Developmental Active Inference: A Computational Framework for Stage-Dependent Self-Regulated Learning

Alex Kwon[1]

[1] Boston Neuromind LLC, Boston, MA, USA

Former Visiting Scholar, Harvard Graduate School of Education

Author Note

Correspondence concerning this article should be addressed to Alex Kwon, Boston Neuromind LLC, Boston, MA, USA. Email: bostonneuromind@gmail.com

Conflicts of Interest. The author is the founder of Boston Neuromind LLC, which develops the clinical implementation (Learning Catcher) referenced in this paper.

Funding. No external funding supported this work.

Abstract

Self-regulated learning (SRL) is widely recognized as central to mental health and educational outcomes, yet existing frameworks remain primarily phenomenological. Active Inference offers a computational account of perception and action under uncertainty, but has rarely been applied to development-sensitive learning contexts. This paper proposes Developmental Active Inference (DAI), a theoretical framework that integrates four traditionally separate accounts—Zimmerman's SRL phase model, Fischer's Dynamic Skill Theory, Polyvagal Theory, and Bayesian personalization—as natural decompositions of expected free energy. We argue that (a) developmental levels correspond to layers in a hierarchical generative model, (b) heart rate variability indexes the precision parameter that gates learning, (c) habits are learned action policies that minimize free energy, and (d) the refinement of a self-generative model is the computational substrate of self-reflection. We derive a single decision rule for adaptive intervention—argmin expected free energy—from which the four mechanisms emerge as structural components, and present a worked example illustrating its application in clinical decision support. We discuss empirical predictions, measurement strategies for each mechanism, and a sequenced research program beginning with single-case experimental designs. The framework provides a principled foundation for adaptive clinical and educational interventions, and a testable theory of why developmental progression in self-regulated learning is computationally tractable.

Keywords: self-regulated learning; active inference; free energy principle; dynamic skill theory; polyvagal theory; developmental psychology; adaptive intervention

Developmental Active Inference: A Computational Framework for Stage-Dependent Self-Regulated Learning

Self-regulated learning (SRL) describes the active, intentional, and goal-directed processes by which learners regulate their cognition, motivation, behavior, and environment in service of learning (Zimmerman, 2000; Pintrich, 2004). It is a robust predictor of academic achievement (Sitzmann & Ely, 2011), clinical outcomes in mental health (Bandura, 2005), and long-term well-being (Boekaerts & Corno, 2005). Despite decades of theoretical and empirical work, however, SRL remains primarily described at the phenomenological level: phases such as forethought, performance, and self-reflection (Zimmerman, 2000) capture what learners do, but not how their nervous systems compute the underlying decisions.

In parallel, the past two decades have produced an increasingly mature computational framework for understanding adaptive behavior. The free energy principle (Friston, 2010) and its process-theoretic instantiation, Active Inference (Friston, FitzGerald, et al., 2017; Parr, Pezzulo, & Friston, 2022), provide a unified account of perception, learning, and action in terms of a single objective: the minimization of expected free energy. This account has been applied to a wide range of clinical and cognitive phenomena, including depression (Stephan et al., 2016), psychosis (Sterzer et al., 2018), autism (Lawson, Rees, & Friston, 2014), and habit formation (Friston, FitzGerald, et al., 2017).

Yet the application of Active Inference to learning has remained largely momentary. Most analyses focus on how an agent acts to minimize prediction error in the immediate sense, often within a static generative model. Development—understood as the structured, stage-dependent transformation of the model itself—has received comparatively little attention (cf. Pezzulo, Rigoli, & Friston, 2018). This omission is consequential. Clinical and educational interventions are inherently developmental: a card or task that is optimal for a learner today may be useless or even harmful three months from now, not because the algorithm has failed, but because the learner's generative model has changed.

The present paper proposes Developmental Active Inference (DAI) as a framework that addresses this gap. The central thesis is that four traditionally separate accounts of human learning and regulation—Zimmerman's SRL model, Fischer's Dynamic Skill Theory (DST; Fischer, 1980; Fischer & Bidell, 2006), Polyvagal Theory (Porges, 2007), and Bayesian models of individual differences—are not arbitrary components to be weighted in some ad hoc fashion, but natural decompositions of expected free energy under a hierarchical generative model. Once this decomposition is made explicit, a single decision rule (select the action minimizing expected free energy) generates each component as a structural feature.

We proceed as follows. Section 2 provides a brief, accessible overview of the Active Inference framework. Section 3 reviews Fischer's Dynamic Skill Theory and identifies why it requires a computational mechanism. Section 4 develops the central integration: how four mechanisms—

hierarchical generative model expansion, precision allocation, empirical prior accumulation, and self-generative model refinement—emerge from expected free energy minimization, and how each maps to one of the existing theoretical accounts. Section 5 illustrates the framework with a worked example. Section 6 discusses empirical predictions and measurement strategies. Section 7 outlines a sequenced research program.

1.1 Why a Framework Rather Than a Model

Before proceeding, a methodological note. The contribution of this paper is theoretical: a conceptual framework that organizes existing constructs under a unifying computational principle. It does not present novel empirical data. We make this choice deliberately, for two reasons. First, the integration we propose has not, to our knowledge, been previously articulated, and developing the theoretical scaffolding is a prerequisite for the empirical work that would test it. Second, several of the mechanisms we propose—particularly the mapping between heart rate variability and the Active Inference precision parameter (Smith et al., 2017)—have substantial independent empirical support, allowing us to draw on existing literature rather than generate new data. The framework yields specific, testable predictions; we end by sketching a research program designed to evaluate them.

2. Active Inference: A Brief Overview

We provide here a non-technical overview of the Active Inference framework sufficient to support the integration developed in subsequent sections. Readers seeking formal treatments are directed to Friston, FitzGerald, et al. (2017) and Parr et al. (2022); Smith, Friston, and Whyte (2022) offer a step-by-step tutorial for empirical researchers.

2.1 The Free Energy Principle

The free energy principle (Friston, 2010) states that any self-organizing system that maintains its existence over time must, in effect, minimize a quantity called variational free energy. Variational free energy is an upper bound on surprise: it quantifies the mismatch between what an agent expects and what it observes. Living systems—from cells to brains to social groups—are characterized by their capacity to bound this mismatch within viable limits. For brains specifically, free energy minimization corresponds to keeping the world predictable: perceiving what is there, acting to make the world conform to expectations, and updating expectations when neither perception nor action is sufficient.

Two ways of minimizing free energy are available to an agent. The first is perceptual: the agent updates its beliefs about the world to better match the sensory input it receives. This corresponds to standard Bayesian inference. The second is active: the agent acts on the world to bring its sensory input into closer alignment with its beliefs and goals. This second mode—acting to confirm one's predictions—is the distinctive feature of Active Inference, and the basis for treating learning and decision-making within the same framework.

2.2 Expected Free Energy and Action Selection

Whereas variational free energy is computed over current observations, expected free energy (EFE) is computed prospectively, over the consequences of possible actions. For each candidate policy π (a sequence of actions), the agent computes:

$$G(\pi) = -\,E[\,I(\pi)\,] - E[\,U(\pi)\,]$$

where $I(\pi)$ is the expected information gain (epistemic value: how much the agent expects to learn) and $U(\pi)$ is the expected pragmatic value (the degree to which the agent expects to achieve preferred outcomes). The agent then selects the policy with the lowest G. Critically, this formulation unifies exploration and exploitation under a single objective: agents prefer actions that are both informative and goal-conducive.

