

The Simulative-Elasticity Model (SEM)

—

*A Systems Theory of Intelligence:
Navigating the Space Between Habit and Adaptation*

Abdul Martinez

Didomi Research
didomiresearch.org

Working Paper — February 2026

v2.4

Abstract

Contemporary measures of intelligence, in both human psychometrics and artificial intelligence benchmarking, remain anchored to static, capacity-based paradigms. Systems are evaluated by how much they know rather than how well they adapt when what they know becomes obsolete. This paper introduces the Simulative-Elasticity Model (SEM), a dynamic systems framework that redefines intelligence as regulatory phase-control: the mechanical ability to shift between automated habit execution and deep, energy-intensive reasoning in response to environmental volatility. The framework identifies five functional states, Architectural, Reconstructive, Simulative, Liquid, and Elastic Intelligence, and proposes that the critical variable governing adaptive behavior is the Delta (Δ), a continuous divergence signal between internal prediction and external reality. We introduce the Script-to-Capability Pipeline showing how Liquidated scripts become encapsulated tools that compose into observable expertise, extend the model to account for gradient Delta states and calibration errors, introduce compound operations, most notably the Inverse Simulation loop for diagnostic and root cause reasoning, draw a strict boundary between Acquired Capacity and Raw Intelligence with attention to individual differences and neurodivergent profiles, outline preliminary conditioning protocols for improving phase-transition speed, propose a SEM-guided path to artificial general intelligence grounded in regulatory architecture rather than parameter scaling, and decompose Emotional Intelligence into Acquired Capacity (social scripts) and Raw Intelligence (emotional phase-control). SEM is positioned against existing frameworks including Kahneman's dual-process theory, Friston's Free Energy Principle, Anderson's ACT-R architecture, and Goleman's Emotional Intelligence model, demonstrating that it offers a more mechanistically complete account of adaptive intelligence across biological and artificial systems.

Keywords: *Simulative-Elasticity Model, adaptive intelligence, phase-control, cognitive elasticity, Delta divergence, predictive processing, habit automation, script-to-capability pipeline, toolchain fragility, dependency collapse, calibration error, inverse simulation, abductive reasoning, artificial general intelligence, AI alignment, emotional intelligence, emotional Delta signals, neurodivergence, individual differences, dual-process theory, free energy principle, working memory, cognitive flexibility*

1. Introduction: The Flaw in How We Measure Minds

For over a century, society has measured intelligence like a bucket. We ask how much data a person can hold, how many facts they can memorize, or how many mathematical rules they can apply. In the world of artificial intelligence, we do the exact same thing, we judge a machine by the billions of parameters it contains.

This creates a dangerous illusion. It leads us to believe that the system with the most data is the smartest. But reality is not a static test. Reality is a volatile, constantly changing environment. This is why highly educated experts often freeze or fail when a crisis hits. It is also why massive AI models completely break down when faced with a problem slightly outside their training data. We call this **Brittle Expertise**: a system might be incredibly fast and efficient in a familiar room, but it shatters the moment the furniture is rearranged.

This observation is not new in isolation. Taleb (2012) has written extensively on the fragility of optimized systems, and Kahneman (2011) has catalogued the failure modes of heuristic cognition. However, what the field lacks is a unified mechanical framework that describes *how* an intelligent system should regulate its own transitions between efficiency and adaptation. The Simulative-Elasticity Model (SEM) offers this paradigm.

Intelligence, under SEM, is not a storage tank of knowledge. Intelligence is a dynamic control system. It is the mechanical ability to shift gears between deep, high-energy reasoning and fast, low-energy habits, and, most critically, to know *when* each gear is appropriate.

2. The Ignition Switch: The Delta (Δ)

Every intelligent system operates on a continuous loop of prediction and observation. You are always guessing what will happen next, and reality is always telling you if you were right. This loop is well-established in computational neuroscience under the banner of predictive processing (Clark, 2013; Friston, 2010). SEM builds on this foundation but reorients the focus from prediction *accuracy* to the *regulatory response* triggered by prediction failure.

The most critical component of the SEM framework is the **divergence monitor**. We call this the **Delta (Δ)**.

The Delta is the mathematical difference between your internal expectation and the external reality. It is the anomaly detector. Formally:

$$\Delta(t) = || E(t) - O(t) ||$$

Where $E(t)$ is the system's expectation at time t , and $O(t)$ is the observed state of the environment at time t . The norm captures the magnitude of divergence across all relevant dimensions of the problem space.

The Foundational Example

Imagine you are typing your password into your computer. You have done this thousands of times. You expect the screen to unlock instantly. But this time, the screen flashes red and says "Access Denied." Your expectation was an unlocked screen. The reality is a locked screen. That sudden gap between expectation and reality is the Delta.

When the Delta is low or zero, the system can relax and rely on automated habits. When the Delta spikes, it is a blaring alarm, it tells the system that the current rules are broken and immediate adaptation is required.

3. The Five States of Intelligence

The SEM framework categorizes intelligence into five functional states. A truly intelligent agent does not just possess one of these traits, it rotates through all five of them as the environment demands. These are not personality types, cognitive styles, or static abilities. They are *operating modes* of a single adaptive engine.

I. Architectural Intelligence (The Mental Workbench)

Architectural Intelligence is the raw capacity of your active workspace, your mental bandwidth. Think of this as a physical workbench in a garage. If you have a tiny workbench, you can only fix one small part at a time. If you have a massive, well-lit workbench, you can lay out an entire engine block and see how every single gear, belt, and piston connects to the others.

Function: Holding multiple, interdependent variables in active cognition simultaneously without representational decay.

This maps closely to the construct of working memory capacity as described by Baddeley (2000), though SEM treats it not as a trait but as a configurable state variable, one whose effective bandwidth is modulated by the current Delta load.

Example: A software developer with high Architectural Intelligence can picture the database structure, the user interface, and the security protocols all at once, understanding how changing a single line of code will ripple through the entire system.

II. Reconstructive Intelligence (The Structural Builder)

Reality rarely gives you all the information you need. Reconstructive Intelligence is the ability to take a few sparse clues and build a complete, high-resolution picture.

Function: Inferring latent structure from partial observational data, connecting the dots to see the hidden topology of a problem.

Example: A veteran mechanic is listening to a car engine. They hear a faint rhythmic ticking and smell a specific type of burning oil. They cannot see inside the steel engine block. However, they use those two tiny clues to completely reconstruct a mental 3D model of a failing valve seal. They built a complete picture from incomplete data.

III. Simulative Intelligence (The Temporal Projector)

Once you have built a model of a problem, you need to test it before you take action in the real world. Simulative Intelligence is your internal physics engine. This is closely related to what cognitive scientists call mental simulation or mental modeling (Hegarty, 2004; Craik, 1943), but SEM embeds it within a regulatory cycle rather than treating it as an isolated competence.

