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

Alex Kwon[1]

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

Author Note

Alex Kwon is a Board Certified Neurofeedback (BCN) practitioner, holds a PhD in Instructional Design and a Master's degree in Mental Health Counseling, and is a former Visiting Scholar at the Harvard Graduate School of Education, where he conducted collaborative research with the late Professor Kurt W. Fischer on dynamic skill theory and educational neuroscience. Correspondence concerning this article should be addressed to Alex Kwon, Boston Neuromind LLC, Canton, MA, USA. Email: bostonneuromind@gmail.com

May 2026 — Working Draft

Abstract

Self-regulated learning (SRL; Zimmerman, 2000) is a foundational construct in educational and clinical psychology, yet existing frameworks describe its phases phenomenologically rather than mechanistically. The Active Inference framework (Friston, 2010; Parr et al., 2022), grounded in the free energy principle, offers a computational account of how organisms minimize prediction error through perception and action. While Active Inference has been applied to learning (Schwartenbeck et al., 2019), psychiatry (Smith et al., 2021), and habit formation (Friston et al., 2016), it has not been integrated with developmental theory in a manner that addresses stage-dependent learning. We propose Developmental Active Inference (DAI) as a computational framework that integrates four traditionally separate domains: (a) Fischer's dynamic skill theory (DST; Fischer, 1980) as hierarchical generative model structure, (b) polyvagal theory (Porges, 2007) and heart rate variability (HRV) as precision (γ) parameter, (c) habit formation theory as learned action policies, and (d) Bayesian personal posterior updating as empirical prior accumulation. We derive the expected free energy (EFE) decomposition for adaptive intervention selection and demonstrate that the four domains constitute natural components of a single computational objective rather than arbitrarily weighted features. We further introduce Observable Model Refinement (OMR) as a clinical operationalization of self-model updating. The framework yields testable predictions about when learners engage, disengage, regulate, and progress. We discuss implementation in Learning Catcher, a clinical decision support system currently in beta, and outline an empirical validation roadmap including single-case experimental designs and prospective cohort studies. The framework contributes to clinical psychology by providing a theory-driven alternative to data-driven recommendation systems, and to developmental science by linking stage-dependent learning to a computational neuroscience foundation.

Keywords: Active Inference; self-regulated learning; dynamic skill theory; heart rate variability; precision; clinical decision support; developmental psychology

1. Introduction

Self-regulated learning (SRL) is among the most influential constructs in contemporary educational psychology. Originally formalized by Zimmerman (1989, 2000), SRL describes a cyclical process in which learners proactively set goals, monitor performance, and reflect on outcomes. Decades of empirical work have established that SRL skills predict academic achievement (Dignath & Büttner, 2008), therapeutic outcomes (Maes & Karoly, 2005), and lifespan adaptation (Heckhausen et al., 2010). Despite this generative success, the field faces a recurring methodological challenge: SRL is described phenomenologically, in terms of phases and components, rather than mechanistically, in terms of computational processes that generate the observed phenomena.

This descriptive-mechanistic gap is consequential for clinical practice. When a clinician seeks to support a patient's regulatory capacity—whether for academic skill, emotion regulation, or therapeutic engagement—the SRL literature provides rich vocabulary (forethought, performance, self-reflection) but limited algorithmic guidance about which intervention to deliver, in what sequence, to which person, at what moment. Practitioners are left to integrate SRL principles with their own clinical judgment, a process that is by nature variable and difficult to communicate, replicate, or improve.

In parallel, computational neuroscience has converged on a candidate unifying framework for adaptive behavior: the free energy principle (FEP; Friston, 2010) and its operationalization as Active Inference (Friston et al., 2017; Parr et al., 2022). Active Inference proposes that organisms minimize the long-run average of variational free energy—a quantity that bounds the surprise of sensory observations under a generative model of the world. From this single principle, perception, action, attention, and learning emerge as different routes to the same end: reducing prediction error. The framework has been productively applied to motor control (Adams et al., 2013), psychiatric phenomena (Schwartenbeck et al., 2015), habit formation (Friston et al., 2016), and decision making (FitzGerald et al., 2015), among many other domains.

Active Inference is, in principle, well-suited to ground SRL. Both frameworks centrally feature the minimization of discrepancy between expected and observed outcomes; both posit hierarchical generative structure; both emphasize the role of action in disambiguating uncertainty. Yet a direct integration has not been advanced. Existing attempts to apply Active Inference to learning (Schwartenbeck et al., 2019; Tschantz et al., 2020) focus on moment-to-moment policy selection within a fixed model, leaving aside the question that most concerns educators and clinicians: how does the model itself develop?

The challenge of developmental change is not new. Dynamic skill theory (DST; Fischer, 1980; Fischer & Bidell, 2006) was developed precisely to characterize how skill emerges, consolidates, and transforms across the lifespan. DST organizes skill development into a series of levels—reflexes, single actions, action mappings, action systems, single representations, representational mappings, representational systems, single abstractions, abstract mappings, abstract systems, and principles—and identifies optimal and functional levels within each domain. While DST has been extensively validated empirically (Fischer & Yan, 2002; Schwartz & Fischer, 2004), it has remained largely outside computational modeling: it specifies the structure of developmental change without specifying the computational mechanism by which that change is realized.

This paper addresses the integration gap. We propose Developmental Active Inference (DAI), a framework that links Active Inference's computational machinery to DST's developmental architecture, with two additional bridges to clinical practice: polyvagal theory (Porges, 2007) as a physiological grounding for precision, and contemporary habit theory (Wood & Rünger, 2016; Clear, 2018) as a structural account of learned policies. We further introduce Observable Model Refinement (OMR), a clinical procedure for externalizing and revising a patient's generative self-model. The result is a unified computational framework in which (a) regulated learning is moment-to-moment expected free energy minimization, and (b) developmental regulated learning is structural transformation of the generative model itself.

The contributions of this paper are five. First, we provide a conceptual mapping between SRL phases (Zimmerman, 2000), Active Inference components (Parr et al., 2022), and DST levels (Fischer & Bidell, 2006). Second, we derive the expected free energy decomposition appropriate for clinical decision support, showing how four traditionally separate clinical inputs—Fischer level, HRV, habit alignment, and personal history—are not arbitrary features to be weighted but natural components of a single computational objective. Third, we propose four developmental mechanisms by which the generative model itself is restructured: hierarchical expansion, precision reallocation, empirical prior accumulation, and self-model refinement. Fourth, we describe the Learning Catcher system, a clinical implementation currently in beta evaluation, and use it as a worked example to illustrate the framework. Fifth, we identify empirical predictions and outline a validation roadmap including single-case experimental designs and prospective cohort studies.

1.1 The Problem of Mechanism in Self-Regulated Learning

Zimmerman's (2000) three-phase model of SRL—forethought, performance, and self-reflection—provides an enormously productive descriptive vocabulary. Within forethought, learners engage in task analysis (goal setting, strategic planning) and self-motivation beliefs (self-efficacy, outcome expectations). Within performance, learners exercise self-control (task strategies, self-instruction) and self-observation (metacognitive monitoring, self-recording). Within self-reflection, learners engage in self-judgment (self-evaluation, causal attribution) and self-reaction (self-satisfaction, adaptive inferences). Each phase has been operationalized through multiple validated instruments, and the model has supported a substantial empirical literature.

Yet a clinician asking, "Why is this patient not progressing?" finds that the model describes the absence (e.g., "poor goal setting," "inadequate self-monitoring") more readily than it explains it. Two patients may exhibit identical SRL profiles on standard measures while differing fundamentally in the computational basis of their regulatory difficulty. One may suffer from a generative model that cannot represent the goal hierarchy required for the task; another may have an intact model but insufficient autonomic regulation to deploy it. The two require different interventions, yet the descriptive framework does not distinguish them.

This problem is not unique to SRL. Many constructs in psychology—working memory, executive function, attention—have moved from purely behavioral description to computational characterization over the past two decades (Marr, 1982; Niv, 2009). The shift typically yields three benefits: (a) mechanistically distinct phenomena that previously appeared similar become distinguishable; (b) interventions can be targeted at specific computational components; and (c) cross-domain integration becomes possible. Applying these benefits to SRL is the principal motivation of the present work.

1.2 Active Inference as a Candidate Mechanism

Active Inference, in its modern formulation (Parr et al., 2022), holds that an organism is best understood as a generative model of its environment. The model implements predictions about sensory inputs, and the organism acts to minimize the divergence between predicted and observed inputs—the prediction error, formalized as variational free energy. Critically, action selection is governed by expected free energy (EFE): the model selects policies (sequences of actions) that, in expectation, will minimize free energy across some horizon.

Expected free energy decomposes into two terms with intuitive interpretations. The first, often termed epistemic value or expected information gain, quantifies how much a policy is expected to reduce uncertainty about hidden states or model parameters. The second, often termed pragmatic value or expected utility, quantifies how much a policy is expected to lead to preferred outcomes. The relative weighting of these two terms is itself modulated by precision parameters that reflect the organism's current confidence in its predictions.

This decomposition maps with remarkable cleanness onto questions a clinician asks at the moment of intervention selection. "What will this patient learn from doing this task?" is a question about epistemic value. "Does this task move the patient toward their stated goal?" is a question about pragmatic value. "Is the patient currently able to act on these predictions with appropriate confidence?" is a question about precision. The clinical reasoning that practitioners describe phenomenologically—weighing learning, goal-progress, and readiness—has a precise computational counterpart.

Crucially, Active Inference is not a recommendation algorithm of the kind familiar from consumer technology (e.g., collaborative filtering, deep learning—based recommenders). It is a theory of how purposive action is generated from a model of the world. This distinction matters for clinical translation: a recommendation system trained on aggregated user data optimizes a population objective (e.g., engagement, click-through), whereas an Active Inference agent optimizes its own informational and pragmatic objectives relative to an explicit generative model. The latter is interpretable, principled, and accountable in a way the former is not—properties that are essential for clinical use.

1.3 The Developmental Gap

Despite its broad applicability, the Active Inference literature has paid limited attention to how generative models themselves develop. Most published applications either assume a fixed model and study policy learning within it, or model parameter updates within a fixed structural form. The structural transformation of the model—the emergence of new representational levels, the differentiation and integration that characterize development—remains underdeveloped (though see Pezzulo et al., 2018, for hierarchical extensions, and Constant et al., 2018, for cultural learning).

This is precisely the gap that DST has spent four decades characterizing. Fischer (1980) proposed that skill development proceeds through a sequence of structural reorganizations, in which new levels of skill emerge from the coordination and abstraction of existing levels. Each level is defined by the type of unit it operates on (action, representation, or abstraction) and by the complexity of the structural form (single elements, mappings between elements, or systems of mappings). The transitions between levels are not gradual quantitative increases but qualitative restructurings.

