

Catcher Navigator Algorithm (CNA) v7:

A Bayesian Differential Diagnosis + Comorbidity Decomposition + Adaptive Re-routing Framework for Adult Neuropsychiatric Disorders

Alex Kwon, BCN, PhD¹

¹ *Boston Neuromind LLC, Canton, MA, USA*

Harvard Graduate School of Education, Visiting Scholar (2018–2020, with the late Prof. Kurt W. Fischer)

Correspondence: bostonneuromind@gmail.com

Manuscript prepared for Nature Mental Health
May 2026

Abstract

Background. Current adult neuropsychiatric diagnosis relies on DSM-5-TR categorical thresholds that ignore comorbidity decomposition. Adult ADHD shows approximately 80% comorbidity with at least one additional psychiatric disorder (Salazar de Pablo et al., 2025), yet identical clinical scores (e.g., attention deficit of 0.39) yield discordant treatment outcomes depending on which primary diagnosis drives the score. McIntosh et al. (2009) demonstrated that selective serotonin reuptake inhibitor monotherapy fails systematically in unrecognized adult ADHD, with a mean diagnostic delay of 6.8 years (Kessler et al., 2023).

Objective. We introduce the Catcher Navigator Algorithm version 7 (CNA v7), a four-module Bayesian framework that performs (i) Differential Diagnosis Engine (DDE) using DSM-5-TR canonical pattern matching, (ii) Comorbidity Decomposition assigning fractional weights to each candidate diagnosis, (iii) Cognitive Simulator generating counterfactual training-outcome predictions, and (iv) Adaptive Re-routing using 4-week longitudinal outcomes to auto-revise diagnostic priors.

Methods. CNA v7 ingests five clinical axes (attention, learning efficiency, peak performance, anxiety regulation, mood stability) integrated with biomarkers (QEEG, HRV) and temporal patterns. Bayesian posterior estimation distributes a single clinical score across multiple diagnostic sources. The adaptive layer compares observed 4-week outcomes against pattern-predicted trajectories and triggers diagnostic reclassification on mismatch. A synthetic cohort ($N = 1,000$) was generated for internal validation, and five preliminary clinical participants completed an 8-week protocol.

Results. Synthetic validation yielded a mean diagnostic AUC of 0.90 (95% CI: 0.87–0.93) across 12 candidate diagnoses. Three clinical case studies demonstrated that comorbidity decomposition distinguished substantively different primary drivers underlying numerically identical attention scores of 0.39. Two of five participants triggered adaptive re-routing events at the 4-week mark, with posterior shifts grounded in explicit Bayesian rationale.

Conclusions. CNA v7 reframes neuropsychiatric assessment as a continuous Bayesian inference problem rather than categorical thresholding, opening a longitudinal adaptive decision layer that complements existing QEEG normative databases.

Keywords: *differential diagnosis, comorbidity, Bayesian inference, adult ADHD, adaptive treatment, Active Inference, Dynamic Skill Theory.*

1.1 The Categorical Diagnosis Problem

The diagnosis of adult neuropsychiatric disorders remains anchored to categorical thresholds defined by the **Diagnostic and Statistical Manual of Mental Disorders, Fifth Edition, Text Revision** (DSM-5-TR; American Psychiatric Association, 2022). Under this framework, a patient meets criteria for attention-deficit/hyperactivity disorder (ADHD), generalized anxiety disorder (GAD), or major depressive disorder (MDD) when a sufficient number of symptoms exceed a binary threshold. This categorical approach, while clinically useful for communication and billing, obscures three critical realities of contemporary neuropsychiatric practice.

First, comorbidity in adult psychiatric populations is not the exception but the rule. A recent meta-analysis of clinical populations (Salazar de Pablo et al., 2025) demonstrated that approximately 80% of adults with ADHD meet criteria for at least one additional psychiatric disorder, with mood and anxiety disorders representing the most common co-occurring conditions. Cortese (2025), in a comprehensive **World Psychiatry** review, reported that the worldwide prevalence of adult ADHD has stabilized at approximately 2.5%, with persistence into adulthood predicted by symptom severity, psychosocial adversity, and—critically—psychiatric comorbidity itself. The World Mental Health Surveys (Kessler et al., 2023) found that nearly 20% of adults with ADHD carry diagnoses spanning three or more additional disorder classes.

Second, identical clinical scores frequently reflect divergent underlying mechanisms. A clinician confronted with an attention deficit score of 0.39 on standardized assessment cannot, from that number alone, distinguish between a patient whose impairment is driven primarily by ADHD versus one whose attention difficulty is secondary to ruminative anxiety, anhedonic depression, or trauma-related hypervigilance. The treatment implications are radically different. McIntosh et al. (2009) demonstrated, in a now-foundational clinical observation, that selective serotonin reuptake inhibitor (SSRI) monotherapy systematically fails in patients with unrecognized adult ADHD even when depressive symptoms are present. The patient does not respond not because the medication is ineffective but because the primary diagnosis driving the clinical picture has been misidentified.

Third, diagnostic decisions in real-world practice are made under uncertainty, and that uncertainty does not resolve at the initial assessment. Kessler et al. (2023) documented that the mean diagnostic delay for adult ADHD in the United States is 6.8 years, during which time patients accumulate iatrogenic complications, mistreatment with inappropriate pharmacotherapy, and substantial psychosocial cost. This delay is not merely an access problem; it reflects a structural inadequacy of the categorical diagnostic system itself, which provides no formal mechanism for revising a diagnostic hypothesis when treatment outcomes diverge from prediction.

1.2 The Comorbidity Decomposition Gap

Existing clinical infrastructure has responded to the comorbidity challenge by deploying parallel single-disorder assessment instruments: the Adult ADHD Self-Report Scale (ASRS) for attention, the Generalized Anxiety Disorder 7-item scale (GAD-7) for anxiety, the Patient Health Questionnaire-9 (PHQ-9) for depression, and so forth. Each instrument returns a score for its target construct. The clinician is then required, informally and without computational support, to integrate these isolated scores into a unified diagnostic formulation.

Quantitative electroencephalography (QEEG) normative databases—including Neuroguide (Thatcher et al., 2003), the Human Brain Index database (Kropotov, 2009), and the Mitsar 19-channel system—have extended this single-snapshot approach into the neurophysiological domain. These tools compare an individual's spectral and event-related potential (ERP) profile against age-matched population norms, returning z-scores indicating deviance. They are powerful within their scope but share a structural limitation: they describe the patient at a moment in time without modeling how the clinical phenotype emerges as a weighted superposition of multiple latent diagnostic sources, nor how that superposition might shift as treatment unfolds.

What is missing, then, is a computational layer that performs three distinct operations: (a) fractional attribution of a single observable clinical score across multiple candidate diagnoses, weighted by Bayesian posterior probability; (b) counterfactual simulation of how interventions on one clinical axis propagate through the multivariate clinical state; and (c) longitudinal revision of diagnostic priors as treatment outcomes accumulate. No existing commercial or open-source system, to our knowledge, integrates these three capacities into a single clinical decision support framework.

1.3 Foundational Frameworks Bridged in the Present Work

The Catcher Navigator Algorithm version 7 (CNA v7), described in this paper, draws on four converging research traditions.

The first is Dynamic Skill Theory (Fischer, 1980; Fischer & Bidell, 2006), which conceptualizes cognitive and emotional capacities as hierarchically nested skills that emerge dynamically through interaction with environmental and physiological context. The late Professor Kurt W. Fischer, with whom the present author trained as a Visiting Scholar at Harvard Graduate School of Education, advanced the position that skill levels are not fixed traits but observable states whose trajectory through development and intervention can be measured and modeled. CNA v7 operationalizes this insight by treating each of five clinical axes as a dynamic skill whose current value reflects the joint contribution of multiple underlying diagnostic processes.

The second is Active Inference and the Free Energy Principle (Friston, 2010; Friston, 2023), which frames neural and behavioral function as Bayesian inference over a hierarchical generative model. Recent extensions of this framework to psychiatry (Smith et al., 2024; Krishnadas et al., 2025) demonstrate that clinical phenomenology can be decomposed into specific computational primitives—precision-weighted prediction error, generative model parameters, hierarchical belief updating—that map onto disorder-specific patterns. Bayesian formulations of mood (Clark, Watson, & Friston, 2018) and depression (Kirchner et al., 2024) have shown that affective psychopathology can be modeled as an inference problem rather than a static categorical state.

The third is computational psychiatry as a clinical-translational discipline (Huys, Maia, & Frank, 2016; Browning et al., 2020; Hess et al., 2025). Hess and colleagues, working within the Translational Neuromodeling Unit at ETH Zurich, recently published a Bayesian workflow for generative modeling in computational psychiatry that formalizes how hierarchical Gaussian filtering and related techniques can be deployed to extract individual-level computational parameters from behavioral and neurophysiological data. CNA v7 builds on this foundation but extends it in the direction of clinician-facing decision support: rather than estimating computational parameters as research endpoints, we use Bayesian inference to drive concrete diagnostic and treatment-selection decisions.

The fourth is dimensional and hierarchical reorganizations of psychopathology, including the Research Domain Criteria (RDoC) framework of the National Institute of Mental Health (Cuthbert & Insel, 2013) and the Hierarchical Taxonomy of Psychopathology (HiTOP; Kotov et al., 2017; Conway et al., 2024). Both traditions reject the assumption that psychiatric disorders are mutually exclusive natural categories and instead propose that observable phenotypes emerge from underlying continuous dimensions of dysfunction. CNA v7 is consonant with this dimensional turn but operates one layer closer to clinical practice: it accepts the categorical labels used by clinicians (ADHD, GAD, MDD, etc.) as practical communication units while decomposing each patient's presentation into fractional contributions from those categories.

1.4 The Present Contribution

CNA v7 integrates four computational modules into a single clinical workflow. The Differential Diagnosis Engine (DDE) performs Bayesian posterior estimation over twelve candidate diagnoses using DSM-5-TR canonical patterns, multimodal biomarker profiles (QEEG, heart rate variability, cognitive testing), symptom self-report, and temporal pattern information. The Comorbidity Decomposition module distributes each observed clinical-axis score across the candidate diagnoses according to their posterior weights, returning a fractional attribution vector.

The Cognitive Simulator generates counterfactual predictions of how clinical-axis scores would shift under specified training interventions. The Adaptive Re-routing module compares observed four-week outcomes against simulator-predicted trajectories and, upon mismatch exceeding a threshold, revises the diagnostic posterior and recommends an updated treatment protocol.

These modules operate over a five-axis clinical architecture spanning attention, learning efficiency, peak performance, anxiety regulation, and mood stability. The architecture is intentionally Layer 2 in design: it complements rather than replaces existing QEEG normative tools and standardized symptom instruments, ingesting their outputs as inputs to the Bayesian engine.

1.5 Hypotheses

We advance three formal hypotheses, to be tested in the present work through synthetic validation and a preliminary clinical beta cohort, and in subsequent prospective trials.

Hypothesis 1. Bayesian posterior decomposition of clinical-axis scores predicts treatment response with greater accuracy than categorical primary diagnosis alone. Specifically, treatment selection based on the highest-posterior diagnosis among decomposed candidates will yield superior four-week outcomes compared to treatment selection based on the symptom-scale instrument with the highest absolute score.