Function: Projecting a constructed model forward in time to evaluate the causal consequences of potential actions before committing resources.

Example: An architect is designing a bridge. Before pouring any concrete, they run a computer simulation to see how the bridge will react to a Category 5 hurricane. Simulative Intelligence is doing that in your head, it is the chess player looking four moves ahead to see the consequences of moving their knight.

IV. Liquid Intelligence (The Optimizer)

Deep reasoning and simulation take a massive amount of metabolic energy. If you had to consciously think about every single action you took all day, your brain would burn out by noon. To survive, a system must automate its successes. Liquid Intelligence is the ability to take a proven solution and turn it into an effortless habit. We call it “Liquid” because it flows without resistance.

Function: Conserving metabolic and computational energy by compiling complex reasoning chains into automated execution scripts.

This process is well-documented in skill acquisition research. Fitts and Posner (1967) described the transition from cognitive to autonomous phases of motor learning. In SEM terminology, Liquidation is the *successful* compilation of a Simulative-Reconstructive output into a low-energy execution routine.

Example: Think about the first time you drove a car. You had to consciously think about the pedals, the steering wheel, and the mirrors. It was exhausting. Ten years later, you can drive all the way to work while having a complex conversation, and you do not even remember the drive. The process of driving has been successfully “liquidated” into a background habit.

Liquidation as Tool-Building: The Script-to-Capability Pipeline

The existing description of Liquid Intelligence captures the compression step, complex reasoning compiled into an automated script, but stops short of what happens next. In practice, Liquidated scripts do not merely save energy. They **become functional tools** that unlock higher-order operations, compound into observable capability, and ultimately constitute what the outside world recognizes as expertise, skill, or competence.

The mechanism operates as a three-stage pipeline:

Stage 1: Compilation (Script). A successful Reconstructive-Simulative cycle produces a solution to a recurring problem. Through repetition and reinforcement, this solution is compiled into a low-energy execution routine, the Liquidated script. This is the stage the paper has already described: the novice driver becoming an automatic driver.

Stage 2: Encapsulation (Tool). Once compiled, the script becomes *encapsulated*, it can be invoked as a single unit without re-deriving its internal logic. At this point, the script is no longer just a habit; it is a **cognitive tool**. The person who has Liquidated driving does not merely “drive automatically.” They now possess a tool called “driving” that can be deployed as a *component* in higher-order operations: commuting to work, delivering goods, navigating an evacuation route, teaching someone else to drive. Each of these operations requires the driving tool as a prerequisite, but none of them *are* the driving tool, they are built *on top of it*. The tool has become infrastructure.

Stage 3: Composition (Capability). Multiple encapsulated tools stack and interlock to produce compound capabilities that no individual tool could achieve alone. A surgeon's expertise is not a single Liquidated script, it is the composition of dozens of encapsulated tools: hand-eye coordination, anatomical spatial reasoning, instrument handling, tissue recognition, hemostatic protocols, patient monitoring. Each was independently compiled through the Reconstructive-Simulative cycle and separately Liquidated. Together, they compose into what we observe as "surgical competence." This is Acquired Capacity, not as a flat library of independent habits, but as a **toolchain**: a structured, interdependent stack of encapsulated scripts where higher-order tools depend on lower-order tools as prerequisites.

The Misattribution Problem

The Script-to-Capability Pipeline explains a persistent error in how intelligence is evaluated. When we observe a surgeon performing a complex procedure with apparent effortlessness, or a programmer building a system with fluid speed, we are observing the *output of a deep toolchain*. The temptation, in education, in hiring, in AI benchmarking, is to mistake this output for intelligence. But toolchain depth is Acquired Capacity. It tells you how many scripts the system has successfully compiled and stacked. It tells you nothing about the system's ability to *disassemble the stack* when the environment changes.

This is where the connection to Section 6 becomes explicit. The distinction between Acquired Capacity and Raw Intelligence is not merely a theoretical boundary, it is a description of *what the toolchain can and cannot do*. A toolchain can execute. It can even execute with breathtaking speed and precision. What it cannot do is *question its own composition*. It cannot recognize when one of its lower-order tools has become invalid due to environmental shift, pull that tool out of the stack, and rebuild the dependent operations on a new foundation. That is the work of the Elastic clutch and the Reconstructive-Simulative cycle. The toolchain is the product; the regulatory architecture is the factory.

Toolchain Fragility

The pipeline also reveals a specific vulnerability: **dependency collapse**. Because higher-order capabilities are built on lower-order tools, a Delta spike that invalidates a foundational tool can cascade through the entire stack. A surgeon who has Liquidated hand-eye coordination into an unconscious tool will experience not merely a skill disruption but a *capability collapse* if that tool is compromised, by injury, by unfamiliar

equipment, or by a novel surgical context that changes the spatial dynamics the tool was calibrated for. The entire toolchain above the compromised layer becomes unreliable simultaneously.

This is the mechanical explanation for why experts sometimes fail more catastrophically than novices in novel environments. The novice has few Liquidated tools and a shallow stack; they are inefficient but have little to collapse. The expert has a deep, highly optimized toolchain that amplifies performance in familiar contexts but creates a larger blast radius when a foundational script breaks. Taleb (2012) described this pattern as the fragility of over-optimization; SEM now provides the specific mechanism, dependency collapse in the Liquidated toolchain, and the specific remedy: an Elastic clutch sensitive enough to detect when a foundational tool has been compromised, paired with sufficient Architectural bandwidth to hold the decomposed stack while the Reconstructive-Simulative cycle rebuilds the affected layer.

V. Elastic Intelligence (The Clutch)

This is the most important concept in the entire framework. Elastic Intelligence is the **regulator**. It is triggered directly by the Delta (Δ) alarm.

To understand this, consider the mechanical metaphor of a car's clutch. In a manual transmission vehicle, the engine is constantly spinning. The wheels need to turn. If the engine is permanently welded to the wheels, you will crash the moment you need to stop. The clutch is the mechanism that physically disconnects the engine from the wheels so you can safely change gears or hit the brakes.

Elastic Intelligence is the **cognitive clutch**. It is the ability to instantly disconnect your automated habits from your actions when the environment changes.

Function: Interrupting automated execution to re-engage high-energy reasoning states when the Delta exceeds the system's threshold for acceptable divergence.

Example: You are driving that familiar route to work on pure autopilot (Liquid state). Suddenly, a truck swerves into your lane (Delta spike). Elastic Intelligence is the mental clutch being pushed in, it instantly kills your autopilot, snaps your attention back to the present moment, and forces you to rapidly simulate a new escape route.

Compound Operations: The Inverse Simulation Loop

The five states described above are presented individually for clarity, but in practice they frequently operate in tightly coupled pairs or chains. The most important compound operation in the SEM framework is **Inverse Simulation**, the process by which a system reasons backward from an observed result to identify its root cause.