DST also identifies a critical empirical regularity: skill performance is highly context-dependent (Fischer & Bidell, 2006). A person who exhibits abstract reasoning in one domain may operate at the representational level in another, and the level achieved depends on environmental support. DST distinguishes the optimal level (the highest level achievable with full contextual support) from the functional level (the level typically exhibited in everyday activity). The gap between these is the developmental space in which scaffolded intervention operates—what Vygotsky (1978) called the zone of proximal development.

Integrating DST with Active Inference yields a coherent computational picture of developmental SRL. The generative model has a hierarchical structure whose depth corresponds to Fischer's developmental levels. Precision parameters allocate the model's confidence across levels, modulating which level is functionally available at a given moment. Learning at a level corresponds to parameter updating; transition between levels corresponds to structural change in the model. Self-regulation, in this view, is not separate from learning—it is the moment-to-moment expression of the same process by which the model develops over the long run.

1.4 Overview of the Framework

The remainder of this paper develops Developmental Active Inference in five parts. Section 2 reviews Active Inference at a level of detail sufficient to support the integration, with emphasis on the EFE decomposition and the role of precision. Section 3 reviews DST, focusing on the hierarchical structure of skill levels and the optimal-functional distinction. Section 4 develops the integration in four mechanisms: (4.1) hierarchical generative model structure, mapping Fischer levels to model depth; (4.2) precision reallocation, linking polyvagal theory and HRV to the γ parameter; (4.3) empirical prior accumulation, linking habit theory to learned action policies; and (4.4) self-model refinement, introducing Observable Model Refinement (OMR) as a clinical procedure. Section 5 provides a worked example using the Learning Catcher system. Section 6 derives empirical predictions and outlines a validation roadmap. Section 7 concludes.

A framework of this scope necessarily simplifies the literature it draws on. We have aimed for theoretical fidelity to each source domain while maintaining computational coherence in the integration. Where mathematical formalism is required, we have provided it; where formalism would obscure rather than clarify, we have relied on natural language and diagrams. The framework is intended to be both rigorous enough to generate testable predictions and accessible enough to guide clinical practice.

2. Active Inference: A Brief Review

This section reviews Active Inference at the level of detail required for the integration that follows. Comprehensive treatments are available in Parr et al. (2022), Smith et al. (2022, tutorial), and Friston et al. (2017). Our aim here is to establish notation and intuition for four constructs that recur throughout the rest of the paper: the generative model, variational free energy, expected free energy and its decomposition, and the precision parameter γ.

2.1 The Free Energy Principle

The free energy principle (FEP; Friston, 2010; Friston et al., 2006) holds that any self-organizing system that maintains a non-equilibrium steady state with its environment must, in effect, minimize a quantity that upper-bounds the surprise of sensory observations. Mathematically, surprise is the negative log-probability of an observation, $-\log p(o)$, under a model of how observations are generated. Because $p(o)$ is generally intractable to compute directly (it requires marginalizing over all possible hidden causes), the FEP proposes that organisms instead minimize a tractable upper bound: variational free energy F.

Variational free energy depends on a generative model $p(o, s)$—the joint probability of observations o and hidden states s under the organism's model of the world—and on an approximate posterior $q(s)$ over hidden states. The defining property of F is that minimizing it with respect to q simultaneously achieves two ends: it makes q a better approximation of the true posterior $p(s|o)$, and it makes the chosen actions more consistent with the model's predictions. These two effects—perceptual inference and action selection—are unified under the single quantity F.

For our purposes, the technical details of variational free energy matter less than the conceptual implication: the brain is engaged, continuously and in every modality, in resolving the discrepancy between what its model predicts and what it observes. Learning, regulation, and action are different routes to the same goal of discrepancy reduction. This unification is what makes Active Inference an attractive candidate for grounding a multi-component theory like

SRL: rather than positing separate mechanisms for cognition, emotion, motivation, and behavior, the framework proposes that all are manifestations of free energy minimization under different aspects of the generative model.

2.2 The Generative Model

The generative model is the central construct in Active Inference. It is the organism's implicit theory of how observations are generated from hidden states, and how hidden states evolve over time given actions. In discrete-state formulations (which are most relevant for our clinical application), the model specifies:

- A likelihood mapping A: $p(o|s)$, specifying which observations are expected under each hidden state.
- A transition mapping B: $p(s'|s, \pi)$, specifying how hidden states evolve under each policy π (sequence of actions).
- Prior preferences C: $p(o)$, specifying which observations the organism prefers (i.e., is biased to predict).
- Initial state priors D: $p(s_0)$, specifying initial beliefs about hidden states.
- Precision parameters γ, controlling the confidence in policy selection.

In hierarchical formulations (Pezzulo et al., 2018; Friston, Parr, & de Vries, 2017), the generative model has multiple levels, each operating at a different temporal scale and level of abstraction. Higher levels generate predictions for lower levels, and lower levels return prediction errors to higher levels. This hierarchical structure is precisely the substrate we will identify, in Section 4, with Fischer's developmental levels.

A critical property of generative models in Active Inference is that they are learned. The parameters of A, B, C, and D are updated through experience, following Bayesian principles. The structure of the model—the number of levels, the dimensionality of hidden states, the partitioning of state space—is also subject to change, though structural learning has received less attention than parameter learning in the published literature. Developmental Active Inference, as we will develop it, gives a central role to structural change.

2.3 Expected Free Energy and Action Selection

Where variational free energy F is concerned with retrospective inference (what was the hidden state that produced this observation?), expected free energy G is concerned with prospective action (which policy should the organism pursue?). G is the free energy that the organism expects to incur if it follows a given policy, marginalized over the hidden states and observations that policy would generate.

Expected free energy decomposes into two terms (Friston et al., 2017; Parr et al., 2022):

$$G(\pi) = -E_{q(o,s|\pi)} [\ln q(s|o,\pi) - \ln q(s|\pi)] - E_{q(o|\pi)} [\ln p(o|C)]$$

The first term is the expected information gain (also called epistemic value): the expected reduction in uncertainty about hidden states from following policy π. Policies that resolve ambiguity—that allow the organism to discriminate between competing hypotheses about the world—have high epistemic value. The second term is the expected utility (also called pragmatic value): the expected log-probability of preferred outcomes under policy π. Policies that lead to preferred observations have high pragmatic value.

The organism selects policies by minimizing expected free energy:

$$\pi^* = \arg\min_{\pi} G(\pi)$$

This formulation has several conceptual virtues. It naturally balances exploration (driven by epistemic value) and exploitation (driven by pragmatic value), without requiring an externally imposed exploration parameter. It treats learning and goal-pursuit as components of a single optimization, rather than competing objectives. And it provides a principled basis for deciding when a system should explore (epistemic value dominates) and when it should commit (pragmatic value dominates).

For clinical decision support, the EFE decomposition translates directly into a question that practitioners ask at every session: of the available interventions, which one will (a) teach the patient something they do not yet know about themselves or their patterns, and (b) move them closer to their stated goals? The first is epistemic value; the second is pragmatic value. Section 4 will show how Fischer level, HRV, habit alignment, and personal history each contribute to one or both components of EFE.

2.4 Precision: The Confidence Parameter

The precision parameter γ is, in our view, the most clinically important construct in Active Inference and the one with the most direct empirical grounding in psychophysiology. Precision is, formally, the inverse variance of a probability distribution—a measure of how concentrated, or confident, that distribution is. In Active Inference, multiple distributions have associated precisions: the precision of likelihoods (how reliable are sensory cues?), the precision of priors (how strongly held are beliefs?), and the precision of policies (how confidently is the next action chosen?).

The policy precision parameter—often denoted γ—plays a particularly central role. It determines the sharpness of the distribution over policies: high γ corresponds to confident, decisive action selection, while low γ corresponds to undecided, exploratory behavior. γ is itself dynamically modulated: it increases when the organism is in conditions of high model confidence and decreases when the organism is uncertain. Several authors have proposed that γ has a specific neurochemical substrate in dopaminergic signaling (FitzGerald et al., 2015), though the broader physiological basis is multifaceted.

Critically for our framework, precision has been linked to autonomic state. Smith, Thayer, Khalsa, and Lane (2017) reviewed evidence that vagal tone—indexed by heart rate variability (HRV)—provides a physiological marker of the precision of beliefs about the body and, by extension, the precision of higher-order beliefs that depend on accurate interoception. When vagal tone is high (high HRV), the autonomic system is in a state that supports flexible, contextually appropriate engagement; when vagal tone is low (low HRV), the system is in a state of mobilization or shutdown that constrains higher-order cognition. This linkage—between vagal tone, precision, and the availability of complex generative models—provides the bridge from polyvagal theory to Active Inference that we develop in Section 4.2.

A clinical implication is immediate. When a patient presents in a state of low HRV (whether due to acute stress, chronic dysregulation, or other causes), their generative model's higher levels are functionally less precise: predictions are noisy, learning is impaired, and policy selection becomes dominated by short-horizon defaults. Intervening to restore vagal tone—through breathing, posture, environment, or coregulation—is not merely "settling the patient down" but, in computational terms, re-establishing the precision conditions under which higher-level learning is possible. This insight, developed informally in clinical practice over many decades, has a precise mathematical home in Active Inference.

2.5 Active Inference and Learning: What Already Connects

Before turning to the developmental gap, it is worth noting what aspects of learning Active Inference already addresses well. Habit formation has been modeled as the consolidation of high-precision policies through repeated execution (Friston et al., 2016): a frequently performed action sequence becomes the default policy under a given state, and the precision associated with that policy increases over time, making alternative policies progressively less likely to be sampled. This account aligns naturally with contemporary habit theory (Wood & Rünger, 2016), which characterizes habits as automaticities cued by contextual features.

Curiosity and exploration have been modeled as expressions of epistemic value (Schwartenbeck et al., 2019). When the organism's model has high uncertainty about a region of state space, policies that visit that region acquire high epistemic value and are preferentially selected. This naturally produces the observation that infants and young learners exhibit pronounced exploratory behavior in novel environments, and that exploration declines as the model becomes well-fit. Whether this account also captures the qualitative restructurings characteristic of developmental transitions is precisely the question that Section 4 will take up.

Affect and emotion, often treated as separate domains from learning, have been integrated with Active Inference through several routes. Joffily and Coricelli (2013) proposed that emotional valence corresponds to the rate of change of free energy: positive affect arises when free energy is decreasing rapidly (learning, mastery, progress), and negative affect arises when free energy is increasing or stuck (confusion, frustration, threat). This account provides a natural connection to the affective experience of self-regulated learning—the satisfaction of solving a problem, the frustration of an impasse, the boredom of an under-challenging task—without requiring separate emotional and cognitive systems.

