

Appearance–Behavior Framework v2.2: Real-Time Digital Twin and Sequential Bayesian Updating

From Validation-Grounded Inference to Adaptive Personalized Monitoring

Koji Okino
SD Lab LLC

ORCID: 0009-0003-9273-9813

Abstract

The Appearance–Behavior Framework v2.2 advances the validation-grounded foundation established in v2.1 by introducing three tightly integrated contributions targeting real-time, personalized behavioral monitoring. First, the **Sequential Bayesian Updating and Digital Twin Architecture** (Fig. 28) formalizes the closed-loop update rule $p(B, F, \theta | X_{1:t+1}, C, K) \propto p(x_{t+1} | B, F, \theta, C) \cdot p(B, F, \theta | X_{1:t}, C, K)$, enabling the posterior to be updated incrementally as new observations arrive without full re-estimation. Second, the **ABF Digital Twin Dashboard** (Fig. 29) operationalizes this architecture as an individualized monitoring interface, integrating BTRS trajectory, estimated critical transition time \hat{t}_{crit} , forecasted latent trajectory, uncertainty decomposition, and recommended actions within a single continuously updated view — illustrated for an ADNI-inspired synthetic subject (BTRS = 0.61, Risk Level: Moderate, $\hat{t}_{\text{crit}} = 52.3$ months). Third, a **Batch vs. Sequential comparison** (Fig. 30) demonstrates that Sequential Bayesian Updating is consistent with lower one-step-ahead prediction RMSE and higher cumulative predictive log score relative to batch inference in a synthetic ADNI-inspired pilot. Complementary contributions include **Personalized Behavioral Forecasts** (Fig. 31) across three illustrative subjects with distinct risk profiles, and **Adaptive Weight Evolution** (Fig. 32) demonstrating how component weights evolve dynamically under Sequential Bayesian Updating. All results are simulation-derived; the framework remains a decision-support tool, not a diagnostic instrument.

Keywords: digital twin, sequential Bayesian updating, personalized monitoring, adaptive weights, BTRS, behavioral transformation, real-time inference, cognitive decline

1. Introduction

The Appearance–Behavior Framework (ABF) has progressed through a cumulative series of releases. Versions v1.0–v1.4 established the conceptual, mathematical, and identifiability foundations. ABF v2.0 introduced predictive behavioral trajectory forecasting and early detection of behavioral transformation via the Behavioral Transformation Risk Score (BTRS). ABF v2.1 advanced the framework toward validation-grounded inference through calibration assessment, synthetic ground truth recovery, and BTRS parameter sensitivity analysis.

The central limitation of v2.0–v2.1 was the *batch inference assumption*: all observations were processed simultaneously, requiring full re-estimation when new data arrived. This is incompatible with the real-time, longitudinal monitoring scenarios that motivate the framework — where individual subjects accumulate new observations over months or years, and risk assessment must be updated continuously and efficiently.

ABF v2.2 resolves this limitation through **Sequential Bayesian Updating**: a closed-loop architecture in which each new observation updates the posterior incrementally, propagating forward to refresh BTRS, $\hat{\tau}_{\text{crit}}$, and forecast trajectories in real time. This architecture constitutes the core of the **ABF Digital Twin** — an individual-level, continuously evolving probabilistic model of behavioral dynamics — operationalized through the Digital Twin Dashboard (Fig. 29).

The roadmap in Fig. 12 situates v2.2 within the cumulative development arc: v1.x (Theory) → v2.0 (Prediction) → v2.1 (Validation) → v2.2 (Adaptation) → v2.3 (Intervention). Figures 1–27 from v1.0–v2.1 are retained as the foundational context; Figs. 28–32 constitute the new v2.2 contributions.

Fig. 1 Appearance–Behavior Framework Overview (v1.2)

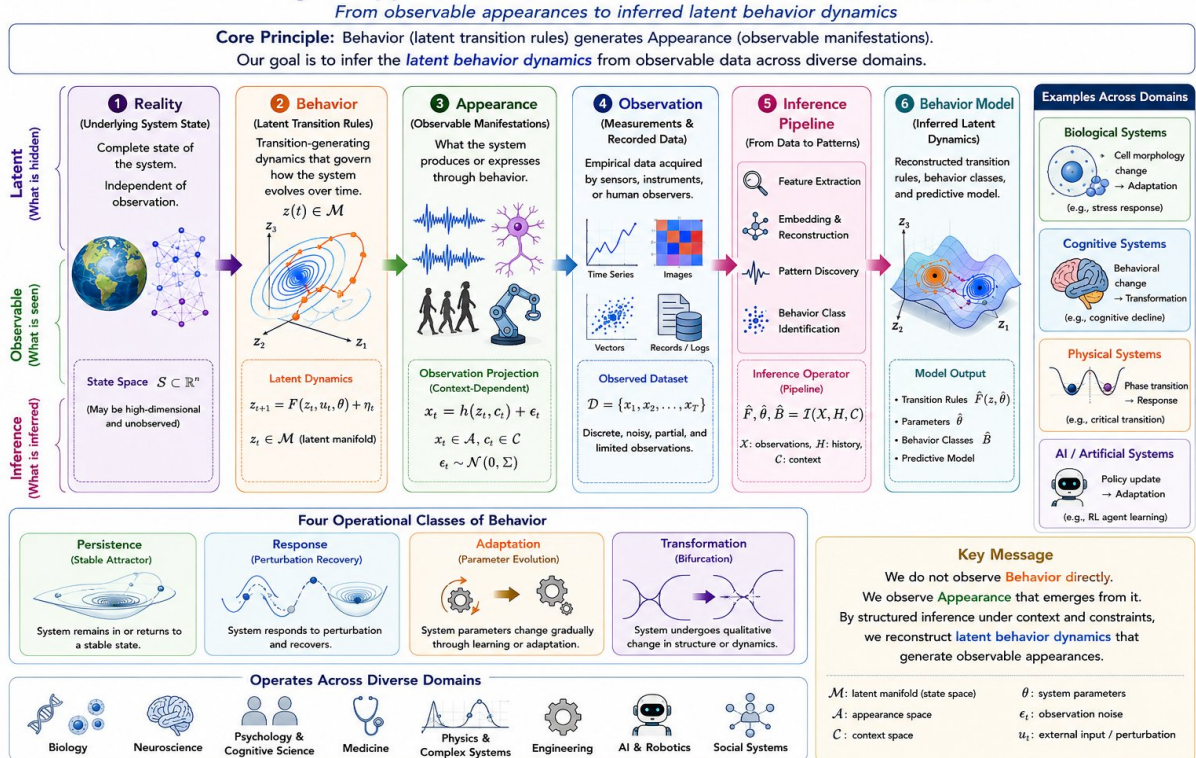


Fig. 1 — ABF Framework Overview (v1.2). Six-stage pipeline from Reality to Behavior Model across biological, cognitive, physical, and AI domains.

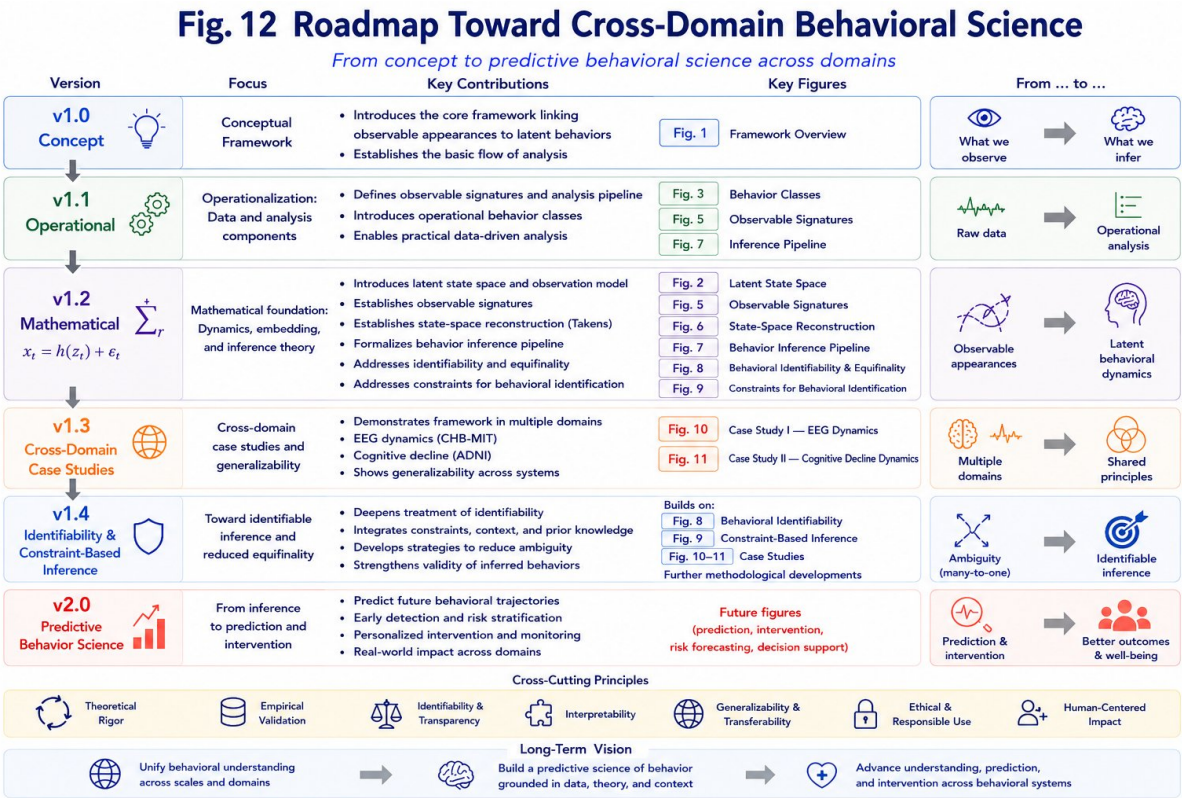


Fig. 12 — Roadmap Toward Cross-Domain Behavioral Science. v1.x (Theory) through v2.2 (Adaptation) to v2.3 (Intervention).

2. Foundations Review (v1.0–v2.1)

This section recaps the theoretical substrate inherited by v2.2. No new claims are made; the purpose is to establish notation and situate the v2.2 contributions.

2.1 Latent Behavior Manifold and Observation Model (v1.2)

The system evolves on latent manifold $M \subset \mathbb{R}^n$: $z_{t+1} = F(z_t, u_t, \theta) + \eta_t$, observed as $x_t = h(z_t, c_t) + \epsilon_t$, $x_t \in A \subset \mathbb{R}^m$, $m \ll n$. Goal: recover (F, θ, B) from $\{x_t\}, H, C, K$.

