placed on it.

Stationary regime. If

$$g_I(t) = \gamma\, g_K(t),$$

then χ is locally stationary. Machine capability can be improving rapidly in absolute terms while merely keeping pace with the reconstructive burden.

Divergent regime. If

$$g_I(t) < \gamma\, g_K(t),$$

then relative interpreter sufficiency falls. Absolute technological progress and declining reconstructive sufficiency can therefore occur simultaneously.

This is the central ACTM result. The relevant growth rate is not a universal law of computation; it is the difference between effective interpreter-capacity growth and complexity-adjusted reconstructive growth. The model therefore makes a stronger claim than "AI is improving quickly" but a weaker and more testable claim than "AI inevitably crosses a universal threshold."

The same result also generates a symmetry that is easy to miss. A system can move toward the convergent regime either by increasing g_I or by reducing g_K. Modularity, standardization, and decomposition can therefore substitute in part for faster machine capability. In ACTM, an organization that simplifies its reconstructive problem can approach the same developmental condition as one that increases interpreter capacity.

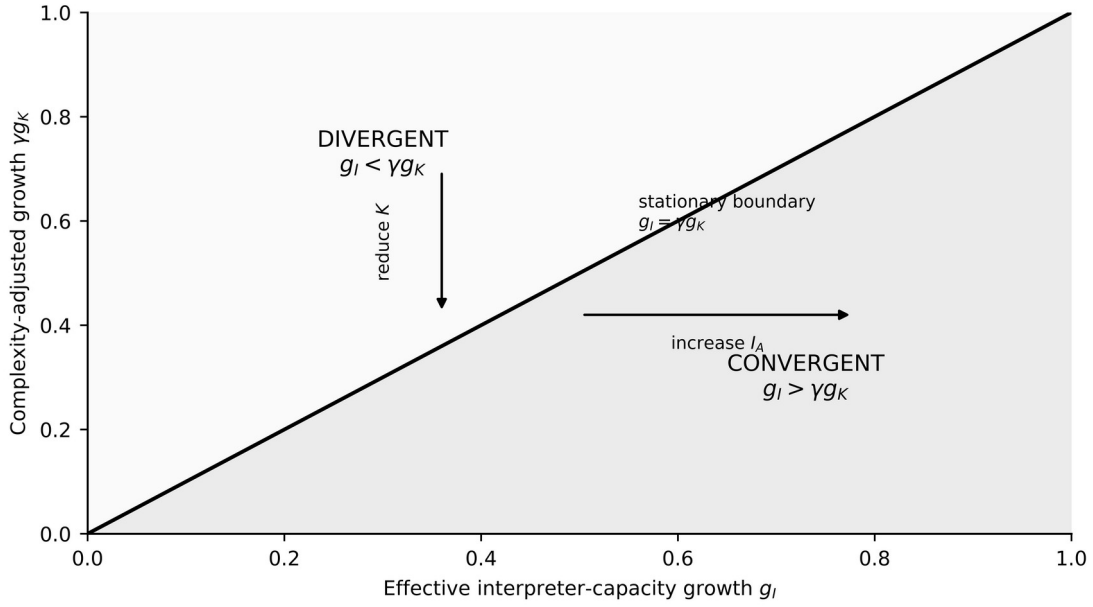


Fig. 2 ACTM phase map. The diagonal $g_I = \gamma\, g_K$ separates convergence from divergence. Systems below the diagonal have interpreter capacity growing faster than complexity-adjusted reconstructive burden; systems above it have complexity growing faster. Movement rightward can arise from increased effective interpreter capacity, while movement downward can arise from reductions in reconstructive complexity

Linking relative sufficiency to executability and heredity

ACTM becomes a theory of developmental transition when relative interpreter sufficiency changes executability. Suppose

$$E(t)=F\left[\chi(t)\right], F'(\chi)>0$$

within the preregistered operating regime. No particular functional form is required for the core directional result. Combining this response with the local heredity relation yields

$$\frac{dH_d}{dt}=\frac{dH_d}{dE}F'(\chi)\chi\left[g_I(t)-\gamma g_K(t)\right].$$

Whenever

$$\frac{dH_d}{dE}>0, F'(\chi)>0, \chi>0,$$

the sign of the temporal change in developmental heredity is

$$sign\left(\frac{dH_d}{dt}\right)=sign\left[g_I(t)-\gamma g_K(t)\right].$$

This is ACTM's principal directional prediction. In the convergent regime, developmental heredity should tend to rise, all else equal. In the divergent regime, increases in raw capability need not close the developmental gap and can coexist with stagnant or declining reconstruction fidelity.