At first glance, root cause analysis appears to require the Simulative engine to run in reverse, to take an outcome and rewind it to its origin. But this is mechanically unnecessary. Inverse Simulation is not a separate state; it is a **convergence loop** between two existing states: Reconstructive Intelligence and Simulative Intelligence, governed by a local Delta signal.

The Mechanism

The process unfolds in three iterative phases:

Phase 1 — Abductive Reconstruction. The system observes an anomalous result, the server crashed, the engine seized, the patient collapsed. Reconstructive Intelligence takes the sparse available evidence (error logs, sounds, symptoms) and generates a set of **candidate causal models**: hypotheses about what hidden configuration of the world could have produced the observed outcome. This is abductive reasoning in the classical Peircean sense (Peirce, 1903), inference to the best explanation from incomplete data.

Phase 2 — Forward Validation. The system takes each candidate root cause and feeds it into the Simulative engine, which runs it *forward* through its internal physics model: “If the valve seal began failing at timestamp X, would the forward propagation of that failure produce the specific ticking pattern, oil burn signature, and thermal profile I am currently observing?” The Simulative engine does not need to run backward. It only ever projects forward. The “backward” character of root cause analysis emerges from the Reconstructive state’s ability to generate starting conditions that precede the observed result.

Phase 3 — Local Delta Evaluation. The system compares the Simulative output against the actual observed result. This generates a **local Delta (Δ_0)**, the divergence between “what my candidate cause would have produced” and “what I actually see.” If Δ_0 is low, confidence in the candidate root cause increases. If Δ_0 remains high, the candidate is discarded and Reconstructive Intelligence generates the next hypothesis. The loop iterates until convergence.

Formally, we can express this as an iterative minimization:

$$\Delta_0(n) = || \text{Sim}(R\Box) - O ||$$

Where $R\Box$ is the n th reconstructive hypothesis, $\text{Sim}(R\Box)$ is the forward simulation of that hypothesis, and O is the observed outcome. The system iterates to minimize Δ_0 across candidate models.

The Chess Analogy

Consider a chess player who arrives at a board mid-game and sees a devastating position. They ask: “What sequence of moves could have led here?” They are not rewinding the game, chess has no rewind function. Instead, they propose candidate move sequences (Reconstructive), play each one forward in their mind (Simulative), and check whether the resulting board matches the position in front of them (Local Delta). The sequence that converges is the diagnosis.

The Engineering Example

A structural engineer is called to a building where a support column has buckled. The column is the observed result. The engineer cannot see the forces that acted on it over time. Reconstructive Intelligence generates candidates: foundation settling, asymmetric load redistribution from a recent renovation, corrosion from water infiltration, or a material defect from original construction. For each candidate, Simulative Intelligence runs a forward stress analysis: “If water infiltration reduced the effective cross-section by 30% over eight years, would the resulting stress concentration produce a buckle at this specific location and angle?” The candidate whose forward projection most closely matches the observed failure mode, including the buckle geometry, crack propagation pattern, and deformation profile, is identified as the root cause.

Implications for the Framework

The Inverse Simulation loop has three significant consequences for SEM.

First, it reveals that the five states are more interdependent than they appear in isolation. Reconstructive and Simulative Intelligence are not merely sequential stages in a pipeline; in diagnostic reasoning, they are locked in a tight feedback cycle that iterates until convergence. This suggests that the SEM framework should be understood not as a linear state machine but as a *recurrent regulatory network* in which states can call each other repeatedly within a single cognitive episode.

Second, it provides a more precise account of the Architectural demands of diagnostic reasoning. The Inverse Simulation loop requires the Architectural

workspace to hold three things simultaneously: the observed result, the current candidate hypothesis, and the Simulative output for comparison. Each iteration may also require retaining traces of previously rejected hypotheses to avoid redundant search. This is why expert diagnosticians, in medicine, engineering, and forensics, require substantial Architectural bandwidth, and why novices with identical Reconstructive and Simulative abilities but smaller workspaces produce slower and less reliable diagnoses.

Third, it sharpens the hallucination diagnosis for artificial systems. A large language model performing “root cause analysis” is running only Phase 1 of the loop, Abductive Reconstruction. It generates causal narratives that are structurally plausible because its training data contains millions of diagnostic patterns. But it never executes Phase 2 (forward simulation to test the candidate) or Phase 3 (local Delta evaluation against observed reality). The result is explanations that *sound* causally coherent but have never been tested against a physics model. The LLM is pattern-completing the abductive step without closing the validation loop. In SEM terms, it has Reconstructive power without Simulative grounding, the same mechanical deficit identified in Section 4, but now visible in its most operationally dangerous form: confident causal explanations that are unfalsified and potentially unfalsifiable within the system’s own architecture.

Relation to Existing Literature

The Inverse Simulation loop maps to several established concepts in the reasoning literature. Harman (1965) described “inference to the best explanation” as the fundamental logic of abductive reasoning. Josephson and Josephson (1994) formalized abductive inference as hypothesis generation followed by evaluation, which corresponds directly to the Reconstructive-Simulative pairing in SEM. In the engineering domain, the loop is closely related to model-based diagnosis (de Kleer & Williams, 1987), where candidate fault models are tested against observed system behavior through forward simulation.

SEM’s contribution is to embed these diagnostic reasoning patterns within a broader regulatory framework. Inverse Simulation is not a standalone cognitive competence; it is a compound operation that emerges from the interaction of states that are themselves governed by the Delta signal and the Elastic clutch. The decision to *initiate* a diagnostic loop, rather than defaulting to a Liquidated heuristic (“it’s probably the same thing as

last time”), is itself an Elastic function, triggered when the observed failure is sufficiently novel that cached explanations are unreliable.

4. Diagnosing Systemic Failure

When we understand these five states, we can precisely diagnose why intelligent systems fail. Each failure mode corresponds to a specific dysfunction in the SEM cycle.

The “Old Dog” Syndrome (Rigidity)

This occurs when a person or organization is locked in the Liquid state with a broken Elastic clutch. They have highly optimized habits, but when the environment changes, they just keep executing the old plan faster, completely unable to adapt. In organizational theory, this is closely related to what Leonard-Barton (1992) called *core rigidities*: the very capabilities that drove past success become the barriers to future adaptation.

Analysis Paralysis

This occurs when a system is stuck in the Architectural and Simulative states. It constantly builds complex models and simulates endless “what-if” scenarios but never pulls the trigger and Liquidates a decision into action. The Elastic clutch is too *sensitive*, it keeps interrupting nascent habits before they can consolidate.

Hallucination

This is increasingly common in modern AI systems. The system has high Reconstructive power but zero Simulative grounding. It connects dots that do not exist because it has no internal physics engine to test whether its reconstructions actually correspond to reality. Bender et al. (2021) have described large language models as “stochastic parrots”; SEM provides a more precise mechanical diagnosis, these systems possess powerful pattern completion (Reconstructive state) without any mechanism for causal verification (Simulative state).