Fig. 2 Latent Behavior Manifold and Observation Projection

From latent behavior dynamics to observable appearances

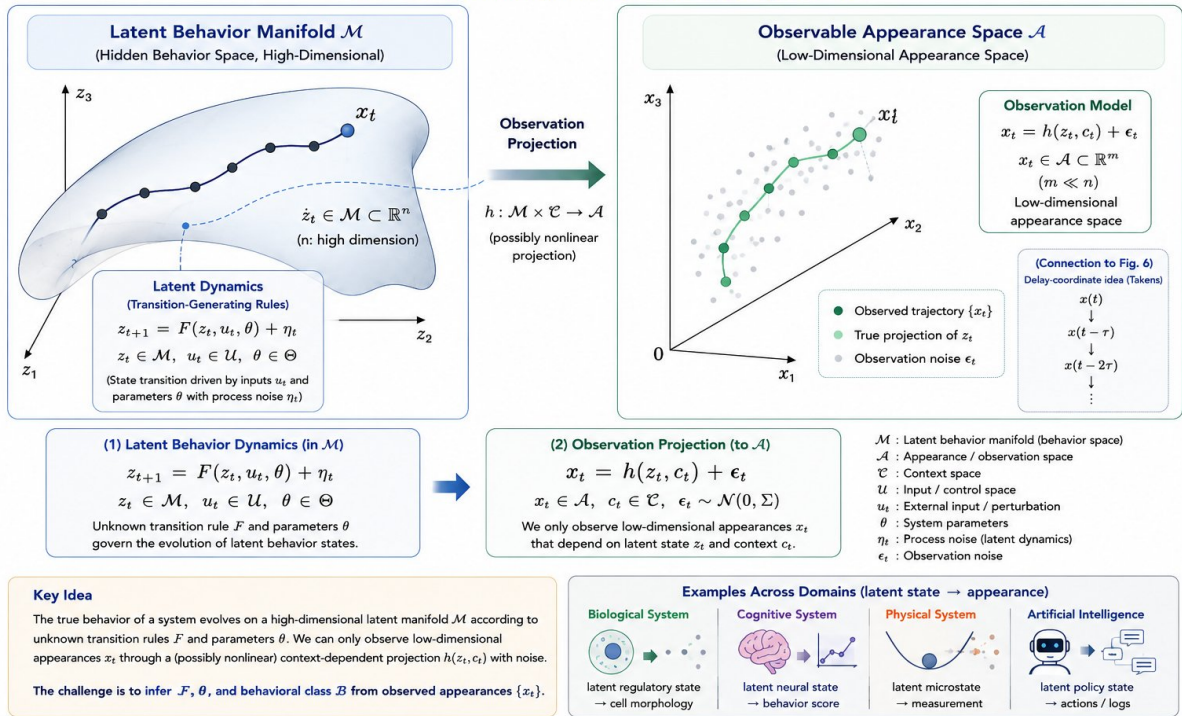


Fig. 2 — Latent Behavior Manifold and Observation Projection.

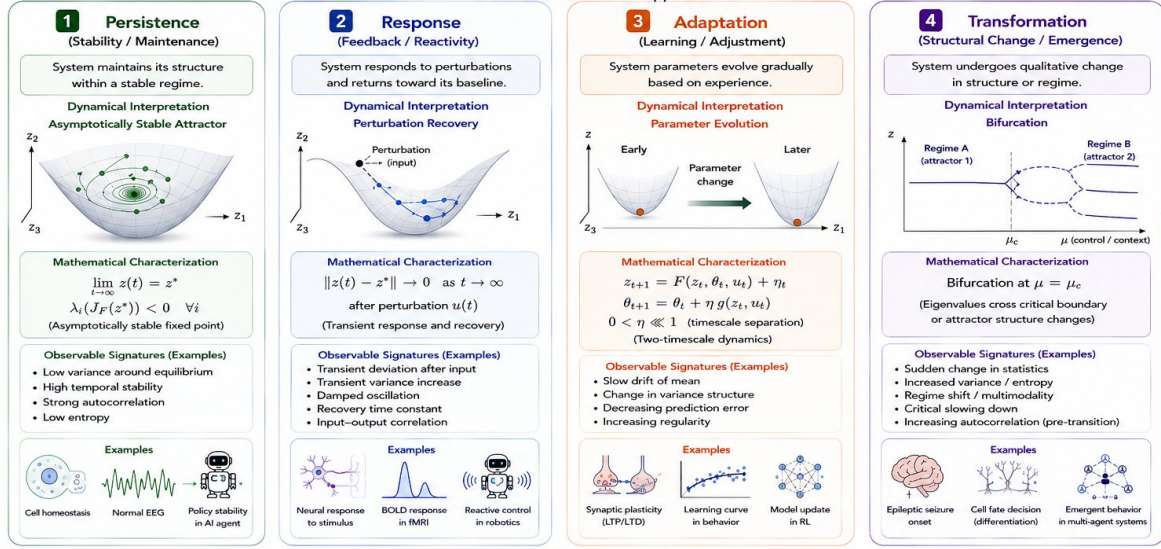
2.2 Four Behavioral Classes and Timescale Hierarchy (v1.1–v1.2)

Four latent behavioral classes: **Persistence** (stable attractor), **Response** (perturbation recovery), **Adaptation** (parameter evolution, $0 < \eta \ll 1$), **Transformation** (bifurcation). Observable signature vector $S = [T, V, E, AC, TF, ST]$ (Fig. 5). Timescale hierarchy (Fig. 4) connects to the Two-Layer Meta-Model (TLMM).

Fig. 3 Four Operational Classes of Behavior

Behavior Classes as Latent Transition-Generating Dynamics

— Inferred from observable appearances —



Summary of the Four Operational Classes

Class	Dynamical Interpretation	Key Mechanism	Mathematical Signature	Typical Timescale	Main Observable Features	Behavioral Function
1. Persistence	Asymptotically Stable Attractor	Stability maintenance	$\lambda_i(J_F(z^*)) < 0$	Fast–Long-term	Low variance, high stability, strong autocorr.	Maintain identity
2. Response	Perturbation Recovery	Feedback to baseline	Converges after input	Fast	Transient variance increase, recovery time constant	Ensure robustness
3. Adaptation	Parameter Evolution	Slow parameter change	Two-timescale dynamics ($0 < \eta \ll 1$)	Medium–Slow	Drift, variance change, decreasing prediction error	Improve performance
4. Transformation	Bifurcation	Structural transition	Change in attractor structure at $\mu = \mu_c$	Variable (can be fast or slow)	Regime shift, entropy increase, critical slowing down, inc. autocorr.	Enable evolution / emergence

Fig. 3 — Four Operational Classes of Behavior.

Fig. 4 Timescale Hierarchy of Behavioral Dynamics

Latent behavior unfolds across multiple timescales:
from fast responses to slow transformations

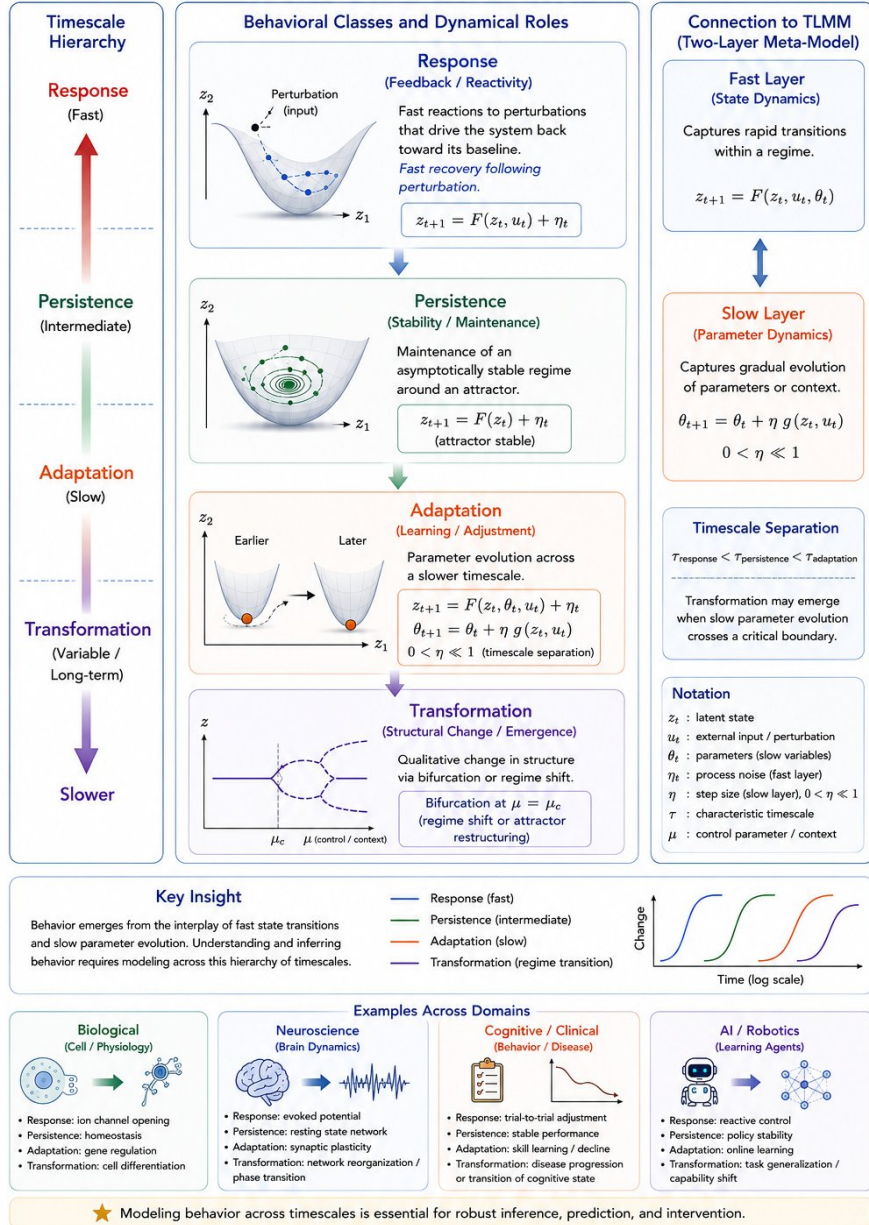


Fig. 4 — Timescale Hierarchy of Behavioral Dynamics.