Hypothesis 2. Comorbidity-aware treatment selection reduces time-to-clinically-meaningful improvement. We predict that patients whose primary diagnosis is correctly identified at intake through CNA v7 will reach clinical improvement thresholds in fewer treatment cycles than patients diagnosed by conventional categorical means.

Hypothesis 3. Adaptive re-routing triggered by four-week outcome divergence prevents diagnostic lock-in and reduces cumulative diagnostic delay. We predict that the adaptive layer will identify cases of initial misdiagnosis at the four-week mark, with a substantial reduction in time-to-correct-diagnosis compared to historical baseline.

The remainder of this paper proceeds as follows. Section 2 details the theoretical framework underlying the four modules. Section 3 describes the methods, including the diagnostic fingerprints database, the Five-Lens AI framework, and the beta study protocol. Section 4 presents results from synthetic validation and preliminary clinical cases. Section 5 discusses implications, limitations, and future directions.

2.1 Five-Axis Clinical Architecture

CNA v7 operates on a five-axis clinical state vector, $\mathbf{s} = (s_1, s_2, s_3, s_4, s_5)$, where each component represents an individual's current functional level on a distinct domain of neuropsychiatric functioning. The selection of these five axes was guided by three converging constraints: clinical interpretability for practicing neurofeedback and counseling clinicians, measurability through validated standardized instruments and noninvasive biomarkers, and theoretical mappability onto Dynamic Skill Theory's hierarchical skill structure (Fischer & Bidell, 2006).

Axis 1: Attention (s_1) captures the patient's capacity for sustained, selective, and divided attention. Operational measures include the Conners Continuous Performance Test (CPT-3), the Adult ADHD Self-Report Scale (ASRS v1.1), and the theta/beta power ratio at frontal-central electrode sites derived from QEEG (Arns et al., 2013; Faraone et al., 2024). Z-score normalization against age-matched healthy population data converts raw measurements into the unit-free attention score.

Axis 2: Learning Efficiency (s_2) captures working memory capacity, encoding-retrieval efficiency, and consolidation. Operational measures include adaptive N-back performance, the Digit Span subtest of the Wechsler Adult Intelligence Scale-IV (WAIS-IV), retention curves on spaced-repetition tasks, and the P300 ERP component, particularly the parietal P3b reflecting working memory updating (Polich, 2007). This axis corresponds operationally to what Fischer (1980) termed the "skill ceiling" for the patient's current developmental level.

Axis 3: Peak Performance (s_3) captures the capacity for sustained high-engagement cognitive work, time perception under load, and autonomic-cognitive coupling. Operational measures include the time estimation task (subjective duration estimation of fixed objective intervals), heart rate variability coherence ratio (high-frequency to low-frequency power ratio under standardized respiration), and flow-state self-report measures.

Axis 4: Anxiety Regulation (s_4) captures the capacity to maintain regulated arousal under stress, recover from acute stress challenges, and prevent ruminative escalation. Operational measures include the GAD-7, HRV in the low-frequency band as an index of sympathetic-parasympathetic balance, and frontal alpha asymmetry, with greater right-frontal alpha indicating withdrawal-related affect (Coan & Allen, 2004).

Axis 5: Mood Stability (s_5) captures the capacity for stable hedonic tone, behavioral activation, and recovery from anhedonic states. Operational measures include the PHQ-9, the daily

Behavioral Activation count, and theta cordance at frontal electrodes, which has shown predictive validity for antidepressant response (Cook et al., 2002).

Table 1 provides the complete mapping from each axis to instruments and biomarkers.

2.2 Bayesian Differential Diagnosis Engine

The Differential Diagnosis Engine (DDE) implements full Bayesian posterior estimation over a diagnostic space $\mathbf{D} = \{D_1, D_2, \dots, D_K\}$ of $K = 12$ candidate diagnoses. The diagnostic space spans the three ADHD subtypes (predominantly inattentive, predominantly hyperactive-impulsive, combined), four anxiety-spectrum conditions (generalized anxiety disorder, panic disorder, social anxiety disorder, and post-traumatic stress disorder), obsessive-compulsive disorder, and four mood-spectrum conditions (major depressive disorder, persistent depressive disorder, bipolar I disorder, and bipolar II disorder).

For an observed clinical feature vector $\mathbf{f} = (f_1, \dots, f_M)$ comprising the five clinical axes augmented with biomarker, symptom, and temporal pattern features, the posterior probability for each candidate diagnosis is computed as

Equation 1:

$$P(D_i | \mathbf{f}) = [P(\mathbf{f} | D_i) \times P(D_i)] / \sum_{j=1}^K P(\mathbf{f} | D_j) \times P(D_j)$$

The prior probabilities $P(D_i)$ are drawn from epidemiological data, with the National Comorbidity Survey Replication (NCS-R; Kessler et al., 2005) and the National Epidemiologic Survey on Alcohol and Related Conditions III (NESARC-III; Grant et al., 2015) providing baseline rates. Priors can be updated patient-specifically based on referral context, age, and history.

The likelihood functions $P(\mathbf{f} | D_i)$ are constructed from the diagnostic fingerprints database described in Section 3.2. For each diagnosis D_i , the fingerprint specifies expected values and variances for the 38 features across 8 modalities. The likelihood is computed under a multivariate Gaussian assumption with diagonal covariance, with feature-wise mean and standard deviation estimated from clinical reference cohorts. Following the workflow of Hess et al. (2025), we evaluate model identifiability through posterior predictive checks before clinical deployment.

2.3 Comorbidity Decomposition

The comorbidity decomposition module performs the operation absent from standard diagnostic instruments: it distributes each observed clinical-axis score across the candidate diagnoses according to their posterior weights. For axis α and observed score s_α , the decomposed contribution from diagnosis D_i is

Equation 2:

$$s_{\alpha}^{(i)} = w_i \times c_{\{\alpha, i\}}$$

where $w_i = P(D_i | \text{***})$ is the posterior weight on diagnosis D_i , and $c_{\{\alpha, i\}}$ is the expected contribution of diagnosis D_i to clinical axis α , drawn from the diagnostic fingerprints database. The constraint $\sum_i s_{\alpha}^{(i)} = s_{\alpha}$ holds by construction.

Consider the canonical illustrative case introduced in Section 1: an observed attention score $s_1 = 0.39$. Under classical categorical diagnosis, this single number admits only one disposition: above or below threshold. Under CNA v7 decomposition, this same score might resolve into

$$s_1^{(ADHD)} = 0.25$$

$$s_1^{(GAD)} = 0.10$$

$$s_1^{(MDD)} = 0.04$$

in one patient, and

$$s_1^{(ADHD)} = 0.06$$

$$s_1^{(GAD)} = 0.28$$

$$s_1^{(MDD)} = 0.03$$

in another. The two decompositions identify radically different primary drivers and direct radically different treatment selections, even though the patients are indistinguishable on the attention score alone.

This decomposition formalizes what experienced clinicians do informally during differential diagnosis: weight the contribution of each candidate diagnosis to the observed clinical picture. CNA v7 makes this weighting explicit, quantitative, and auditable.

2.4 Cognitive Simulator: Counterfactual Inference

The Cognitive Simulator module addresses the question: "If we intervene to improve clinical axis α by Δ , what is the predicted change in axes β , γ , δ , ϵ ?" This counterfactual inference draws on Pearl's causal hierarchy (Pearl, 2009; Pearl & Mackenzie, 2018), specifically the third rung concerning interventions and their downstream effects.

The simulator implements a directed acyclic causal graph $G = (V, E)$, where vertices V correspond to the five clinical axes augmented with intermediate cognitive constructs (working memory capacity, executive control, emotion regulation strategy use, behavioral activation), and edges E encode causal influences. Edge weights are derived from the empirical literature: for

instance, the meta-analytic effect of working memory training on attention ($r = 0.43$ in Sala & Gobet, 2017, though contested by Melby-Lervåg et al., 2016), and the meta-analytic effect of emotion regulation training on anxiety symptoms ($d = 0.605$; Aldao, Nolen-Hoeksema, & Schweizer, 2010).

For a proposed intervention I that targets axis α with magnitude Δ_α , the simulator computes the predicted post-intervention state vector via the do-operator:

Equation 3:

$$E[s_\beta \mid \text{do}(s_\alpha := s_\alpha + \Delta_\alpha)] = s_\beta + \sum_{\{\text{paths } \alpha \rightarrow \beta\}} (\text{product of edge weights along path}) \times \Delta_\alpha$$

The counterfactual prediction serves two purposes. Clinically, it allows the clinician to evaluate alternative treatment options against their projected multivariate effects before committing to a protocol. Methodologically, it generates the predicted trajectory against which observed four-week outcomes will be compared by the Adaptive Re-routing module (Section 2.5).

2.5 Adaptive Re-routing: Outcome-Driven Posterior Revision

The Adaptive Re-routing module implements Bayesian update of the diagnostic posterior in response to observed treatment outcomes. At intake (time $t = 0$), the DDE returns posterior probabilities $P_0(D_i \mid \mathbf{f}_0)$ over the diagnostic space. The patient is enrolled in a treatment protocol informed by the primary posterior diagnosis. At $t = 4$ weeks, post-intervention measurements yield a new feature vector $\mathbf{f}_4$. The Cognitive Simulator had earlier generated a predicted feature vector $\mathbf{f}_4^{\{\text{predicted} \mid D_i\}}$ for each candidate diagnosis under the assumption that D_i was the true primary.

The posterior is updated according to the likelihood of the observed outcome under each diagnostic hypothesis:

Equation 4:

$$P_4(D_i \mid \mathbf{f}_0, \mathbf{f}_4) \propto P_0(D_i \mid \mathbf{f}_0) \times P(\mathbf{f}_4 \mid \mathbf{f}_4^{\{\text{predicted} \mid D_i\}})$$

Diagnoses whose predicted outcomes match the observed four-week trajectory gain posterior weight; diagnoses whose predictions diverge lose weight. If the updated posterior P_4 differs sufficiently from P_0 (defined operationally as a shift of the maximum-a-posteriori diagnosis or a Kullback-Leibler divergence exceeding a threshold), the module triggers a diagnostic re-routing event. The clinician is presented with the revised posterior, the rationale (which predictions matched, which diverged), and a revised treatment recommendation.

This adaptive layer is the architectural feature distinguishing CNA v7 from snapshot diagnostic tools. The same framework can be applied iteratively at 8 weeks, 12 weeks, and beyond, with each cycle providing additional Bayesian evidence to refine the diagnostic hypothesis.

2.6 Connection to Active Inference and Computational Psychiatry

CNA v7's architecture aligns with the active inference framework articulated by Friston (2010, 2023) and recently extended to clinical applications by Smith, Friston, and Whyte (2024) and Krishnadas et al. (2025). Under active inference, the brain is modeled as a hierarchical generative model that maintains beliefs about latent states (in this case, latent diagnostic categories and skill levels) and updates those beliefs through precision-weighted prediction error.

The DDE corresponds to inference over latent diagnostic states. The Cognitive Simulator corresponds to generative-model-based prediction of sensory and behavioral consequences of actions (treatments). The Adaptive Re-routing module corresponds to belief updating when prediction errors accumulate beyond expected variance—the computational signature of model revision under active inference. The five clinical axes correspond to observable consequences at the lowest hierarchical level, while the diagnostic posterior corresponds to beliefs at a higher hierarchical level.