2.3 Hierarchical Generative Models

A central technical feature of Active Inference is its use of hierarchical generative models (Friston, 2008; Pezzulo et al., 2018). Such models represent the world at multiple temporal and abstract scales. Lower levels encode fast, concrete signals (e.g., visual edges, autonomic state); higher levels encode slow, abstract regularities (e.g., social rules, personal goals). Inference and learning proceed in both directions: higher levels send predictions to lower levels, and lower levels return prediction errors to higher levels. The depth and structure of this hierarchy are not given a priori but develop over an organism's lifetime—a feature we will return to repeatedly.

2.4 Precision

A subtle but consequential feature of Active Inference is the role of precision, denoted γ. Precision quantifies the confidence the agent assigns to particular beliefs or sensory channels. High precision on a belief means the agent treats that belief as reliable; low precision means it is provisional. Precision allocation is itself learned, and dysregulated precision has been implicated in several clinical phenomena (Lawson et al., 2014; Sterzer et al., 2018). For the present framework, the most important fact about precision is that it is modulated by autonomic state—a connection we elaborate in Section 4.

3. Dynamic Skill Theory and the Computational Gap

Fischer’s Dynamic Skill Theory (DST; Fischer, 1980; Fischer & Bidell, 2006) is a theory of cognitive and emotional development that has been influential in education, clinical psychology, and the Mind, Brain, and Education field. DST holds that human development proceeds through a series of hierarchically nested levels of skill, organized into four broad tiers: reflexes, actions, representations, and abstractions. Within each tier, learners progress through three increasingly complex types of skills: single sets, mappings, and systems. Across the lifespan, an adult may operate at levels 7 through 13 across different domains, with substantial variation depending on the specific skill, context, and supports available.

3.1 Key Features of DST

Four features of DST are central for our purposes. First, DST is domain-specific: a learner may be at one level for one skill and a different level for another, even within the same individual at the same moment (Fischer & Bidell, 2006). Second, DST is context-sensitive: under high support and emotional safety, learners operate at their optimal level; under stress or low support, they regress to a functional level (Fischer & Bidell, 2006; Lamb, Mascolo, & Fischer, 2002). Third, DST emphasizes that development is non-linear—learners spiral through levels rather than progressing in a strictly monotonic fashion (Granott & Parziale, 2002). Fourth, DST holds that the optimal challenge for learning occurs not within a learner’s current level, but at the next level up, which can only be reached with scaffolding (Fischer & Bidell, 2006; Vygotsky, 1978).

3.2 The Computational Gap

Despite its theoretical richness and empirical support, DST has a notable limitation: it offers no computational mechanism for the transitions it describes. Why is a learner ready to move from one level to the next? What internal signal indicates readiness? Why does development proceed in this specific hierarchical pattern? DST describes the structure of development but does not explain how the brain implements that structure.

This gap has practical consequences. Without a computational account, the recommendation to “teach at one level up” cannot be operationalized in adaptive systems beyond rough heuristic matching. There is no principled way to identify when a learner has consolidated their current level sufficiently to benefit from new material, no principled way to integrate developmental level with momentary state (e.g., fatigue, emotion), and no principled way to model how learners differ from one another in their developmental trajectories.

Active Inference, we propose, fills this gap. The hierarchical generative model of Active Inference provides a computational substrate for DST: each Fischer level corresponds to a layer in the hierarchy, and developmental transitions correspond to the elaboration of these layers over time. We turn now to the central integration.

4. The Integration: Four Mechanisms

We propose that four mechanisms—each grounded in an existing literature—emerge as natural decompositions of expected free energy under a hierarchical generative model. Figure 1 contrasts the current landscape of fragmented theories with the proposed integration.

Figure 1

Conceptual gap (left) and proposed integration (right).

Note. Existing theoretical accounts (Zimmerman SRL, Fischer DST, Polyvagal Theory, Active Inference) operate independently. Developmental Active Inference (DAI) treats them as decompositions of a single quantity, expected free energy.

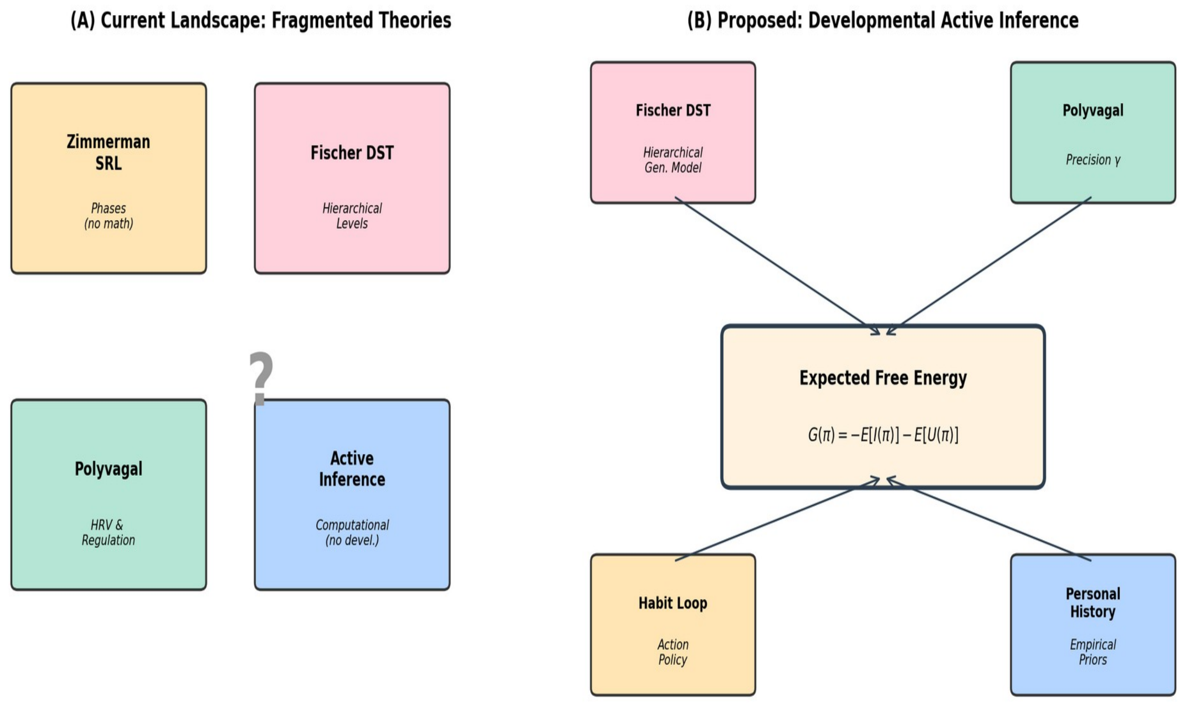

We describe each mechanism in turn, identifying its empirical support and its specific role within the integrated framework. Table 1 summarizes the mapping.

Table 1

Mapping Between Existing Theoretical Accounts and Active Inference Components

Note. Each row identifies a traditional theoretical account, the construct it describes, the Active Inference component it maps onto, and the primary empirical literature supporting the mapping.