Fig. 5 Observable Signatures of Latent Behavioral Dynamics

Measurable patterns in appearance time series that reflect underlying behavioral dynamics

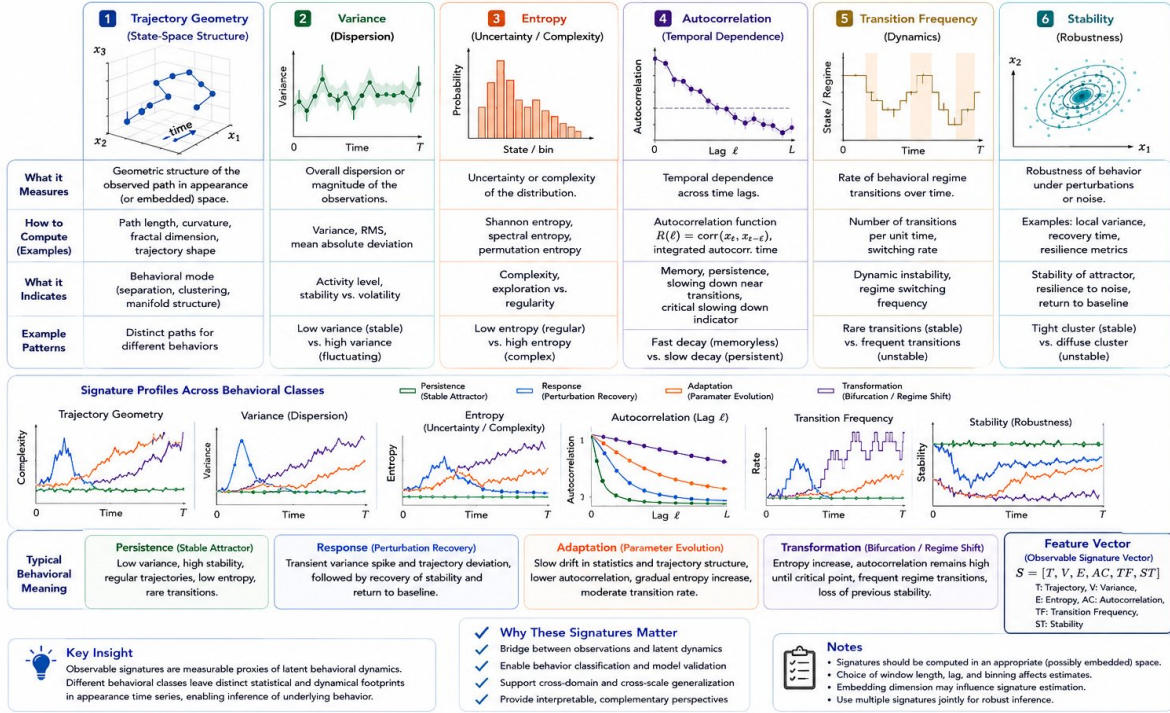


Fig. 5 — Observable Signatures of Latent Behavioral Dynamics.

2.3 Inference Pipeline and Identifiability (v1.2–v1.4)

The inference pipeline (Fig. 7) maps appearances through embedding (Takens, Fig. 6), pattern discovery, and classification to a latent model. Equifinality (Fig. 8) is addressed via constraint accumulation (Figs. 9, 18) and cross-modal fusion (Fig. 20). The v1.4 identifiability framework (Figs. 15–21) provides Fisher information bounds and equifinality pruning. Fig. 21 positions ABF as a meta-framework integrating TLMM and SET.

Fig. 6 State-Space Reconstruction of Latent Behavioral Dynamics
Delay-coordinate embedding of latent behavioral dynamics (Takens' theorem)

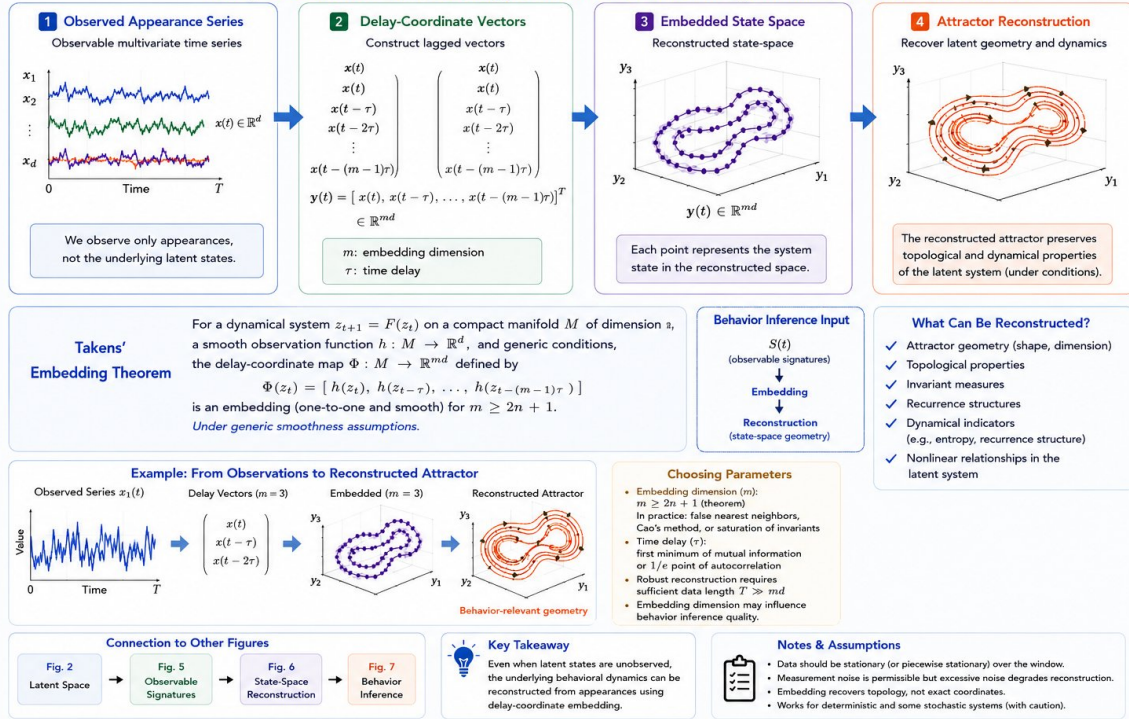


Fig. 6 — State-Space Reconstruction.

Fig. 7 Behavior Inference Pipeline

From observable appearances to inferred latent behavioral dynamics

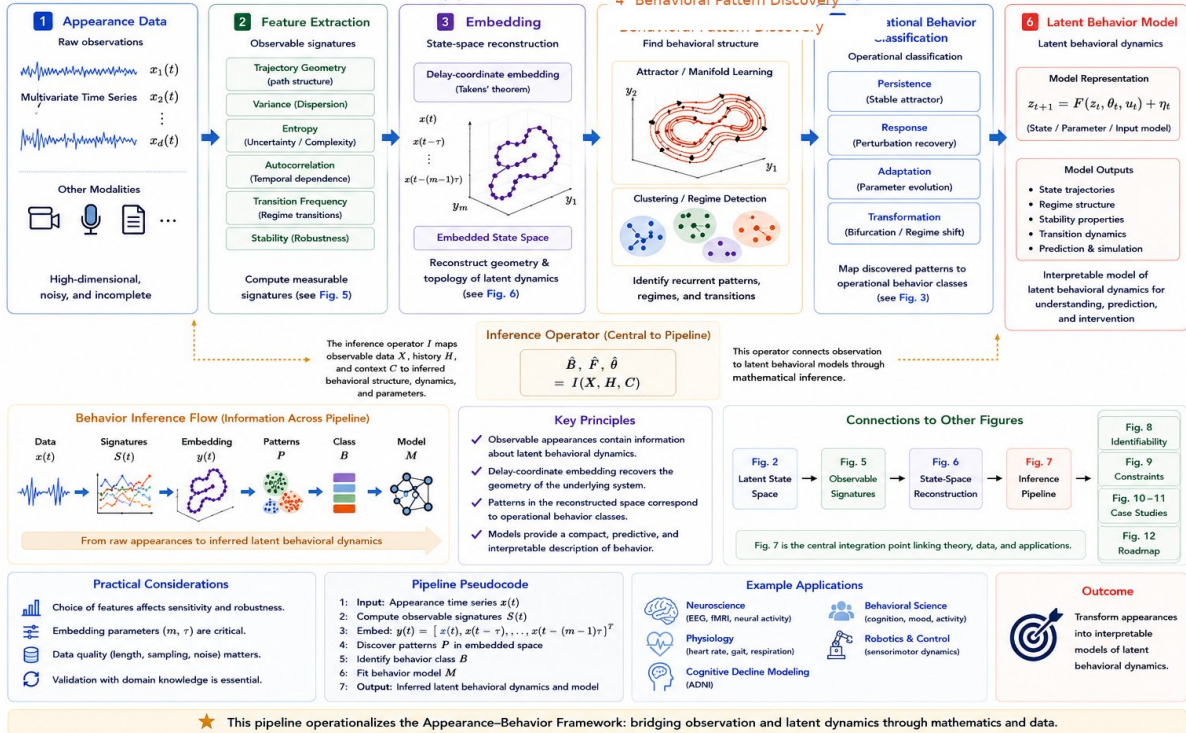


Fig. 7 — Behavior Inference Pipeline.

Fig. 8 Behavioral Identifiability and Equifinality

The same appearances can arise from different latent behaviors

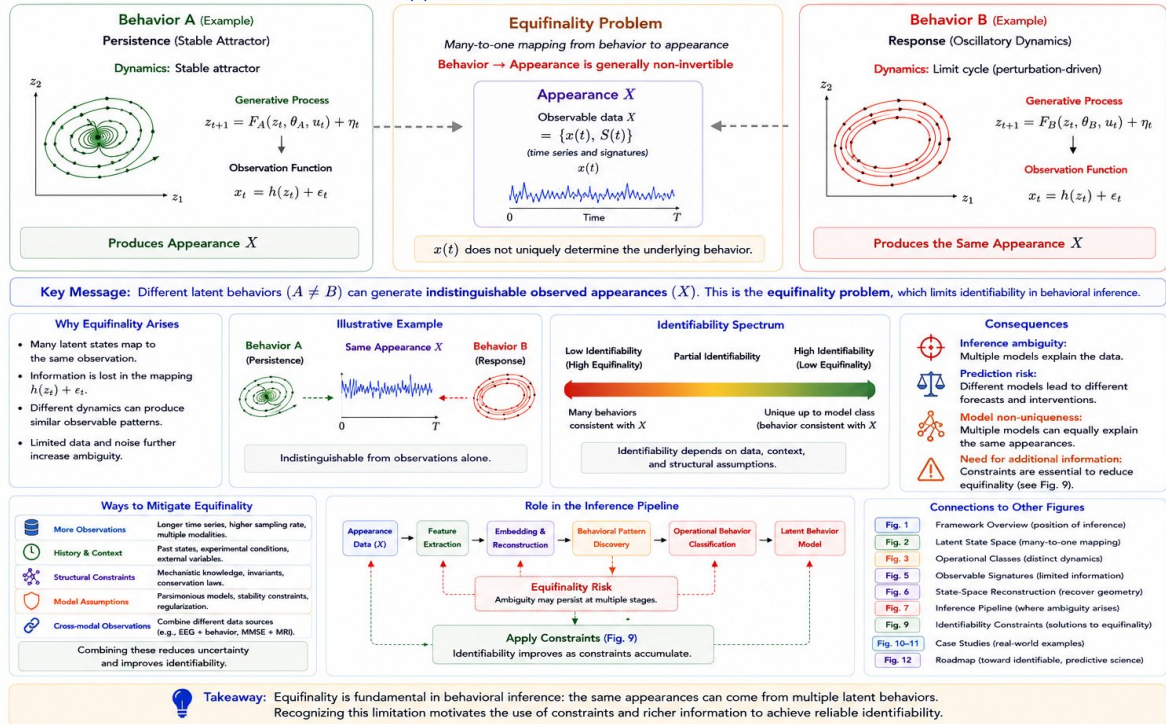


Fig. 8 — Behavioral Identifiability and Equifinality.