This alignment is not incidental. We argue that the categorical diagnostic systems of DSM-5-TR and ICD-11 implicitly assume a particular generative model—one in which a patient instantiates one and only one disorder, and clinical features are conditionally independent given that disorder. The comorbidity literature, the dimensional turn of RDoC (Cuthbert & Insel, 2013) and HiTOP (Conway et al., 2024), and the present framework all converge on the view that this implicit generative model is empirically inadequate for adult neuropsychiatric populations. A generative model permitting mixed-diagnosis states and dynamic revision of those states under treatment is required.

CNA v7 is one operationalization of such a generative model, specifically engineered for clinician-facing decision support. It does not claim to displace experienced clinical judgment; rather, it makes the implicit Bayesian inference that experienced clinicians perform mentally into an explicit, auditable computational process, with the additional capacity for systematic four-week outcome feedback that exceeds the cognitive capacity of unaided clinical reasoning.

2.7 Architectural Position: Layer 2

A clarifying note on positioning. CNA v7 is designed as a Layer 2 system, situated on top of existing diagnostic infrastructure rather than competing with it. The Layer 1 tools—standardized

symptom instruments (ASRS, GAD-7, PHQ-9), QEEG normative databases (Neuroguide, HBI, Mitsar), and clinician interview-based diagnostic assessment—remain essential. They provide the input features over which CNA v7 performs Bayesian inference. The contribution of CNA v7 is the inferential layer above these inputs: posterior decomposition, counterfactual simulation, and adaptive revision.

This architectural choice has practical implications. CNA v7 can be deployed in clinical settings without requiring replacement of existing QEEG hardware or symptom-scale workflows. A clinic running Neuroguide, BioGraph Infiniti, or a Mitsar 19-channel system can integrate CNA v7 by exporting their existing outputs into the engine's input pipeline. The longitudinal adaptive decision layer constitutes the genuinely novel functional niche that, to our knowledge, no existing clinical tool occupies.

Figure 1 presents the complete four-module architecture as a flow diagram, with the Layer 2 positioning explicitly visualized.




Figure 1. CNA v7 four-module architecture. Layer 1 inputs feed the four computational modules (DDE → Comorbidity Decomposition → Cognitive Simulator → Adaptive Re-routing). The 4-week outcome creates an adaptive feedback loop (marigold arrow).

3.1 System Architecture and Implementation

CNA v7 is implemented as a three-tier web application with strict separation between the clinical interface, the inference engine, and the data persistence layer. The clinical interface (`decision.neurocatchers.com/clinician.html`) is a single-page application built with HTML5 and vanilla JavaScript, providing clinicians with intake forms, real-time inference visualization, five-axis radar charts, and longitudinal outcome tracking. Access to the clinical interface is restricted to authenticated clinical users via a parent admin portal (`neurocatchers.com/admin.html`) with session-based authentication.

The inference engine is implemented in Python 3.11, deployed as a serverless function on the Vercel platform. The engine encapsulates the four computational modules described in Section 2: the Differential Diagnosis Engine, the Comorbidity Decomposition module, the Cognitive Simulator, and the Adaptive Re-routing module. All Bayesian computations use NumPy 1.26 and SciPy 1.13 for numerical stability, with custom diagnostic priors loaded at runtime from a JSON-formatted database (Section 3.2).

The data persistence layer uses Supabase PostgreSQL (Pro tier), with automated nightly backups, Row-Level Security policies isolating clinical data by clinician account, and standardized REST API endpoints. The schema separates intake data (`patients``, `intake_sessions``), measurement data (`qeeq_sessions``, `hrv_sessions``, `cognitive_sessions``, `symptom_sessions``), inference outputs (`diagnostic_posteriors``, `decomposition_results``, `simulator_predictions``), and longitudinal outcomes (`outcome_sessions``, `adaptive_routing_events``, `clinical_notes``). Use of patient-identifying terms in user-facing labels is restricted to "client" in keeping with the scope of practice of the operating clinician (Board Certified in Neurofeedback, PhD), while legacy database column names retain the term "patient" for backward compatibility.

A separate beta enrollment subsystem manages prospective participant recruitment, informed consent collection, and parental consent for minor participants. This subsystem supports four consent pathways: adult in-person signature, adult email-delivered consent, minor with parent present, and minor with parent email consent. Consent signatures are captured as canvas-based images and stored with cryptographic timestamps for audit purposes.

3.2 Diagnostic Fingerprints Database

The diagnostic fingerprints database is a structured JSON document specifying expected feature profiles for each of the twelve candidate diagnoses (Section 2.2). For each diagnosis D_i , the database stores:

- Mean and standard deviation for each of 38 features spanning 8 modalities
- Feature-level base rates and conditional probabilities for symptom presence
- Temporal pattern signatures (onset age, episode duration, recurrence pattern)
- Documented effect sizes for first-line interventions on each axis

The 38 features span eight measurement modalities: (a) attention-domain symptoms (ASRS items 1–18), (b) anxiety-domain symptoms (GAD-7 items, panic frequency, social avoidance ratings), (c) mood-domain symptoms (PHQ-9 items, behavioral activation count, suicidality indicator), (d) trauma/OCD-domain symptoms (PTSD Checklist items, Y-BOCS dimensions), (e) cognitive testing (CPT omission errors, CPT commission errors, N-back accuracy, N-back reaction time), (f) QEEG biomarkers (theta/beta ratio frontal-central, frontal alpha asymmetry, theta cordance frontal), (g) HRV biomarkers (RMSSD, HF power, LF/HF ratio), and (h) temporal patterns (symptom onset age, episode frequency, comorbidity history).

The fingerprints database is treated as a living document. Initial values were populated from published normative data: Faraone et al. (2024) and Cortese (2025) for ADHD; Bandelow, Michaelis, and Wedekind (2017) for anxiety disorders; Hasin et al. (2018) for major depressive disorder; Smoller, Andreassen, et al. (2019) for bipolar disorders. Empirical refinement of these values is anticipated as beta and prospective cohort data accumulate.

Table 2 presents an excerpt of the diagnostic fingerprints database for the five most clinically common diagnoses (ADHD-Combined, GAD, MDD, Bipolar II, PTSD), showing mean $\pm$ SD values for the 38 features.

3.3 The Five-Lens AI Framework

The Five-Lens AI Framework provides multiple complementary perspectives on each patient's clinical profile, addressing the well-documented problem that no single normative comparison captures the full clinical picture (Hess et al., 2025).

Lens 1: General Population Norms. The patient's measurements are z-scored against a healthy general adult population database drawn from published normative studies. This lens answers the question: "How does this patient deviate from the typical adult on each measure?"

Lens 2: Age-Gender Cohort. The patient is compared against a cohort matched on age band (decadal: 20s, 30s, 40s, 50s, 60s+) and gender. This lens controls for the well-documented age and gender effects on QEEG, cognitive, and symptom measures (Cortese, 2025).

Lens 3: Symptom-Profile Cohort. The patient is compared against a cohort of individuals presenting with similar high-level symptom profiles, allowing identification of subtle deviations

within an apparently homogeneous presentation. This lens implements the dimensional approach of HiTOP (Conway et al., 2024) within CNA v7's categorical-Bayesian architecture.

Lens 4: Personal Baseline. Once a patient has at least one prior assessment, subsequent measurements are compared against their own baseline. This lens detects within-individual change, the variable on which treatment response is ultimately judged.

Lens 5: Item-Level Variance. Within a single symptom instrument, the variance pattern across items is examined for signatures of specific diagnoses. For instance, an ASRS profile elevated specifically on items reflecting executive dysfunction (e.g., items 1, 2, 3) versus items reflecting hyperactivity (e.g., items 11, 12) carries diagnostic information beyond the total ASRS score.

Each lens produces an independent posterior contribution. The five posteriors are combined through Bayesian model averaging (Hoeting et al., 1999) weighted by lens-specific reliability estimates.

3.4 The Fingerprint Engine: Z-Score Weighted Cosine Similarity

The fingerprint engine implements the multimodal similarity computation that translates a patient's measurement profile into posterior probabilities over the diagnostic space. Each diagnosis D_i is represented as a 38-dimensional fingerprint vector F_i with components corresponding to expected z-scores on each feature. The patient's measurements are similarly converted to a 38-dimensional observed vector O of z-scores.

The weighted cosine similarity between O and F_i is computed as

Equation 5:

$$\text{sim}(O, F_i) = (\sum_k w_k \times O_k \times F_{i,k}) / (\sqrt{\sum_k w_k \times O_k^2} \times \sqrt{\sum_k w_k \times F_{i,k}^2})$$

where w_k is the reliability weight for feature k , estimated from the test-retest reliability of the underlying instrument. The weights ensure that more reliable features contribute more to the similarity computation.

The similarity scores are converted to posterior probabilities via softmax normalization with temperature parameter τ :

Equation 6:

$$P(D_i | O) = \exp(\text{sim}(O, F_i) / \tau) / \sum_j \exp(\text{sim}(O, F_j) / \tau)$$

The temperature τ is tuned via cross-validation on synthetic data and recalibrated against expert-rated cases in the empirical validation phase. Lower τ produces sharper posteriors (more confident single-diagnosis attribution); higher τ produces flatter posteriors reflecting more equivocal evidence. Default $\tau = 0.5$ yielded interpretable posteriors in pilot testing.

3.5 Beta Study Protocol

A preliminary clinical validation cohort ($N = 5$) has been recruited from the operating clinician's existing client base, consistent with the operating scope of practice (Board Certified in Neurofeedback, PhD in Instructional Design with focus on Mind, Brain, and Education). All participants provided informed consent following local ethical standards, with parent or legal guardian consent additionally obtained for any participant under 18 years of age. The study protocol was reviewed by an institutional ethics committee equivalent and registered internally at Boston Neuromind LLC.

Eligibility criteria included adult presentation with at least one diagnostic question that conventional categorical instruments had been unable to resolve, capacity to complete eight weeks of weekly engagement, and absence of exclusion conditions. Exclusion criteria included acute psychosis, active suicidal crisis, uncontrolled bipolar disorder, recent (within six months) suicide attempt, and current substance dependence.

The protocol comprises three phases. ****Phase A (Intake, 60–90 minutes)**** includes the Entry Protocol 5 battery: time estimation task (approximately 1 minute), heart rate variability recording during paced respiration (approximately 5 minutes), Continuous Performance Test (approximately 8 minutes), N-back task at 1- and 2-back loads (approximately 5 minutes), and Pre-loaded Stem Reading (VSR) acoustic-emotion mapping (approximately 3 minutes). Symptom scales (ASRS, GAD-7, PHQ-9, PCL-5 as indicated) are administered. QEEG is recorded for participants where available; in the present preliminary cohort, QEEG was unavailable for two of five participants.

Phase B (Training Protocol, 4–8 weeks) involves daily app-based engagement with a small task ("Catcher card"), weekly clinician check-ins, and bi-weekly five-axis radar tracking. The specific training protocol (working memory training, anxiety regulation training, behavioral activation, or composite) is selected by the clinician based on the CNA v7 posterior decomposition and recommendation.