5. The Delta Gradient: Calibration, Drift, and the Messy Middle

The preceding sections have described the Delta primarily in binary terms, low or spiking. In practice, however, the Delta operates along a continuous gradient, and the most consequential failures of intelligence occur not at the extremes but in the

intermediate zone where novelty is moderate and the correct regulatory response is ambiguous.

5.1 The Continuous Delta Spectrum

We can formalize three broad regimes of the Delta signal:

Delta Regime	Signal Character	Optimal Response	Failure Mode
Low ($\Delta \approx 0$)	Stable, predictable environment	Liquid execution	Complacency if sustained too long
Moderate ($0 < \Delta < \theta$)	Subtle drift, ambiguous signals	Heightened monitoring; selective habit audit	False confidence or premature override
High ($\Delta \geq \theta$)	Acute novelty or crisis	Full Elastic disengage; Reconstructive-Simulative cycle	Panic, freeze, or rigid habit repetition

Here, θ represents the system's **clutch threshold**, the Delta value at which the Elastic mechanism fully engages and overrides Liquid execution. This threshold is not universal; it is a tunable system parameter, and its calibration is a primary determinant of adaptive success.

5.2 The Moderate-Delta Problem

The moderate Delta zone is where most real-world intelligent failure occurs. This is the zone of slow drift, gradual environmental change, and ambiguous signals, the market that is declining at 2% per quarter rather than crashing overnight, the relationship that is eroding rather than exploding, the codebase that is accumulating technical debt rather than throwing errors.

In this zone, the system faces a genuine **regulatory dilemma**: the Delta signal is nonzero but does not clearly exceed θ . The system must decide whether to trust its Liquidated habits (which are still *mostly* working) or to incur the metabolic cost of engaging Reconstructive and Simulative states to audit whether the environment has fundamentally shifted.

This dilemma maps to what Simon (1956) described as *bounded rationality*, the system does not have infinite resources to monitor every parameter simultaneously, and so it must allocate its attention under uncertainty. SEM provides a mechanical account of *how* that allocation is governed: through the continuous monitoring of $\Delta(t)$ against the threshold θ .

5.3 Calibration Errors

If θ is the clutch threshold, then **calibration** is the process by which the system sets and adjusts θ over time. Two systematic miscalibrations are particularly destructive:

Hypersensitive Calibration (θ too low)

When the threshold is set too low, the system engages the Elastic clutch at the slightest perturbation. Every minor fluctuation triggers a full Reconstructive-Simulative cycle. In biological systems, this manifests as chronic anxiety, hypervigilance, or the organizational equivalent, the company that pivots strategy every quarter in response to normal market noise. The system burns enormous metabolic resources re-examining threats that would be better handled by Liquid execution, and it never consolidates habits long enough to achieve efficiency.

Hyposensitive Calibration (θ too high)

When the threshold is set too high, the system requires catastrophic divergence before it will interrupt its habits. This is complacency at the mechanical level. The Delta accumulates incrementally, but the clutch does not engage until the environment has shifted so far that the Liquidated scripts are producing catastrophic output. In human terms, this is the experienced professional who dismisses early warning signs because “I’ve seen this before”, until the situation has deteriorated beyond recovery. In organizational terms, it is Kodak continuing to optimize film production while digital photography consumed the market (Lucas & Goh, 2009).

5.4 Drift Detection and the Monitoring Duty Cycle

A well-calibrated system does not merely react to Delta spikes; it maintains a **background monitoring duty cycle** that continuously samples $\Delta(t)$ even during Liquid execution. This is analogous to the experienced driver who is on autopilot but still glancing at mirrors every few seconds, a low-energy environmental audit that does not require full Elastic disengage but ensures the system is not accumulating undetected drift.

We propose that the monitoring duty cycle is itself a trainable parameter. Systems with high adaptive fitness maintain a higher sampling rate even in stable environments, creating a *prophylactic buffer* against slow drift. This is metabolically expensive in the short term but catastrophe-preventive in the long term, and it represents a key mechanism by which Raw Intelligence can be systematically conditioned (see Section 7).

6. Raw Intelligence vs. Acquired Capacity

A persistent flaw in both human psychometrics and computer science is the conflation of “knowing” with “thinking.” Society and engineers frequently mistake a highly optimized system for a highly intelligent one. To resolve this, the SEM framework draws a strict mechanical boundary between **Acquired Capacity** and **Raw Intelligence**.

These are not separate types of minds, but rather a description of how a system allocates its resources relative to environmental volatility, the Delta (Δ).

Acquired Capacity (The Liquidated Library)

Acquired Capacity is the accumulation of successful Liquid states, the vast library of automated scripts, hardened habits, and compiled policies a system has built over time.

Mechanism: When an environment is stable ($\Delta \approx 0$), the system relies almost entirely on Acquired Capacity. It executes stored solutions with maximum metabolic and computational efficiency.

Limitation: Acquired Capacity is inherently brittle. Because it is optimized for a specific, known environment, it possesses no internal mechanism for self-correction. It is the vehicle driving in overdrive on a perfectly paved, straight highway.

Raw Intelligence (Adaptive Phase-Control)

Raw Intelligence is the mechanical capacity of the system to navigate the full SEM cycle when the environment fractures. It is the system’s proficiency in **state arbitration**, specifically, the sensitivity of the Elastic clutch to break a habit, the volumetric bandwidth of the Architectural workspace to hold novel data, and the power of the Reconstructive and Simulative engines to build a new causal model.

Mechanism: When the environment introduces severe novelty or systemic shocks (High Δ), Acquired Capacity fails. Raw Intelligence must take over. It is the off-road suspension and manual gearbox required when the highway abruptly ends.

Applied Profiles

To observe this distinction in practice, we can map systems based on their ratio of Raw Intelligence to Acquired Capacity:

The Human Child (High Raw Intelligence, Low Acquired Capacity)

A child operates in a state of near-constant Δ spikes because everything is novel. They possess very few Liquidated scripts. Consequently, they must constantly engage their Reconstructive and Simulative states, building causal models from scratch. This makes their cognitive process highly adaptive but metabolically exhausting and painfully slow. They are pure phase-control engines with an empty library. Gopnik (2009) has described children as operating in a state of “lantern consciousness”, broadly attentive and exploratory, which in SEM terms reflects an agent whose θ is extremely low and whose monitoring duty cycle is near-continuous.

The Legacy Corporation (Low Raw Intelligence, High Acquired Capacity)

A century-old bureaucracy possesses massive Acquired Capacity in the form of Standard Operating Procedures (Liquid states). As long as the market remains stable ($\Delta \approx 0$), the corporation is highly profitable and efficient. However, when a disruptive technology emerges (High Δ), the corporation routinely collapses. Its Elastic clutch is rusted shut; it lacks the Raw Intelligence to halt its own procedures and reconstruct a new business model. Christensen (1997) documented this pattern as the “innovator’s dilemma”, SEM provides the mechanical explanation for why this pattern is so predictable.