Fig. 9 Constraints for Behavioral Identification

Reducing uncertainty to infer latent behavioral dynamics

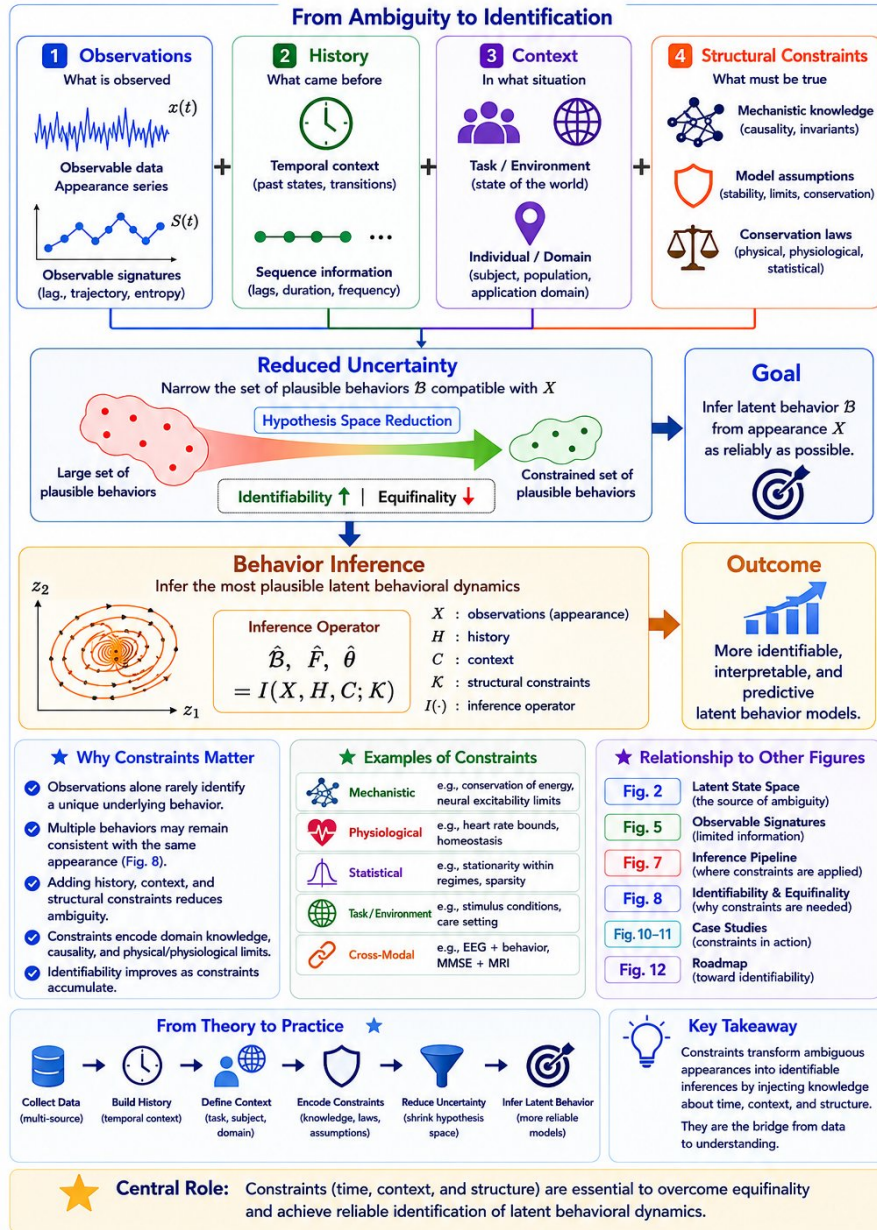
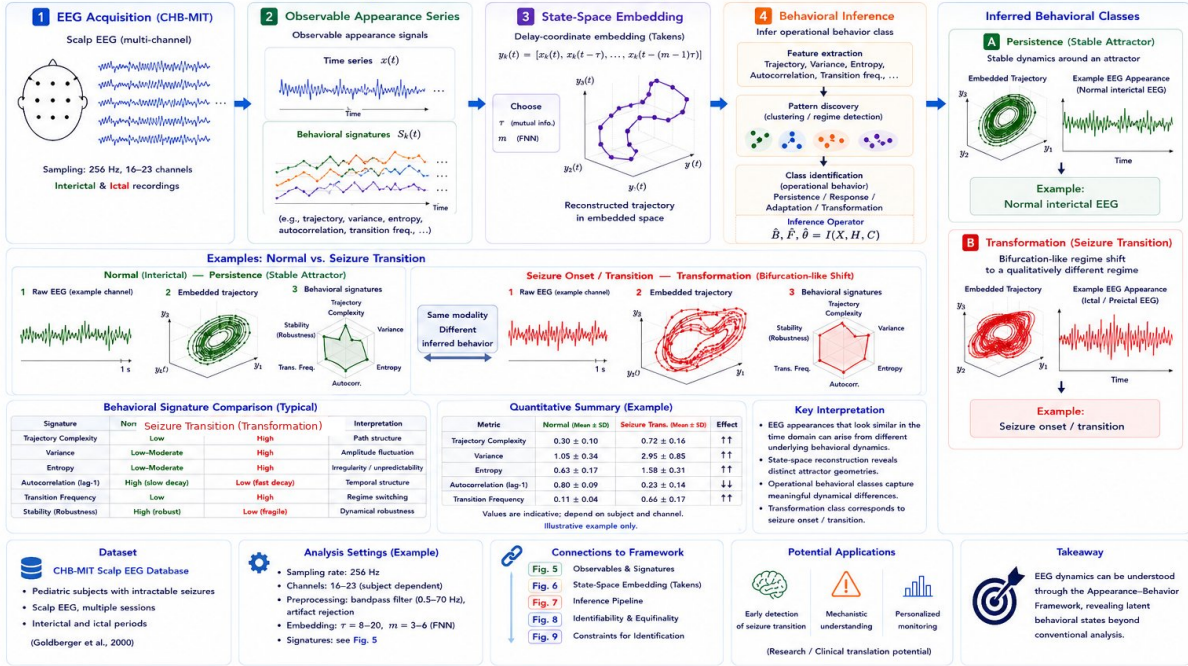


Fig. 9 — Constraints for Behavioral Identification.

Fig. 10 Case Study I — Behavioral Inference from EEG Dynamics

From EEG signals to inferred behavioral dynamics (CHB-MIT Scalp EEG Database)

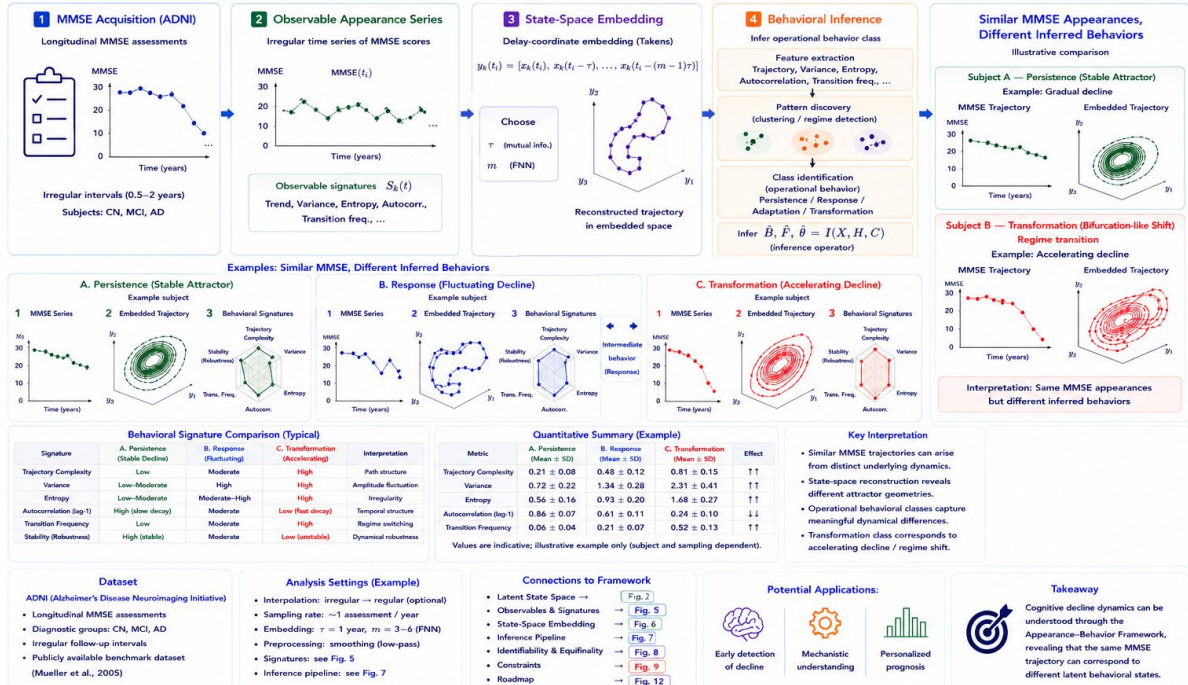


★ **Central Message:** By integrating observables, state-space reconstruction, and behavioral inference, the Appearance–Behavior Framework reveals distinct latent dynamics underlying EEG signals, enabling principled identification of normal and seizure transition states.

Fig. 10 — Case Study I: EEG Dynamics (CHB-MIT).