The result creates several discriminations that a raw-capability theory cannot make. First, two organizations using the same model can occupy different regimes because their integration architectures differ. Second, a less capable model can generate greater developmental change if it is more deeply integrated into authoritative and persistent workflows. Third, reducing reconstructive complexity can raise developmental heredity even when machine capability is unchanged. Fourth, a sequence of model upgrades can fail to improve heredity when organizational complexity rises even faster.

The monotonicity assumption is itself testable. High-bandwidth machine coupling can amplify correlated error, reduce local autonomy, or make an organization more dependent on opaque vendor-side state. In those cases $F'(\chi)$ may fail to be positive over the relevant range, or I_A may have been mismeasured. ACTM should then be rejected or revised in that regime rather than rescued by redefining every integration as beneficial.

Thresholds and the transition surface

Directional change is weaker than threshold prediction. ACTM can make a prospective threshold claim only when a developmental criterion is independently justified. Let $H_{d,crit}$ denote a preregistered developmental-fidelity level relevant to the empirical system. From

$$H_d(E)=\exp\left\{-\lambda_d\left[a\sigma_s^2+b(1-E)^2\sigma_d^2+c\sigma_\varepsilon^2\right]\right\},$$

the minimum executability compatible with $H_d \geq H_{d,crit}$ is

$$E_{H,crit}=1-\sqrt{\frac{-\ln\left(H_{d,crit}\right)/\lambda_d-a\sigma_s^2-c\sigma_\varepsilon^2}{b\sigma_d^2}},$$

provided the target lies within the attainable fidelity range. If F is invertible, the corresponding relative-sufficiency threshold is

$$\chi_{H,crit}=F^{-1}\left(E_{H,crit}\right).$$

A logistic response is one testable special case,

$$E(\chi)=\frac{1}{1+\exp\left[-\beta\left(\ln\chi-\ln\chi_{50}\right)\right)},$$

which implies

$$\chi_{H,crit}=\chi_{50}\left(\frac{E_{H,crit}}{1-E_{H,crit}}\right)^{1/\beta}.$$

The logistic form is not a premise of ACTM. It is a parametric candidate that makes explicit what confirmatory work would have to estimate. Other monotone or nonmonotone response families can be compared out of sample.

The threshold should be understood as a **transition surface**, not a universal constant. Different systems can have different σ_s^2, σ_d^2, environmental floors, sensitivity coefficients, architectures, and complexity exponents. The same nominal model capability can therefore place one organization above its developmental threshold and another below it. This is not a weakness of ACTM; it is one of its central predictions.

Exponential growth as a benchmark, not a premise

ACTM does not require a Kurzweil-style or Moore-style law. Any observed paths for $I_A(t)$ and $K(t)$ can be inserted into the general equation. Exponential growth is analytically useful because it yields a closed-form benchmark.

Suppose

$$I_A(t)=I_{A0}2^{(t-t_0)/\tau_I},K(t)=K_0 2^{(t-t_0)/\tau_K}.$$

Then

$$\chi(t)=\chi_0 2^{(t-t_0)(1/\tau_I-\gamma/\tau_K)}.$$

If

$$\frac{1}{\tau_I}-\frac{\gamma}{\tau_K}>0,$$

the earliest modeled crossing of a registered relative-sufficiency threshold is

$$t_{H,crit}=t_0+\frac{\log_2\left(\chi_{H,crit}/\chi_0\right)}{1/\tau_I-\gamma/\tau_K}.$$

If the denominator is zero, relative sufficiency remains stationary. If it is negative, the threshold recedes despite absolute capability growth.

The more general condition is obtained by integration:

$$\int_{t_0}^{t_{H,crit}}\left[g_I(s)-\gamma\, g_K(s)\right)ds=\ln\left(\frac{\chi_{H,crit}}{\chi_0}\right).$$

This expression is the most general temporal form of ACTM. It allows punctuated model improvements, architectural discontinuities, periods of organizational simplification, or rising complexity burdens. A threshold date should therefore be reported only as a conditional model output with uncertainty, not as a calendar prediction of technological inevitability.