The Large Language Model (High Acquired Capacity, Zero Raw Intelligence)

Modern LLMs represent the ultimate extreme of Acquired Capacity. They possess billions of parameters, frozen, Liquidated weights optimized during training. When operating within the distribution of their training data ($\Delta \approx 0$), their performance mimics elite expertise. However, because they lack an active Architectural workspace and an Elastic clutch, they shatter when faced with novel logic puzzles (High Δ). They cannot pause, clear the board, and run a causal simulation. They simply execute the closest available habit, resulting in hallucination. This diagnosis aligns with the empirical findings of Marcus (2020), who has extensively documented the brittleness of LLM reasoning under distributional shift.

Individual Differences: Traits, Storage, and the Integration Problem

The distinction between Acquired Capacity and Raw Intelligence also illuminates a persistent confusion in how we evaluate individual cognitive differences, particularly in neurodivergent populations where specific capacities may be extraordinarily developed while integrative phase-control remains uneven.

Consider exceptional memory. A system, biological or artificial, may possess remarkable raw storage: the ability to encode, retain, and retrieve vast quantities of information with high fidelity. In conventional assessment, this capacity is often mistaken for intelligence itself. A child who can recite entire books verbatim, recall dates and sequences with perfect accuracy, or absorb encyclopedic knowledge at extraordinary speed is frequently labeled “gifted.”

SEM reframes this. Exceptional memory is a form of **Acquired Capacity amplification**, a larger-than-typical Liquidated library. It is raw storage, not raw intelligence. The critical question is not *how much* the system can store, but *how well it can integrate, regulate, and deploy* that storage under environmental volatility. A system with massive storage but limited Architectural bandwidth cannot hold its own knowledge in productive relation, it can retrieve individual facts but may struggle to synthesize them into novel causal models. A system with massive storage but a rigid Elastic clutch may have encyclopedic knowledge locked in Liquid scripts that cannot be interrupted or recombined when the context shifts.

This pattern is well-documented in certain neurodivergent profiles. Research on autism spectrum conditions has identified individuals with exceptional domain-specific abilities, sometimes termed “islands of ability” or savant skills, coexisting with significant challenges in flexible adaptation to novel social or environmental demands (Happé & Frith, 2006; Baron-Cohen, 2008). In SEM terms, these profiles are not paradoxical. They represent a specific mismatch: high Acquired Capacity (often in targeted domains) with atypical calibration of the Elastic clutch and Delta sensitivity. The storage is exceptional; the phase-control architecture processes signals differently.

This has important implications for both assessment and intervention. Traditional metrics that conflate storage with intelligence will systematically misevaluate individuals whose storage capacity dramatically exceeds their integrative phase-control, or, conversely, individuals whose phase-control is highly adaptive but whose Liquidated library is sparse. A child with extraordinary memory who struggles with novel problem-solving is not “smart but lazy” or “gifted but underperforming.” In SEM’s language, their Acquired Capacity is deep while their Architectural workspace, Elastic sensitivity, or Reconstructive-Simulative coordination may require targeted conditioning, conditioning that current educational frameworks rarely provide because they do not distinguish between storage and regulation.

The trainability thesis (Section 7) applies directly here. If Raw Intelligence is not fixed but conditionable, then the mismatch between storage and integration is not a permanent diagnostic category but a *starting point for targeted intervention*. Delta sensitivity training, phase-transition drills, and progressive complexity loading could be specifically calibrated for profiles where the Liquidated library is rich but the regulatory architecture needs development, an approach that traditional cognitive training paradigms, which do not distinguish between these components, are poorly equipped to deliver.

The Integrative Standard

A system cannot be judged solely by its performance in a static room. A massive database of habits is only useful until the rules of the environment change. The true measure of a system's resilience and capability is not defined by how many optimized scripts it has accumulated, but by its **regulatory phase-control**. Intelligence is the speed at which a system can safely abandon its own habits.

7. Trainability: Conditioning the Elastic Clutch

A central claim of SEM, and one that distinguishes it from trait-based intelligence models, is that Raw Intelligence is not a fixed, deterministic, or innate property. It is a systemic function that can be conditioned. Just as a physical clutch can be calibrated and a digital workspace can be expanded, a system can be trained to increase its Delta sensitivity and accelerate its phase-transitions.

This claim is supported by a growing body of evidence in cognitive neuroscience. Bavelier et al. (2012) demonstrated that action video game training improves attentional control and task-switching speed. Diamond and Ling (2016) reviewed evidence for executive function training, finding that interventions targeting inhibitory control and cognitive flexibility, both core components of the Elastic state, can produce measurable and transferable improvements.

SEM proposes three primary conditioning axes:

7.1 Delta Sensitivity Training

The objective is to lower the system's detection threshold so that moderate-Delta signals are recognized earlier, before accumulated drift produces crisis-level divergence.

Protocol sketch: Structured exposure to environments with gradually increasing ambiguity, where the system receives delayed feedback about whether its initial “stay in Liquid” or “engage Elastic” decision was correct. Over time, this calibrates θ downward without producing hypersensitivity. In human contexts, this is analogous to scenario-based training in fields like aviation and emergency medicine, where professionals practice recognizing subtle precursors to catastrophic failure (Klein, 1998).

7.2 Phase-Transition Speed Training

The objective is to reduce the latency between Delta detection and full engagement of the Reconstructive-Simulative cycle. Even a well-calibrated θ is useless if the transition takes too long, the crisis may evolve faster than the system can adapt.

Protocol sketch: Rapid-switch drills in which the system alternates between Liquid execution and Elastic override under time pressure. This is effectively the principle behind interval training in athletics applied to cognition, building the “muscle” of state-switching through repeated, time-constrained practice.

7.3 Architectural Bandwidth Expansion

The objective is to increase the effective capacity of the active workspace so that more variables can be held in simultaneous relation during the Reconstructive and Simulative phases.

Protocol sketch: Progressive complexity loading, in which the system is tasked with problems that require holding an increasing number of interdependent variables. Research on working memory training (Klingberg, 2010) suggests that such capacity is modestly expandable, though the degree of transfer to novel domains remains debated. SEM predicts that the greatest gains will come not from expanding raw capacity in isolation but from training the *coordination* between Architectural capacity and the other four states, a prediction that current working memory training paradigms do not typically test.