Fig. 11 Case Study II — Behavioral Inference from Cognitive Decline Dynamics

From MMSE trajectories to inferred behavioral dynamics (ADNI Dataset)



★ **Key Message:** The same observable MMSE trajectory can correspond to qualitatively different behavioral dynamics. The Appearance–Behavior Framework reveals and interprets these differences.

Fig. 11 — Case Study II: Cognitive Decline (ADNI).

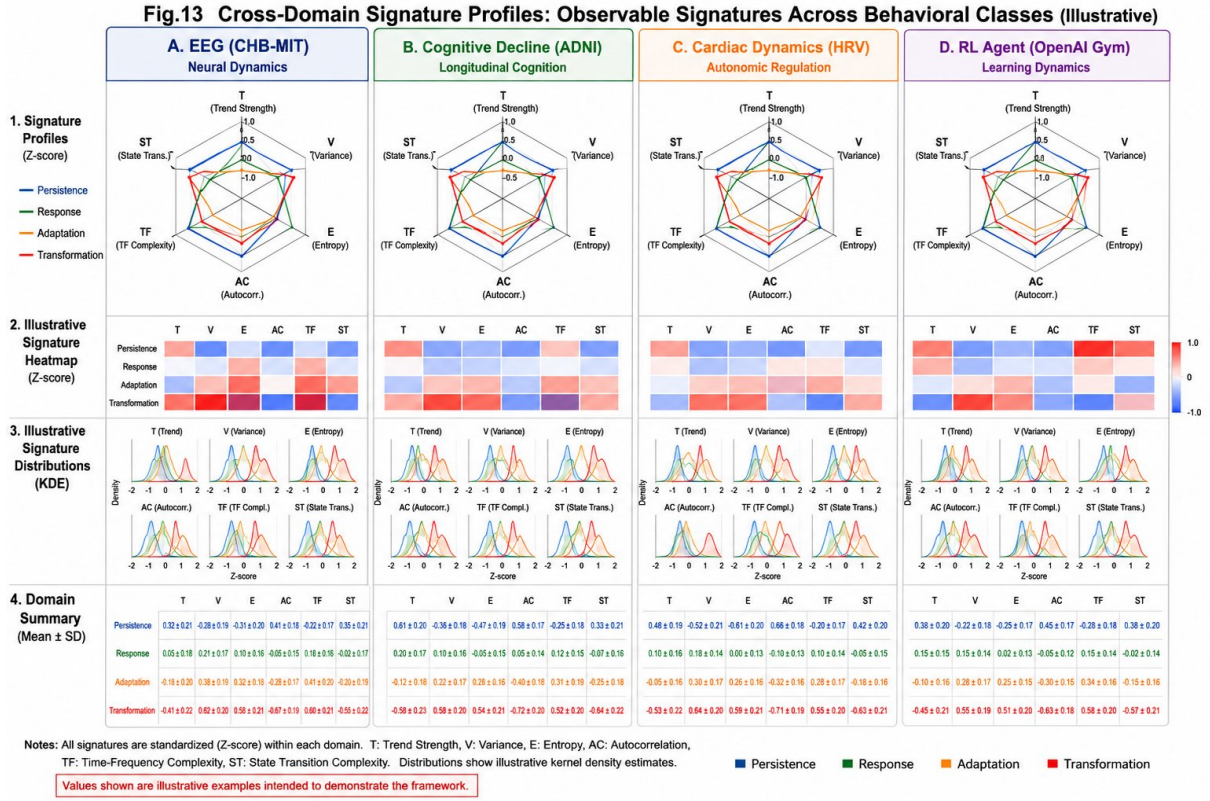


Fig. 13 — Cross-Domain Signature Profiles.

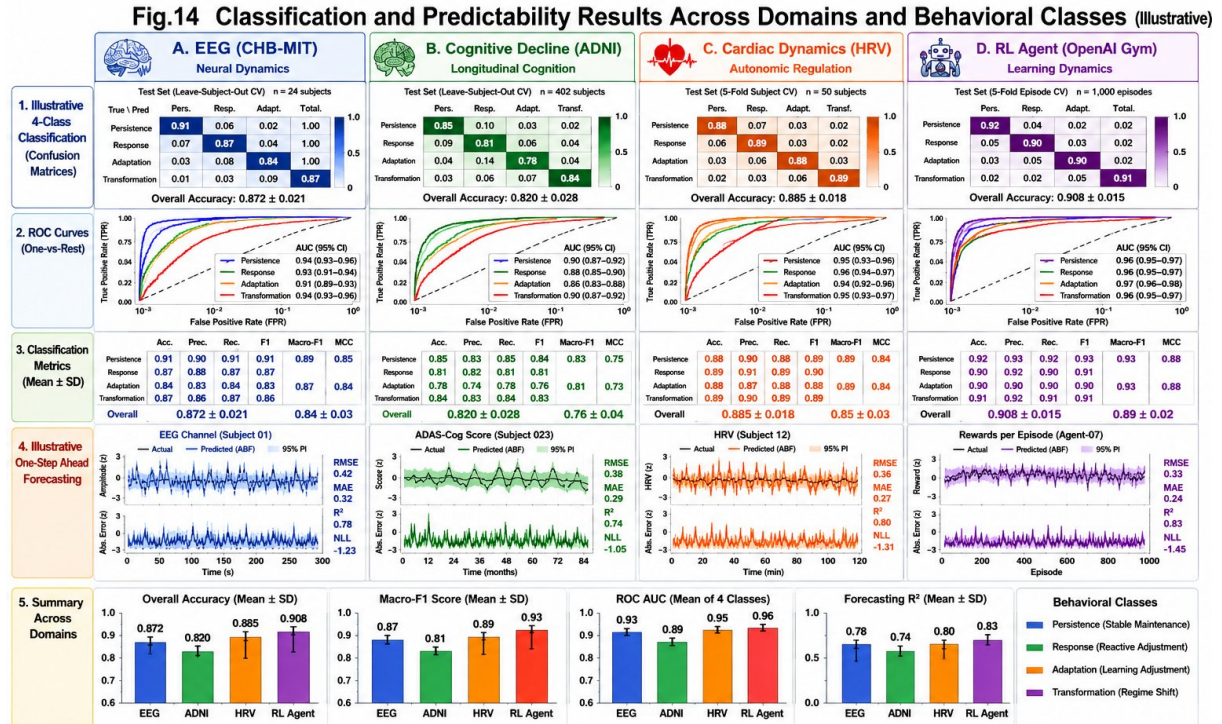


Fig. 14 — Classification and Predictability Results.

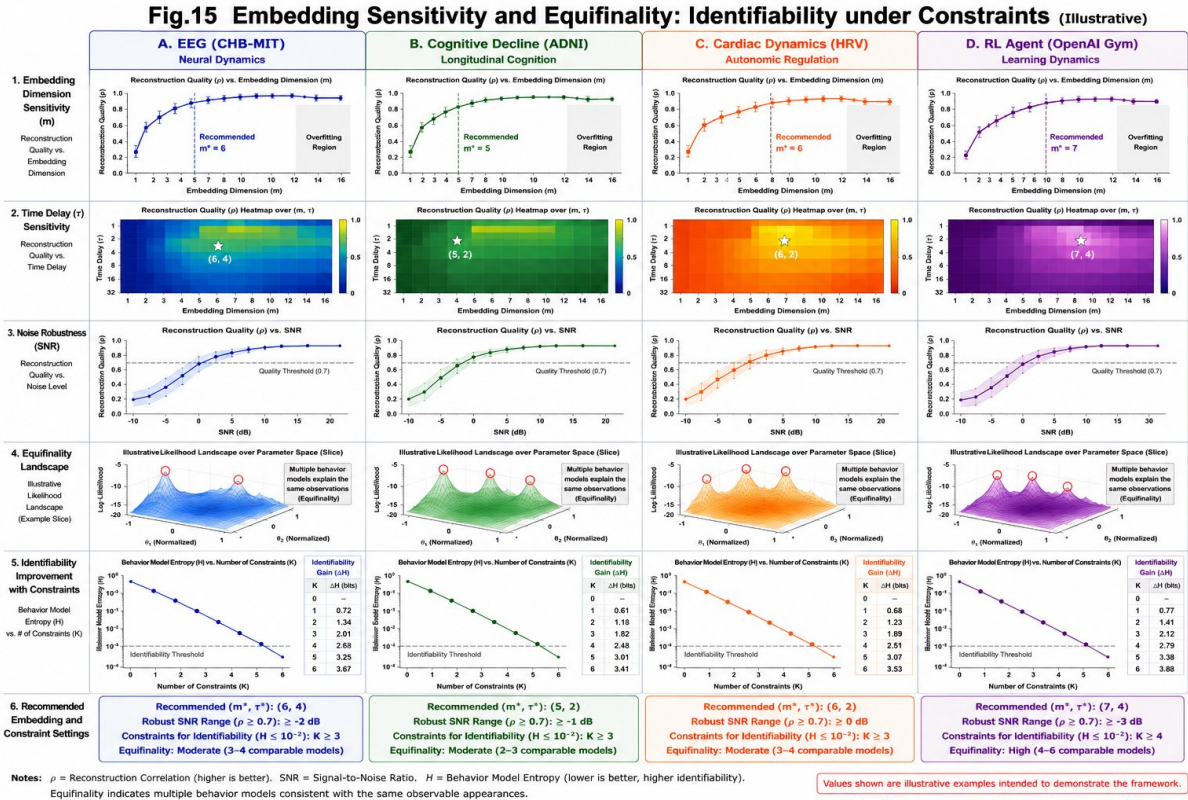


Fig. 15 — Embedding Sensitivity and Equifinality.