In sum, Active Inference offers a unified computational substrate for many of the phenomena central to SRL: learning, habit, exploration, affect, and adaptive action. What it has not yet offered is an account of how the generative model itself develops across the lifespan. Section 3 turns to dynamic skill theory as the developmental complement, and Section 4 develops the integration.

Table 1

Conceptual Correspondences Among Self-Regulated Learning, Active Inference, and Dynamic Skill Theory

Construct	**SRL (Zimmerman, 2000)**	**Active Inference (Parr et al., 2022)**	**DST (Fischer, 1980)**
Goal-setting	Forethought: task analysis	Prior preferences C	Optimal level under support
Monitoring	Performance: self-observation	Posterior $q(s\|o)$	Functional level assessment
Self-evaluation	Self-reflection: judgment	Model evidence $\ln p(o)$	Skill assessment across contexts
Strategy selection	Forethought: strategic planning	Policy π via min $G(\pi)$	Skill selection in context
Self-efficacy	Forethought: motivation beliefs	Precision γ on policies	

Confidence at functional level
Adaptive inference
Self-reflection: causal attribution
Model update via Bayes
Level transition

Note. Correspondences are conceptual rather than strict identities; see Section 4 for detailed mapping. SRL = self-regulated learning; DST = dynamic skill theory; G = expected free energy; γ = precision parameter.

3. Dynamic Skill Theory: A Developmental Architecture

Dynamic skill theory (DST; Fischer, 1980; Fischer & Bidell, 2006) is a developmental framework that has, over four decades, accumulated substantial empirical support while remaining underutilized in computational modeling. This section reviews DST at the level needed for the integration in Section 4, with emphasis on three features that map naturally onto Active Inference constructs: hierarchical levels of skill, domain-specificity, and the optimal-functional distinction.

3.1 Skill as Coordinated Action

DST defines a skill as a person's capacity to act in an organized way in a specific context. This definition is deliberately concrete: a skill is not an abstract capacity but a way of operating on the world that requires both the person's internal organization and the supportive structure of a context. This contextualism distinguishes DST from approaches that treat development as the unfolding of internal stages independent of environmental support (Piaget, 1972), and aligns it more closely with sociocultural perspectives (Vygotsky, 1978; Bruner, 1985).

Importantly, skill in this sense includes both what is traditionally called cognition and what is traditionally called emotion. Fischer and Bidell (2006) explicitly extend skill analysis to affective skills—a person's capacity to organize feelings and emotional responses in context-appropriate ways. The implication for our framework is that the same theoretical machinery describes academic skill, clinical regulation, and therapeutic insight. This is not a metaphorical extension but a structural claim: skills of all kinds exhibit the same patterns of acquisition, consolidation, and transformation.

3.2 The Hierarchy of Skill Levels

DST organizes skill development into a sequence of levels, each defined by the type of unit it operates on and the complexity of structural organization. Three tiers, each containing four levels, span the developmental range from infancy to adulthood:

The reflex tier (Level 1, single reflexes through Level 4, systems of reflexes) is characteristic of the first months of life. The action tier (Level 5, single actions through Level 8, action systems) emerges across the second half of the first year and consolidates through early childhood. The representational tier (Level 9, single representations through Level 12, representational systems) emerges between approximately 18 months and middle childhood. The abstract tier (Level 13, single abstractions through Level 16, principles) emerges from late childhood through adulthood, with the highest level rarely fully realized.

Within each tier, the four levels reflect a recurring pattern of structural complexity: single units, mappings between two units, systems combining multiple units, and systems of systems. The transition from one level to the next is not a gradual quantitative increase but a qualitative restructuring: at each level, units that were previously coordinated become the elements of new, higher-order units. The cycle then repeats at the next tier of abstraction.

For our framework, two features of this hierarchy are essential. First, the structure is recursive: the same organizational principle (units → mappings → systems → tier transition) applies at every level. This recursiveness allows DST to characterize development with relatively few principles while accommodating the diversity of skills people acquire. Second, higher levels are constructed from lower levels: a representational system (Level 12) is built from coordinations of representational mappings (Level 11), which are themselves coordinations of single representations (Level 9-10). This compositional structure parallels the hierarchical generative models of Active Inference (Pezzulo et al., 2018; Friston et al., 2017) in a way we will exploit in Section 4.

3.3 Domain-Specificity and the Web Model

Perhaps DST's most distinctive empirical contribution is the demonstration that skill development is highly domain-specific. A person who reasons abstractly about mathematics may operate at the representational level about social relationships; a clinician who functions at the system level diagnostically may regulate emotion at the action level. Skill levels are not properties of a person globally but of a person operating in a specific domain under specific contextual conditions (Fischer & Yan, 2002).

Fischer and Bidell (2006) capture this domain-specificity with what they call the web model: development proceeds along multiple, partially independent strands, each strand corresponding to a domain of skill. Strands can be partially synchronized—for example, when conceptual development in mathematics supports problem-solving in physics—but they are not lockstep. The pattern of synchronization across strands is itself a developmental phenomenon, capturing how the integration of skills across domains constitutes a major part of mature cognition.

For clinical practice, the web model has direct implications. A standard psychometric instrument that yields a single composite score risks obscuring exactly the cross-domain pattern that is most clinically informative. Two patients with identical composite scores may have entirely different web profiles—different patterns of strength across domains—and thus require different interventions. The Learning Catcher system (Section 5) adopts a five-dimensional skill profile (sustained attention, working memory, emotional regulation, time awareness, self-awareness) specifically to honor the web structure that DST has documented.

3.4 Optimal Level and Functional Level

The single most clinically actionable construct in DST is the distinction between optimal level and functional level. The optimal level is the highest level a person can achieve in a domain when full contextual support is provided—an attentive coach, well-designed materials, time to think, low emotional load. The functional level is the level the person typically exhibits in everyday activity, with whatever support is available. The gap between the two is the developmental space in which scaffolded intervention operates, conceptually equivalent to Vygotsky's (1978) zone of proximal development but operationalized in terms of skill levels rather than general age stages.

Empirically, the gap between optimal and functional level can be large—often two or more levels (Fischer & Yan, 2002). A person who, under supportive conditions, can construct an abstract system about a topic may, in everyday discourse, deploy only single abstractions about the same topic. This gap is not a failure of the person but a feature of how skill is distributed across the person-context system: support raises functional level toward optimal, withdrawal of support lowers it.

The optimal-functional distinction is especially important for clinical work because it provides a principled answer to the question, "What level should I be intervening at?" Intervention at the functional level provides little leverage—the person is already operating there. Intervention at or beyond the optimal level produces frustration and disengagement. Intervention in the gap—at one level above functional, scaffolded by support—produces what we will, in Section 4.2, formally identify as conditions of high expected information gain.

3.5 What DST Specifies, and What It Leaves Open

DST specifies the structure of developmental change but leaves the computational mechanism by which that change occurs largely open. The theory describes how skills are organized at each level, how transitions between levels are structured, and how skills are distributed across contexts. It does not specify what process determines whether a given intervention will be effective for a given learner at a given moment, beyond appeals to general principles of optimal challenge and contextual support.

This is, of course, not a criticism of DST: every theory has a chosen level of explanation, and DST's level—the structural organization of skill across development—is empirically rich and conceptually distinctive. But it does identify a gap that Active Inference is uniquely positioned to fill. Where DST specifies that the generative model has hierarchical structure of a particular form, Active Inference specifies how that hierarchical structure produces moment-to-moment perception and action. Where DST specifies that there exists an optimal challenge zone, Active Inference specifies the computational quantity (expected information gain) that defines that zone. The integration is not a coincidence: both frameworks have converged on hierarchical predictive structures, from different directions, over the same several decades.

Section 4 develops the integration in detail.

Table 2

Fischer's Dynamic Skill Theory: Tiers, Levels, and Approximate Age Ranges

Tier	**Level**	**Structural Form**	**Approx. Age of Emergence**
Reflex	1–4	Single reflexes → reflex mappings → reflex systems → systems of systems	Birth to 4 months
Action	5–8	Single actions → action mappings → action systems → systems of systems	4 months to 2 years
Representation	9–12	Single representations → mappings → systems → systems of systems	2 to 11 years
Abstract	13–16	Single abstractions → mappings → systems → principles	11 years to adulthood

Note. The abstract tier (highlighted) corresponds to most clinical adult work. Within each tier, levels follow the same recursive pattern: single elements, mappings between elements, systems combining elements, and systems of systems (which become elements at the next tier). Adapted from Fischer (1980) and Fischer and Bidell (2006). Ages are approximate optimal-level emergence under supportive conditions; functional levels in everyday life are typically one to two levels lower.

4. Developmental Active Inference: Four Integrative Mechanisms

This section is the conceptual core of the paper. We develop four mechanisms by which Active Inference and dynamic skill theory can be integrated to yield a computational framework for developmental self-regulated learning. The mechanisms are not separate hypotheses to be tested independently; they are facets of a single proposal that the generative model in Active Inference can be identified, with appropriate elaboration, with the developmental architecture in DST. Two further mechanisms—precision allocation and self-model refinement—connect the integrated framework to clinically actionable inputs (HRV) and clinically actionable procedures (Observable Model Refinement).

We organize the discussion around four propositions:

1. Hierarchical generative model structure. Fischer's developmental levels correspond to the depth of the generative model. Development is structural expansion of the model.
2. Precision allocation as autonomic state. The precision parameter γ that governs policy confidence in Active Inference has a physiological substrate in vagal tone, indexed by HRV. Regulation is precision availability.
3. Empirical prior accumulation as habit formation. Habits, in DST terms repeated skills, are learned policies in Active Inference, instantiated as high-precision priors. Habit is consolidated policy under precision.
4. Self-model refinement as Observable Model Refinement (OMR). The generative model of self develops through externalized hypothesis-testing in clinical conversation, formalizable as Bayesian update with explicit observation.

Each proposition is developed in a subsection below. Throughout, we use the notation established in Section 2, with extensions where needed.

4.1 Mechanism 1: Hierarchical Generative Model Structure

Active Inference's hierarchical generative models and Fischer's developmental hierarchy share a key structural property: both are recursive compositions in which higher levels are constructed from lower levels. In Active Inference, a hierarchical model has multiple levels, each generating predictions for the level below and receiving prediction errors from below. Higher levels operate at slower temporal scales and represent more abstract content; lower levels operate at faster temporal scales and represent more concrete content (Friston, Parr, & de Vries, 2017; Pezzulo et al., 2018).