Theoretical Account	**Construct**	**Active Inference Component**	**Key Empirical Support**
Fischer DST	Hierarchical developmental levels	Hierarchical generative model layers	Pezzulo et al. (2018); Friston (2008)
Polyvagal Theory	HRV / autonomic regulation	Precision parameter γ	Smith et al. (2017); Thayer & Lane (2009)
Habit Theory	Cue–routine–reward–craving	Learned action policy	Friston, FitzGerald, et al. (2017); Wood & Rünger (2016)
Bayesian Individual Difference Models	Personal baseline / history	Empirical priors	Tschantz et al. (2020); Stephan et al. (2016)
Self-Awareness / Metacognition	OMR-style self-monitoring	Self-generative model refinement	Apps & Tsakiris (2014); Fleming & Lau (2014)

4.1 Mechanism 1: Hierarchical Generative Model Expansion

The first mechanism concerns the structural relationship between Fischer developmental levels and the layers of a hierarchical generative model. We propose that each Fischer level corresponds to a level in the generative hierarchy, with developmental progression understood as the elaboration—and, ultimately, addition—of higher layers. Figure 2 illustrates this correspondence.

Figure 2

Hierarchical generative model with Fischer developmental levels.

Note. Each level in the Active Inference hierarchy corresponds to a Fischer tier. Higher levels operate at slower timescales and encode more abstract regularities. Top-down predictions and bottom-up prediction errors couple the levels.

Figure 2. Hierarchical Generative Model with Fischer Developmental Levels

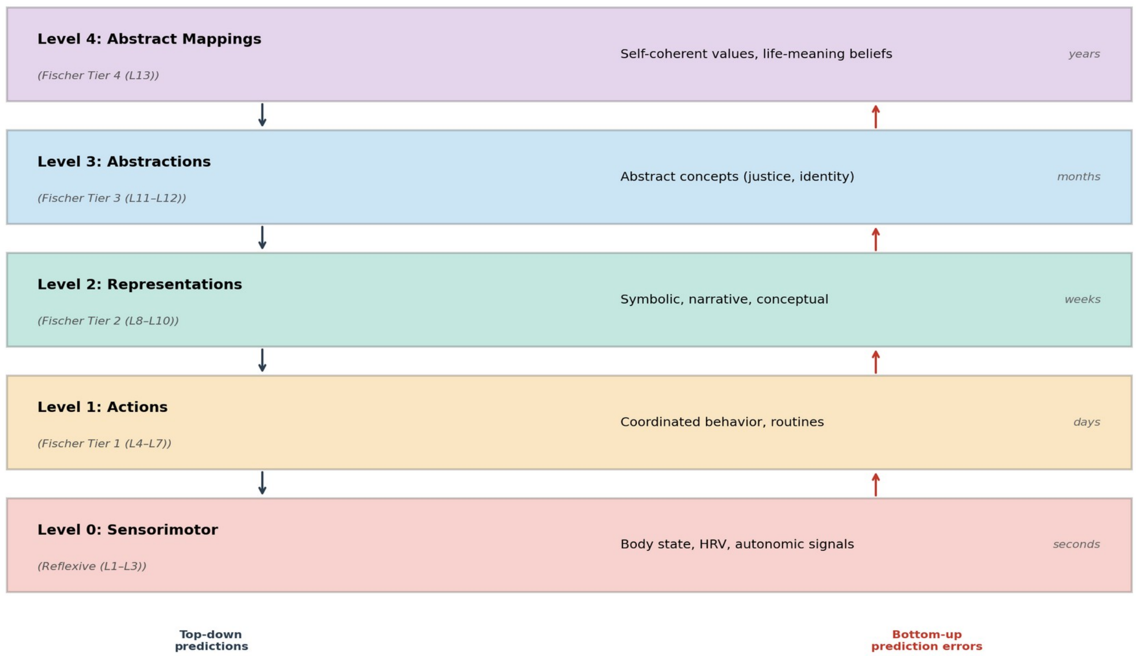

Note. *Development entails the emergence of new hierarchical layers, each operating at a slower timescale.*

This mapping is not merely descriptive. It generates predictions. First, learners should show evidence of stable activity at lower levels (e.g., autonomic, sensorimotor) preceding the emergence of stable activity at higher levels (e.g., abstract conceptual). Second, developmental progression should be marked by the emergence of slower, more abstract patterns of prediction, not merely by increased speed or accuracy. Third, regression under stress—a hallmark of DST—should manifest as the breakdown of upper layers and reliance on lower-layer predictions, a pattern consistent with clinical observations in trauma and dissociation (van der Kolk, 2014; Porges, 2007).

Pezzulo et al. (2018) have developed a hierarchical Active Inference framework that aligns closely with these predictions, though without explicit reference to Fischer. The present framework extends this work by drawing on DST to specify the empirical content of the higher levels (i.e., the specific abstractions characteristic of human adult development) and the developmental sequence by which they emerge.

4.2 Mechanism 2: Precision Allocation via Autonomic State

The second mechanism concerns how autonomic state, indexed primarily by heart rate variability (HRV), modulates the precision parameter γ of the generative model. This mapping has the strongest direct empirical support of the four mechanisms.

Smith, Thayer, Khalsa, and Lane (2017) provide a synthetic review arguing that vagal tone (as indexed by HRV) reflects the precision of beliefs about the body and about the contingencies that link action to outcome. Higher HRV corresponds to higher precision on prefrontal predictions, supporting executive function, working memory, and flexible behavior. Lower HRV corresponds to a shift toward lower-level, more reflexive processing. Thayer and Lane's (2009) neurovisceral integration model converges on a similar conclusion from independent evidence.

Figure 3 illustrates two consequences of this mapping. Panel A shows the relationship between HRV and γ, with thresholds defining a dysregulated zone (in which top-down learning predictions are insufficient to support meaningful information gain), a marginal zone (in which only scaffolded engagement is feasible), and a regulated zone (in which challenge becomes productive). Panel B illustrates the inverted-U shape of expected information gain across Fischer levels: too easy generates no learning, too hard produces frustration, and optimal challenge—Fischer's L+1—falls in between.

Figure 3

HRV–precision mapping (A) and information gain across Fischer levels (B).

Note. Panel A: HRV operationalizes the precision parameter γ via a sigmoid function. Thresholds define regulatory zones. Panel B: Expected information gain follows an inverted-U shape with peak at L+1 relative to the learner's current level.