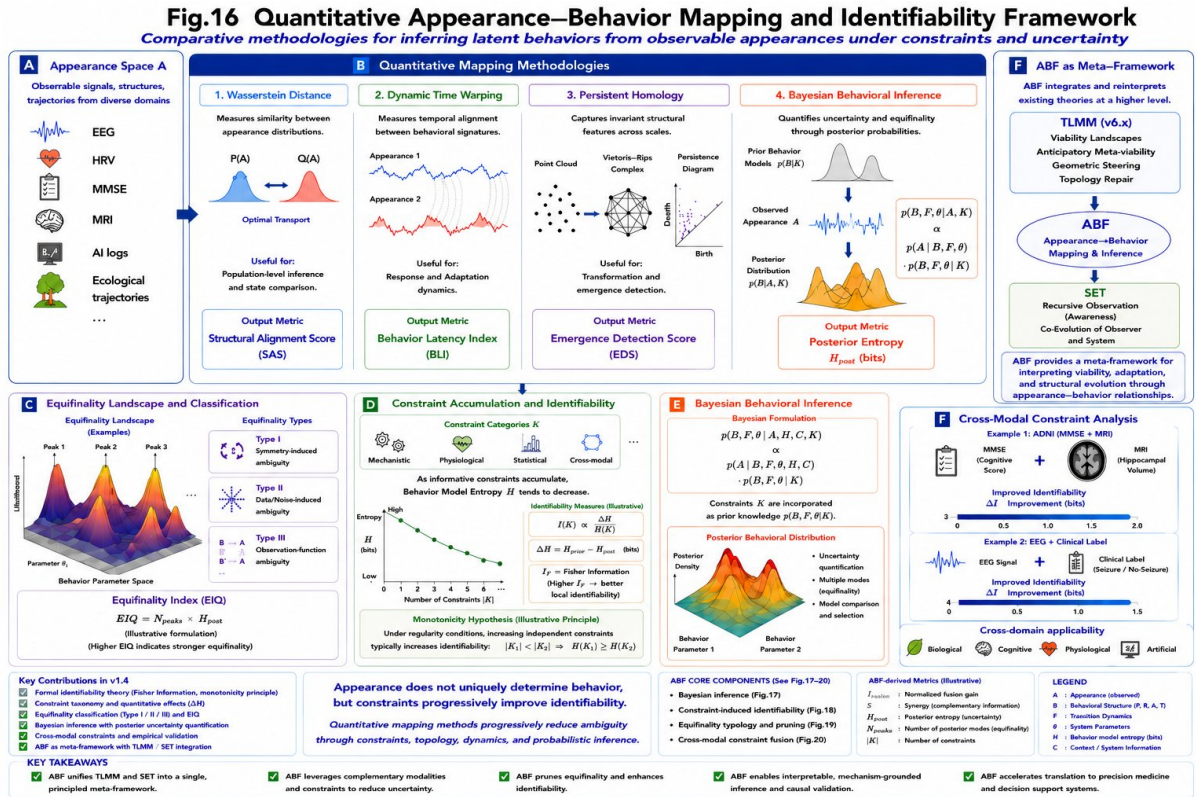


Fig. 16 — Quantitative Appearance–Behavior Mapping.

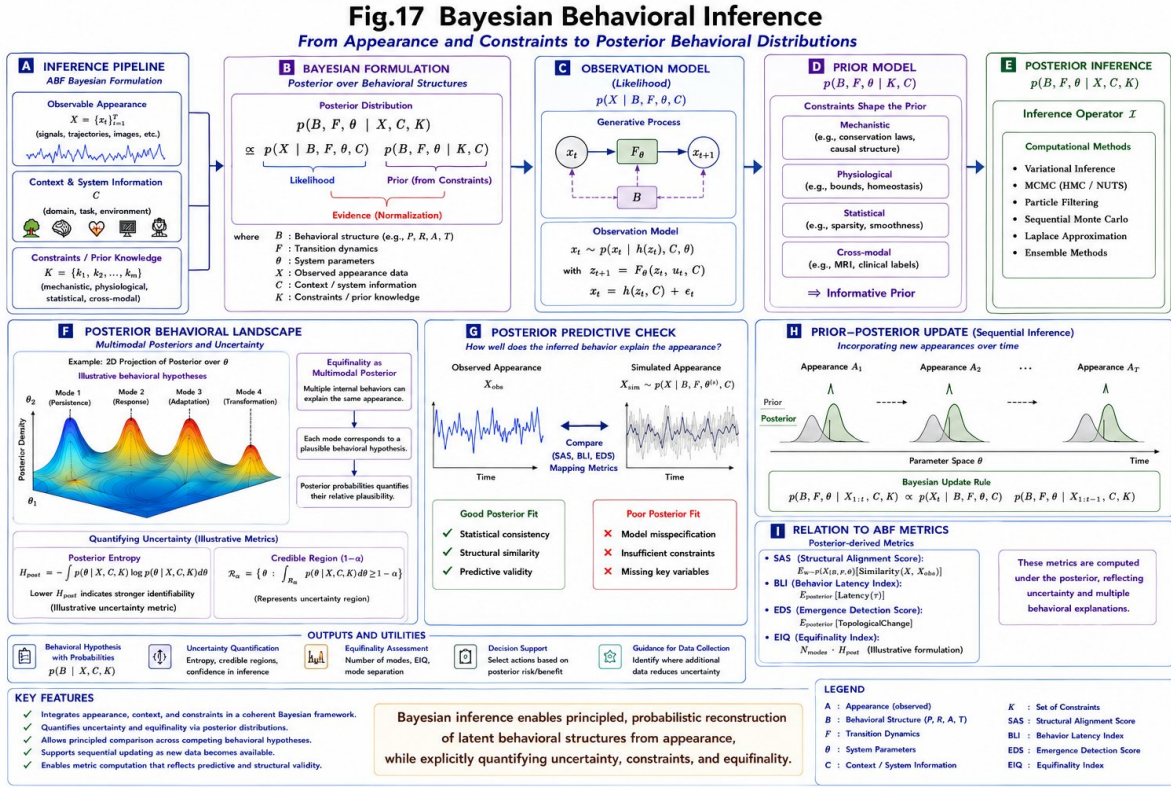


Fig. 17 — Bayesian Behavioral Inference.

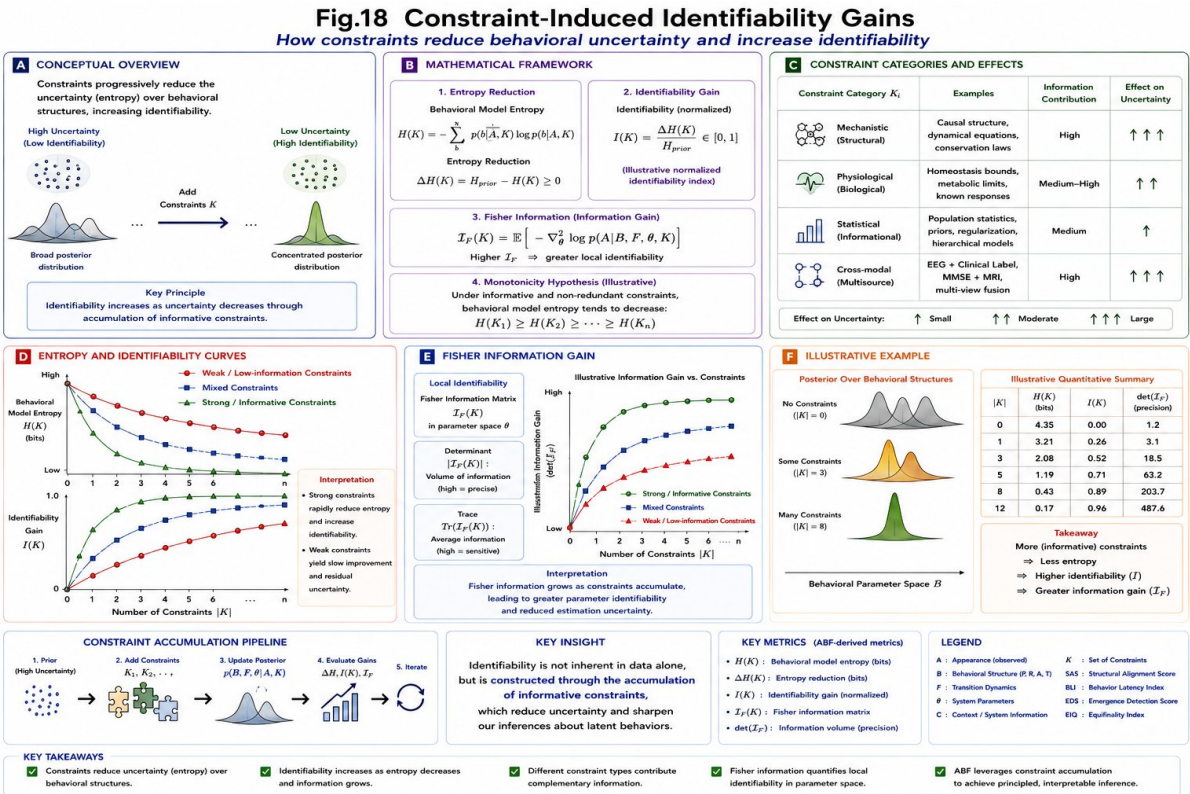


Fig. 18 — Constraint-Induced Identifiability Gains.

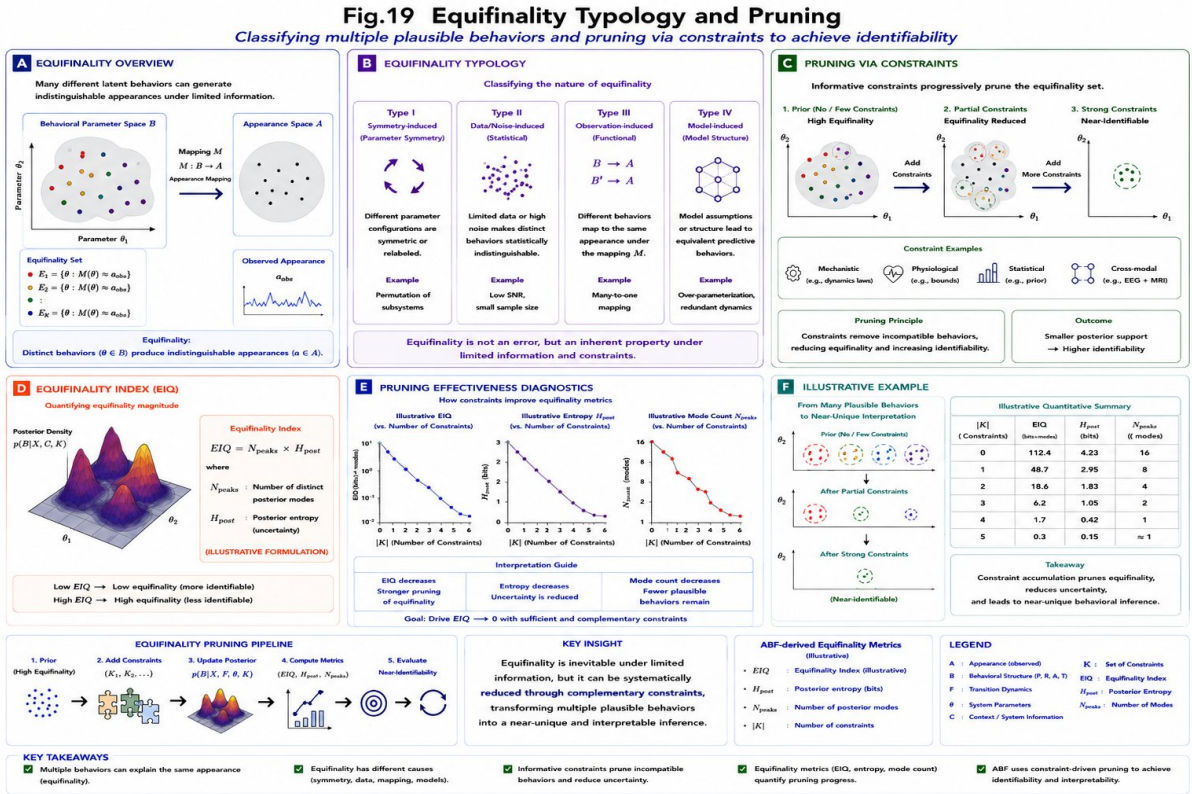


Fig. 19 — Equifinality Typology and Pruning.