Testable propositions

ACTM is useful only if its components can fail separately. The following propositions distinguish the core model from stronger optional claims.

P1. Architecture mediation. Increases in raw capability C should increase effective interpreter capacity I_A more strongly when persistent access, authority, memory, portability, reproducibility, and auditability are present. If model upgrades predict reconstruction equally well regardless of architecture, the proposed mediation is weakened.

P2. Relative sufficiency over raw capability. χ should explain registered executability or developmental-heredity outcomes beyond raw C alone. If raw capability consistently predicts the outcomes while I_A / K^{γ} adds no out-of-sample information, ACTM's central relational claim is weakened.

P3. Capability-complexity race. Conditional on the stated assumptions, developmental heredity should tend to rise when $g_I > \gamma g_K$, remain approximately stationary near equality, and fail to rise or decline when the inequality reverses. Persistent violations reject the proposed dynamic mapping.

P4. Complexity-reduction equivalence. Interventions that reduce reconstructive complexity should increase χ and, where $F'(\chi) > 0$, should increase executability even when machine capability is held approximately constant. Strong modularization or standardization that lowers measured K without changing E challenges the operationalization of complexity or the response function.

P5. Prospective threshold. A preregistered $\chi_{H,crit}$ estimated on a training period or sample should predict later crossing of the associated developmental-heredity criterion out of sample. A threshold that fits retrospectively but fails prospective replication does not support the stronger ACTM claim.

P6. AI necessity is optional. At sufficiently high reconstructive complexity, machine-mediated interpreter capacity may become necessary for some registered developmental criterion. This is not a core ACTM requirement. Durable high-complexity human-only or non-AI systems that satisfy the same criteria through modularity, hierarchy, markets, standardization, or other architectures reject the necessity conjecture without rejecting ACTM.

Measurement and empirical strategy

The strongest initial tests should use bounded sociotechnical lineages rather than whole national or global economies. Software projects, protocol ecosystems, platforms, franchise-like systems, replicated operating systems, or other collectives with documented state, governance, turnover, automation, and lineage events make the relevant quantities more observable.

Measure capability independently of integration. C should be estimated from task-relevant machine capability before architectural variables are introduced. The measure can be multidimensional, but it must not be defined by whether organizational reconstruction succeeds. Otherwise the mediation claim becomes circular.

Measure effective interpreter capacity from causal access and persistence. I_A should reflect what the machine layer can actually interpret and execute within the organization. Candidate dimensions include persistent access to canonical state, authority to execute or veto, stable memory, machine-readable rules, reproducibility, portability, audit logs, identity continuity, and independence from private prompt histories or inaccessible vendor state. Validation should include intervention or reconstruction tasks rather than usage intensity alone.

Measure reconstructive complexity as burden, not size. A preregistered K index should capture the amount and structure of information and coordination that must be recovered to recreate the validated phenotype. Candidate measures can include dependency-graph complexity, state dimensionality after justified compression, role interdependence, exception density, coupling across modules, turnover sensitivity, and the amount of context that must be supplied externally. Multiple candidate constructions should be compared using held-out prediction rather than selected after observing the ACTM outcome.

Measure developmental heredity on phenotype. H_d should concern reconstruction of organizational form, not repository or document similarity. Domain-specific phenotypes can include governance topology, release discipline, role structure, permission architecture, control rules, conflict-resolution procedures, or other validated operational states. Syntactic fidelity should remain a separate manipulation check.

Exploit interventions on both sides of the ratio. ACTM is unusually testable because I_A and K can sometimes be manipulated independently. Interpreter-capacity interventions include persistent agents, policy-as-code, authority

changes, memory portability, or machine-readable governance. Complexity interventions include modularization, interface standardization, removal of dependencies, or decomposition of coupled processes. A two-sided intervention design can test whether raising the numerator and lowering the denominator have the predicted directional equivalence.