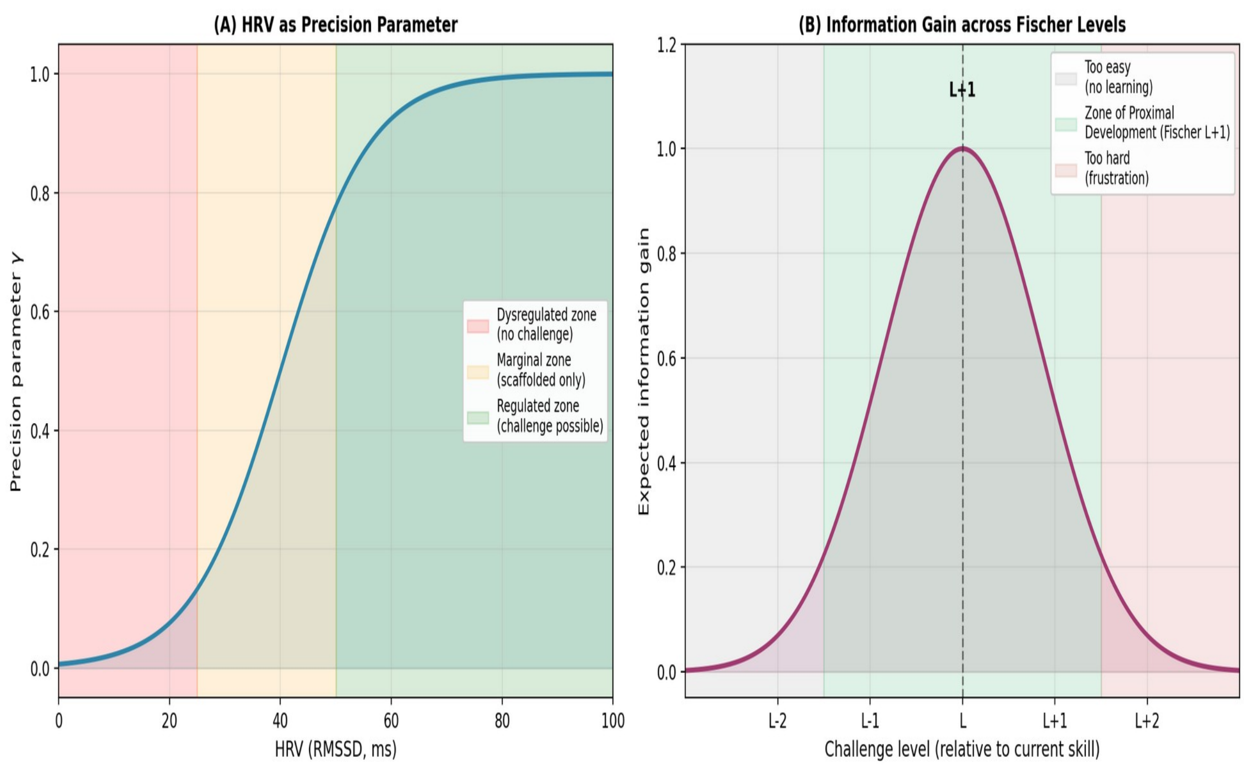

A practical implication of this mechanism is what we term *precision gating*: any adaptive intervention should evaluate the learner's regulatory state before selecting an action. When γ is low, the optimal action is one that increases γ (e.g., a regulation card, a breathing exercise, a body-anchored intervention), not one that maximizes information gain at the cognitive level. This is the computational expression of the clinical principle “body first, mind second,” a principle widely invoked in trauma-informed practice (van der Kolk, 2014) but rarely articulated computationally.

4.3 Mechanism 3: Empirical Prior Accumulation via Habit

The third mechanism concerns the role of habits in self-regulated learning. Habits have been variously characterized in cognitive psychology (Wood & Rünger, 2016), behavioral economics (Duhigg, 2012), and clinical theory (Marlatt & Gordon, 1985). Within Active Inference, habits correspond to learned policies—stable mappings from situations to actions that have, through repetition, accumulated high precision and low cost (Friston, FitzGerald, et al., 2017; Tschantz et al., 2020).

Three implications follow. First, established habits effectively constitute strong empirical priors: the agent expects to perform familiar action sequences in familiar contexts, and this expectation is rarely re-evaluated. This is computationally efficient, but creates a learning bottleneck: new patterns are difficult to install where existing habits compete. Second, the canonical habit loop (cue–routine–reward) maps onto the Active Inference cycle (context-sensitive policy selection → action → outcome-driven belief update); the “craving” component sometimes added to this loop (Duhigg, 2012) corresponds to expected pragmatic value. Third, habit formation has its own developmental trajectory: early in development, habits are formed primarily around bodily and emotional regulation; later, around social and abstract goals.

For an adaptive intervention, the role of habit-aware design is to align with the learner’s natural cue structure and existing routines rather than against them. Asking a learner to engage with a card at a time of day inconsistent with their existing rhythms is, from this perspective, a request for the learner to override a strong empirical prior—an expensive computation that typically fails. This is consistent with the empirical literature on behavior change, which finds that interventions exploiting natural cues are substantially more effective than those that do not (Lally, van Jaarsveld, Potts, & Wardle, 2010).

4.4 Mechanism 4: Self-Generative Model Refinement

The fourth mechanism concerns the refinement of the learner’s model of themselves. Apps and Tsakiris (2014) have argued, within the free energy framework, that the self is itself a generative model: a prior over the agent’s own states, dispositions, and patterns of response. Self-awareness, in this formulation, is the agent’s ability to perform inference about its own model. Self-reflection, in Zimmerman’s SRL framework, corresponds closely to this construct.

A central observation, however, is that the self-model is initially low-precision and often inaccurate. Learners frequently misattribute their patterns (“I procrastinate because I am lazy”) and underestimate the contextual contingencies that drive their behavior. Refining the self-model—making it more accurate, more granular, and better calibrated to actual contingencies—is the developmental task of self-reflection.

This is the computational role we assign to what we will call *observable model refinement* (OMR). In an adaptive intervention context, OMR consists of explicit, testable hypotheses about the learner's patterns, externalized so that the learner can confirm, deny, or refine them. From the perspective of Active Inference, OMR functions as a structured Bayesian update on the self-model: hypothesis (prior) + evidence (the learner's response) → posterior (revised self-model). When the learner engages with this loop over time, their self-model becomes increasingly accurate and predictive—the formal expression of what Fischer and colleagues described as the construction of higher-tier abstractions about the self (Fischer, 2009).

4.5 Integration: A Single Decision Rule

The four mechanisms above are not independent. They are facets of a single computational problem: which action best minimizes expected free energy, given the agent's current hierarchical generative model, current precision state, available policies, and history-conditioned priors? Formally, we can write:

$$\pi^* = argmin\_\pi\ G(\pi) = argmin\_\pi\ [-\gamma(HRV) \cdot (E[I(\pi \mid hierarchy, history)] + E[U(\pi \mid goals)])]$$

Where γ depends on autonomic state (Mechanism 2), the hierarchy is structured by Fischer levels (Mechanism 1), the priors over actions are shaped by habit (Mechanism 3), and the model being updated includes the self-model (Mechanism 4). The decision rule itself remains the same in all cases: select the policy that minimizes expected free energy. The four mechanisms are not additional terms to be weighted—they are constraints on the structure of the model over which G is computed.

This formulation has a critical methodological implication: weights in the integration are not free parameters to be fit ad hoc. They are determined by the structure of the underlying model. To a substantial degree, the choice between alternative actions is not a matter of how we choose to combine the four mechanisms, but of how the four mechanisms are implemented in the brain in the first place.

5. Worked Example: A Single Decision Cycle

To illustrate how the four mechanisms combine in practice, we walk through a single decision cycle for a hypothetical patient. The example is drawn from a clinical decision support context (the Learning Catcher platform; see Section 7), but the principles apply more generally.

5.1 Patient Description

The patient is a 35-year-old adult presenting with longstanding attention difficulties consistent with adult ADHD. Initial assessment placed the patient at Fischer Level 9 (single abstractions) for working memory and Level 10 (abstract mappings) for sustained attention. Resting HRV (RMSSD) was 35 ms (within typical adult range). The patient has used the intervention daily for three weeks, generating sufficient personal history to constitute meaningful empirical priors.

5.2 The Decision

On a Wednesday afternoon, the patient opens the app. The system gathers current observations and runs through the five steps illustrated in Figure 6.

Figure 6

Worked example: single decision cycle.

Note. Sequential application of the four mechanisms in a single intervention decision. Sensing yields current state; locating identifies position in the hierarchical model; deciding involves hard-filtering, soft-scoring, and reranking; deliver enacts the policy; update revises the model.