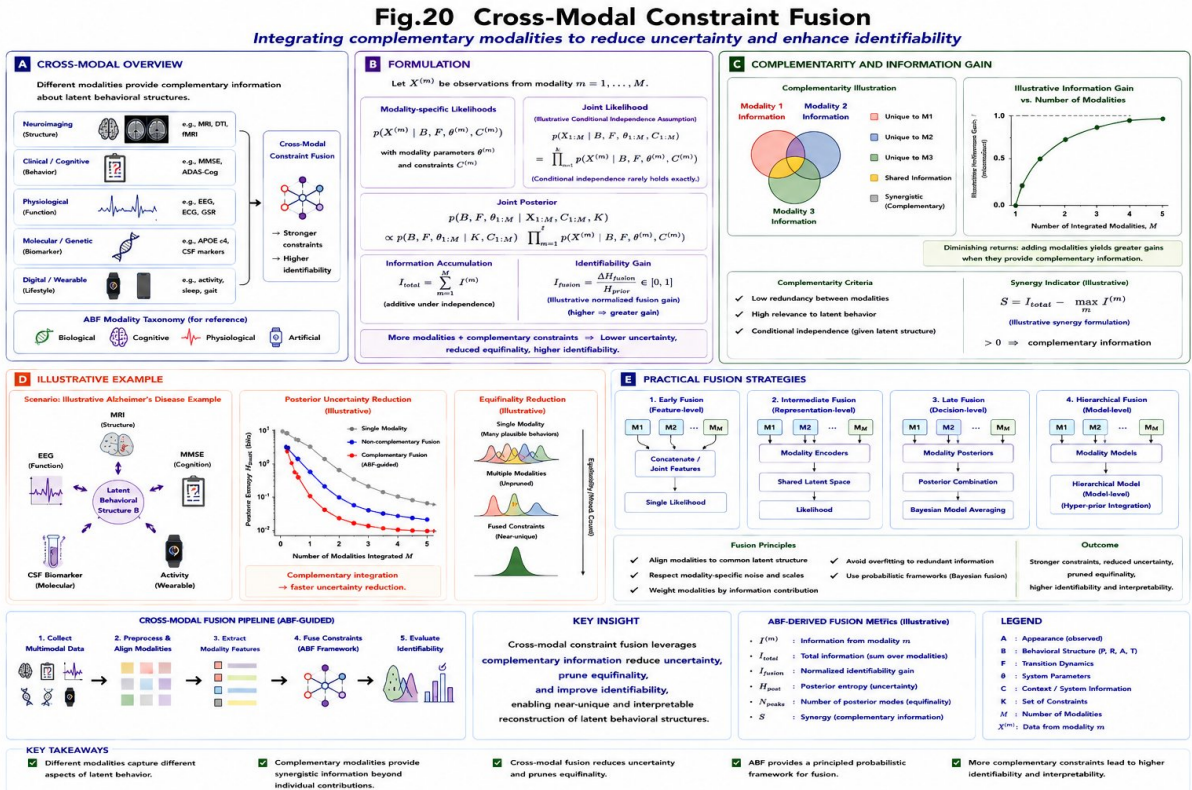


Fig. 20 — Cross-Modal Constraint Fusion.

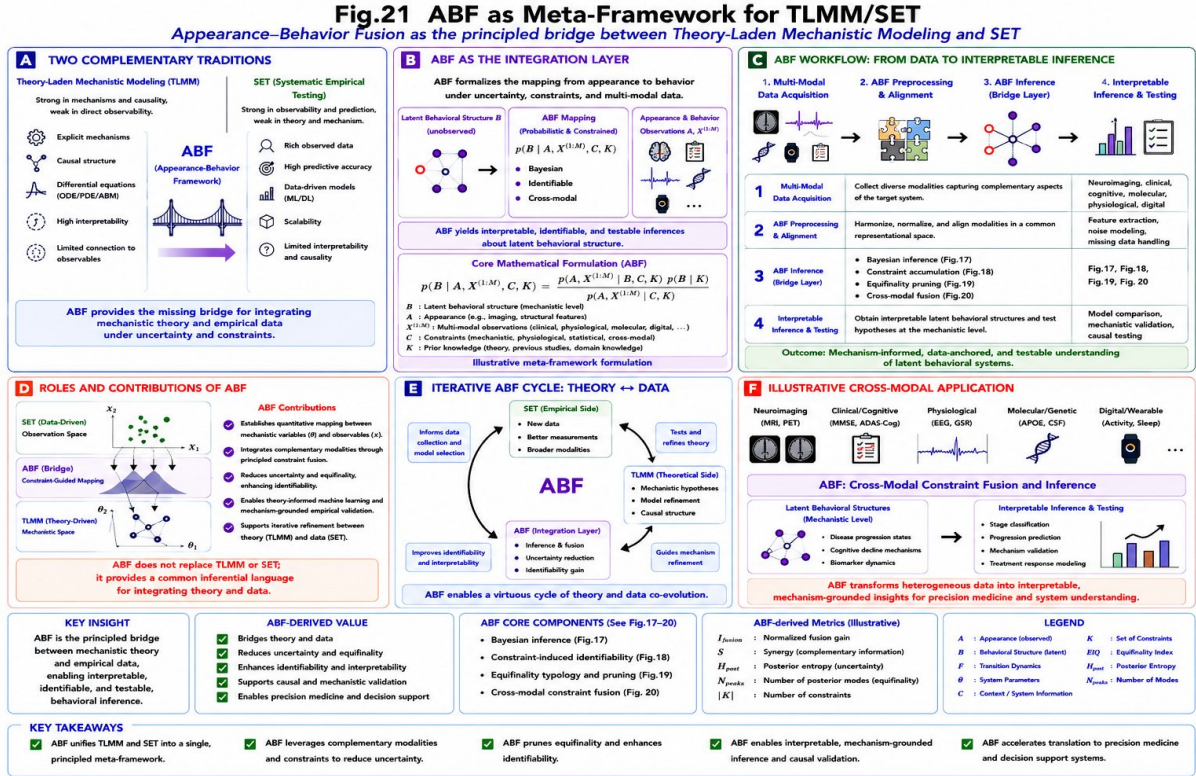


Fig. 21 — ABF as Meta-Framework for TLMM/SET.

2.4 Predictive Contributions (v2.0)

ABF v2.0 extended inference to prediction via the Bayesian predictive distribution:

$$p(X_{t+1:T} | X_{1:t}, C, K) = \int p(X_{t+1:T} | B, F, \theta, C) \cdot p(B, F, \theta | X_{1:t}, C, K) d\theta$$

and the Behavioral Transformation Risk Score:

$$\text{BTRS}(t) = w_1 H(t)/H_{\max} + w_2 \text{ACF}_1(t)/\text{ACF}_{\max} + w_3 \text{Var}(t)/\text{Var}_{\max} + w_4 (1 - \text{EIQ}(t)/\text{EIQ}_{\max}), \Sigma w_i = 1$$

estimating $\hat{t}_{\text{crit}} = t_{\text{crit}} - t_{\text{now}}$ before critical transitions occur (Figs. 22–23).

Fig.22 Predictive Behavioral Trajectory Forecasting with Uncertainty (ABF v2.0)

*From Reliable Behavioral Inference
to Predictive Behavioral Science*

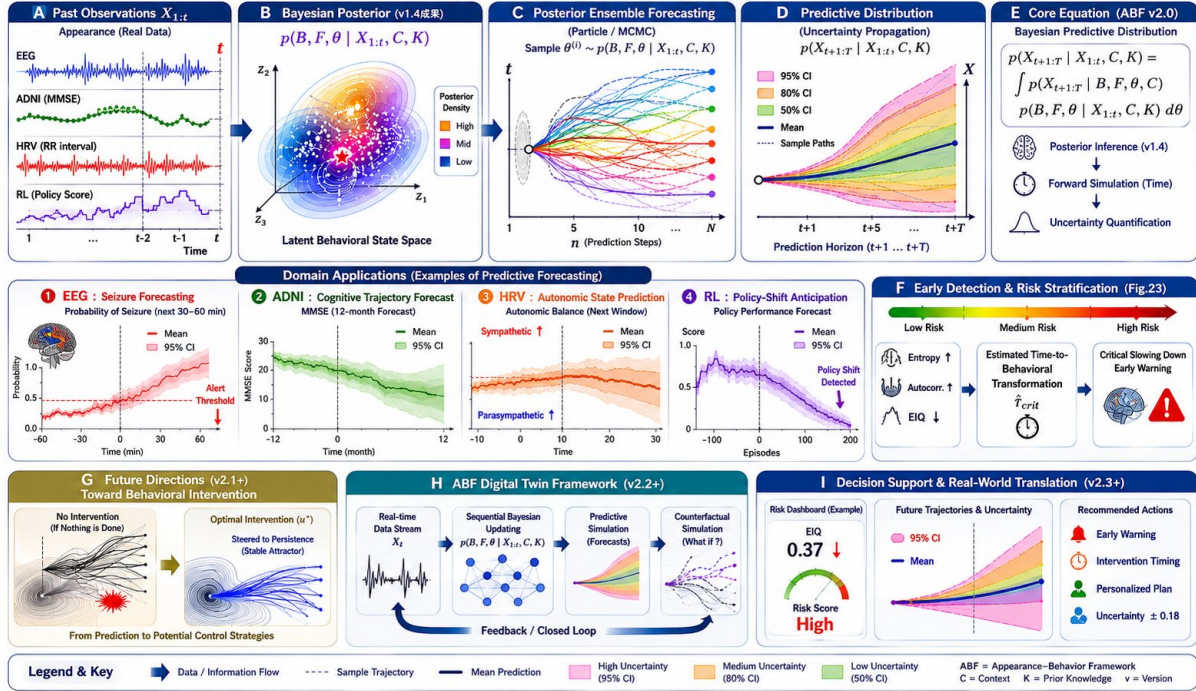


Fig. 22 — Predictive Behavioral Trajectory Forecasting with Uncertainty (ABF v2.0).