Fischer's hierarchy has the same structural form. The reflex tier coordinates immediate sensorimotor patterns at the fastest temporal scale. The action tier coordinates reflexes into goal-directed sequences at a longer scale. The representational tier abstracts actions into mental tokens that can be manipulated independently of immediate sensorimotor presence. The abstract tier abstracts representations into concepts that span multiple contexts. Each tier operates at a slower characteristic temporal scale and a more abstract level of organization than the preceding tier.

We propose that the levels of Fischer's hierarchy correspond, mechanistically, to the layers of a hierarchical generative model. The reflex tier corresponds to the sensorimotor coupling level: predictions about immediate input from the body and environment. The action tier corresponds to the policy level: sequences of action that produce desired sensory outcomes. The representational tier corresponds to an intermediate latent level: predictions about states of the world that persist independently of immediate sensory input. The abstract tier corresponds to higher latent levels: predictions about general principles, identities, values, and long-horizon goals.

This mapping is more than analogical. Both frameworks were developed to account for similar phenomena—the emergence of organized adaptive behavior—and have converged on similar structural commitments because the underlying problem (achieving robust adaptive control in a partially observable world) admits a limited number of solutions. Pezzulo et al. (2018) showed that hierarchical Active Inference models can reproduce key features of motivated control, including the integration of immediate sensory input with long-horizon goals. Constant et al. (2018) extended hierarchical Active Inference to cultural learning, in which higher levels of the hierarchy encode socially transmitted priors. We propose that the developmental sequence Fischer described corresponds to the historical emergence, in each individual's lifetime, of progressively higher levels in their hierarchical generative model.

This correspondence yields a precise computational interpretation of developmental transitions. A transition from one Fischer level to the next is not merely the addition of more parameters to an existing model. It is the emergence of a new layer in the hierarchical generative model—a new representational level capable of generating predictions about a new class of content. Such structural changes are formally distinct from parameter updates within an existing structure: they require what has been called Bayesian model selection (MacKay, 2003), in which the data are explained by a structurally distinct model that is not reachable through parameter learning alone.

The clinical implication is consequential. A patient who has not yet achieved a given Fischer level cannot be brought to operate at that level by repetition, reinforcement, or fluency training within the level below. They require the emergence of a new structural level in their model, which depends on (a) sufficient consolidation of the level below to provide the constituents from which the new level is built, and (b) supportive context that scaffolds the structural transition. This is precisely the situation that Fischer characterized as transitions between optimal levels: contextual support enables a person to operate at a level that, without support, they would not yet exhibit.

4.1.1 Structural Learning Versus Parameter Learning

Most Active Inference applications focus on parameter learning: given a fixed model structure, how are the parameters of the likelihood, transition, and preference distributions updated through experience? Parameter learning is well-understood, mathematically tractable, and empirically supported. Structural learning—how new layers emerge, how the dimensionality of latent state spaces grows, how new partitions of the state space are introduced—is less developed. Most theoretical work assumes that structural learning is a slow background process operating on developmental timescales, distinct from the moment-to-moment perceptual and motor inference that occupies most of Active Inference modeling.

We adopt this distinction with one important addition: developmental SRL operates at the intersection of moment-to-moment and structural timescales. A clinical encounter—a single session, a single card delivered in a learning system, a single therapeutic exchange—occupies the moment-to-moment timescale. The developmental trajectory of a patient over months and years occupies the structural timescale. The framework we develop holds that these timescales are coupled: the moment-to-moment selection of interventions, when patterned over many sessions, scaffolds the structural changes that constitute development.

4.2 Mechanism 2: Precision Allocation and the Autonomic Bridge

If hierarchical structure provides the developmental architecture, precision allocation provides the moment-to-moment dynamics. As reviewed in Section 2.4, precision γ in Active Inference governs the sharpness of distributions over policies and, more broadly, the confidence assigned to different elements of the generative model. High precision corresponds to confident, committed action; low precision corresponds to undecided, exploratory behavior. Precision is itself dynamically modulated, increasing under conditions of model confidence and decreasing under uncertainty.

Smith, Thayer, Khalsa, and Lane (2017) reviewed evidence linking vagal tone—operationalized as heart rate variability (HRV)—to the precision of beliefs about the body. The argument runs as follows. The brain maintains a model of the body's internal state (interoception), with predictions descending from cortical to subcortical and peripheral structures and prediction errors ascending in the reverse direction. Vagal tone provides a key channel for the precise transmission of bodily signals, and individual differences in vagal tone reflect, in part, individual differences in the precision of these interoceptive predictions.

Several extensions of this argument are relevant for our framework. First, interoceptive precision is not isolated from higher-order cognitive precision: there are well-documented connections between autonomic state and cognitive performance (Thayer et al., 2009; Hansen et al., 2003), with low vagal tone associated with impaired executive function. Second, the relationship is not one-directional. Cognitive load can suppress vagal tone, and emotional regulation strategies can elevate it. Third, vagal tone has been linked to therapeutic engagement and outcome (Porges, 2011): patients in low vagal states tend to exhibit reduced capacity for the kind of integrative cognitive-emotional work that therapy requires.

Taken together, these observations support the following proposal: the precision parameter γ in Active Inference can be operationalized clinically by measuring HRV. When HRV is high, the patient's generative model is operating with elevated precision: predictions can be made confidently, prediction errors can be informatively integrated, and learning can occur. When HRV is low, the patient's generative model is operating at reduced precision: predictions are noisy, prediction errors are not reliably distinguished from noise, and learning is impaired.

This proposal has direct clinical implications. Intervening to elevate HRV—through breathing, posture, environmental adjustment, coregulation with the clinician—is not merely "settling the patient down." It is restoring the precision conditions under which higher-order learning is possible. Conversely, attempting to deliver learning content to a patient in low HRV is not merely difficult; it is, in computational terms, attempting to update parameters using observations weighted at precision near zero. The data may pass through but cannot revise beliefs.

4.2.1 The Step-Zero Regulator

The clinical operationalization of this insight is straightforward: every intervention should be preceded by a precision assessment. If precision (HRV) is below a domain-specific threshold, the indicated intervention is regulation—an action chosen for its expected effect on autonomic state rather than for its informational or pragmatic content per se. Only when precision is restored above threshold should the learning-oriented intervention be delivered. We refer to this procedural element as the Step-Zero Regulator, by analogy with engineering systems in which a regulator establishes operating conditions before functional processing begins.

Implementing the Step-Zero Regulator computationally requires (a) a measurement protocol for HRV that is feasible in the clinical setting (e.g., 1–5 minute resting recordings, or continuous monitoring during intervention); (b) a mapping from HRV values to γ values; (c) a threshold for γ below which only regulation interventions are eligible; and (d) a library of regulation interventions matched to specific patterns of autonomic dysregulation. The Learning Catcher system (Section 5) implements each of these components in beta form.

It is worth emphasizing that this proposal does not reduce SRL to autonomic regulation. The framework holds that SRL requires both regulatory and learning components, and that the two are computationally distinct: regulation operates on precision, learning operates on the content of the generative model. What is novel in the framework is the explicit recognition that learning is precision-conditional, and the corresponding clinical priority on establishing precision conditions before delivering learning content.

4.3 Mechanism 3: Empirical Priors and Habit Formation

The third mechanism concerns how repeated experience consolidates into stable patterns of action—what colloquial language calls habit, what DST analyzes as skill consolidation at a level, and what Active Inference formalizes as policy precision and empirical prior strengthening.

Friston et al. (2016) modeled habit formation as the consolidation of high-precision policies. Initially, a novel context elicits broad exploration: multiple policies are evaluated, each with low precision, and the choice between them depends on subtle differences in expected free energy. With repetition under similar contexts, certain policies acquire higher precision: they become the default response, executed with reduced deliberation. The transition from controlled to automatic processing, long documented in cognitive psychology, has in this account a precise computational characterization.

Contemporary habit theory in psychology (Wood & Rünger, 2016; Lally et al., 2010; Clear, 2018) describes habits as automaticities cued by contextual features, executed without explicit goal pursuit, and resistant to outcome devaluation. These properties map cleanly onto the Active Inference account. Contextual cueing corresponds to context-conditional policy selection.

Execution without explicit goal pursuit corresponds to action under high policy precision (the policy is selected almost deterministically, so the search across alternatives is minimal). Resistance to outcome devaluation corresponds to the property that high-precision policies require substantial counter-evidence to be displaced.

The Cue-Routine-Reward-Craving framework popularized in clinical and self-help contexts (Duhigg, 2012; Clear, 2018) has its own Active Inference reading. The cue is the context that conditions policy selection (sets the prior over policies). The routine is the policy itself (the sequence of actions selected). The reward is the observation that confirms the policy's pragmatic value. The craving—the felt anticipation that precedes habit execution—corresponds to the predicted observation under the selected policy, instantiated as a perceptual representation in advance of action.

For our developmental framework, what is essential is that habit formation operates at the level of policy precision while skill development operates at the level of model structure. A habit is a consolidated way of acting within a given level of skill. A developmental transition is the emergence of a new level. These are computationally distinct: a person can have well-formed habits at a given Fischer level without exhibiting the structural conditions for transition to the next level. The clinical implication is that habit work and developmental work are complementary but not interchangeable. Strengthening a habit at the functional level may improve consistency without producing developmental progression; producing developmental progression requires structural conditions that habit consolidation alone does not provide.

4.3.1 Habit Loop Components and Active Inference Counterparts

The four-element habit loop has a direct computational correspondence summarized in Table 3 (presented after Section 4.4). Each element has a specific role in the policy precision dynamics that constitute habit formation. The therapeutic implications of this correspondence are significant: interventions that target one component (e.g., reward manipulation) without addressing others (e.g., cue restructuring, craving regulation) operate on only a portion of the relevant dynamics, with predictable limits on effectiveness.

4.4 Mechanism 4: Self-Model Refinement and Observable Model Refinement

The fourth mechanism is the one most distinctive to our framework and the one with the most direct clinical operationalization. Active Inference holds that organisms maintain generative models of themselves, not only of the external environment. The self-model encodes predictions about one's own dispositions, preferences, capabilities, and trajectories—what Apps and Tsakiris (2014) call the "free-energy self." Self-knowledge, in this view, is not a static introspective access to inner states; it is a continuous inferential process by which the self-model is refined against the evidence provided by one's own behavior, internal states, and feedback from others.