Phase C (Outcome Assessment, 60 minutes) repeats the Entry Protocol 5 battery and symptom scales at 4 weeks (mid-protocol) and 8 weeks (final). The 4-week outcome triggers the

Adaptive Re-routing module; the 8-week outcome provides the primary endpoint for the present validation work.

Primary outcome measures include change in the five-axis composite scores, change in standardized symptom scales (clinically meaningful change defined as $\geq 50\%$ reduction from baseline), and Clinical Global Impression-Improvement (CGI-I) rated by the operating clinician.

3.6 Synthetic Validation Approach

In parallel with the preliminary clinical cohort, synthetic validation provides a complementary route to assess CNA v7's inferential properties. Following the synthetic data generation paradigm articulated by Goodfellow et al. (2014) and adapted to psychiatry by recent work (e.g., Wang et al., 2024), we generate populations of synthetic patient profiles by sampling from the diagnostic fingerprint distributions.

For each synthetic patient, a primary diagnosis D_{primary} is sampled with probability proportional to epidemiological base rates. Comorbid diagnoses are sampled with probabilities reflecting documented comorbidity rates (Salazar de Pablo et al., 2025). The patient's 38-feature vector is then generated by mixing samples from the fingerprints of the primary and comorbid diagnoses, weighted by the comorbidity proportions. Measurement noise is added at empirically realistic levels.

A synthetic cohort of $N = 1,000$ patients is generated. CNA v7's DDE is then applied to each synthetic patient. Sensitivity and specificity for primary diagnosis identification are computed against the known ground-truth diagnosis. Receiver Operating Characteristic (ROC) curves are constructed via the standard one-vs-rest approach. Bootstrap confidence intervals (1,000 iterations) provide uncertainty quantification for all reported metrics.

A second synthetic cohort generation paradigm employs multi-agent self-play: pairs of CNA v7 instances generate and evaluate synthetic cases adversarially, with one instance generating challenging mixed-presentation cases and the other attempting accurate decomposition. This paradigm, inspired by AlphaGo's self-play training (Silver et al., 2017), generates synthetic test cases that target the inferential weak points of the current diagnostic fingerprints database. The technique is particularly useful for identifying diagnostic boundary regions where the current fingerprints provide insufficient discrimination.

3.7 Ethical and Regulatory Framework

CNA v7 is positioned as a clinical decision support tool within the scope of practice of Board Certified Neurofeedback (BCN) practice and Instructional Design / Mind, Brain, and Education (MBE) doctoral-level research. The system does not prescribe medication, does not issue DSM diagnoses as primary clinical determinations, and does not replace clinician judgment. All outputs are presented to the clinician as decision-support recommendations to be considered alongside clinical interview, physical examination as indicated, and the clinician's own integrated assessment.

The operating clinician (the first author) holds the BCN credential and the PhD in Instructional Design with concentration in cognitive development under the late Professor Kurt W. Fischer at Harvard Graduate School of Education. The clinician operates under supervisor licensure consistent with Massachusetts statutory requirements at the time of writing. Where psychiatric diagnoses with prescribing implications are required, referral to a licensed psychiatrist is made. The CNA v7 framework is offered as a complement to, not a substitute for, psychiatric medical care.

Data security and patient privacy follow Health Insurance Portability and Accountability Act (HIPAA)-aligned practices: PostgreSQL Row-Level Security at the database level, encrypted backups, restricted access through authenticated clinician portals, and explicit consent for all data collection and use. Business Associate Agreements with cloud service providers (Supabase, Vercel, Resend for email) are obtained prior to admission of any non-beta clinical participant.

Figure 2 presents the complete data flow from patient intake through inference to outcome assessment, including the security boundaries and consent gates at each step.




Figure 2. Data flow with HIPAA-aligned security boundaries. Three layers (client → API/inference → persistence) with TLS 1.3, OAuth 2.0, and Row-Level Security (RLS) at appropriate boundaries.

4.1 Synthetic Validation of the Differential Diagnosis Engine

The Differential Diagnosis Engine was first evaluated on the $N = 1,000$ synthetic patient cohort generated as described in Section 3.6. For each synthetic patient, the engine produced posterior probabilities over the 12 candidate diagnoses, and the maximum-a-posteriori diagnosis was compared to the known ground-truth primary diagnosis.

Table 4 presents sensitivity, specificity, positive predictive value, and area under the receiver operating characteristic curve (AUC) for each of the 12 diagnoses, with bootstrap-derived 95% confidence intervals. Overall, the engine achieved high diagnostic accuracy on the synthetic cohort, with mean sensitivity of 0.84 (95% CI: 0.81–0.87) and mean specificity of 0.93 (95% CI: 0.91–0.95) across diagnoses. Per-diagnosis AUC ranged from 0.79 (Bipolar II, the diagnosis with the most overlapping symptom signature with MDD) to 0.96 (PTSD, the diagnosis with the most distinctive temporal pattern and biomarker signature).

Figure 3 presents ROC curves for the five most clinically common diagnoses (ADHD-Combined, GAD, MDD, Bipolar II, PTSD), illustrating that the engine maintains high specificity at clinically useful sensitivity thresholds across the diagnostic space.

The sensitivity and specificity values reported here reflect performance on the synthetic cohort, which was generated from the same diagnostic fingerprints database that the engine uses for inference. As such, these values represent an upper bound on real-world performance. The values are most appropriately interpreted as a check on internal consistency of the inferential framework rather than as a predictor of out-of-sample clinical performance. Prospective validation on independent clinical samples is required before generalizing these performance estimates to clinical practice.

4.2 Comorbidity Decomposition: Three Illustrative Cases

We illustrate the comorbidity decomposition module through three case examples drawn from the preliminary clinical cohort. All cases are anonymized and clinical details modified to prevent identification while preserving the substantive structure of the decomposition.

Case A presented with chief complaint of "trouble focusing at work" and a primary care referral for evaluation of possible depression. Standardized symptom scales returned PHQ-9 = 12 (moderate depression) and GAD-7 = 8 (mild anxiety). CPT performance revealed elevated omission errors ($z = 1.6$) and elevated commission errors ($z = 1.4$). QEEG was available and showed an elevated theta/beta ratio at frontal-central sites ($z = 1.9$). Time estimation revealed substantial overestimation ($z = 1.7$), and HRV showed reduced high-frequency power ($z = -0.9$).

The CNA v7 attention axis composite for Case A was 0.39. The Comorbidity Decomposition module returned:

$$s_1^{\{(ADHD-C)\}} = 0.25$$

$$s_1^{\{(GAD)\}} = 0.10$$

$$s_1^{\{(MDD)\}} = 0.04$$

with posterior weights $P(ADHD-C) = 0.64$, $P(GAD) = 0.26$, $P(MDD) = 0.10$. The system flagged ADHD-Combined as the primary diagnostic candidate despite the referral for depression, motivated by the CPT, QEEG, and time-estimation pattern. The clinician confirmed the ADHD-C primary diagnosis on follow-up interview and initiated an ADHD-targeted training protocol with concurrent behavioral activation support for the depressive symptoms.

Case B presented with the same chief complaint of "trouble focusing at work" and an attention axis composite of 0.39—numerically identical to Case A. The decomposition returned:

$$s_1^{\{(ADHD-C)\}} = 0.06$$

$$s_1^{\{(GAD)\}} = 0.28$$

$$s_1^{\{(MDD)\}} = 0.03$$

with posterior weights $P(GAD) = 0.71$, $P(ADHD-C) = 0.15$, $P(MDD) = 0.08$. CPT performance in Case B was within normal limits despite the subjective attention complaint. HRV revealed a marked reduction in high-frequency power ($z = -1.8$), and the symptom scale GAD-7 = 14 (moderate-severe anxiety). The temporal pattern of symptom onset was consistent with anxiety-driven attentional disruption rather than developmental ADHD: symptoms emerged in the patient's late twenties in association with workplace responsibility expansion.

The treatment selection consequence was decisive. Case A's training protocol emphasized executive function training, working memory training, and consideration of stimulant evaluation through psychiatric referral. Case B's training protocol emphasized HRV biofeedback, cognitive behavioral therapy for generalized anxiety, and explicit attentional training contingent on prior anxiety regulation. Despite identical attention scores at intake, the two patients received substantively different treatment plans grounded in the decomposition.

Case C, also presenting with attention complaint and attention axis composite of 0.39, decomposed as:

$$s_1^{\{(MDD)\}} = 0.30$$

$$s_1^{\{(ADHD-I)\}} = 0.06$$

$$s_1^{\{(GAD)\}} = 0.03$$

with posterior weights $P(\text{MDD}) = 0.76$, $P(\text{ADHD-I}) = 0.14$, $P(\text{GAD}) = 0.08$. Case C's PHQ-9 = 18 (moderately severe depression), behavioral activation count was 1.2 activities per day (vs. healthy norm 4.6), and QEEG showed elevated frontal theta cordance consistent with antidepressant response prediction. Time estimation was within normal limits. The attention complaint reflected attentional disengagement secondary to depressive psychomotor slowing and anhedonia rather than primary attentional dysfunction.

These three cases illustrate the central claim: an identical attention score of 0.39 can reflect three distinct primary diagnostic processes, each requiring a distinct treatment approach. Conventional categorical diagnosis is silent on which process predominates. CNA v7 makes the decomposition explicit and clinically actionable.

Figure 4 presents the three decompositions side-by-side as stacked bar charts, with the diagnostic fingerprint contributions to each axis displayed for direct comparison.

4.3 Cognitive Simulator Outputs

For Case A, the Cognitive Simulator was queried with the proposed intervention: "Working memory training to produce a $\Delta = +0.20$ improvement in axis 2 (Learning Efficiency) over 4 weeks." The simulator returned predicted changes in the other four axes:

$$\Delta s_1 (\text{Attention}) = +0.12$$

$$\Delta s_2 (\text{Learning Efficiency}) = +0.20 (\text{intervened})$$

$$\Delta s_3 (\text{Peak Performance}) = +0.07$$

$$\Delta s_4 (\text{Anxiety Regulation}) = -0.03$$

$$\Delta s_5 (\text{Mood Stability}) = +0.04$$

The propagation pattern reflected the empirical literature on working memory training effects: substantial positive transfer to attention (consistent with Sala & Gobet, 2017), small positive transfer to peak performance and mood, and a small negative effect on anxiety regulation (reflecting the cognitive demand of the training itself).

For comparison, the simulator was queried with an alternative intervention: "HRV biofeedback to produce $\Delta = +0.20$ improvement in axis 4 (Anxiety Regulation)." The propagation pattern was substantially different:

$$\Delta s_1 (\text{Attention}) = +0.05$$

Δs_2 (Learning Efficiency) = +0.02

Δs_3 (Peak Performance) = +0.14

Δs_4 (Anxiety Regulation) = +0.20 (intervened)

Δs_5 (Mood Stability) = +0.06

The HRV-biofeedback intervention produced larger gains in peak performance and mood through reduced sympathetic arousal, but smaller gains in attention than the working-memory training intervention. For Case A, with attention as the primary clinical concern, the working-memory training intervention was selected as the optimal first-line approach based on the simulator's projection.

Figure 5 presents the two propagation patterns as before-after radar charts with the predicted four-week endpoint overlaid.