Prefer within-parent or sibling comparisons. Deep AI integration is endogenous. Organizations that are already formalized, successful, or easier to reconstruct may adopt machine systems more effectively. Stronger designs exploit staged rollouts, externally timed access changes, mandated tooling transitions, interpreter substitution, or sibling descendants from a common parent.

Test prospectively. The strongest ACTM evidence is not retrospective curve fitting. Estimate I_A, K, γ, and the F relation in a training window; preregister the regime or threshold; then test later lineages or later periods. A particularly strong result would identify a system below $\chi_{H,crit}$, observe an exogenous interpreter-capacity increase or complexity reduction, and correctly predict subsequent crossing of the registered developmental-heredity criterion.

Alternative explanations and failure modes

Several mechanisms can generate apparent ACTM-like patterns without supporting the theory.

Productivity without reconstruction. Artificial intelligence can increase output, decision quality, or labor substitution without changing the machinery of organizational heredity. Human-AI organization research already emphasizes the allocation of tasks and decision rights and the automation-augmentation tension (Puranam 2021; Raisch and Krakowski 2021). ACTM must explain registered reconstruction outcomes beyond ordinary productivity.

Algorithmic control without persistence. Algorithms can record, recommend, restrict, reward, and replace workers while remaining dependent on current personnel, permissions, or vendors. Present-day control is not equivalent to reconstructive persistence. A system can therefore have high algorithmic control and low I_A under ACTM.

Opacity and vendor dependence. More capable models can make organizational reconstruction harder when behavior depends on inaccessible weights, hidden instructions, private memory, unstable APIs, or changing provider policy. Such cases can produce $C \uparrow$ with little or negative change in I_A.

Control instability. Increased coupling can synchronize failure or amplify positive feedback. If deeper integration raises fragility rather than executability, $F'(\chi)$ may be nonmonotone. ACTM should then be restricted to the empirically supported regime rather than assuming that maximal integration is optimal.

Selection on adopters. Formalized organizations may both integrate AI more deeply and reproduce more faithfully. Cross-sectional associations among C, I_A, and H_d are therefore insufficient. Within-unit interventions, instrumental variation, or matched lineage designs are needed for causal interpretation.

Misidentified complexity. If K is treated as headcount or raw information volume, the theory can be made trivially wrong. Reconstructive complexity is the burden relevant to rebuilding the phenotype. A larger but highly modular system can legitimately have lower K than a smaller entangled one.

Relation to evolutionary individuality

ACTM is not itself a criterion of individuality. Crossing a developmental threshold establishes, at most, that a specified organization can be reconstructed with sufficient fidelity under a declared model. It does not establish a privileged system boundary, causal self-maintenance, recursive reproduction, or selection among lineages.

This separation matters because information-theoretic, causal, and evolutionary accounts of individuality impose additional requirements (Krakauer et al. 2020; Bourrat 2024; Clarke 2013, 2025). Major-transition theories further require evidence about dependence, specialization, conflict regulation, reproduction, and the level at which fitness differences are realized (Maynard Smith and Szathmáry 1995; Michod 1999; Godfrey-Smith 2009; West et al. 2015).

Within AEIH, ACTM provides the dynamic layer. AEIH asks whether persistent machine interpretation can increase developmental heredity by carrying interpreter function that otherwise depends on transient human routines. ACTM asks when the available machine interpreter is becoming more or less sufficient relative to reconstructive complexity. A system can therefore satisfy ACTM's directional prediction without qualifying as a higher-order evolutionary individual.

For a stronger Darwinian-individuality claim, developmental heredity would have to be combined with independent evidence. One compact bookkeeping expression is

$$I_D = 1(\Delta_B > \delta_B)\,1(H_d \geq H_{d,crit})\,Q\,1(S > S_{crit}),$$

where Δ_B is a registered boundary contrast, $Q \in \{0,1\}$ indicates recursive reproductive capacity, and S captures evidence that inherited collective traits predict differential lineage success. The expression is not proposed as a universal definition of life or individuality. It records the logical separation between a computational-developmental threshold and an evolutionary transition.

Bourrat's account of scaffold endogenization provides a useful neighboring idea: evolutionary transitions can involve formerly external scaffolds becoming internalized into a collective's developmental organization (Bourrat 2025). Machine-mediated interpretation is a candidate technological analogue, not an established instance. ACTM supplies a quantitative question for such cases: is the capacity of the increasingly internalized interpreter actually outrunning the complexity of what it must reconstruct?