Developmentally, the self-model becomes progressively more articulated. The infant's implicit self-model is largely sensorimotor: predictions about what bodily movements produce what sensory consequences. The young child's self-model becomes representational: predictions about social roles, capabilities in specific contexts, and emotional patterns. The adolescent and adult self-model becomes abstract: predictions about identity, values, life trajectories, and the kinds of person one is or wishes to become. This developmental progression follows the same hierarchical pattern identified in Section 4.1.

Self-model refinement is, we propose, the developmental core of self-regulated learning. The phases Zimmerman identifies—forethought, performance, self-reflection—each contribute to self-model refinement. Forethought operationalizes predictions ahead of action. Performance generates the observations that test those predictions. Self-reflection updates the self-model on the basis of the observation-prediction comparison. Across many such cycles, the self-model becomes a more accurate and articulate predictor of the person's own behavior—what is colloquially called self-knowledge.

4.4.1 Observable Model Refinement (OMR) as Clinical Procedure

Observable Model Refinement (OMR) is the procedural operationalization of self-model refinement in clinical practice. The procedure has three components. First, the clinician (or, in computational form, the clinical decision support system) externalizes a hypothesis about a pattern in the patient's behavior or experience. The hypothesis is articulated in concrete terms: "It seems that when X happens, you tend to do Y, and afterwards Z follows." Second, the patient is invited to confirm, qualify, or disconfirm the hypothesis in light of their experience. Third, the hypothesis is updated based on the patient's response, and the updated version becomes the patient's revised self-model entry.

The procedure draws on long-established clinical traditions—reflective listening (Rogers, 1957), Socratic questioning (Padesky, 1993), motivational interviewing (Miller & Rollnick, 2012)—and gives them a computational reading. The clinician's hypothesis is a prediction generated by an external generative model (the clinician's model of the patient). The patient's response is an observation. The update is Bayesian: the prior model is updated to a posterior that better accounts for the observation. The patient internalizes the updated model as a revised self-prediction.

What is novel in OMR is the explicit treatment of this process as a developmental intervention—an intervention specifically targeted at the structure of the patient's self-model rather than at any particular content of that model. The hypothesis is not designed to be "correct" in the sense of accurately describing the patient; it is designed to be a candidate self-prediction that the patient can evaluate against their own experience. The value of the intervention lies not in the accuracy of the initial hypothesis but in the patient's engagement with the comparison: the act of evaluating one's behavior against an externalized prediction is itself the self-model refinement, regardless of whether the initial prediction proves correct.

Empirically, OMR produces observable outputs that can be tracked: the rate at which patients confirm hypotheses (high rates may suggest agreeable response bias or accurate hypothesis generation), the rate of qualified responses (suggesting active engagement with the procedure), the rate of disconfirmation (suggesting accurate self-assessment and tolerance of inconsistency with the clinician's view). Over time, the pattern of these responses constitutes a window into the patient's developing self-model. We consider this empirical accessibility a significant advantage of OMR over alternative operationalizations of self-knowledge that rely entirely on retrospective self-report.

4.4.2 OMR and the Generative Model of Self

OMR has a precise interpretation within hierarchical Active Inference. The patient's self-model is a layer (or layers) of their generative model, encoding predictions about their own behavior. The clinician's externalized hypothesis is an observation: a candidate prediction about the patient that is exposed for testing. The patient's response provides the data for Bayesian update of the relevant layer. Over repeated trials, the precision of self-predictions—the patient's confidence in what they will do under specified conditions—is calibrated against actual evidence. Self-knowledge, in this reading, is the accumulated outcome of many such updates.

A developmental property is implicit in this account. The level at which OMR operates depends on the developmental level of the patient's self-model. With a patient operating at the action level, OMR may concern behavioral sequences ("when you are tired you tend to skip breakfast"). With a patient operating at the representational level, OMR may concern emotional patterns ("when you feel criticized you tend to withdraw"). With a patient operating at the abstract level, OMR may concern identity-level patterns ("you describe yourself as someone who avoids conflict but you also describe situations in which you initiate conflict to clarify boundaries"). The same procedural form operates across levels; the content of the hypothesis is calibrated to the level.

This level-sensitivity is important clinically. A common failure mode of intervention is the delivery of content at the wrong developmental level: representational interventions to a patient operating at the action level, abstract interventions to a patient operating at the representational level. OMR provides a procedural check: hypotheses that are confirmed or productively disconfirmed are at an appropriate level; hypotheses that produce blank confusion or rote agreement are at the wrong level. The clinician can recalibrate within the same procedure.

Table 3

Habit Loop Components and Active Inference

Counterparts

Habit Component	**Phenomenology (Clear, 2018)**	**Active Inference Construct**	**Clinical Lever**
Cue	Contextual trigger that initiates the response	Context-conditional prior over policies $p(\pi\|c)$	Environment restructuring
Routine	The action sequence executed	Policy π itself; high precision after consolidation	Practice; substitution
Reward	Outcome that reinforces the routine	Observation confirming pragmatic value (C-distribution match)	Outcome shaping; reflection on natural reinforcers
Craving	Anticipatory pull toward the routine	Predicted observation under selected policy, instantiated perceptually in advance of action	Mindful awareness; precision modulation

Note. Each habit-loop component has both a phenomenological description (familiar from contemporary self-help and habit psychology) and a computational counterpart in Active Inference. Clinical interventions target one or more components and have predictable effects on the corresponding computational quantities.

Table 4

Developmental Active Inference: Summary of Four Integrative Mechanisms

Mechanism	**Active Inference Construct**	**Source Theory**	**Clinical Operationalization**
1. Hierarchical structure	Hierarchical generative model; depth corresponds to representational tier	Fischer's Dynamic Skill Theory (1980)	

Fischer level assessment per domain
2. Precision allocation
Precision γ on policies; modulated by autonomic state
Polyvagal Theory (Porges, 2007); precision/HRV (Smith et al., 2017)
HRV measurement; Step-Zero Regulator
3. Empirical priors
Learned policy precision; context-conditional priors
Habit theory (Wood & Rünger, 2016; Clear, 2018)
Habit-loop assessment and intervention
4. Self-model refinement
Bayesian update of self-related layers of generative model
Reflective listening; motivational interviewing
Observable Model Refinement (OMR) procedure
Note. The four mechanisms are facets of a single proposal: the hierarchical generative model in Active Inference can be identified with the developmental skill hierarchy in DST, with precision providing the moment-to-moment dynamics and habit and OMR providing the consolidation and update routes, respectively. None of the mechanisms is a separate hypothesis; together they constitute the framework's structure. **5. A Worked Example: The Learning Catcher System**

This section illustrates the Developmental Active Inference framework through a worked example drawn from Learning Catcher, a clinical decision support system currently in beta evaluation. We describe the system at the level of detail needed to illustrate how the four mechanisms developed in Section 4 are operationalized, and we walk through a single session to show how the framework determines intervention selection. The system is not the framework; many alternative implementations are conceivable. Our purpose here is to make the framework's components concrete enough that their implications for clinical practice are evident.

5.1 System Overview

Learning Catcher is a web-based application designed to support brief, daily learning activities ("cards") between in-person clinical sessions. Each card is a single, completable activity—typically requiring three to five minutes—designed to advance the patient on one of five skill dimensions: sustained attention, working memory, emotional regulation, time awareness, and self-awareness. These dimensions correspond to commonly assessed cognitive-affective functions in clinical practice and were selected on the basis of (a) feasibility of in-session measurement using standard clinical instruments and (b) developmental relevance across the adult Fischer tiers (levels 7 through 13).

The system architecture comprises three layers. The measurement layer collects in-session and between-session data: heart rate variability (HRV) recordings, behavioral measures (continuous performance test, N-back, time estimation, voiced sentence reading), and self-report responses to cards and to Observable Model Refinement hypotheses. The decision layer computes expected free energy for each candidate card given the current state and selects the card that minimizes expected free energy subject to a precision threshold (the Step-Zero Regulator). The delivery layer presents the selected card to the patient through a brief interactive interface.

The patient-facing interface emphasizes simplicity: the patient sees a single card per day, with options to complete, skip, or request alternative content. The clinician-facing dashboard provides longitudinal views of the five-dimensional skill profile, HRV trends, card completion patterns, and OMR response patterns. Decision-relevant computation occurs in the background; neither patient nor clinician interacts with the EFE calculation directly. This design follows what is known about effective clinical decision support: the computation should support rather than replace clinical judgment, and the interface should present interpretable summaries rather than raw outputs (Bates et al., 2003).

5.2 The Five-Dimensional Skill Profile

Each of the five skill dimensions is assessed at intake using a domain-appropriate clinical measure, and re-assessed at intervals during ongoing treatment. The assessment protocol is intentionally brief—total in-session time approximately 25 minutes—to be feasible within typical clinical session lengths. The five dimensions and their primary measures are summarized in Table 5.

The selection of these five dimensions and their associated measures reflects several design considerations. First, all five can be assessed by a Board Certified in Neurofeedback (BCN) practitioner within scope of practice, without requiring physician supervision for routine administration. Second, all five have established population norms permitting comparison with reference distributions. Third, all five have been linked, in independent literatures, to clinically meaningful outcomes: sustained attention to academic and occupational functioning (Conners, 2014), working memory to executive function and emotion regulation (Schmeichel & Tang, 2015), emotional regulation to therapeutic outcomes (Berking & Wupperman, 2012), time awareness to ADHD and procrastination (Barkley, 2012), and self-awareness to therapeutic alliance and outcome (Hill, 2009).

Importantly, the five dimensions are not assumed to be statistically independent. Empirically, they show partial correlations consistent with shared executive-attentional substrates. The dimensional approach is intended not to claim that the five are independent factors but to honor the web structure (Section 3.3) by which skill develops in partially independent domains. Two patients with similar global executive function may have notably different five-dimensional profiles, and those profiles may indicate different intervention targets.

5.3 Mapping Skill Dimensions to Fischer Levels

For each dimension, the system maintains an estimate of the patient's optimal and functional Fischer level. The estimates are initialized at intake based on clinical observation and refined over subsequent sessions through performance on cards calibrated to specific levels. The system maintains, for each level, a library of cards designed for that level; the level estimate determines the eligible card library for the next selection.

To make this concrete: a patient whose functional emotional-regulation level is estimated at L9 (single representations) is offered cards that involve recognizing single emotional states ("notice when you feel anxious today"). A patient at L10 (representational mappings) is offered cards that involve recognizing relationships between two states ("notice when feeling rejected leads to feeling angry"). A patient at L11 (representational systems) is offered cards involving systems of three or more states ("notice the pattern of feelings that follows criticism: defensive, then guilty, then withdrawn"). And so on. Cards at the functional level provide consolidation; cards one level above functional level provide developmental scaffolding.