Figure 6. Worked Example: One Decision Cycle for a Patient (Adult, mid-development, ADHD presentation)

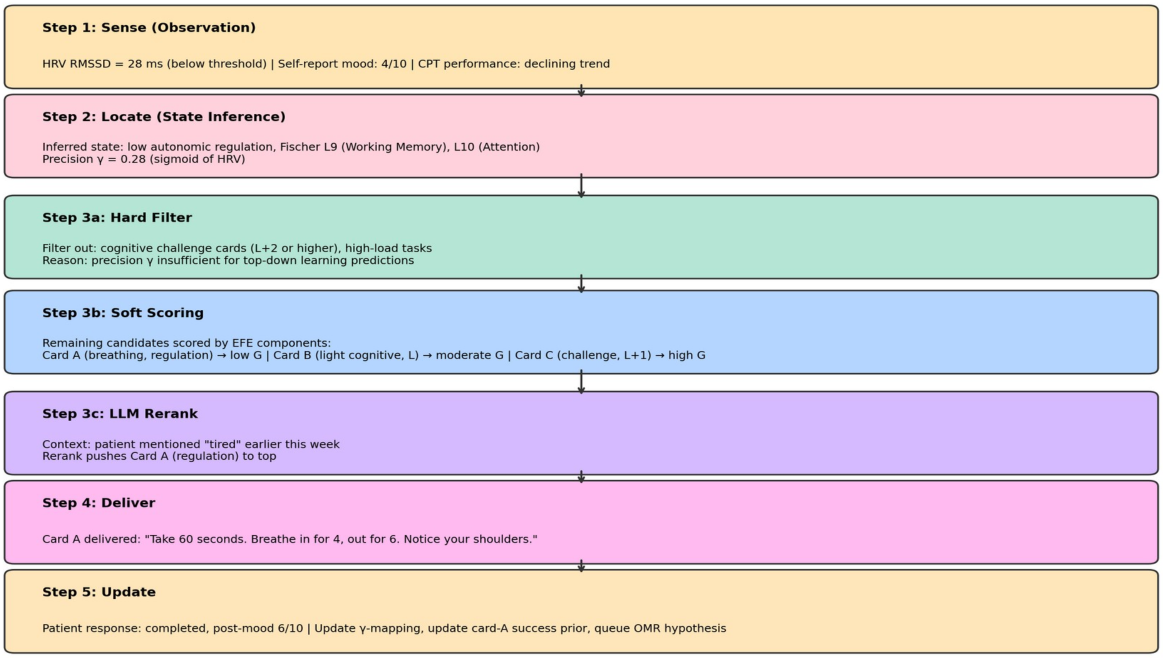

5.3 Interpretation

Note that the system did not choose Card A because of a hard rule (“if HRV < 30, give regulation”). It chose Card A because, given the current state of the generative model, Card A had the lowest expected free energy. With sufficient HRV (γ), Card C might have had the lowest G, because its high expected information gain would outweigh its initial cost. With ample history and an established habit of breathing exercises being unhelpful at this time of day (Mechanism 3), the LLM rerank might have suppressed Card A in favor of a different regulation card. The decision rule remained constant; what changed was the structure of the model over which it was computed.

This is the practical sense in which the framework is integrated rather than merely additive. Each mechanism shapes the model in a different way; the decision rule is invariant.

6. Self-Regulated Learning as Free Energy Minimization in Time

Having developed the four-mechanism integration, we return to Zimmerman's SRL framework and show how its three phases—forethought, performance, and self-reflection—correspond to natural stages of the Active Inference cycle. Figure 7 illustrates the mapping.

Figure 7

Zimmerman SRL phases mapped to Active Inference mechanisms.

Note. Forethought, performance, and self-reflection correspond to anticipated EFE computation, precision-gated action selection, and posterior model update, respectively. The three phases form a unified inference cycle.

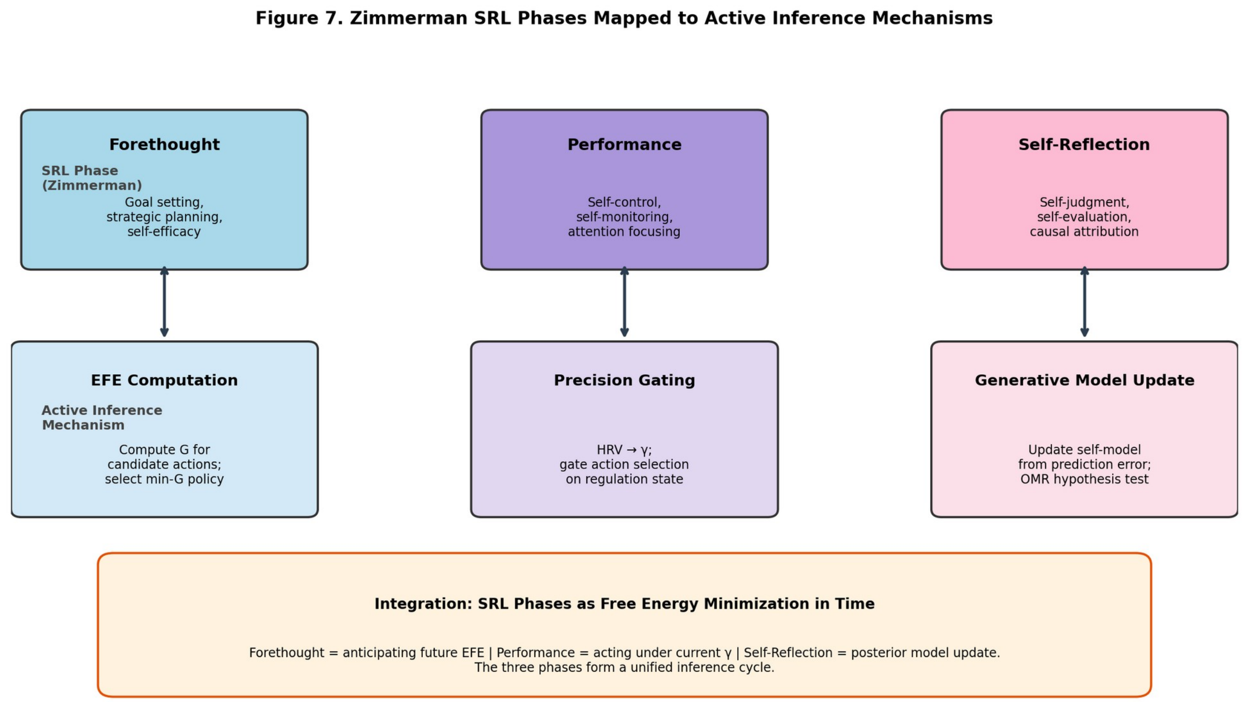

6.1 Forethought as Expected Free Energy Computation

In Zimmerman's framework, forethought encompasses goal setting, strategic planning, and motivational beliefs. From the standpoint of Active Inference, these constitute the agent's explicit representation of pragmatic value—what outcomes are preferred, and what policies are expected to achieve them. The computational task of forethought is to enumerate candidate policies and evaluate their expected free energy.

6.2 Performance as Precision-Gated Action

Performance—self-control, self-monitoring, and attention focusing—corresponds to the actual selection and execution of policies under the current precision state. Failures of performance often reflect not a deficiency of will, but a mismatch between the policies the agent has committed to in forethought and the precision γ available in the moment. Interventions targeting performance, on this view, should target γ as well as the policy itself.