Discussion

ACTM reframes a common forecasting question. Instead of asking when artificial intelligence becomes "powerful enough," it asks powerful enough relative to what. The answer depends on the organization, its architecture, and the rate at which its reconstructive burden changes. This turns a one-dimensional capability forecast into a relational dynamical model.

The key variable χ has three advantages. First, it prevents raw computational capability from being mistaken for organizational executability. Second, it allows complexity-reduction strategies to enter the same model as capability gains. Third, it makes negative cases informative. A decade of rapidly improving models can coexist with stagnant relative sufficiency if integration remains shallow or reconstructive complexity rises equally quickly.

The model also suggests a different interpretation of technological thresholds. Thresholds need not be global moments at which a technology suddenly transforms every organization. They can be system-specific surfaces in a state space. A highly standardized software platform may cross its reconstructive threshold early, while a loosely documented institution with opaque dependencies remains below it despite access to the same machine capability. The transition can therefore be distributed across organizations rather than synchronized across an economy.

Several limitations are substantial. I_A and K are latent constructs and may be difficult to identify cleanly. The exponent γ need not be constant across regimes. The mapping F can be nonlinear, hysteretic, or nonmonotone. Organizational reconstruction may exhibit path dependence and phase changes that invalidate local first-order approximations. Environmental variation can dominate both machine capacity and organizational complexity. These limitations argue for preregistered domain-specific measurement rather than against the model itself: ACTM is intended to specify what must be measured for the transition claim to become testable.

A further limitation is that the current model treats relative sufficiency as effectively scalar. Real organizations contain multiple interpreter subsystems and multiple complexity bottlenecks. One subsystem can be above threshold while another remains below it. A natural extension is therefore a vector or network ACTM in which I_A and K are matched across functional domains and system-level reconstruction depends on bottlenecks, redundancy, and coupling among them. The scalar model should be treated as the minimal theory required to generate falsifiable first predictions.

The most important empirical advance would be prospective intervention. If a preregistered estimate of χ correctly predicts which systems will increase developmental heredity after an interpreter-capacity intervention, and if complexity-reduction interventions produce the corresponding predicted effect, ACTM would explain a dimension of organizational change that raw model capability does not. If those predictions fail, the theory should be revised or rejected without changing the fact that artificial intelligence may remain economically transformative.

Conclusion

The Computational Transition Model proposes that the relevant technological race is not computation against time but interpreter capacity against reconstructive complexity. Raw machine capability matters only insofar as architecture

converts it into persistent capacity to interpret and execute organizational state. Organizational complexity matters only insofar as it increases the burden of reconstructing the validated organizational phenotype.

The model compresses this relation into relative interpreter sufficiency,

$$\chi(t) = \chi_0 \frac{\tilde{I}_A(t)}{\tilde{K}(t)^\gamma},$$

with dynamics

$$\frac{d \ln \chi}{dt} = g_I(t) - \gamma\, g_K(t).$$

When relative sufficiency monotonically increases executability and executability increases developmental heredity, the same difference determines the predicted direction of change in reconstruction fidelity:

$$sign\left(\frac{d H_d}{dt}\right) = sign\left[g_I(t) - \gamma\, g_K(t)\right).$$

The result is deliberately conditional. ACTM does not predict a universal singularity, a fixed amount of compute at which organizations become individuals, or a calendar date at which an economy becomes an organism. It predicts something narrower and empirically stronger: systems should move toward a specified developmental threshold when effective interpreter capacity persistently outruns complexity-adjusted reconstructive burden, and away from it when the inequality reverses.

That prediction can be tested before any stronger claim about evolutionary individuality is entertained. In this sense ACTM is a path model: it specifies how technological capability can become organizational developmental capacity, and the conditions under which that capacity becomes increasingly sufficient for reconstructing complex systems through time.

Statements and Declarations

Funding. No external funding was received for this work.

Competing interests. The author is Founder and CEO of Clockwise Capital, LLC, an SEC-registered investment adviser. The firm had no role in the design, analysis, or writing of this article. The author declares no other competing interests.