The system favors cards one level above the current functional level, consistent with the optimal challenge principle developed in Section 4.1. However, the actual selection is not deterministic: it depends on the expected free energy computation, which also depends on the patient's current precision state and history. Section 5.5 walks through the computation in detail.

5.4 HRV and the Step-Zero Regulator

HRV is measured both in-session (using a chest-strap sensor during a 5-minute resting baseline) and, optionally, before each card via a brief home recording. The system maintains a longitudinal HRV trajectory and a current estimated precision γ derived from recent HRV values relative to the patient's own baseline.

The Step-Zero Regulator operates on this precision estimate. Before selecting a learning card, the system checks whether γ is above the patient-specific threshold for that day. If γ is below threshold, the system selects from a separate library of regulation cards designed to support autonomic recovery: paced breathing, brief body awareness, environmental adjustment, or contact with the clinician. The regulation library is independent of the Fischer-level estimate; regulation is treated as a precondition for learning rather than as learning content itself.

This separation of regulation from learning is an explicit design commitment. It would be possible to fold regulation into a single ranked list of cards, with regulation cards competing against learning cards on a unified EFE score. We have chosen not to do so because the empirical and conceptual literatures we reviewed in Section 4.2 are unambiguous: at low precision, learning content is poorly metabolized regardless of how appropriate it would be at higher precision. The Step-Zero Regulator implements this conditionality structurally rather than parametrically.

5.5 A Session Walk-Through

To illustrate, consider a hypothetical patient: a 35-year-old in the second month of treatment for symptoms of inattention, mild anxiety, and difficulties with self-direction. After two prior sessions of intake assessment, the system has the following estimates for this patient:

- Sustained attention: functional level 10, optimal level 11 (representational mappings); CPT performance moderately below age-matched norms.
- Working memory: functional level 10, optimal level 11; N-back performance within normal range.
- Emotional regulation: functional level 9, optimal level 11 (large gap); HRV baseline RMSSD = 28 ms, below age-gender norms.
- Time awareness: functional level 10, optimal level 11; time estimation errors moderately elevated.
- Self-awareness: functional level 11, optimal level 12; engagement with prior OMR hypotheses high.

The patient logs in for the day. The system first reads the most recent HRV value, recorded by the patient that morning using a home monitor: RMSSD = 22 ms, notably below their own baseline. The Step-Zero Regulator computes γ as below threshold. The learning libraries are bypassed; the system selects from the regulation library. The selection within regulation is itself an EFE computation: among regulation candidates, the system selects the one expected to provide both the highest information gain ("learning something about my own regulation") and the highest pragmatic value ("actually elevating HRV").

The selected card is a 4-minute paced-breathing exercise with a brief framing: "Let's start with something gentle. Your body is asking for a moment to settle." After completion, the patient is offered a brief OMR hypothesis: "It seems that some mornings your body is in a slower gear. When that happens, the smallest cards feel like the right starting place." The patient confirms with a one-line addition: "Especially after nights I don't sleep well." The system records the confirmation and the qualification, which will refine the precision-state model: the patient's report links low morning HRV to prior-night sleep, a relationship the system can now incorporate.

Two days later, the patient logs in with a morning HRV of RMSSD = 41 ms—elevated relative to baseline, perhaps due to a restorative weekend. The Step-Zero Regulator passes. The system now considers learning cards. Across the five dimensions, the system computes expected free energy for each candidate, considering the current functional level, the candidate's targeted level, and the patient's history of engagement with cards in each dimension.

For sustained attention, the highest-EFE candidate (lowest G) is a 5-minute focused-reading card at level 11 (mapping reading to mood reflection)—a developmental step above the functional level, in a dimension where the patient has been consolidating. For emotional regulation, the highest-EFE candidate is a level-10 card (mapping a recent emotional event to a familiar pattern). For self-awareness, the highest-EFE candidate is a level-12 OMR engagement ("how does your week's pattern compare to last week's?").

The system selects across dimensions probabilistically, weighted by EFE values and by recent dimension coverage. Today, the selected card is the emotional-regulation card—the dimension with the largest functional-optimal gap and the strongest current evidence of progress. The card is delivered with a brief framing: "Your body is in a good place today. Let's look at something a little more demanding." The patient completes the card and provides an emotion rating. The response data update the system's estimates: the dimension-specific Fischer level posterior is sharpened, the policy precision for cards of this type is increased, the patient's history of engagement is extended.

Over weeks and months, the cumulative effect of these daily selections is to scaffold developmental progression in the domains where it is most clinically indicated, while respecting the precision conditions that make learning possible. The clinician, viewing the longitudinal dashboard, observes which dimensions are progressing, which are stable, and where the gaps between functional and optimal levels are closing. The in-person sessions can then be devoted to integration and to the OMR work that scaffolds the highest level of the system's intervention—the patient's developing self-model.

5.6 What This Worked Example Demonstrates

The worked example illustrates four properties of the framework that we wish to emphasize. First, no element of the daily intervention is arbitrary: every choice—regulation versus learning, dimension selection, level targeting—is determined by a transparent computation grounded in the four mechanisms of Section 4. Second, the system does not replace clinical judgment but supports it: the longitudinal patterns it makes visible inform the in-person work, and the in-person work informs the system's parameters (especially through OMR). Third, the framework is naturally interpretable: at any point, the clinician can ask why a specific card was selected and receive an answer in terms of precision state, current levels, and EFE components. Fourth, the framework generates testable predictions about patient trajectories that can be evaluated against actual outcomes (Section 6).

We do not claim that this implementation is optimal. Many design decisions—the choice of five dimensions, the specific HRV threshold, the prior weight on dimension coverage—are empirical questions that the framework cannot answer from first principles. What the framework provides is a principled basis for stating those decisions explicitly, evaluating them against outcomes, and revising them when evidence warrants. The framework's value lies in this discipline, not in the specific parameters of any particular implementation.

Table 5

Five Skill Dimensions in Learning Catcher:

Measures, Norms, and Fischer Level Range

Dimension	**Primary Measure**	**Norm Source**	**Adult Fischer Range**
Sustained Attention	Continuous Performance Test (CPT)	Conners CPT (Conners, 2014)	L9–L13
Working Memory	N-back (typically 2-back)	Kirchner (1958); meta-analyses	L9–L13
Emotional Regulation	Resting HRV (RMSSD, 5 min)	Nunan et al. (2010); HRV norms	L7–L13
Time Awareness	Time estimation tasks	Zakay & Block (1997); ADHD literature	L9–L13
Self-Awareness	OMR response patterns	Internal calibration (longitudinal)	L11–L16

Note. Emotional regulation (highlighted) doubles as the precision input for the Step-Zero Regulator. RMSSD = root mean square of successive differences; OMR = Observable Model Refinement; L = Fischer level. Adult Fischer range indicates levels at which cards in each library are calibrated; not all levels are populated for every dimension.

6. Empirical Predictions and Validation Roadmap

A theoretical framework's value depends on its capacity to generate predictions that can be empirically evaluated. This section identifies five testable predictions of Developmental Active Inference and outlines a phased validation roadmap. Predictions are stated at progressively narrower levels of resolution, from broad architectural claims testable with available data to specific quantitative claims that will require dedicated empirical work. The validation roadmap is structured to allow each phase to be conducted within the resource constraints of a small clinical practice, while remaining capable of supporting cumulative evidence appropriate for peer review.

6.1 Five Empirical Predictions

6.1.1 Prediction 1: HRV as Precision (Step-Zero Validation)

If HRV operationalizes the precision parameter γ that gates learning under Active Inference, then card-completion probability should increase monotonically with pre-card HRV, within-subject, controlling for time-of-day and prior-night sleep. Quantitative formulation: in a logistic regression with completion as outcome and HRV as predictor, the within-subject slope should be positive and significantly greater than zero. Effect sizes are not strongly constrained by theory, but the qualitative prediction is robust: at extremely low HRV, completion probability should approach zero for cognitively demanding cards; at typical HRV, completion probability should be substantially higher.

This prediction can be evaluated with relatively modest data: a single patient over several weeks generates the within-subject variation needed. A more rigorous version of the prediction adds a discontinuity at the Step-Zero threshold: completion probability for learning cards should drop sharply when HRV crosses below the threshold, since the system itself bypasses learning cards under that condition. This is, of course, a property of the system rather than of the underlying biology, but it provides a quality-control check that the system is behaving as designed.

6.1.2 Prediction 2: Optimal Challenge (Information Gain Validation)

If cards calibrated one level above functional level provide higher expected information gain than cards at functional level or two levels above, then completion rates for level-appropriate cards should be highest, with a curvilinear (inverted-U) relationship between difficulty and completion. Specifically, completion rate should be maximal for cards at functional+1, lower for cards at functional level (too easy), and lowest for cards at functional+2 or higher (too hard).

This prediction can be evaluated by occasionally delivering cards at functional level and at functional+2 (with patient consent and clinical oversight) and comparing completion rates with those at the standard functional+1. A practical variant uses naturally occurring variation: when the system selects cards at slightly different levels because of dimension-coverage or recency considerations, the resulting completion-rate differences can be analyzed. Adequately powered analyses will likely require pooled within-subject data from several patients.

6.1.3 Prediction 3: Developmental Progression

If sustained engagement with cards at functional+1 scaffolds developmental progression, then over time the functional level for engaged dimensions should rise. The prediction is dose-dependent: dimensions in which the patient engages with more cards should show larger functional-level gains than dimensions in which the patient engages with fewer cards, holding intake level constant. Across patients, this prediction yields a positive correlation between engagement and progression.

This prediction operates on longer timescales (months to a year) than the prior two, and requires more careful measurement of functional level. The Learning Catcher implementation currently uses card-performance patterns to update level estimates; future implementations will need to incorporate independent assessment (e.g., periodic standardized testing) to ensure that progression estimates are not artifactually inflated by within-system improvements.

6.1.4 Prediction 4: OMR Engagement and Self-Awareness Trajectory

If OMR is the procedural operationalization of self-model refinement, then the pattern of OMR responses—rate of confirmation, qualification, disconfirmation, and elaboration—should track the patient's self-awareness Fischer level. Specifically, patients at lower self-awareness levels should show higher rates of either unqualified confirmation (agreeable response bias) or blank confusion (level too high). Patients at higher self-awareness levels should show more qualified responses, more disconfirmation, and more spontaneous elaboration.

This prediction is testable through coded analysis of OMR response transcripts. It requires inter-rater reliability work on the coding scheme but does not require additional clinical measurement. The validation can proceed retrospectively on existing OMR data, with the prediction operationalized as a correlation between coded response complexity and self-awareness level estimated independently by clinical interview.