6.3 Self-Reflection as Posterior Update

Self-reflection—self-judgment, self-evaluation, and causal attribution—corresponds to the agent updating its generative model on the basis of observed outcomes. Critically, this update includes the self-model itself: not only “did that action work,” but “what does it say about me that I responded that way?” Effective self-reflection is precisely the structured updating of the self-model in light of evidence, and OMR (Section 4.4) is the operational form of this update in an adaptive intervention context.

6.4 The Time Horizon Problem

A subtle but important consequence of this mapping concerns the time horizon over which free energy is minimized. The agent can minimize immediate G by resting, by avoiding challenge, by retreating to familiar routines. But the agent can also project G into the future and recognize that immediate minimization—at sufficient extension in time—produces long-term G accumulation (decline, stagnation, loss of skill). Figure 5 illustrates this dynamic.

Figure 5

Time horizon and self-regulated learning.

Note. Immediate vs. long-term G trajectories for resting and learning policies. Mature SRL involves selecting actions with higher immediate G to minimize long-term G — the computational basis of developmental progression.

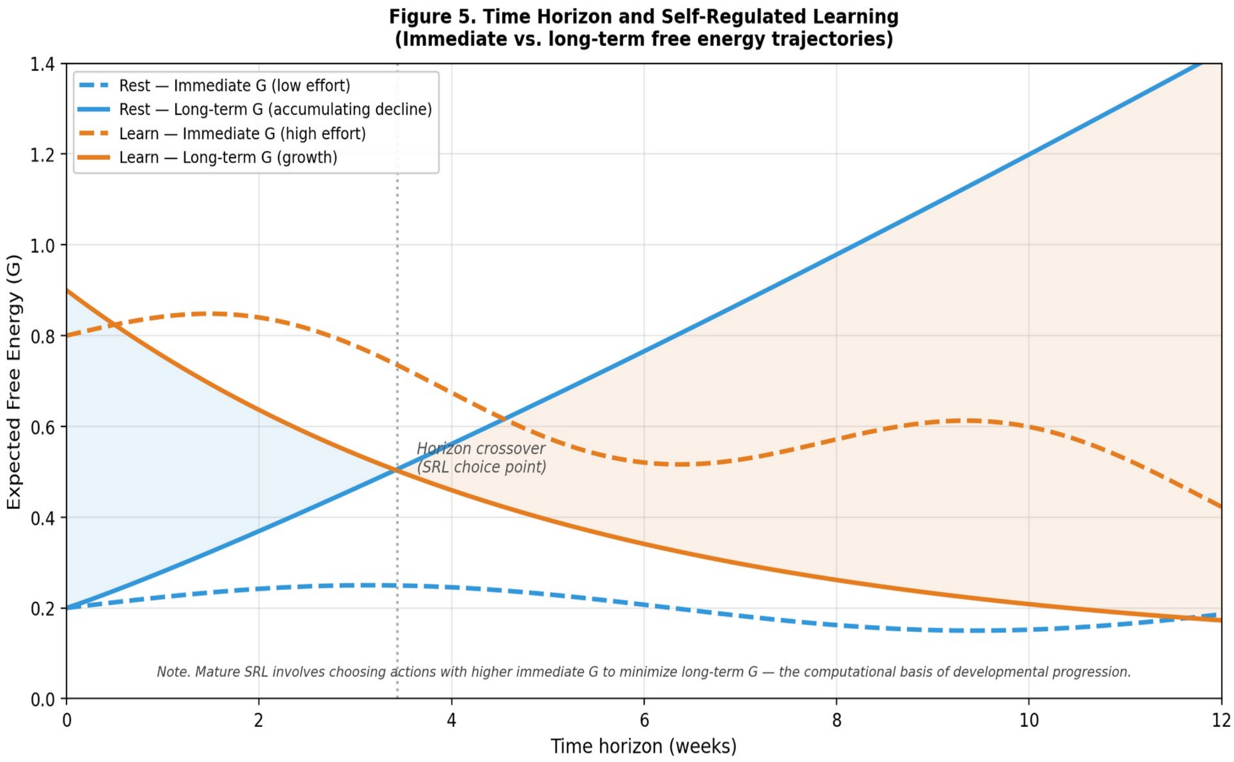

Self-regulated learning, in this view, is the capacity to act on the basis of long-horizon G rather than immediate G. This is not merely a psychological observation but a computational one: it requires hierarchical generative models with sufficient depth and stability to represent the future self and to discount its predictions appropriately. Development, in turn, can be characterized as the gradual extension of the effective time horizon over which the agent computes G.

7. Implementation and Empirical Predictions

The framework we have developed yields specific, testable predictions and provides a structure for an implementation. In this section we sketch both, with reference to the Learning Catcher platform under development at Boston Neuromind LLC.

7.1 The Nine-Loop Architecture

The Learning Catcher platform implements the framework through a nine-loop architecture, shown in Figure 4. Each loop corresponds to a component of the Active Inference cycle: perception, inference, action selection, action, learning, and self-model refinement. Loops 1–2 (Sense, Locate) constitute perception; Loop 3 (Decide) performs the EFE-based action selection; Loop 4 (Deliver) enacts the policy; Loop 5 (Update) revises beliefs; Loops 6–9 implement clinician-side review, feedback, audit, and engagement features.

Figure 4

Nine-loop architecture mapped to Active Inference components.

Note. Each loop in the implementation corresponds to a component of the Active Inference cycle. The architecture is iterative; loops cycle continuously rather than executing once.

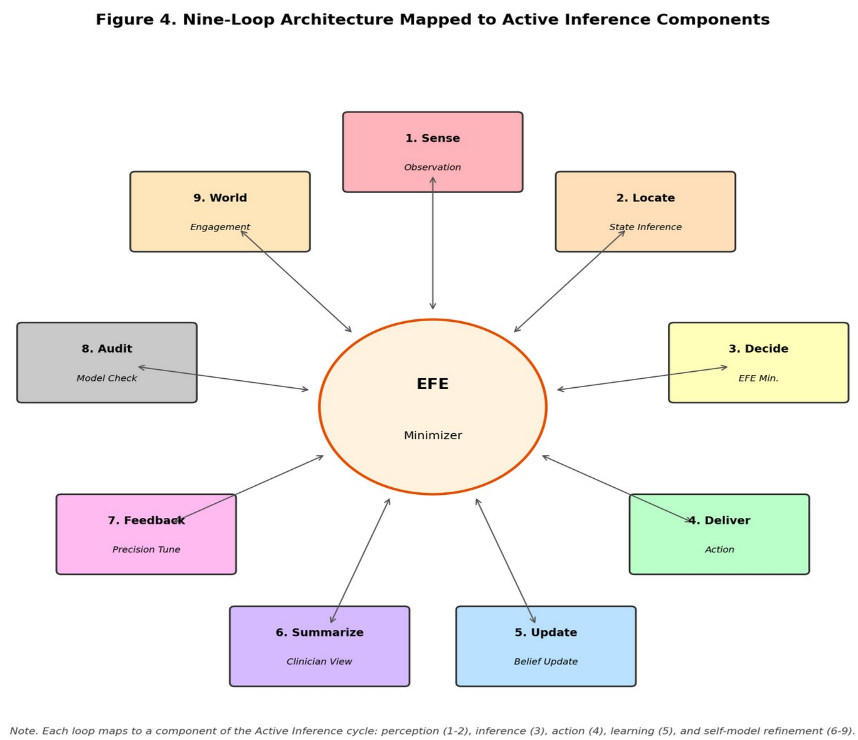

7.2 Measurement Strategy

Table 2 specifies a measurement strategy for each of the four mechanisms, along with the recommended instrument and frequency of measurement.