Author contributions. The author conceived the Computational Transition Model, developed the formal derivations and theoretical predictions, prepared the figures, and wrote and revised the manuscript.

Ethics approval. Not applicable. No human participants, human data, animals, or animal data were used.

Data and code availability. No empirical datasets were generated or analyzed for this theoretical article. The figures contain conceptual and analytic model output only; no empirical data are plotted.

Generative AI disclosure. OpenAI ChatGPT and Anthropic Claude were used as drafting, restructuring, review, and code-assistance tools during manuscript development. The author independently reviewed and revised all substantive claims and takes full responsibility for the final manuscript. No generative-AI image model was used to create figures or data; figures were generated deterministically from the stated theoretical model and conceptual layout.

References

Ashby WR (1956) *An Introduction to Cybernetics*. Chapman & Hall, London

Beer S (1972) *Brain of the Firm: A Development in Management Cybernetics*. Allen Lane, London

Bourrat P (2024) Adding causality to the information-theoretic perspective on individuality. *Eur J Philos Sci* 14:9. https://doi.org/10.1007/s13194-023-00566-1

Bourrat P (2025) Evolutionary transitions in individuality by endogenization of scaffolded properties. *Br J Philos Sci* 76(2):269–294. https://doi.org/10.1086/719118

Clarke E (2013) The multiple realizability of biological individuals. *J Philos* 110(8):413–435. https://doi.org/10.5840/jphil2013110817

Clarke E (2025) *The Units of Life: Kinds of Individual in Biology*. Oxford University Press, Oxford. https://doi.org/10.1093/9780191948008.001.0001

Conant RC, Ashby WR (1970) Every good regulator of a system must be a model of that system. *Int J Syst Sci* 1(2):89–97. https://doi.org/10.1080/00207727008920220

Farmer JD, Lafond F (2016) How predictable is technological progress? *Res Policy* 45(3):647–665. https://doi.org/10.1016/j.respol.2015.11.001

Godfrey-Smith P (2009) *Darwinian Populations and Natural Selection*. Oxford University Press, Oxford. https://doi.org/10.1093/acprof:osobl/9780199552047.001.0001

Hodgson GM, Knudsen T (2008) Information, complexity and generative replication. *Biol Philos* 23(1):47–65. https://doi.org/10.1007/s10539-007-9073-y

Hodgson GM, Knudsen T (2010) *Darwin's Conjecture: The Search for General Principles of Social and Economic Evolution*. University of Chicago Press, Chicago. https://doi.org/10.7208/chicago/9780226346922.001.0001

Krakauer DC, Bertschinger N, Olbrich E, Flack JC, Ay N (2020) The information theory of individuality. *Theory Biosci* 139(2):209–223. https://doi.org/10.1007/s12064-020-00313-7

Maynard Smith J, Szathmáry E (1995) *The Major Transitions in Evolution*. W. H. Freeman, Oxford

Michod RE (1999) *Darwinian Dynamics: Evolutionary Transitions in Fitness and Individuality*. Princeton University Press, Princeton, NJ

Nagy B, Farmer JD, Bui QM, Trancik JE (2013) Statistical basis for predicting technological progress. *PLoS ONE* 8(2):e52669. https://doi.org/10.1371/journal.pone.0052669

Nelson RR, Winter SG (1982) *An Evolutionary Theory of Economic Change*. Belknap Press of Harvard University Press, Cambridge, MA

Puranam P (2021) Human-AI collaborative decision-making as an organization design problem. *J Organ Des* 10:75–80. https://doi.org/10.1007/s41469-021-00095-2

Raisch S, Krakowski S (2021) Artificial intelligence and management: The automation-augmentation paradox. *Acad Manag Rev* 46(1):192–210. https://doi.org/10.5465/amr.2018.0072

Sevilla J, Heim L, Ho A, Besiroglu T, Hobbhahn M, Villalobos P (2022) Compute trends across three eras of machine learning. In: *2022 International Joint Conference on Neural Networks (IJCNN)*, pp 1–8. https://doi.org/10.1109/IJCNN55064.2022.9891914

West SA, Fisher RM, Gardner A, Kiers ET (2015) Major evolutionary transitions in individuality. *Proc Natl Acad Sci USA* 112(33):10112–10119. https://doi.org/10.1073/pnas.1421402112