6.1.5 Prediction 5: Five-Dimensional Profile Heterogeneity

If skill development is genuinely web-structured rather than globally synchronized, then within diagnostic categories (e.g., adult ADHD, generalized anxiety, depression) the five-dimensional profiles should show substantial heterogeneity. Specifically, patients with identical DSM diagnoses should differ from one another in five-dimensional profile by more than the difference between mean profiles across diagnostic categories. Quantitatively: within-diagnosis profile variance should be larger than between-diagnosis profile variance.

This prediction is the most resource-intensive to test, requiring larger samples than the others. However, even with modest samples (10–20 patients per diagnostic category), the within-versus-between comparison can be informative. If profiles are indeed heterogeneous within diagnosis, the practical implication is consequential: clinical decision support based on five-dimensional profile may yield treatment-relevant distinctions that DSM categories do not.

6.2 Phased Validation Roadmap

We propose a four-phase validation program corresponding to the natural growth of the practice and the increasing maturity of the framework. Each phase is designed to be feasible within the resources of a small clinical practice while supporting cumulative evidence. Phases are listed in Table 6 with expected timelines and primary outcomes.

6.2.1 Phase 1: Single-Case Series (Months 1–6)

With the beta cohort of five patients, a single-case experimental series can evaluate Predictions 1 and 2 (HRV-precision, optimal challenge). The design uses each patient as their own control, with within-subject variation in HRV providing the natural experiment for Prediction 1 and occasional cross-level card delivery (with consent and clinical oversight) for Prediction 2.

Sample size constraints preclude strong inferential claims, but effect-size estimation and pre-registration provide meaningful inferential discipline. Expected output: a methodology paper and a brief empirical companion.

6.2.2 Phase 2: Extended Cohort and OMR Analysis (Months 6–18)

As the practice scales to 15–25 patients, Predictions 3 and 4 become accessible. Developmental progression (Prediction 3) becomes measurable on longer trajectories; OMR response analysis (Prediction 4) requires sufficient corpus volume that 15–25 patients with several months of data each is roughly the lower bound. Expected output: empirical papers on each prediction, with the OMR coding scheme formalized and inter-rater reliability established.

6.2.3 Phase 3: Multi-Site Replication and Profile Heterogeneity (Months 18–36)

Prediction 5 (profile heterogeneity) requires samples larger than typical small practices can generate. Phase 3 envisions collaboration with one or more allied practices to pool data on five-dimensional profiles within diagnostic categories. Multi-site collaboration introduces methodological complexity (training, fidelity monitoring, shared infrastructure) but is also where the framework's claims about web structure become genuinely testable. Expected output: a multi-site validation paper, ideally in a high-impact clinical journal.

6.2.4 Phase 4: Mechanism Studies (Months 36+)

With the architectural predictions established or refined, Phase 4 turns to mechanism: studies that probe whether the framework's specific computational commitments are necessary, or whether simpler alternative architectures yield comparable performance. This includes ablation studies (e.g., removing the Step-Zero Regulator and observing whether learning is degraded), parameter studies (e.g., systematically varying the precision threshold), and comparison with alternative decision rules (e.g., simple weighted-feature recommenders without explicit EFE). Phase 4 work is the natural setting for theoretical refinement and for engagement with the broader computational psychiatry literature.

6.3 Limits of the Framework

It is important to be explicit about what the framework does not and cannot do, both to maintain epistemic honesty and to direct attention to where further development is needed. Three limitations bear specific mention.

First, the framework is silent on the substantive content of clinical care. It specifies how to select interventions given a generative model and current state, but it does not specify what should be in the card library, how to design effective cards, or how to recognize specific clinical presentations. These remain matters of clinical expertise, accumulated case knowledge, and the broader therapeutic literature. The framework supports rather than supplants this substantive knowledge.

Second, the framework's developmental commitments are largely structural rather than substantive. We have proposed that Fischer levels correspond to depths of the generative model and that developmental progression corresponds to the emergence of new levels. We have not specified what those levels look like in functional or neural terms, nor have we modeled the mechanisms of structural change in the kind of detail that would support direct neuroscientific testing. These remain open questions for collaboration with computational neuroscience.

Third, the precision-HRV linkage, while supported by published evidence (Smith et al., 2017; Thayer et al., 2009), is a domain of active inquiry rather than settled science. The framework's clinical utility does not depend on the precision-HRV mapping being precisely correct in all details; it depends on HRV being a useful clinical index of an autonomic-cognitive state that gates learning, which is empirically robust. Future work will need to refine the mapping, particularly for clinical populations in which HRV interpretation is complicated by medication, comorbidity, or atypical autonomic patterns.

Table 6

Phased Validation Roadmap for Developmental Active Inference

Phase	**Timeline**	**Predictions Tested**	**Sample and Method**
1. Single-Case Series	Months 1–6	P1 (HRV-precision); P2 (optimal challenge)	N = 5 beta patients; within-subject ABA designs
2. Extended Cohort	Months 6–18	P3 (progression); P4 (OMR-self-awareness)	N = 15–25; longitudinal; coded OMR responses
3. Multi-Site Replication	Months 18–36	P5 (profile heterogeneity); P1–P4 replication	Pooled multi-practice; N ≥ 60
4. Mechanism Studies	Months 36+		

Ablation; parameter sensitivity; alt. architectures
Comparative system designs
Note. Each phase is designed to be feasible within the resources of a small clinical practice (Phases 1–2) or modest multi-site collaboration (Phase 3). Phases 1 and 2 yield empirical papers individually; Phase 3 supports a higher-impact validation; Phase 4 supports theoretical refinement and engagement with computational psychiatry. P1–P5 refer to Predictions 1–5 in Section 6.1.

7. Conclusion

Self-regulated learning has long been recognized as central to educational and clinical practice, but the field has lacked a computational mechanism by which to ground its phenomenology. We have proposed Developmental Active Inference (DAI) as such a mechanism, integrating four traditionally separate theoretical traditions: Active Inference for the computational machinery, Fischer's dynamic skill theory for the developmental architecture, polyvagal theory for the autonomic grounding of precision, and habit theory for the consolidation of learned policies. We introduced Observable Model Refinement (OMR) as a clinical operationalization of self-model refinement, complementing the three other mechanisms with an explicit procedural intervention.

The framework's central proposal is that the four traditionally separate components are not arbitrary features to be combined by weighted summation but natural components of a single computational objective—expected free energy—decomposable as the sum of expected information gain and expected utility, gated by precision. This integration has three immediate consequences. First, it provides a theoretically principled basis for adaptive intervention selection in clinical settings, replacing ad hoc or data-driven recommendation with a transparent computation grounded in established neuroscience. Second, it links short-timescale phenomena (moment-to-moment regulation, single-card learning) with long-timescale phenomena (developmental progression, self-model refinement) through a single mechanism. Third, it generates testable predictions at multiple levels of resolution, allowing the framework to be evaluated by phased empirical work appropriate to small clinical practice.

We have illustrated the framework through the Learning Catcher system, currently in beta evaluation with a small cohort. The empirical claims developed in Section 6 will be evaluated as the cohort matures and as the practice scales. We expect those evaluations to revise the framework in some respects; specifications of the precision-HRV mapping, the level-card calibration procedures, and the OMR response coding scheme are all likely to be refined in light of accumulating data. The framework's theoretical commitments are stated at a level intended to remain useful even as these specifications evolve.

Three broader contributions warrant brief mention. First, for clinical psychology, the framework offers an alternative to data-driven recommendation systems that has the virtues of interpretability and theoretical grounding. The choice between data-driven and theory-driven adaptive systems is not, in our view, a competition: data-driven systems have produced extraordinary advances in many domains, and there is no general argument against them. But in clinical settings where transparency, accountability, and the explicit modeling of developmental architecture are valuable, a theory-driven framework has specific advantages. We have aimed to identify a framework that is rigorous enough to support clinical decision-making while remaining accessible enough to be implemented by clinical practitioners without specialized computational training.

Second, for developmental science, the framework provides a computational complement to the structural account of skill development that dynamic skill theory has provided over four decades. Fischer's framework specified what developmental change looks like; Active Inference offers an account of how that change is computationally realized. The integration is not the only such bridge that could be built—Vygotskyan and Piagetian frameworks each admit Active Inference readings, and other developmental theories will doubtless yield productive integrations as well. But the specific bridge we have constructed, between Fischer's web-structured developmental hierarchy and the hierarchical generative models of Active Inference, has, we believe, particular fidelity to both source frameworks.

Third, for the computational neuroscience of psychiatric and educational intervention, the framework illustrates how a unified computational principle can support practically useful clinical decision support without sacrificing theoretical rigor. The free energy principle has been criticized at times for being so general as to be empirically vacuous (Klein, 2018; Andrews, 2021). The framework presented here is, we hope, a productive response to that concern: by specifying which clinical quantities map onto which computational components, and by stating predictions that can be empirically evaluated, the framework demonstrates that the FEP can yield specific, falsifiable claims when integrated with appropriate developmental and physiological theory.

Future work will pursue the validation roadmap outlined in Section 6, with particular attention to the OMR coding scheme and the multi-site collaboration that Phase 3 will require. We also intend to develop the framework's connections to neighboring literatures—particularly to mindfulness-based approaches, which share with DAI a commitment to present-moment awareness and to the regulatory significance of autonomic state, and to the emerging literature on enactivism in clinical psychology, which shares with DAI a commitment to skill as situated activity rather than abstract competence.

The framework, like the patients it is intended to serve, is in a process of development. Its initial articulation, as presented in this paper, reflects our current best understanding; revisions and extensions are expected and welcomed. What we hope to have established is that a theory-driven, developmentally informed, computationally grounded approach to self-regulated learning is feasible, useful, and capable of supporting cumulative empirical work. The path from here to a mature framework will be long. It begins, we suggest, with the explicit articulation of the framework's commitments and the public statement of its predictions, both of which the present paper has attempted to provide. **Figures**

This section describes the four primary figures of the paper. In the final submitted version, each figure would be provided as a high-resolution image file (PDF or PNG); here we provide a text description of each, suitable for sketching and refinement before final figure preparation. Figures are referenced from the corresponding sections of the main text.

Figure 1. The Generative Model in Active Inference

Schematic of a hierarchical generative model. Three vertical levels are shown, from bottom (sensorimotor, fastest temporal scale, most concrete) to top (abstract, slowest temporal scale, most abstract). Each level contains a node representing hidden states (s) and an arrow to the level below indicating top-down prediction (red). Arrows from the level below upward represent prediction error (blue). The bottom level receives sensory observations (o); the top level connects to prior preferences (C). The diagram emphasizes the bidirectional flow of information and the role of each level in generating predictions for the level below. Annotations identify the components named in Section 2.2.