7.4 Implications for Artificial Systems

The trainability thesis carries significant implications for AI architecture. If SEM is correct that intelligence requires a tunable Elastic mechanism, then systems designed without one cannot be patched into intelligence through scale alone. The current paradigm of training ever-larger static models is, in SEM terms, an exercise in

expanding Acquired Capacity while leaving Raw Intelligence at zero. A more promising direction, one that several research groups are beginning to explore (e.g., Goyal & Bengio, 2022; LeCun, 2022), involves architectures with explicit world models, active inference loops, and the ability to *halt and recompute* when prediction error exceeds a dynamic threshold.

7.5 A SEM-Guided Path to Artificial General Intelligence

If SEM's central thesis is correct, that intelligence is regulatory phase-control rather than accumulated capacity, then the path to true Artificial General Intelligence (AGI) cannot be achieved through parameter scaling alone. Current frontier models, regardless of size, operate as massive Liquidated libraries: they excel in low- Δ regimes where the problem distribution matches their training data, and they fail catastrophically when the environment fractures. SEM predicts that genuine AGI requires an architecture that embeds the full regulatory cycle, a system capable of self-governing across arbitrary Delta gradients, including domains it has never encountered.

Architectural Requirements

SEM implies a minimum set of structural components for any AGI-class system:

A tunable Delta monitor. The system must maintain a continuous divergence signal between its predictions and observed reality, not as a static loss function computed during training, but as a *live runtime signal* that governs behavior in real time. Active inference frameworks (Friston, 2010) provide a starting point, but SEM requires that this monitor feeds into an explicit regulatory mechanism rather than simply updating model weights.

An Elastic meta-controller. This is the architectural clutch. When the Delta monitor exceeds a dynamically calibrated threshold θ , the meta-controller must be able to halt Liquid execution, clear the active workspace, and redirect computational resources to Reconstructive and Simulative engines. No current production system possesses this capability. Chain-of-thought prompting and self-reflection techniques approximate it, but they operate within the Liquid state, they are cached reasoning patterns, not genuine phase-control.

A Simulative world model. The system requires an internal engine capable of forward causal projection, not pattern completion, but physics-like simulation of how candidate actions will propagate through the environment. LeCun's (2022) Joint

Embedding Predictive Architecture (JEPA) represents one promising direction, as does the broader push toward world models in reinforcement learning (Ha & Schmidhuber, 2018). SEM adds the requirement that this engine must also support *inverse operation*, Reconstructive hypothesis generation followed by forward validation, to enable root cause reasoning (see Section 3, Compound Operations).

An expandable Architectural workspace. The system must be able to hold multiple interdependent variables in active relation during high- Δ episodes. Current transformer architectures have fixed context windows that serve as a static approximation, but genuine Architectural Intelligence requires dynamic allocation, expanding the workspace when the problem demands it and contracting it during Liquid execution to conserve resources.

Training Paradigm: Volatility Exposure

SEM's trainability thesis implies that an AGI system must be *trained on volatility itself*, not merely on data. Current training paradigms optimize for prediction accuracy across a fixed distribution. SEM-informed training would instead prioritize **progressive Delta ramps**: beginning with stable environments to build Liquid efficiency, then introducing controlled drift to condition monitoring duty cycles, then escalating to acute novelty to stress-test the Elastic mechanism and Reconstructive-Simulative loops. Meta-learning approaches (Finn et al., 2017) provide relevant infrastructure, but would need to be extended to explicitly reward phase-transition speed and calibration accuracy rather than task performance alone.

Milestones and Benchmarks

Current AI benchmarks, including HELM, MMLU, and domain-specific suites, overwhelmingly test Acquired Capacity in low- Δ conditions. SEM suggests that meaningful AGI benchmarks should evaluate performance *across the Delta gradient*: stable-environment efficiency (Liquid), moderate-drift detection (monitoring duty cycle), crisis adaptation (Elastic disengage and Reconstructive-Simulative recovery), and post-crisis consolidation (re-Liquidation of successful new strategies). A system that scores well on static benchmarks but fails moderate- Δ adaptation tests is, by SEM's definition, not generally intelligent, it is merely extensively trained.

Alignment Through Calibration

SEM's calibration framework has direct implications for AI alignment and safety. A hyposensitive AGI (θ too high) would fail to detect distributional shifts in its operating

environment, including shifts in human values, social norms, or safety constraints, until catastrophic misalignment had already occurred. A hypersensitive AGI (θ too low) would over-adapt to noise, producing erratic behavior as it constantly abandons stable policies in response to minor perturbations. Alignment, in SEM terms, is not a static property baked in during training; it is a *calibration problem*, maintaining a θ that is sensitive enough to detect genuine environmental shifts (including ethical ones) while stable enough to sustain coherent, predictable behavior. This reframes alignment research from “how do we constrain the system?” to “how do we calibrate its regulatory sensitivity?”, a question SEM is uniquely equipped to formalize.

Limitations and Open Questions

SEM provides a regulatory scaffold, not a complete AGI architecture. It does not address embodiment, multi-agent coordination, or the substrate-level engineering required to implement these mechanisms at scale. The computational cost of maintaining live Simulative engines and Inverse Simulation loops during runtime remains an open challenge. Furthermore, the interaction between SEM’s phase-control model and emergent properties of large-scale neural networks, where Liquid and Elastic behaviors may arise implicitly rather than through explicit architectural design, requires further investigation. These limitations do not diminish SEM’s contribution; they define the research frontier that a systems theory of intelligence must eventually address.

8. Positioning Against Existing Frameworks

SEM does not emerge in a vacuum. It draws on and extends several influential frameworks in cognitive science and AI. To clarify its contribution, we map SEM explicitly against three major reference points.

8.1 Kahneman’s System 1 / System 2 (Dual-Process Theory)

Kahneman (2011) proposed that cognition operates through two systems: System 1 (fast, automatic, heuristic) and System 2 (slow, deliberate, analytical). SEM shares the fundamental insight that cognition alternates between automatic and deliberate modes.

What SEM adds: Dual-process theory gives you two gears. SEM gives you a full manual transmission with a clutch. Specifically, SEM contributes three things that dual-process theory lacks. First, it decomposes System 2 into three mechanistically distinct states (Architectural, Reconstructive, Simulative) rather than treating “slow thinking” as a monolith. Second, it specifies the *regulatory mechanism* that governs

transitions between modes, the Elastic clutch responding to the Delta signal, whereas Kahneman's framework leaves the transition dynamics largely implicit. Third, it identifies the Delta gradient and calibration errors as a systematic source of cognitive failure, explaining not just *that* System 1 sometimes overrides System 2 (as Kahneman documents), but *why* the override mechanism fails in predictable ways.

8.2 Friston's Free Energy Principle

Friston (2010) proposed that all self-organizing biological systems minimize variational free energy, a quantity closely related to prediction error. Under the Free Energy Principle (FEP), organisms act to minimize surprise (free energy) by either updating their internal models or acting on the environment.