4.4 Adaptive Re-routing: Two Cases

Of the five preliminary cohort participants, two completed at least one full Adaptive Re-routing cycle (4-week intermediate outcome assessment). We present these two cases.

Case A re-routing event. At intake, the DDE assigned $P_0(\text{ADHD-C}) = 0.64$ and $P_0(\text{GAD}) = 0.26$. The protocol was working-memory and attention training. At 4 weeks, the observed change pattern was $\Delta s_1 = +0.09$ (predicted +0.12, observed slightly below), $\Delta s_2 = +0.16$ (predicted +0.20, observed below), and notably $\Delta s_4 = -0.11$ (predicted -0.03, observed substantially below predicted).

The Adaptive Re-routing module computed the updated posterior, with the observed greater-than-predicted anxiety dysregulation providing evidence against the pure-ADHD hypothesis and in favor of comorbid GAD as a substantive contributor. $P_4(\text{ADHD-C}) = 0.51$, $P_4(\text{GAD}) = 0.40$. While the maximum-a-posteriori diagnosis remained ADHD-C, the substantial increase in $P(\text{GAD})$ triggered a recommendation to add HRV biofeedback to the existing protocol, addressing the comorbid anxiety component that the initial decomposition had partially recognized but the initial protocol had not directly targeted.

Case D re-routing event. Case D presented with PHQ-9 = 16 and elevated theta cordance, leading to an initial posterior $P_0(\text{MDD}) = 0.72$ and a protocol of behavioral activation training. At 4 weeks, behavioral activation count had increased from 1.4 to 2.8 activities per day (predicted: 3.2). However, attention and learning efficiency scores had also improved more than the depression-focused protocol predicted. The Adaptive Re-routing module identified this pattern as

inconsistent with pure MDD and shifted weight toward MDD-with-ADHD-comorbidity. $P_4(\text{MDD}) = 0.55$, $P_4(\text{ADHD-I}) = 0.27$. The clinician was alerted that the patient may benefit from additional ADHD-targeted assessment, and an ADHD-specific symptom inventory was administered, returning supporting evidence.

These two cases illustrate the adaptive layer's intended function: detection of patterns in outcome data that warrant diagnostic revision, surfaced to the clinician with explicit Bayesian rationale.

Figure 6 presents the two re-routing events as Sankey diagrams showing the posterior shift from intake to 4 weeks.

4.5 Five-Axis Radar Visualization

Figure 7 presents pre- and post-protocol five-axis radar charts for all five preliminary cohort participants. The radar visualization provides clinicians and participants with an at-a-glance view of multidimensional change. Mean change across the cohort was $\Delta s_1 = +0.14$, $\Delta s_2 = +0.12$, $\Delta s_3 = +0.09$, $\Delta s_4 = +0.07$, $\Delta s_5 = +0.10$, reflecting modest improvements across all five axes over the 4–8 week protocol period. Individual variation was substantial, as expected for an $N = 5$ preliminary cohort.

Table 5 presents the preliminary cohort demographics (age range, gender distribution, primary diagnosis as identified by CNA v7), intake five-axis scores, post-protocol five-axis scores, and Clinical Global Impression-Improvement (CGI-I) ratings. Three of five participants received CGI-I ratings of 2 ("much improved") at 8 weeks; one received CGI-I = 3 ("minimally improved"); one (Case D, which experienced the re-routing event) received CGI-I = 2 after the re-routing-driven protocol adjustment but was rated CGI-I = 4 ("no change") at the 4-week pre-adjustment timepoint.

4.6 Summary of Preliminary Findings

The synthetic validation supports the internal consistency of the inferential framework. The three illustrative decomposition cases demonstrate that the system can distinguish substantively different diagnostic processes underlying numerically identical symptom scores. The two re-routing events demonstrate that the adaptive layer can detect outcome patterns inconsistent with the initial diagnostic hypothesis and surface revision recommendations.

These results are explicitly preliminary. The preliminary cohort is small ($N = 5$), the synthetic validation circular in the sense that the same fingerprints database generates both the test cases

and the inference engine's expectations, and the case studies are illustrative rather than confirmatory. The intended trajectory is to use these preliminary results to motivate prospective multi-site validation with independent clinical samples. Sections 5.5 and 5.6 discuss limitations and future directions in detail.

Figure 3. ROC Curves for Five Common Diagnoses (DDE Synthetic Validation, N=1,000)




Figure 3. Receiver Operating Characteristic (ROC) curves for the five most common diagnoses on the synthetic validation cohort (N = 1,000). PTSD shows the highest AUC (0.96) reflecting its distinctive temporal pattern; Bipolar II shows the lowest AUC among displayed diagnoses (0.79; not pictured) due to overlap with MDD.


Figure 4. Comorbidity Decomposition: Three Cases with Identical Attention Score

Figure 4. Comorbidity decomposition of an identical attention axis score ($s_1 = 0.39$) into three substantively different diagnostic profiles. Case A: ADHD-primary. Case B: GAD-primary. Case C: MDD-primary. Treatment selection differs accordingly.

Figure 5. Cognitive Simulator: Counterfactual Predictions for Case A




Figure 5. Cognitive Simulator counterfactual predictions for Case A under two alternative interventions. Panel A: working memory training ($\Delta s_2 = +0.20$) propagates substantial gains to attention but a small cost to anxiety. Panel B: HRV biofeedback ($\Delta s_4 = +0.20$) produces gains in peak performance and mood but smaller transfer to attention.

Figure 6. Adaptive Re-routing: Posterior Shifts at 4-Week Outcome Assessment




Figure 6. Adaptive re-routing posterior shifts at the 4-week outcome assessment. Panel A: Case A — increased GAD weight after observed anxiety dysregulation. Panel B: Case D — increased ADHD-I weight after attention/learning gains beyond MDD-only prediction.

Figure 7. Pre- and Post-Protocol 5-Axis Radar Charts (Preliminary Cohort, N=5)




Figure 7. Pre- and post-protocol 5-axis radar charts for all five preliminary cohort participants. Light blue: pre-protocol baseline. Dark blue: post-protocol (8 weeks).

5.1 Principal Findings

The present work introduces and provides preliminary evidence for the Catcher Navigator Algorithm version 7 (CNA v7), a four-module Bayesian framework for adult neuropsychiatric assessment. Three findings warrant emphasis.

First, the comorbidity decomposition module resolves a clinical phenomenon that conventional categorical instruments cannot address: the systematic ambiguity of identical clinical scores arising from divergent underlying processes. The three case examples in Section 4.2, in which an attention axis composite of exactly 0.39 decomposed into ADHD-primary, GAD-primary, and MDD-primary patterns respectively, illustrate the practical significance of this resolution for treatment selection. The clinician confronted with the raw 0.39 score alone is constrained to clinical interview-based judgment with whatever supplementary information that interview elicits; the clinician equipped with the decomposition has a Bayesian-quantified rationale for differentiated treatment selection.

Second, the adaptive re-routing module provides a structured mechanism for diagnostic revision in response to four-week outcome evidence. The two re-routing events documented in Section 4.4—Case A's posterior shift toward greater GAD comorbidity, and Case D's posterior shift toward MDD-with-ADHD comorbidity—illustrate the layer's intended function. Critically, this revision occurs through an explicit Bayesian computation with auditable rationale: the clinician sees which predictions matched, which diverged, and how the posterior shifted as a result. This contrasts with the implicit and undocumented form of diagnostic revision that occurs in conventional practice, where treatment non-response may or may not prompt the clinician to reconsider the initial diagnosis.

Third, the five-axis architecture provides a measurement framework that captures multidimensional change in a form interpretable both to clinicians and to participants themselves. The radar visualization (Figure 7) renders abstract psychometric change into a perceptually salient representation. Pilot participants in the preliminary cohort consistently reported that the radar visualization clarified their own understanding of what was changing and what was not over the course of the protocol.

5.2 Comparison to Existing Frameworks

CNA v7 occupies a specific functional niche distinct from existing diagnostic and decision-support tools. **Table 6** presents a structured comparison across five reference frameworks: QEEG normative databases (Neuroguide, Thatcher et al., 2003), the Human Brain Index (HBI;

Kropotov, 2009), the Research Domain Criteria (RDoC; Cuthbert & Insel, 2013), the Hierarchical Taxonomy of Psychopathology (HiTOP; Conway et al., 2024), and CNA v7.

The QEEG normative databases provide rigorous snapshot comparison of an individual's neurophysiological profile to age-matched population norms but offer no longitudinal adaptive layer and no formal mechanism for comorbidity decomposition. The HBI system extends this snapshot approach into the ERP domain with comparable advantages and limitations. RDoC reorganizes psychopathology dimensionally but remains a research framework rather than a clinician-facing decision support tool; its operationalization for routine clinical practice is still emerging (Cuthbert, 2022). HiTOP offers a hierarchical dimensional taxonomy with strong empirical support (Conway et al., 2024) but, like RDoC, has not yet been operationalized as a clinical decision-support system with adaptive treatment-response feedback.

CNA v7 differs from these frameworks principally in its Layer 2 positioning and its longitudinal adaptive architecture. It does not aspire to replace any of the Layer 1 instruments; rather, it ingests them as inputs and adds the Bayesian inferential layer that those instruments do not natively provide. The combination of comorbidity decomposition, counterfactual simulation, and adaptive re-routing constitutes, to our knowledge, a genuinely novel functional niche.

The closest existing work in spirit is the Bayesian Workflow for Generative Modeling in Computational Psychiatry articulated by Hess et al. (2025) at the ETH Zurich Translational Neuromodeling Unit. The Hess et al. framework uses Hierarchical Gaussian Filtering to extract individual-level computational parameters from behavioral data. CNA v7 shares the Bayesian commitment of that work but differs in clinical orientation: Hess et al. target computational psychiatry as a research enterprise generating mechanistic insight; CNA v7 targets clinician-facing decision support with directly actionable treatment recommendations. The two enterprises are complementary rather than competing.

5.3 Clinical Implications

If the preliminary findings generalize through prospective validation, three clinical implications follow.

First, the time-to-correct-diagnosis problem identified by Kessler et al. (2023) for adult ADHD—an average 6.8-year delay—may be substantially reducible through systematic four-week outcome-driven posterior revision. A patient initially treated for the wrong primary diagnosis would be flagged at the four-week mark when treatment response deviates from prediction, prompting diagnostic reconsideration. The diagnostic delay would compress from years to weeks.

Second, the category of "treatment-resistant" presentations—particularly treatment-resistant depression, which is well-documented to harbor substantial unrecognized ADHD (McIntosh et al., 2009)—may be partially re-classified as cases of initial misidentification of primary diagnosis. The decomposition module makes the comorbid contribution to the clinical picture explicit at intake, before the patient has accumulated weeks or months of inadequate response to first-line treatment.

Third, personalized protocol selection becomes grounded in posterior weights rather than categorical thresholds. The clinician can transparently explain to the patient: "Your attention difficulty appears to be driven 64% by ADHD-like patterns and 26% by anxiety-driven attentional disruption. The recommended protocol addresses both, weighted toward the larger contributor." This explanation is both more accurate and more therapeutically engaging than the binary categorical communication.

5.4 Theoretical Implications

CNA v7 contributes to three theoretical conversations.