Table 2

Measurement Strategy for the Four Mechanisms

Note. For each mechanism, the corresponding construct, recommended instrument, and measurement frequency. RMSSD = root mean square of successive differences in interbeat intervals; CPT = continuous performance test; OMR = observable model refinement.

Mechanism	Construct	Instrument	Frequency
Mechanism 1 (Hierarchy)	Fischer level per domain	Clinician-rated dimensional ladder; behavioral tasks (CPT, N-back)	Monthly (initial), then every 3 months
Mechanism 2 (Precision)	HRV (RMSSD)	Photoplethysmography or ECG; 5-min resting	Each session (clinician-administered); daily (consumer device)
Mechanism 3 (Habit)	Cue–routine–reward alignment	Self-report timing log; completion rate by time of day	Daily (in-app); reviewed weekly
Mechanism 4 (Self-Model)	OMR confirmation rate; pattern accuracy	In-app OMR prompts; clinician interview	Weekly (in-app); monthly review

7.3 Empirical Predictions

The framework yields several predictions that can be tested with relatively modest sample sizes, using single-case experimental designs or pilot studies. Table 3 lists five predictions, each with an associated test.

Table 3

Empirical Predictions of the Framework

Note. For each prediction, the proposed empirical test and minimum design requirement. SCED = single-case experimental design; AB = baseline–intervention; ABAB = baseline–intervention–baseline–intervention.

Prediction	Test	Minimum Design
P1: HRV predicts card completion within-subject	Within-subject correlation of pre-session RMSSD with completion rate	$N \geq 5$, daily measurement, 4–8 weeks

Prediction	Test	Minimum Design
P2: Optimal challenge (L+1) maximizes information gain	Crossover design with cards at L−1, L, L+1, L+2; measure post-session skill change	SCED, ABAB, N ≥ 5
P3: OMR confirmation rate increases over time	Track OMR confirmation/refinement rate weekly; expect logistic-growth pattern	Time-series, N ≥ 5, ≥12 weeks
P4: Habit-aligned timing improves completion	Compare completion at habit-aligned vs. arbitrary times	Within-subject AB, N ≥ 5
P5: Long-horizon framing increases challenge acceptance	Manipulate framing of card (immediate vs. long-term consequence) and measure acceptance	Vignette experiment or within-subject design

7.4 A Sequenced Research Program

We propose a four-stage research program. Stage 1 (months 0–6): pilot single-case experimental designs with N = 5 to test P1 (HRV–completion link), as this prediction has the strongest existing literature and the lowest technical barrier. Stage 2 (months 6–18): expand to N = 20–30 across multiple clinics, allowing tests of P2 and P3 with adequate statistical power. Stage 3 (months 18–36): longitudinal evaluation of developmental progression (Mechanism 1), with measurement quarterly across 12+ months. Stage 4 (months 36+): randomized controlled trial comparing the integrated framework against active controls (e.g., a habit-only intervention, a regulation-only intervention) to isolate the contribution of integration.

8. Discussion

8.1 Theoretical Contributions

The primary theoretical contribution of this framework is the demonstration that four traditionally separate accounts of human regulation and development can be unified under a single computational principle. This integration is not merely an exercise in conceptual housekeeping. It has substantive consequences for how we think about intervention design, measurement, and validation. In particular, it provides a principled rationale for why interventions should integrate physiological monitoring, developmental assessment, habit-aware design, and self-model refinement, rather than treating these as separate tracks.

A second contribution is the explicit treatment of development within the Active Inference framework. Most existing applications of Active Inference treat the agent's generative model as static, or as updated only at the level of beliefs about specific contingencies. The present framework treats the structure of the model itself as developmentally elaborated, with Fischer levels providing empirical specification of the resulting layers.

A third contribution is the operationalization of self-reflection as Bayesian updating on the self-model, mediated by structured externalized hypotheses (OMR). This bridges Zimmerman's phenomenological account of self-reflection with the computational account of inference on a generative model, and suggests specific design principles for interventions that aim to support self-reflection.

8.2 Limitations

Several limitations should be acknowledged. First, the framework as presented is theoretical; the empirical work proposed in Section 7 has not yet been conducted, and the predictions are accordingly tentative. Second, the mapping between HRV and γ, while supported by independent literature (Smith et al., 2017; Thayer & Lane, 2009), is itself a hypothesis rather than an established equivalence; the specific functional form of the mapping requires empirical determination. Third, our characterization of Fischer levels as layers of a hierarchical generative model is provisional; in particular, the relationship between Fischer's domain-specificity and the more modular structure of generative models in machine learning implementations requires further development. Fourth, the framework as presented assumes a single agent; group-level dynamics (peer learning, family systems, classroom effects) are not addressed.

8.3 A Note on Falsifiability

The free energy principle has been criticized as unfalsifiable on the grounds that any behavior can be reconstructed as free energy minimization given an appropriately chosen generative model (Williams, 2018; cf. Friston et al., 2023 for response). We take this concern seriously, but note two things. First, the present framework specifies concrete constraints on the generative model (Fischer hierarchical structure, HRV-precision mapping, habit-policy correspondence, OMR-mediated self-model update) that are themselves testable. The framework is falsifiable to the extent that these mappings can be tested. Second, the predictions in Table 3 are independent of the broader question of whether the free energy principle is itself a complete theory of brain function. P1 (HRV predicts completion) does not require commitment to Active Inference; it follows from independent empirical literature. The proposed integration is one in which Active Inference provides organization, but does not bear the full empirical weight of the framework.

8.4 Implications for Clinical and Educational Practice

The framework has several implications for practice. First, interventions should be precision-gated: regulation should precede challenge, not because of clinical heuristic, but because the computational substrate of learning requires it. Second, developmental assessment should be domain-specific and multidimensional; a single global score cannot capture the structure of an individual's generative model. Third, habit alignment is not a peripheral consideration but a primary constraint on intervention feasibility. Fourth, self-reflection requires scaffolding; expecting learners to construct accurate self-models without structured input is unrealistic.

A final implication concerns the role of the clinician or instructor. In the framework we have developed, the clinician serves as an external precision regulator and a co-author of OMR hypotheses—not as a deliverer of content. This is consistent with longstanding traditions in psychotherapy (e.g., the therapeutic alliance literature; Wampold, 2015) and education (Vygotsky, 1978), but provides a computational rationale for these traditions.

9. Conclusion

We have proposed Developmental Active Inference (DAI) as a framework for understanding self-regulated learning as a developmentally structured form of free energy minimization. The framework integrates four traditionally separate accounts—Fischer's Dynamic Skill Theory, Polyvagal Theory, habit theory, and Bayesian models of individual difference—as natural decompositions of expected free energy under a hierarchical generative model. The framework yields a single decision rule for adaptive intervention, specific empirical predictions, and a structured implementation in a clinical decision support platform.

Three claims summarize the contribution. First, Active Inference provides a computational substrate that has been missing in developmental theories of self-regulated learning. Second, Fischer's Dynamic Skill Theory, supplemented by precision-based regulation and habit-aware design, provides empirical content that is missing in computational treatments of Active Inference. Third, the integration of these is not merely additive: it generates new predictions (particularly about the time horizon of free energy minimization) that neither tradition produces on its own.