What SEM adds: SEM is *compatible* with FEP but operates at a different level of description. Where FEP provides a unified variational objective (minimize free energy), SEM provides a *state-architecture* for how that minimization is actually implemented in practice. The Delta is a specific operationalization of prediction error. The five states describe the mechanical repertoire available for reducing it. Crucially, SEM adds the concept of the clutch threshold (θ) and calibration errors, which FEP does not address, FEP tells you the system *should* minimize surprise, but not what happens when the regulatory mechanism that governs the system's response to surprise is itself miscalibrated. SEM is, in this sense, a *failure theory* of adaptive intelligence in a way that FEP is not.

8.3 Anderson's ACT-R

ACT-R (Anderson, 2007) is a cognitive architecture that distinguishes between declarative memory (factual knowledge) and procedural memory (production rules). It models cognition as the activation and execution of production rules that compete for access to limited processing resources.

What SEM adds: ACT-R provides a detailed computational account of how knowledge is stored and retrieved, but it does not have a built-in theory of adaptive phase-control under environmental volatility. SEM's Liquid state is roughly analogous to ACT-R's procedural execution, and the Architectural state maps partially to ACT-R's buffers. However, SEM's contribution is the explicit regulatory layer, the Delta monitor and Elastic clutch, that governs when the system should *abandon* its production rules rather than simply selecting among them. ACT-R models cognition within a stable task environment; SEM is designed for environments that break.

8.4 Emotional Intelligence and Goleman's Framework

No discussion of intelligence frameworks is complete without addressing the construct of Emotional Intelligence (EI), popularized by Goleman (1995) and formalized in the ability model by Mayer and Salovey (1997). EI encompasses the perception, understanding, management, and facilitation of emotions, both one's own and others'. It has been widely adopted in organizational psychology, education, and leadership development, and is often presented as a distinct "type" of intelligence parallel to cognitive ability.

SEM challenges this framing. Under the model's lens, most of what is called Emotional Intelligence decomposes into two components that are already accounted for in the framework's existing architecture:

Emotional regulation as Acquired Capacity. The ability to manage one's emotional responses in socially appropriate ways, remaining calm under pressure, suppressing an impulsive reaction, displaying empathy at the right moment, is, in SEM terms, a set of **Liquidated scripts** compiled through socialization. These are learned behavioral routines, automated through repeated practice in social environments. A person who has been trained from childhood to "count to ten before responding" has Liquidated an impulse-control protocol. A leader who "reads the room" before speaking has Liquidated a social-perception heuristic. These are valuable skills, but they are Acquired Capacity, optimized for familiar social environments and brittle when the social context shifts dramatically (e.g., cross-cultural encounters, unprecedented crises, or neurodivergent social dynamics that violate learned scripts).

Emotional adaptation as phase-control. The *genuinely intelligent* component of EI, under SEM, is not the possession of emotional scripts but the ability to *override* them when the social environment produces a Delta spike. This is Elastic Intelligence applied to the socio-emotional domain. The person who can recognize that their trained empathy response is misreading the situation, that the other person is not upset but excited, or that the cultural context demands a different emotional register, and who can halt their automated response to reconstruct and simulate a more appropriate one, is exercising Raw Intelligence in the emotional domain. The clutch, not the library, is doing the real work.

This decomposition resolves a longstanding tension in the EI literature. Critics have argued that EI is merely personality plus general intelligence repackaged (Locke, 2005; Waterhouse, 2006). Defenders insist it captures something distinct. SEM suggests that

both sides are partially correct. The scripted, stable components of EI, social habits, impulse control protocols, learned empathy responses, are indeed closer to personality traits (Acquired Capacity in a social domain). But the adaptive, flexible components, the ability to detect emotional Delta signals, override cached social scripts, and reconstruct appropriate responses in novel interpersonal contexts, are genuine expressions of Raw Intelligence operating on socio-emotional inputs rather than cognitive or physical ones.

Emotional Delta Signals

If SEM's Delta framework extends to the emotional domain, it implies the existence of **emotional Delta signals**: divergences between expected and observed social or emotional states. You expect gratitude; you observe resentment. You expect calm; you observe panic. These emotional Deltas trigger the same regulatory cascade as cognitive Deltas, Elastic disengage, Reconstructive hypothesis generation ("Why are they angry? What am I missing?"), Simulative forward testing ("If I respond with X, will it de-escalate or inflame?"), and eventual re-Liquidation of a new social script if the pattern recurs.

This extension has direct trainability implications. Just as cognitive Delta sensitivity can be conditioned (Section 7.1), emotional Delta sensitivity can be systematically trained, through structured exposure to emotionally ambiguous scenarios, perspective-taking exercises with delayed feedback, and practice in social contexts that deliberately violate cached expectations. Mindfulness-based interventions, which have shown efficacy in improving emotional regulation (Chambers et al., 2009), may work precisely because they train the emotional Elastic clutch, building the capacity to *pause* an automated emotional response and observe the actual emotional landscape before reacting.

Is Emotional Intelligence "True Intelligence"?

SEM's answer is nuanced. The *label* "Emotional Intelligence" conflates two mechanistically distinct phenomena. Emotional scripts and learned social behaviors are Acquired Capacity, valuable, trainable, but not intelligence in SEM's sense. Emotional phase-control, the ability to detect socio-emotional Deltas, interrupt cached responses, and reconstruct adaptive behavior in novel interpersonal contexts, *is* intelligence, applied to a specific input domain. The confusion arises because Goleman's framework does not distinguish between the library and the clutch. SEM does.

8.5 Summary of Contributions

Framework	Shared Ground with SEM	SEM's Specific Contribution
-----------	------------------------	-----------------------------

Dual-Process Theory (Kahneman)	Fast vs. slow cognition	Multi-state decomposition; explicit clutch mechanism; Delta gradient and calibration error theory
Free Energy Principle (Friston)	Prediction error as driver; model updating	State-architecture for minimization; clutch threshold (θ); failure theory of miscalibrated regulation
ACT-R (Anderson)	Declarative/procedural distinction; limited resources	Regulatory layer for environmental volatility; explicit rule-abandonment mechanism; trainability of phase-control
AGI Architectures (LeCun, Ha & Schmidhuber)	World models; active inference; meta-learning	Elastic regulation for volatility; calibration-based alignment; training on Δ gradients rather than static distributions
Emotional Intelligence (Goleman, Mayer & Salovey)	Socio-emotional adaptation; impulse regulation	Decomposition into Acquired Capacity (scripts) vs. Raw Intelligence (emotional phase-control); emotional Δ signals; trainability of emotional Elastic clutch

9. Conclusion: The True Definition of Adaptation

The Simulative-Elasticity Model proposes that true intelligence is not a static library of facts. It is a highly regulated engine.