Figure 2. Fischer's Hierarchy of Skill Levels

Diagram of the four tiers and 13 levels of Fischer's dynamic skill theory. The vertical axis represents level (1 at bottom through 13 at top, with 14–16 indicated as more rarely attained). Within each tier, the four levels are shown as nested boxes (single, mapping, system, system of systems), illustrating the recursive structure. The horizontal axis represents approximate age of optimal-level emergence under supportive conditions. Within the adult range (right side of figure), the levels relevant to clinical work are highlighted. The diagram is adapted from Fischer and Bidell (2006) with structural simplifications appropriate to introductory presentation.

Figure 3. The Four Mechanisms of Developmental Active Inference

Conceptual diagram showing how the four mechanisms of Section 4 combine into a single computational framework. At the center, a hierarchical generative model is shown (cf. Figure 1) with Fischer levels annotated alongside the levels of the model (Mechanism 1: hierarchical structure). To the left, an arrow labeled γ (precision) modulates the model, with a sub-arrow indicating that γ is informed by HRV (Mechanism 2: precision allocation). To the right, an arrow labeled "learned policies" indicates the consolidation of action sequences (Mechanism 3: empirical priors). At the top, an arrow labeled "OMR" indicates the externalized hypothesis-update loop (Mechanism 4: self-model refinement). The four mechanisms together determine expected free energy G, computed for each candidate intervention; the system selects the intervention minimizing G subject to a precision threshold.

Figure 4. Learning Catcher Session Flow

Flowchart of a single session in the Learning Catcher system. Entry node: patient logs in. First decision node: HRV measurement (or recent home recording). Step-Zero Regulator branch: if HRV below threshold, route to regulation library; if HRV above threshold, route to learning library. Within learning library: compute EFE across five dimensions; select dimension-and-card combination minimizing G; deliver card. After card completion: collect response and emotion rating; (optional) deliver OMR hypothesis; collect OMR response; update system parameters. Exit node: session complete. Dotted-line arrows indicate parameter-update flows back to the central state representation. The diagram emphasizes that no element of session flow is arbitrary: every branch is determined by a transparent computation.

References

Adams, R. A., Shipp, S., & Friston, K. J. (2013). Predictions not commands: Active inference in the motor system. Brain Structure and Function, 218(3), 611–643. https://doi.org/10.1007/s00429-012-0475-5

Andrews, M. (2021). The math is not the territory: Navigating the free energy principle. Biology & Philosophy, 36(3), 30. https://doi.org/10.1007/s10539-021-09807-0

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

Barkley, R. A. (2012). Executive functions: What they are, how they work, and why they evolved. Guilford Press.

Bates, D. W., Kuperman, G. J., Wang, S., Gandhi, T., Kittler, A., Volk, L., Spurr, C., Khorasani, R., Tanasijevic, M., & Middleton, B. (2003). Ten commandments for effective clinical decision support: Making the practice of evidence-based medicine a reality. Journal of the American Medical Informatics Association, 10(6), 523–530. https://doi.org/10.1197/jamia.M1370

Berking, M., & Wupperman, P. (2012). Emotion regulation and mental health: Recent findings, current challenges, and future directions. Current Opinion in Psychiatry, 25(2), 128–134. https://doi.org/10.1097/YCO.0b013e3283503669

Bruner, J. (1985). Vygotsky: A historical and conceptual perspective. In J. V. Wertsch (Ed.), Culture, communication, and cognition: Vygotskian perspectives (pp. 21–34). Cambridge University Press.

Clear, J. (2018). Atomic habits: An easy and proven way to build good habits and break bad ones. Avery.

Conners, C. K. (2014). Conners Continuous Performance Test—Third Edition (Conners CPT 3). Multi-Health Systems.

Constant, A., Ramstead, M. J. D., Veissière, S. P. L., Campbell, J. O., & Friston, K. J. (2018). A variational approach to niche construction. Journal of the Royal Society Interface, 15(141), 20170685. https://doi.org/10.1098/rsif.2017.0685

Dignath, C., & Büttner, G. (2008). Components of fostering self-regulated learning among students. A meta-analysis on intervention studies at primary and secondary school level. Metacognition and Learning, 3(3), 231–264. https://doi.org/10.1007/s11409-008-9029-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., & 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.

Fischer, K. W., & Yan, Z. (2002). Darwin's construction of the theory of evolution: Microdevelopment of explanations of variation and change in species. In N. Granott & J. Parziale (Eds.), Microdevelopment: Transition processes in development and learning (pp. 294–318). Cambridge University Press.

FitzGerald, T. H. B., Schwartenbeck, P., Moutoussis, M., Dolan, R. J., & Friston, K. (2015). Active inference, evidence accumulation, and the urn task. Neural Computation, 27(2), 306–328. https://doi.org/10.1162/NECO_a_00699

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., 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

Friston, K., Kilner, J., & Harrison, L. (2006). A free energy principle for the brain. Journal of Physiology-Paris, 100(1–3), 70–87. https://doi.org/10.1016/j.jphysparis.2006.10.001

Friston, K. J., Parr, T., & de Vries, B. (2017). The graphical brain: Belief propagation and active inference. Network Neuroscience, 1(4), 381–414. https://doi.org/10.1162/NETN_a_00018

Friston, K., Schwartenbeck, P., FitzGerald, T., Moutoussis, M., Behrens, T., & Dolan, R. J. (2016). The anatomy of choice: Active inference and agency. Frontiers in Human Neuroscience, 7, 598. https://doi.org/10.3389/fnhum.2013.00598

Hansen, A. L., Johnsen, B. H., & Thayer, J. F. (2003). Vagal influence on working memory and attention. International Journal of Psychophysiology, 48(3), 263–274. https://doi.org/10.1016/S0167-8760(03)00073-4

Heckhausen, J., Wrosch, C., & Schulz, R. (2010). A motivational theory of life-span development. Psychological Review, 117(1), 32–60. https://doi.org/10.1037/a0017668

Hill, C. E. (2009). Helping skills: Facilitating exploration, insight, and action (3rd ed.). American Psychological Association.

Joffily, M., & Coricelli, G. (2013). Emotional valence and the free-energy principle. PLoS Computational Biology, 9(6), e1003094. https://doi.org/10.1371/journal.pcbi.1003094

Kirchner, W. K. (1958). Age differences in short-term retention of rapidly changing information. Journal of Experimental Psychology, 55(4), 352–358. https://doi.org/10.1037/h0043688

Klein, C. (2018). What do predictive coders want? Synthese, 195(6), 2541–2557. https://doi.org/10.1007/s11229-016-1250-6

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

MacKay, D. J. C. (2003). Information theory, inference, and learning algorithms. Cambridge University Press.

Maes, S., & Karoly, P. (2005). Self-regulation assessment and intervention in physical health and illness: A review. Applied Psychology, 54(2), 267–299. https://doi.org/10.1111/j.1464-0597.2005.00210.x

Marr, D. (1982). Vision: A computational investigation into the human representation and processing of visual information. W. H. Freeman.

Miller, W. R., & Rollnick, S. (2012). Motivational interviewing: Helping people change (3rd ed.). Guilford Press.

Niv, Y. (2009). Reinforcement learning in the brain. Journal of Mathematical Psychology, 53(3), 139–154. https://doi.org/10.1016/j.jmp.2008.12.005

Nunan, D., Sandercock, G. R. H., & Brodie, D. A. (2010). A quantitative systematic review of normal values for short-term heart rate variability in healthy adults. Pacing and Clinical Electrophysiology, 33(11), 1407–1417. https://doi.org/10.1111/j.1540-8159.2010.02841.x

Padesky, C. A. (1993, September). Socratic questioning: Changing minds or guiding discovery? Keynote address at the European Congress of Behaviour and Cognitive Therapies, London.

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

Piaget, J. (1972). Intellectual evolution from adolescence to adulthood. Human Development, 15(1), 1–12. https://doi.org/10.1159/000271225

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

Porges, S. W. (2011). The polyvagal theory: Neurophysiological foundations of emotions, attachment, communication, and self-regulation. W. W. Norton.

Rogers, C. R. (1957). The necessary and sufficient conditions of therapeutic personality change. Journal of Consulting Psychology, 21(2), 95–103. https://doi.org/10.1037/h0045357

Schmeichel, B. J., & Tang, D. (2015). Individual differences in executive functioning and their relationship to emotional processes and responses. Current Directions in Psychological Science, 24(2), 93–98. https://doi.org/10.1177/0963721414555178

Schwartenbeck, P., FitzGerald, T. H. B., Mathys, C., Dolan, R., & Friston, K. (2015). The dopaminergic midbrain encodes the expected certainty about desired outcomes. Cerebral Cortex, 25(10), 3434–3445. https://doi.org/10.1093/cercor/bhu159

Schwartenbeck, P., Passecker, J., Hauser, T. U., FitzGerald, T. H. B., Kronbichler, M., & Friston, K. J. (2019). Computational mechanisms of curiosity and goal-directed exploration. eLife, 8, e41703. https://doi.org/10.7554/eLife.41703

Schwartz, M. S., & Fischer, K. W. (2004). Building general knowledge and skill: Cognition and microdevelopment in science learning. In A. Demetriou & A. Raftopoulos (Eds.), Cognitive developmental change: Theories, models and measurement (pp. 157–185). Cambridge University Press.

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., Khalsa, S. S., & Paulus, M. P. (2021). An active inference approach to dissecting reasons for nonadherence to antidepressants. Biological Psychiatry: Cognitive Neuroscience and Neuroimaging, 6(9), 919–934. https://doi.org/10.1016/j.bpsc.2019.11.012

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

Thayer, J. F., Hansen, A. L., Saus-Rose, E., & Johnsen, B. H. (2009). Heart rate variability, prefrontal neural function, and cognitive performance: The neurovisceral integration perspective on self-regulation, adaptation, and health. Annals of Behavioral Medicine, 37(2), 141–153. https://doi.org/10.1007/s12160-009-9101-z

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

Vygotsky, L. S. (1978). Mind in society: The development of higher psychological processes (M. Cole, V. John-Steiner, S. Scribner, & E. Souberman, Eds.). Harvard University Press.

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

Zakay, D., & Block, R. A. (1997). Temporal cognition. Current Directions in Psychological Science, 6(1), 12–16. https://doi.org/10.1111/1467-8721.ep11512604

Zimmerman, B. J. (1989). A social cognitive view of self-regulated academic learning. Journal of Educational Psychology, 81(3), 329–339. https://doi.org/10.1037/0022-0663.81.3.329

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.