The framework is in its infancy. The next decade of empirical work will determine whether the integration we have proposed is productive. We have sketched a sequenced research program to evaluate it. In the meantime, the framework offers practitioners a coherent theoretical foundation for the integrative interventions many already deliver intuitively, and offers researchers a structure within which apparently disparate findings can be related.

References

Apps, M. A. J., & Tsakiris, M. (2014). The free-energy self: A predictive coding account of self-recognition. Neuroscience & Biobehavioral Reviews, 41, 85–97. https://doi.org/10.1016/j.neubiorev.2013.01.029

Bandura, A. (2005). The primacy of self-regulation in health promotion. Applied Psychology, 54(2), 245–254. https://doi.org/10.1111/j.1464-0597.2005.00208.x

Boekaerts, M., & Corno, L. (2005). Self-regulation in the classroom: A perspective on assessment and intervention. Applied Psychology, 54(2), 199–231. https://doi.org/10.1111/j.1464-0597.2005.00205.x

Duhigg, C. (2012). The power of habit: Why we do what we do in life and business. Random House.

Fischer, K. W. (1980). A theory of cognitive development: The control and construction of hierarchies of skills. Psychological Review, 87(6), 477–531. https://doi.org/10.1037/0033-295X.87.6.477

Fischer, K. W. (2009). Mind, brain, and education: Building a scientific groundwork for learning and teaching. Mind, Brain, and Education, 3(1), 3–16. https://doi.org/10.1111/j.1751-228X.2008.01054.x

Fischer, K. W., & Bidell, T. R. (2006). Dynamic development of action and thought. In W. Damon & R. M. Lerner (Eds.), Handbook of child psychology: Vol. 1. Theoretical models of human development (6th ed., pp. 313–399). Wiley.

Fleming, S. M., & Lau, H. C. (2014). How to measure metacognition. Frontiers in Human Neuroscience, 8, 443. https://doi.org/10.3389/fnhum.2014.00443

Friston, K. (2008). Hierarchical models in the brain. PLoS Computational Biology, 4(11), e1000211. https://doi.org/10.1371/journal.pcbi.1000211

Friston, K. (2010). The free-energy principle: A unified brain theory? Nature Reviews Neuroscience, 11(2), 127–138. https://doi.org/10.1038/nrn2787

Friston, K., Da Costa, L., Sajid, N., Heins, C., Ueltzhöffer, K., Pavliotis, G. A., & Parr, T. (2023). The free energy principle made simpler but not too simple. Physics Reports, 1024, 1–29. https://doi.org/10.1016/j.physrep.2023.07.001

Friston, K. J., FitzGerald, T., Rigoli, F., Schwartenbeck, P., & Pezzulo, G. (2017). Active inference: A process theory. Neural Computation, 29(1), 1–49. https://doi.org/10.1162/NECO_a_00912

Granott, N., & Parziale, J. (Eds.). (2002). Microdevelopment: Transition processes in development and learning. Cambridge University Press.

Lally, P., van Jaarsveld, C. H. M., Potts, H. W. W., & Wardle, J. (2010). How are habits formed: Modelling habit formation in the real world. European Journal of Social Psychology, 40(6), 998–1009. https://doi.org/10.1002/ejsp.674

Lamb, M. E., Mascolo, M. F., & Fischer, K. W. (2002). The development of emotion regulation. In Handbook of dynamics in parent-child relations (pp. 31–56). Sage.

Lawson, R. P., Rees, G., & Friston, K. J. (2014). An aberrant precision account of autism. Frontiers in Human Neuroscience, 8, 302. https://doi.org/10.3389/fnhum.2014.00302

Marlatt, G. A., & Gordon, J. R. (Eds.). (1985). Relapse prevention: Maintenance strategies in the treatment of addictive behaviors. Guilford Press.

Parr, T., Pezzulo, G., & Friston, K. J. (2022). Active inference: The free energy principle in mind, brain, and behavior. MIT Press.

Pezzulo, G., Rigoli, F., & Friston, K. J. (2018). Hierarchical active inference: A theory of motivated control. Trends in Cognitive Sciences, 22(4), 294–306. https://doi.org/10.1016/j.tics.2018.01.009

Pintrich, P. R. (2004). A conceptual framework for assessing motivation and self-regulated learning in college students. Educational Psychology Review, 16(4), 385–407. https://doi.org/10.1007/s10648-004-0006-x

Porges, S. W. (2007). The polyvagal perspective. Biological Psychology, 74(2), 116–143. https://doi.org/10.1016/j.biopsycho.2006.06.009

Sitzmann, T., & Ely, K. (2011). A meta-analysis of self-regulated learning in work-related training and educational attainment: What we know and where we need to go. Psychological Bulletin, 137(3), 421–442. https://doi.org/10.1037/a0022777

Smith, R., Friston, K. J., & Whyte, C. J. (2022). A step-by-step tutorial on active inference and its application to empirical data. Journal of Mathematical Psychology, 107, 102632. https://doi.org/10.1016/j.jmp.2021.102632

Smith, R., Thayer, J. F., Khalsa, S. S., & Lane, R. D. (2017). The hierarchical basis of neurovisceral integration. Neuroscience & Biobehavioral Reviews, 75, 274–296. https://doi.org/10.1016/j.neubiorev.2017.02.003

Stephan, K. E., Manjaly, Z. M., Mathys, C. D., Weber, L. A. E., Paliwal, S., Gard, T., Tittgemeyer, M., Fleming, S. M., Haker, H., Seth, A. K., & Petzschner, F. H. (2016). Allostatic self-efficacy: A metacognitive theory of dyshomeostasis-induced fatigue and depression. Frontiers in Human Neuroscience, 10, 550. https://doi.org/10.3389/fnhum.2016.00550

Sterzer, P., Adams, R. A., Fletcher, P., Frith, C., Lawrie, S. M., Muckli, L., Petrovic, P., Uhlhaas, P., Voss, M., & Corlett, P. R. (2018). The predictive coding account of psychosis. Biological Psychiatry, 84(9), 634–643. https://doi.org/10.1016/j.biopsych.2018.05.015

Thayer, J. F., & Lane, R. D. (2009). Claude Bernard and the heart-brain connection: Further elaboration of a model of neurovisceral integration. Neuroscience & Biobehavioral Reviews, 33(2), 81–88. https://doi.org/10.1016/j.neubiorev.2008.08.004

Tschantz, A., Seth, A. K., & Buckley, C. L. (2020). Learning action-oriented models through active inference. PLoS Computational Biology, 16(4), e1007805. https://doi.org/10.1371/journal.pcbi.1007805

van der Kolk, B. A. (2014). The body keeps the score: Brain, mind, and body in the healing of trauma. Viking.

Vygotsky, L. S. (1978). Mind in society: The development of higher psychological processes. Harvard University Press.

Wampold, B. E. (2015). How important are the common factors in psychotherapy? An update. World Psychiatry, 14(3), 270–277. https://doi.org/10.1002/wps.20238

Williams, D. (2018). Predictive coding and thought. Synthese, 197, 1749–1775. https://doi.org/10.1007/s11229-018-1768-x

Wood, W., & Rünger, D. (2016). Psychology of habit. Annual Review of Psychology, 67, 289–314. https://doi.org/10.1146/annurev-psych-122414-033417

Zimmerman, B. J. (2000). Attaining self-regulation: A social cognitive perspective. In M. Boekaerts, P. R. Pintrich, & M. Zeidner (Eds.), Handbook of self-regulation (pp. 13–39). Academic Press.