The first concerns the application of Active Inference (Friston, 2010, 2023) to longitudinal clinical practice. Most active inference work in psychiatry to date has addressed cross-sectional questions: what computational signature characterizes a given disorder, what model parameters distinguish patient groups. CNA v7 demonstrates that the active inference framework extends naturally to longitudinal clinical reasoning, with the diagnostic posterior playing the role of a hierarchically-superordinate belief that is updated through precision-weighted prediction errors arising from four-week outcome observations.

The second concerns the operationalization of Dynamic Skill Theory (Fischer, 1980; Fischer & Bidell, 2006) in adult neuropsychiatry. Fischer's framework, developed primarily in the context of cognitive development, has had less penetration into clinical adult assessment than its empirical and theoretical robustness warrants. The five-axis CNA v7 architecture treats each clinical axis as a dynamic skill in Fischer's sense—a capacity that emerges through interaction with environmental and physiological context, that can be measured at its current level, and that exhibits a trajectory under intervention. This operationalization, the present author hopes, would have pleased Professor Fischer, with whom this work began conceptually during Harvard Graduate School of Education research from 2018 to 2020.

The third concerns the bridge between computational psychiatry as a research enterprise and clinical practice. Browning et al. (2020), in the Banbury Center meeting report, identified the persistent gap between computational psychiatric research and clinical translation as the field's

central unmet challenge. CNA v7 is one attempt to bridge that gap by deploying computational inference in a clinician-facing decision-support tool, with the explicit constraint of remaining within the scope of practice of non-prescribing clinical professionals.

5.5 Limitations

Several limitations bound the interpretation of the present findings.

The preliminary clinical cohort is small ($N = 5$), drawn from the operating clinician's existing client base, and not blinded. The case examples in Section 4 are illustrative rather than confirmatory. Prospective validation in larger, independent samples by clinicians without conflict of interest with respect to the framework is required before clinical generalization. The intended next step is a multi-site prospective trial with $N = 200$ target enrollment.

The synthetic validation in Section 4.1 is internally consistent but circular: the same diagnostic fingerprints database that generates the synthetic patient population also generates the engine's expectations. Performance metrics derived from this validation overestimate real-world clinical performance. The synthetic validation should be understood as a sanity check on inferential mechanics rather than as evidence of clinical accuracy.

The diagnostic fingerprints database is itself a hypothesis under refinement. Initial values were drawn from published normative literature, but the specific feature-by-diagnosis expected values and variances will require empirical refinement as data accumulate. The framework is designed to support this iterative refinement, but the current fingerprints may misspecify certain diagnostic boundaries, particularly for less-studied conditions and for atypical presentations.

The Cognitive Simulator's causal graph encodes empirical effect sizes drawn from heterogeneous literature, some of which is contested (e.g., the working memory training transfer literature, Sala & Gobet, 2017 vs. Melby-Lervåg et al., 2016). Where the literature is divided, the simulator's predictions inherit that uncertainty.

The operating scope of the framework is restricted to non-prescribing clinical practice. CNA v7 does not provide medication recommendations and does not aspire to replace psychiatric medical care. For patients whose treatment plan involves pharmacotherapy, psychiatric referral is required and the CNA v7 outputs serve as one input among many to the psychiatrist's clinical reasoning.

Finally, the data architecture relies on cloud service providers (Supabase, Vercel, Resend). Business Associate Agreements with these providers are required before admission of non-beta clinical participants. Self-hosted infrastructure may be required for some clinical deployment contexts.

5.6 Future Directions

Five extensions are anticipated.

A prospective multi-site validation trial with $N = 200$ target enrollment will provide the central empirical test of the framework's clinical utility. This trial will involve blinded clinician raters at multiple sites, pre-registered primary outcomes, and an active-comparator design comparing CNA v7-guided treatment selection to treatment as usual.

Integration with neurofeedback adaptive protocols. The operating clinician is Board Certified in Neurofeedback, and the longer-term aim of CNA v7 is integration with neurofeedback training protocol selection, where the framework's posterior decomposition would directly inform protocol parameters (frequency bands, electrode sites, training thresholds) for individual patients.

Pediatric extension. The current framework is restricted to adult populations. Adaptation to pediatric ADHD, anxiety, and mood presentations would require both fingerprint database extension for developmental presentations and methodological adaptation for age-appropriate measurement instruments.

USPTO patent portfolio. Four provisional patent applications have been filed covering the Differential Diagnosis Engine, the Comorbidity Decomposition module, the Multi-Source Diagnostic Algorithm bridging behavioral and biological data, and the Adaptive Re-routing layer. The intellectual property strategy aims to preserve open clinical use while protecting commercial replication.

Open framework for collaborative refinement. The present author intends to release the diagnostic fingerprints database schema, the synthetic validation dataset, and the inference engine code under permissive license on the project GitHub repository, enabling other research groups to contribute fingerprint refinements, alternative inferential modules, and validation in additional clinical populations.

5.7 Conclusion

CNA v7 reframes adult neuropsychiatric assessment from a categorical thresholding problem to a continuous Bayesian inference problem with explicit comorbidity decomposition, counterfactual simulation, and adaptive longitudinal revision. The framework occupies a distinct Layer 2 functional niche above existing QEEG normative databases and symptom-scale instruments, ingesting their outputs as inputs to its Bayesian engine.

Preliminary clinical evidence ($N = 5$) and synthetic validation ($N = 1,000$) support the internal consistency of the inferential framework. The three illustrative decomposition cases demonstrate

substantive treatment-selection implications of comorbidity-aware diagnosis. The two adaptive re-routing events demonstrate the layer's capacity to surface diagnostic revision recommendations grounded in four-week outcome evidence.

Prospective multi-site validation is the necessary next step. The framework is offered as a contribution to the longer-running enterprise of bringing computational psychiatry into routine clinical practice, complementing rather than replacing existing tools and clinician judgment. The late Professor Kurt W. Fischer's commitment to dynamic skill measurement, with whom this work began conceptually at Harvard Graduate School of Education, animates the framework's commitment to longitudinal observation of human capacities under intervention.

Adult neuropsychiatric care has long been hampered by categorical thinking, single-snapshot assessment, and the resulting diagnostic delays and treatment-resistant presentations that follow from initial misidentification of primary diagnosis. CNA v7 offers one structured response to that long-standing problem. The empirical work to follow will determine whether the response holds up under prospective scrutiny.

End of Section 5.

Tables

Table 1. Five-Axis Clinical Architecture

Axis	Domain	Symptom Instrument	Cognitive Test	Biomarker
S ₁	Attention	ASRS v1.1 (18 items)	CPT-3 (omission/commission)	QEEG θ/β ratio (Fz–Cz)
S ₂	Learning Efficiency	—	N-back, WAIS-IV Digit Span	P3b ERP (Pz)
S ₃	Peak Performance	Flow State Scale	Time estimation	HRV HF/LF coherence
S ₄	Anxiety Regulation	GAD-7	Stroop interference	HRV RMSSD, frontal α asym.
S ₅	Mood Stability	PHQ-9, BA count	Affective Go/No-Go	Theta cordance (F3–F4)

Note. ASRS = Adult ADHD Self-Report Scale; CPT-3 = Conners Continuous Performance Test 3rd ed.; HRV = heart rate variability; RMSSD = root mean square of successive differences; QEEG = quantitative electroencephalography; BA = behavioral activation.

Table 4. DDE Performance on Synthetic Validation Cohort (N = 1,000)

Diagnosis	Sensitivity	Specificity	PPV	AUC (95% CI)
ADHD-Inattentive	0.82	0.94	0.79	0.91 (0.88–0.94)
ADHD-Hyperactive	0.79	0.95	0.76	0.89 (0.86–0.92)
ADHD-Combined	0.87	0.93	0.81	0.92 (0.89–0.95)
GAD	0.85	0.92	0.78	0.91 (0.88–0.94)
Panic Disorder	0.81	0.96	0.80	0.90 (0.87–0.93)
Social Anxiety	0.83	0.94	0.77	0.90 (0.87–0.93)
OCD	0.86	0.95	0.81	0.92 (0.89–0.95)
PTSD	0.91	0.96	0.85	0.96 (0.93–0.98)

MDD	0.84	0.91	0.74	0.91 (0.88–0.94)
Persistent Depressive	0.78	0.93	0.71	0.87 (0.84–0.90)
Bipolar I	0.83	0.97	0.82	0.92 (0.89–0.95)
Bipolar II	0.76	0.93	0.69	0.79 (0.75–0.83)
Overall (macro)	0.84	0.93	0.78	0.90 (0.87–0.93)

Note. PPV = positive predictive value; AUC = area under the receiver operating characteristic curve; 95% CI = bootstrap confidence interval (1,000 iterations).

Table 5. Preliminary Beta Cohort (N = 5)

Case	Age	Sex	Referral Concern	CNA v7 Primary Dx	4-wk Re-route?	8-wk CGI-I
A	34	F	"Trouble focusing at work"	ADHD-C (0.64)	Yes — add HRV	2
B	41	M	"Trouble focusing at work"	GAD (0.71)	No	2
C	29	F	"Brain fog, low energy"	MDD (0.76)	No	3
D	47	M	"Low motivation, concentration"	MDD → MDD + ADHD-I	Yes — add EF training	2
E	38	F	"Anxiety, insomnia"	GAD (0.68)	No	2

Note. CGI-I = Clinical Global Impression-Improvement (1 = very much improved, 7 = very much worse).

Table 6. CNA v7 vs Existing Frameworks

Feature	Neuroguide	HBI	RDoC	HiTOP	CNA v7
Comorbidity decomposition	No	No	Conceptual	Conceptual	Yes (quantitative)
Counterfactual simulation	No	No	No	No	Yes
Longitudinal adaptive layer	No	No	No	No	Yes (4-wk)
Clinician-facing interface	Yes	Yes	No	No	Yes

Architectural layer	L1	L1	Conceptual	Conceptual	L2
---------------------	----	----	------------	------------	----

References

- Aldao, A., Nolen-Hoeksema, S., & Schweizer, S. (2010). Emotion-regulation strategies across psychopathology: A meta-analytic review. *Clinical Psychology Review*, *30*(2), 217–237. <https://doi.org/10.1016/j.cpr.2009.11.004>
- American Psychiatric Association. (2022). *Diagnostic and statistical manual of mental disorders* (5th ed., text rev.). <https://doi.org/10.1176/appi.books.9780890425787>
- Arns, M., Conners, C. K., & Kraemer, H. C. (2013). A decade of EEG theta/beta ratio research in ADHD: A meta-analysis. *Journal of Attention Disorders*, *17*(5), 374–383. <https://doi.org/10.1177/1087054712460087>
- Arns, M., de Ridder, S., Strehl, U., Breteler, M., & Coenen, A. (2009). Efficacy of neurofeedback treatment in ADHD: The effects on inattention, impulsivity and hyperactivity—A meta-analysis. *Clinical EEG and Neuroscience*, *40*(3), 180–189. <https://doi.org/10.1177/155005940904000311>
- Aschbacher, K., Mather, M., Lehrer, P., Gevirtz, R., Epel, E., & Peiper, N. C. (2024). Real-time heart rate variability biofeedback amplitude during a large-scale digital mental health intervention differed by age, gender, and mental and physical health. *Psychophysiology*, *61*(7), e14533. <https://doi.org/10.1111/psyp.14533>
- Bandelow, B., Michaelis, S., & Wedekind, D. (2017). Treatment of anxiety disorders. *Dialogues in Clinical Neuroscience*, *19*(2), 93–107. <https://doi.org/10.31887/DCNS.2017.19.2/bbandelow>
- Barrett, L. F., Quigley, K. S., & Hamilton, P. (2016). An active inference theory of allostasis and interoception in depression. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *371*(1708), 20160011. <https://doi.org/10.1098/rstb.2016.0011>
- Birk, J. L., Cumella, R., Lopez-Veneros, D., Agarwal, S., & Kronish, I. M. (2024). Feasibility of a remote heart rate variability biofeedback intervention for reducing anxiety in cardiac arrest survivors: A pilot trial. *Contemporary Clinical Trials Communications*, *37*, 101251.
- Bottemanne, H., Mouchabac, S., & Gauld, C. (2025). Reshaping computational neuropsychiatry beyond synaptopathy. *Brain*, *148*(5), 1526–1530.
- Browning, M., Carter, C. S., Chatham, C., Den Ouden, H., Gillan, C. M., Baker, J. T., Chekroud, A. M., Cools, R., Dayan, P., Gold, J., Goldstein, R. Z., Hartley, C. A., Kepecs, A., Krystal, J. H., Lee, D., Levine, S., London, E. D., Murphy, S. E., Nahum, M., ... Paulus, M. P.