A brilliant system must have a large workbench to hold complexity (Architectural Intelligence). It must be able to reconstruct missing data (Reconstructive Intelligence) and simulate future outcomes (Simulative Intelligence). It must automate its successes to conserve energy (Liquid Intelligence). And most importantly, it must maintain a highly sensitive clutch, ready to abandon its own rules the exact millisecond the world changes (Elastic Intelligence).

What this revised framework contributes beyond the initial formulation is eightfold. First, the Script-to-Capability Pipeline (Section 3) makes explicit how Liquidated scripts become encapsulated tools that compose into observable expertise, revealing that what the world calls “capability” is a dependent toolchain vulnerable to cascading collapse when foundational scripts are invalidated. Second, the concept of compound operations, specifically the Inverse Simulation loop (Section 3), demonstrates that the five states are not a linear pipeline but a recurrent regulatory network, capable of locking into iterative feedback cycles for diagnostic reasoning and root cause analysis. Third, the Delta gradient and calibration theory (Section 5) demonstrate that the most dangerous

failures of intelligence occur not in obvious crises but in the slow, ambiguous drift of moderate- Δ environments. Fourth, the analysis of individual differences and neurodivergence (Section 6) shows that exceptional storage capacity does not constitute intelligence without integrative phase-control, reframing common assessment mismatches as differences in regulatory architecture rather than deficits. Fifth, the trainability thesis (Section 7) establishes that the Elastic clutch is not a fixed property but a tunable parameter, opening the door to targeted conditioning protocols for both cognitive and emotional domains. Sixth, the SEM-guided path to AGI (Section 7.5) argues that general intelligence cannot emerge from parameter scaling alone, requiring instead explicit regulatory architecture trained on volatility itself rather than static data distributions. Seventh, the decomposition of Emotional Intelligence (Section 8.4) reveals that most of what is called EI is Acquired Capacity, Liquidated social scripts, while the genuinely intelligent component is emotional phase-control: the ability to detect socio-emotional Delta signals and override cached responses in novel interpersonal contexts. Eighth, the explicit positioning against existing frameworks clarifies that SEM is not a replacement for prior theories but a regulatory layer that addresses what they leave unspecified: the mechanics of when and how a system should abandon its own compiled knowledge.

Intelligence is not the accumulation of perfect habits. It is the speed at which you can break them, and, crucially, the wisdom to know when breaking them is warranted and when it is premature.

References

- Anderson, J. R. (2007). *How Can the Human Mind Occur in the Physical Universe?* Oxford University Press.
- Baddeley, A. D. (2000). The episodic buffer: A new component of working memory? *Trends in Cognitive Sciences*, 4(11), 417–423.
- Baron-Cohen, S. (2008). *Autism and Asperger Syndrome*. Oxford University Press.
- Bavelier, D., Green, C. S., Pouget, A., & Schrater, P. (2012). Brain plasticity through the life span: Learning to learn and action video games. *Annual Review of Neuroscience*, 35, 391–416.
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? *Proceedings of FAccT '21*, 610–623.
- Christensen, C. M. (1997). *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*. Harvard Business Review Press.
- Chambers, R., Gullone, E., & Allen, N. B. (2009). Mindful emotion regulation: An integrative review. *Clinical Psychology Review*, 29(6), 560–572.
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(3), 181–204.
- Craik, K. J. W. (1943). *The Nature of Explanation*. Cambridge University Press.
- de Kleer, J., & Williams, B. C. (1987). Diagnosing multiple faults. *Artificial Intelligence*, 32(1), 97–130.
- Diamond, A., & Ling, D. S. (2016). Conclusions about interventions, programs, and approaches for improving executive functions that appear justified and those that, despite much hype, do not. *Developmental Cognitive Neuroscience*, 18, 34–48.
- Fitts, P. M., & Posner, M. I. (1967). *Human Performance*. Brooks/Cole.
- Finn, C., Abbeel, P., & Levine, S. (2017). Model-agnostic meta-learning for fast adaptation of deep networks. *Proceedings of the 34th International Conference on Machine Learning*, 1126–1135.
- Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11(2), 127–138.
- Gopnik, A. (2009). *The Philosophical Baby: What Children's Minds Tell Us About Truth, Love, and the Meaning of Life*. Farrar, Straus and Giroux.

- Goleman, D. (1995). *Emotional Intelligence: Why It Can Matter More Than IQ*. Bantam Books.
- Goyal, A., & Bengio, Y. (2022). Inductive biases for deep learning of higher-level cognition. *Proceedings of the Royal Society A*, 478(2266), 20210068.
- Harman, G. H. (1965). The inference to the best explanation. *The Philosophical Review*, 74(1), 88–95.
- Ha, D., & Schmidhuber, J. (2018). World models. arXiv preprint arXiv:1803.10122.
- Happé, F., & Frith, U. (2006). The weak coherence account: Detail-focused cognitive style in autism spectrum disorders. *Journal of Autism and Developmental Disorders*, 36(1), 5–25.
- Hegarty, M. (2004). Mechanical reasoning by mental simulation. *Trends in Cognitive Sciences*, 8(6), 280–285.
- Josephson, J. R., & Josephson, S. G. (1994). *Abductive Inference: Computation, Philosophy, Technology*. Cambridge University Press.
- Kahneman, D. (2011). *Thinking, Fast and Slow*. Farrar, Straus and Giroux.
- Klein, G. (1998). *Sources of Power: How People Make Decisions*. MIT Press.
- Klingberg, T. (2010). Training and plasticity of working memory. *Trends in Cognitive Sciences*, 14(7), 317–324.
- LeCun, Y. (2022). A path towards autonomous machine intelligence (Version 0.9.2). OpenReview preprint.
- Leonard-Barton, D. (1992). Core capabilities and core rigidities: A paradox in managing new product development. *Strategic Management Journal*, 13(S1), 111–125.
- Locke, E. A. (2005). Why emotional intelligence is an invalid concept. *Journal of Organizational Behavior*, 26(4), 425–431.
- Lucas, H. C., & Goh, J. M. (2009). Disruptive technology: How Kodak missed the digital photography revolution. *Journal of Strategic Information Systems*, 18(1), 46–55.
- Marcus, G. (2020). The next decade in AI: Four steps towards robust artificial intelligence. arXiv preprint arXiv:2002.06177.
- Mayer, J. D., & Salovey, P. (1997). What is emotional intelligence? In P. Salovey & D. J. Sluyter (Eds.), *Emotional Development and Emotional Intelligence* (pp. 3–31). Basic Books.

Peirce, C. S. (1903). Pragmatism as a principle and method of right thinking. In P. A. Turrisi (Ed.), *The 1903 Harvard Lectures on Pragmatism*. State University of New York Press, 1997.

Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, 63(2), 129–138.

Taleb, N. N. (2012). *Antifragile: Things That Gain from Disorder*. Random House.

Waterhouse, L. (2006). Multiple intelligences, the Mozart effect, and emotional intelligence: A critical review. *Educational Psychologist*, 41(4), 207–225.