- (2020). Realizing the clinical potential of computational psychiatry: Report from the Banbury Center meeting, February 2019. **Biological Psychiatry**, **88**(2), e5–e10. <https://doi.org/10.1016/j.biopsych.2019.12.026>
- Bussalib, A., Congedo, M., Barthélemy, Q., Ojeda, D., Acquaviva, E., Delorme, R., & Mayaud, L. (2019). Clinical and experimental factors influencing the efficacy of neurofeedback in ADHD: A meta-analysis. **Frontiers in Psychiatry**, **10**, 35. <https://doi.org/10.3389/fpsyt.2019.00035>
- Carcone, D., & Ruocco, A. C. (2017). Six years of research on the National Institute of Mental Health's Research Domain Criteria (RDoC) initiative: A systematic review. **Frontiers in Cellular Neuroscience**, **11**, 46. <https://doi.org/10.3389/fncel.2017.00046>
- Choi, W.-S., Woo, Y. S., Wang, S.-M., Lim, H. K., & Bahk, W.-M. (2022). The prevalence of psychiatric comorbidities in adult ADHD compared with non-ADHD populations: A systematic literature review. **PLOS ONE**, **17**(11), e0277175. <https://doi.org/10.1371/journal.pone.0277175>
- Clark, J. E., Watson, S., & Friston, K. J. (2018). What is mood? A computational perspective. **Psychological Medicine**, **48**(14), 2277–2284. <https://doi.org/10.1017/S0033291718000430>
- Coan, J. A., & Allen, J. J. B. (2004). Frontal EEG asymmetry as a moderator and mediator of emotion. **Biological Psychology**, **67**(1–2), 7–49. <https://doi.org/10.1016/j.biopsycho.2004.03.002>
- Cook, I. A., Leuchter, A. F., Morgan, M., Witte, E., Stubbeman, W. F., Abrams, M., Rosenberg, S., & Uijtdehaage, S. H. J. (2002). Early changes in prefrontal activity characterize clinical responders to antidepressants. **Neuropsychopharmacology**, **27**(1), 120–131. [https://doi.org/10.1016/S0893-133X\(02\)00294-4](https://doi.org/10.1016/S0893-133X(02)00294-4)
- Conway, C. C., Forbes, M. K., & South, S. C. (2024). The Hierarchical Taxonomy of Psychopathology (HiTOP): Current state of the science and a research agenda for the next decade. **Annual Review of Clinical Psychology**, **20**, 161–186.
- Cortese, S. (2025). Attention-deficit/hyperactivity disorder (ADHD) in adults: Evidence base, uncertainties and controversies. **World Psychiatry**, **24**(3), 366–391. <https://doi.org/10.1002/wps.21374>
- Cuthbert, B. N. (2014). The RDoC framework: Facilitating transition from ICD/DSM to dimensional approaches that integrate neuroscience and psychopathology. **World Psychiatry**, **13**(1), 28–35. <https://doi.org/10.1002/wps.20087>

- Cuthbert, B. N. (2020). The role of RDoC in future classification of mental disorders. **Dialogues in Clinical Neuroscience**, **22**(1), 81–85.
<https://doi.org/10.31887/DCNS.2020.22.1/bcuthbert>
- Cuthbert, B. N. (2022). Research Domain Criteria (RDoC): Progress and potential. **Current Directions in Psychological Science**, **31**(2), 107–114.
- Cuthbert, B. N., & Insel, T. R. (2013). Toward the future of psychiatric diagnosis: The seven pillars of RDoC. **BMC Medicine**, **11**, 126. <https://doi.org/10.1186/1741-7015-11-126>
- Diamond, A. (2013). Executive functions. **Annual Review of Psychology**, **64**, 135–168.
<https://doi.org/10.1146/annurev-psych-113011-143750>
- Donnelly, D., Georgiadis, E., & Stavrou, N. (2023). A meta-analysis investigating the outcomes and correlation between heart rate variability biofeedback training on depressive symptoms and heart rate variability outcomes versus standard treatment in comorbid adult populations. **Acta Biomedica Atenei Parmensis**, **94**(4), e2023214.
<https://doi.org/10.23750/abm.v94i4.14305>
- Faraone, S. V., Bellgrove, M. A., Brikell, I., Cortese, S., Hartman, C. A., Hollis, C., Newcorn, J. H., Philipsen, A., Polanczyk, G. V., Rubia, K., Sibley, M. H., & Buitelaar, J. K. (2024). Attention-deficit/hyperactivity disorder. **Nature Reviews Disease Primers**, **10**, 11.
<https://doi.org/10.1038/s41572-024-00495-z>
- Farías Venegas, C., Varela Espinoza, L. F., Ramírez Matta, C., Barra Cañas, B., Araneda Maldonado, M., & Sánchez Oñate, A. (2025). New developments in psychiatric classification: A transdiagnostic approach. **Cureus**, **17**(5), e84580.
<https://doi.org/10.7759/cureus.84580>
- 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** (6th ed., Vol. 1, pp. 313–399). Wiley.
- Fournié, C., Chouchou, F., Dalleau, G., Caderby, T., Cabrera, Q., & Verkindt, C. (2021). Heart rate variability biofeedback in chronic disease management: A systematic review. **Complementary Therapies in Medicine**, **60**, 102750.

- 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. (2023). Computational psychiatry: From synapses to sentience. *Molecular Psychiatry*, *28*(1), 256–268. <https://doi.org/10.1038/s41380-022-01743-z>
- Friston, K. J., Stephan, K. E., Montague, R., & Dolan, R. J. (2014). Computational psychiatry: The brain as a phantastic organ. *The Lancet Psychiatry*, *1*(2), 148–158. [https://doi.org/10.1016/S2215-0366\(14\)70275-5](https://doi.org/10.1016/S2215-0366(14)70275-5)
- Geladé, K., Janssen, T. W. P., Bink, M., Twisk, J. W. R., van Mourik, R., Maras, A., & Oosterlaan, J. (2018). A 6-month follow-up of an RCT on behavioral and neurocognitive effects of neurofeedback in children with ADHD. *European Child & Adolescent Psychiatry*, *27*(5), 581–593.
- Goessl, V. C., Curtiss, J. E., & Hofmann, S. G. (2017). The effect of heart rate variability biofeedback training on stress and anxiety: A meta-analysis. *Psychological Medicine*, *47*(15), 2578–2586.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. *Advances in Neural Information Processing Systems*, *27*, 2672–2680.
- Grant, B. F., Goldstein, R. B., Saha, T. D., Chou, S. P., Jung, J., Zhang, H., Pickering, R. P., Ruan, W. J., Smith, S. M., Huang, B., & Hasin, D. S. (2015). Epidemiology of DSM-5 alcohol use disorder: Results from the National Epidemiologic Survey on Alcohol and Related Conditions III. *JAMA Psychiatry*, *72*(8), 757–766. <https://doi.org/10.1001/jamapsychiatry.2015.0584>
- Hasin, D. S., Sarvet, A. L., Meyers, J. L., Saha, T. D., Ruan, W. J., Stohl, M., & Grant, B. F. (2018). Epidemiology of adult DSM-5 major depressive disorder and its specifiers in the United States. *JAMA Psychiatry*, *75*(4), 336–346. <https://doi.org/10.1001/jamapsychiatry.2017.4602>
- Hess, A. J., Iglesias, S., Köchli, L., Marino, S., Müller-Schrader, M., Rigoux, L., Mathys, C., Harrison, O. K., Heinzle, J., Frässle, S., & Stephan, K. E. (2025). Bayesian workflow for generative modeling in computational psychiatry. *Computational Psychiatry*, *9*(1), 76–99. <https://doi.org/10.5334/cpsy.116>
- Hoeting, J. A., Madigan, D., Raftery, A. E., & Volinsky, C. T. (1999). Bayesian model averaging: A tutorial. *Statistical Science*, *14*(4), 382–417.

- Huys, Q. J. M., Maia, T. V., & Frank, M. J. (2016). Computational psychiatry as a bridge from neuroscience to clinical applications. *Nature Neuroscience*, *19*(3), 404–413. <https://doi.org/10.1038/nn.4238>
- Kessler, R. C., Berglund, P., Demler, O., Jin, R., Merikangas, K. R., & Walters, E. E. (2005). Lifetime prevalence and age-of-onset distributions of DSM-IV disorders in the National Comorbidity Survey Replication. *Archives of General Psychiatry*, *62*(6), 593–602. <https://doi.org/10.1001/archpsyc.62.6.593>
- Kessler, R. C., Adler, L. A., Berglund, P., Green, J. G., McLaughlin, K. A., Fayyad, J., Russo, L. J., Sampson, N. A., Shahly, V., & Zaslavsky, A. M. (2023). The effects of temporally secondary co-morbid mental disorders on the associations of DSM-IV ADHD with adverse outcomes in adulthood. *Psychological Medicine*, *53*(7), 2873–2885.
- Kirchner, L., Eckert, A.-L., Berg, M., Endres, D., Straube, B., & Rief, W. (2024). An active inference approach to interpersonal differences in depression. *New Ideas in Psychology*, *74*, 101092. <https://doi.org/10.1016/j.newideapsych.2024.101092>
- Kirchner, L., Schummers, J., & Rief, W. (2025). Disorder at the synapse: How the active inference framework unifies competing perspectives on depression. *Entropy*, *27*(9), 970.
- Kotov, R., Krueger, R. F., Watson, D., Achenbach, T. M., Althoff, R. R., Bagby, R. M., Brown, T. A., Carpenter, W. T., Caspi, A., Clark, L. A., Eaton, N. R., Forbes, M. K., Forbush, K. T., Goldberg, D., Hasin, D., Hyman, S. E., Ivanova, M. Y., Lynam, D. R., Markon, K., ... Zimmerman, M. (2017). The Hierarchical Taxonomy of Psychopathology (HiTOP): A dimensional alternative to traditional nosologies. *Journal of Abnormal Psychology*, *126*(4), 454–477. <https://doi.org/10.1037/abn0000258>
- Krishnadas, R., Friston, K. J., & Hutchinson, P. J. (2025). Active inference applied to clinical neuropsychiatry: A computational framework for individualized treatment selection. *Translational Psychiatry*, *15*, in press.
- Kropotov, J. D. (2009). *Quantitative EEG, event-related potentials and neurotherapy*. Academic Press.
- Lehrer, P. M., Vaschillo, E., & Vaschillo, B. (2000). Resonant frequency biofeedback training to increase cardiac variability: Rationale and manual for training. *Applied Psychophysiology and Biofeedback*, *25*(3), 177–191.
- Lehrer, P., Kaur, K., Sharma, A., Shah, K., Huseby, R., Bhavsar, J., Sgobba, P., Zhang, Y., & Lehrer, P. (2020). Heart rate variability biofeedback improves emotional and physical

- health and performance: A systematic review and meta-analysis. **Applied Psychophysiology and Biofeedback**, *45*(3), 109–129.
- McIntosh, D., Kutcher, S., Binder, C., Levitt, A., Fallu, A., & Rosenbluth, M. (2009). Adult ADHD and comorbid depression: A consensus-derived diagnostic algorithm for ADHD. **Neuropsychiatric Disease and Treatment**, *5*, 137–150.
- Melby-Lervåg, M., Redick, T. S., & Hulme, C. (2016). Working memory training does not improve performance on measures of intelligence or other measures of "far transfer": Evidence from a meta-analytic review. **Perspectives on Psychological Science**, *11*(4), 512–534. <https://doi.org/10.1177/1745691616635612>
- Mittal, V. A., & Wakschlag, L. S. (2017). Research domain criteria (RDoC) grows up: Strengthening neurodevelopment investigation within the RDoC framework. **Journal of Affective Disorders**, *216*, 30–35.
- Montague, P. R., Dolan, R. J., Friston, K. J., & Dayan, P. (2012). Computational psychiatry. **Trends in Cognitive Sciences**, *16*(1), 72–80. <https://doi.org/10.1016/j.tics.2011.11.018>
- Moradi, N., Rajabi, S., & Mansouri Nejad, A. (2024). The effect of neurofeedback training combined with computer cognitive games on the time perception, attention, and working memory in children with ADHD. **Applied Neuropsychology: Child**, *13*(1), 24–36.
- Moreno-García, I., Meneres-Sancho, S., Camacho-Vara de Rey, C., & Servera, M. (2019). A randomized controlled trial to examine the posttreatment efficacy of neurofeedback, behavior therapy, and pharmacology on ADHD measures. **Journal of Attention Disorders**, *23*(4), 374–383.
- Morris, S. E., Sanislow, C. A., Pacheco, J., Vaidyanathan, U., Gordon, J. A., & Cuthbert, B. N. (2022). Revisiting the seven pillars of RDoC. **BMC Medicine**, *20*(1), 220. <https://doi.org/10.1186/s12916-022-02414-0>
- Pearl, J. (2009). **Causality: Models, reasoning, and inference** (2nd ed.). Cambridge University Press.
- Pearl, J., & Mackenzie, D. (2018). **The book of why: The new science of cause and effect**. Basic Books.
- Petzschner, F. H., Weber, L. A., Gard, T., & Stephan, K. E. (2017). Computational psychosomatics and computational psychiatry: Toward a joint framework for differential

- diagnosis. **Biological Psychiatry**, **82**(6), 421–430.
<https://doi.org/10.1016/j.biopsych.2017.05.012>
- Pizzoli, S. F. M., Marzorati, C., Gatti, D., Monzani, D., Mazzocco, K., & Pravettoni, G. (2021). A meta-analysis on heart rate variability biofeedback and depressive symptoms. **Scientific Reports**, **11**(1), 6650. <https://doi.org/10.1038/s41598-021-86149-7>
- Polich, J. (2007). Updating P300: An integrative theory of P3a and P3b. **Clinical Neurophysiology**, **118**(10), 2128–2148. <https://doi.org/10.1016/j.clinph.2007.04.019>
- Sala, G., & Gobet, F. (2017). Working memory training in typically developing children: A meta-analysis of the available evidence. **Developmental Psychology**, **53**(4), 671–685.
<https://doi.org/10.1037/dev0000265>
- Salazar de Pablo, G., Iniesta, R., Bellato, A., Caye, A., Dobrosavljevic, M., Parlatini, V., Garcia-Argibay, M., Li, L., Cabras, A., Haider Ali, M., Archer, L., Meehan, A. J., Suleiman, H., Solmi, M., Fusar-Poli, P., Chang, Z., Larsson, H., Faraone, S. V., & Cortese, S. (2025). Prevalence of attention deficit hyperactivity disorder/hyperkinetic disorder of pediatric and adult populations in clinical settings: A systematic review, meta-analysis and meta-regression. **Molecular Psychiatry**, **30**, 1–15. <https://doi.org/10.1038/s41380-025-03178-8>
- Sanislow, C. (2025). The research domain criteria (RDoC) project as a biological framework for research on mental disorders. **European Psychiatry**, **68**(S1), S46–S47.
<https://doi.org/10.1192/j.eurpsy.2025.69>
- Santamaría-García, H., Migeot, J., Medel, V., Hazelton, J. L., Teckentrup, V., Romero-Ortuno, R., Piguet, O., Lawor, B., Northoff, G., & Ibanez, A. (2025). Interoception in mental health: Toward a computational psychiatry framework. **Biological Psychiatry**, **97**(1), 28–40. <https://doi.org/10.1016/j.biopsych.2024.06.024>
- Shereena, E. A., Gupta, R. K., Bennett, C. N., Sagar, K. J. V., & Rajeswaran, J. (2019). EEG neurofeedback training in children with attention deficit/hyperactivity disorder: A cognitive and behavioral outcome study. **Clinical EEG and Neuroscience**, **50**(4), 242–255.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., Chen, Y., Lillicrap, T., Hui, F., Sifre, L., van den Driessche, G., Graepel, T., & Hassabis, D. (2017). Mastering the game of Go without human knowledge. **Nature**, **550**(7676), 354–359. <https://doi.org/10.1038/nature24270>

- Smith, R., Friston, K. J., & Whyte, C. J. (2024). 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>
- Smoller, J. W., Andreassen, O. A., Edenberg, H. J., Faraone, S. V., Glatt, S. J., & Kendler, K. S. (2019). Psychiatric genetics and the structure of psychopathology. **Molecular Psychiatry**, **24**(3), 409–420. <https://doi.org/10.1038/s41380-017-0010-4>
- Sobanski, E. (2006). Psychiatric comorbidity in adults with attention-deficit/hyperactivity disorder (ADHD). **European Archives of Psychiatry and Clinical Neuroscience**, **256**(Suppl 1), i26–i31.
- Steiner, N. J., Frenette, E. C., Rene, K. M., Brennan, R. T., & Perrin, E. C. (2014). Neurofeedback and cognitive attention training for children with attention-deficit hyperactivity disorder in schools. **Journal of Developmental & Behavioral Pediatrics**, **35**(1), 18–27.
- Stephan, K. E., & Mathys, C. (2014). Computational approaches to psychiatry. **Current Opinion in Neurobiology**, **25**, 85–92. <https://doi.org/10.1016/j.conb.2013.12.007>
- Stephan, K. E., Bach, D. R., Fletcher, P. C., Flint, J., Frank, M. J., Friston, K. J., Heinz, A., Huys, Q. J. M., Owen, M. J., Binder, E. B., Dayan, P., Johnstone, E. C., Meyer-Lindenberg, A., Montague, P. R., Schnyder, U., Wang, X.-J., & Breakspear, M. (2017). Charting the landscape of priority problems in psychiatry, part 1: Classification and diagnosis. **The Lancet Psychiatry**, **4**(1), 77–83. [https://doi.org/10.1016/S2215-0366\(16\)30339-4](https://doi.org/10.1016/S2215-0366(16)30339-4)
- Thatcher, R. W., Walker, R. A., Biver, C. J., North, D. N., & Curtin, R. (2003). Quantitative EEG normative databases: Validation and clinical correlation. **Journal of Neurotherapy**, **7**(3–4), 87–121.
- Van Doren, J., Arns, M., Heinrich, H., Vollebregt, M. A., Strehl, U., & Loo, S. K. (2019). Sustained effects of neurofeedback in ADHD: A systematic review and meta-analysis. **European Child & Adolescent Psychiatry**, **28**(3), 293–305.
- Wu, G., He, Q., Li, D., Zhang, Z., Miao, J., & Shu, Y. (2024). Comparative efficacy of neurofeedback interventions for attention-deficit/hyperactivity disorder in children: A network meta-analysis. **Brain and Behavior**, **14**(12), e70194. <https://doi.org/10.1002/brb3.70194>

Yeh, W.-H., Hsueh, J.-J., & Shaw, F.-Z. (2022). A pilot meta-analysis on self-reported efficacy of neurofeedback for adolescents and adults with ADHD. *Scientific Reports*, *12*(1), 9931. <https://doi.org/10.1038/s41598-022-14220-y>

Zhong, X., Yuan, X., Dai, Y., Liu, X., Wang, J., & Chen, M. (2025). Neurofeedback training for executive function in ADHD children: A systematic review and meta-analysis. *Scientific Reports*, *15*, 28148. <https://doi.org/10.1038/s41598-025-94242-4>

Total: 70 references (peer-reviewed journals 60+, books/chapters 7, in-press/preprint 3).

Target for final Nature submission: 80–120 references. Additional 10–50 to be added in final manuscript preparation covering bipolar disorder differential diagnosis, PTSD computational models, precision medicine in psychiatry, and OCD biomarkers.

Acknowledgments

The author dedicates this work to the late Professor Kurt W. Fischer (1943–2020) of Harvard Graduate School of Education, whose Dynamic Skill Theory and personal mentorship during the author's 2018–2020 Visiting Scholar appointment provided the foundational framework for the present operationalization of five-axis longitudinal skill measurement in adult neuropsychiatric assessment.

The author thanks the preliminary cohort participants for their engagement and trust during the 8-week protocol. All clinical data are anonymized in this manuscript.

Author Contributions

A.K.: conceptualization, methodology, software, validation, formal analysis, writing — original draft, writing — review and editing, supervision.

Competing Interests

A.K. is the founder and principal of Boston Neuromind LLC, the commercial entity developing CNA v7. Four U.S. Patent and Trademark Office provisional patent applications have been filed covering the Differential Diagnosis Engine, the Comorbidity Decomposition module, the Multi-Source Diagnostic Algorithm, and the Adaptive Re-routing layer. The framework code and synthetic validation datasets are released under permissive license at github.com/bostonneuromind.

Data Availability

Synthetic validation dataset ($N = 1,000$) and diagnostic fingerprints database (anonymized JSON schema) are available at the project GitHub repository. Clinical data from the preliminary cohort are not publicly available due to participant privacy considerations but may be made available, in anonymized form, upon reasonable request to the corresponding author subject to institutional ethics approval.

Code Availability

The CNA v7 inference engine, including the DDE Bayesian module, comorbidity decomposition, cognitive simulator, and adaptive re-routing components, is released under permissive open-source license at github.com/bostonneuromind/bnm-decision-engine.