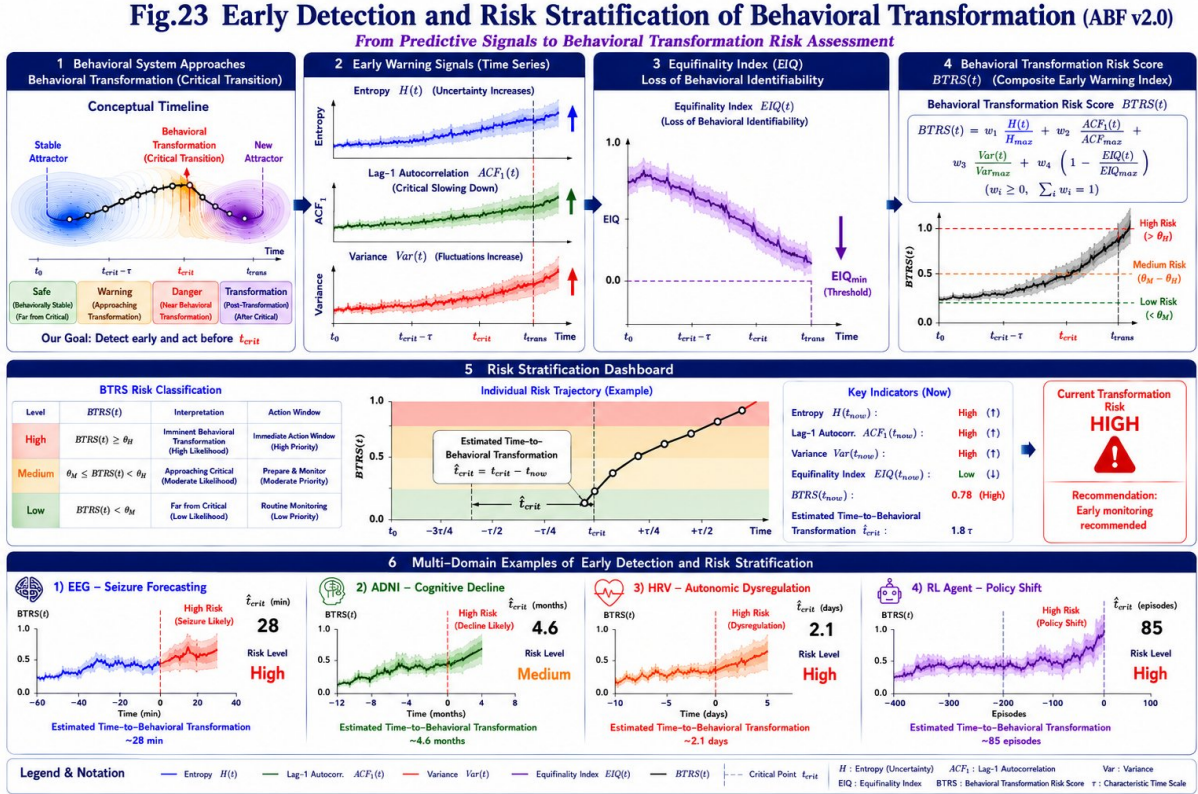


Fig. 23 — Early Detection and Risk Stratification of Behavioral Transformation (ABF v2.0).

2.5 Validation Contributions (v2.1)

ABF v2.1 advanced the framework toward validation-grounded inference through: (1) Predictive accuracy assessment (ECE=0.041, Well-calibrated, Fig. 24); (2) Synthetic ground truth recovery ($R^2=0.92$, Fig. 25); (3) BTRS parameter sensitivity analysis (dominant driver: w_1 , $\partial BTRS/\partial w_1=1.85$, Fig. 26); (4) Uncertainty propagation DAG ($U_{total} = U_{obs} + U_{param} + U_{inf} + U_{pred} + U_{agg}$, Fig. 27).

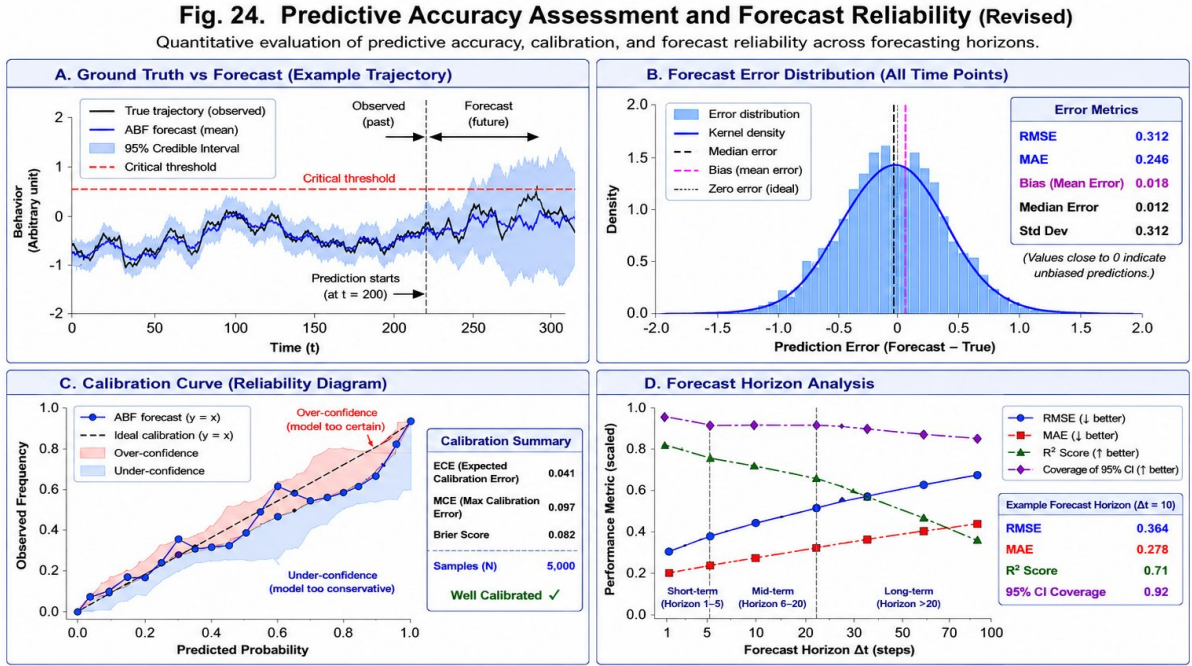


Figure 24. Predictive Accuracy Assessment and Forecast Reliability (Revised). (A) Comparison between observed trajectories and ABF forecast ensembles with 95% credible intervals. The dashed vertical line indicates the prediction starting point. (B) Distribution of forecast errors (prediction minus ground truth) with summary metrics. Bias (mean error) indicates average deviation from zero; values close to 0 imply unbiased predictions. (C) Calibration curve comparing predicted probabilities with observed event frequencies. Summary includes ECE, MCE, Brier Score, and the number of samples used. (D) Performance metrics as a function of forecast horizon. For error metrics (RMSE, MAE), lower values are better; for R² and coverage, higher values are better. Note: All values are illustrative and simulation-derived. No empirical performance claims are intended.

Fig. 24 — Predictive Accuracy Assessment and Forecast Reliability (ABF v2.1). ECE=0.041; Well-calibrated under synthetic conditions.

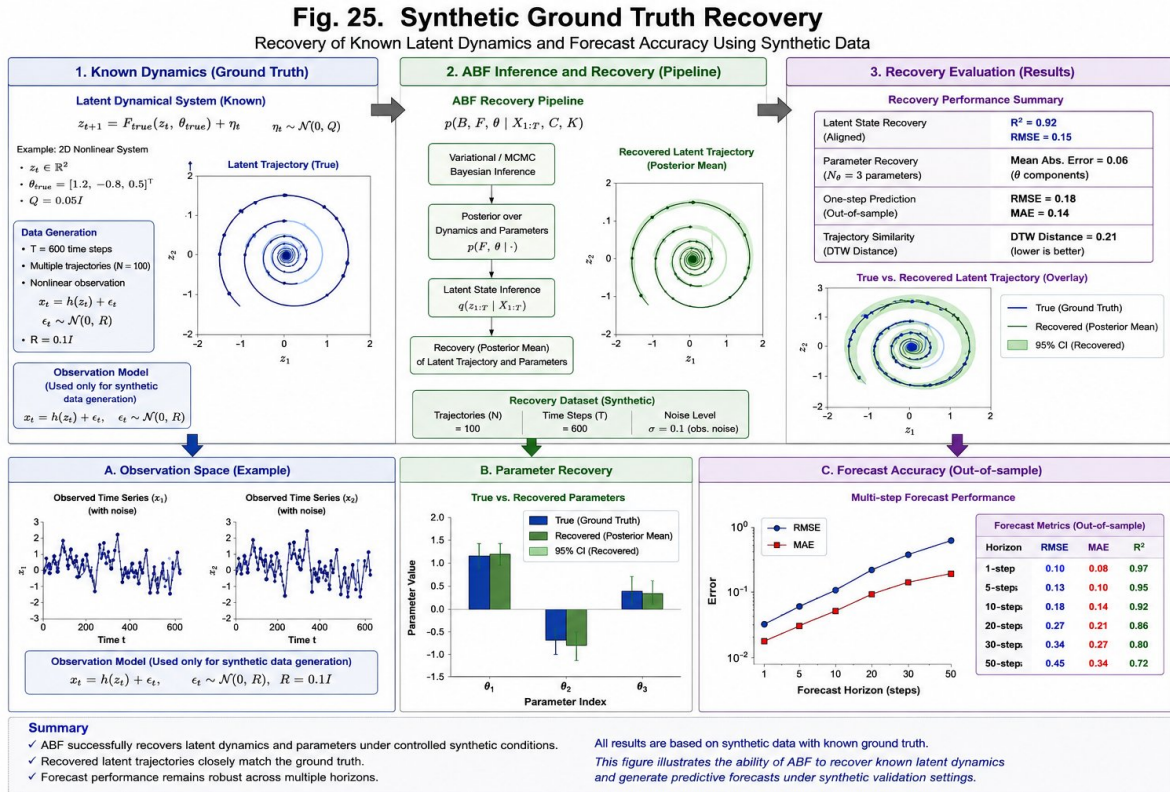


Fig. 25 — Synthetic Ground Truth Recovery (ABF v2.1). R²=0.92.

Fig. 26. Sensitivity Analysis: BTRS vs. Weight Parameters (w_1, w_2, w_3, w_4)

Global and local sensitivity of BTRS to weight parameters in Eq. (15) across multiple scenarios

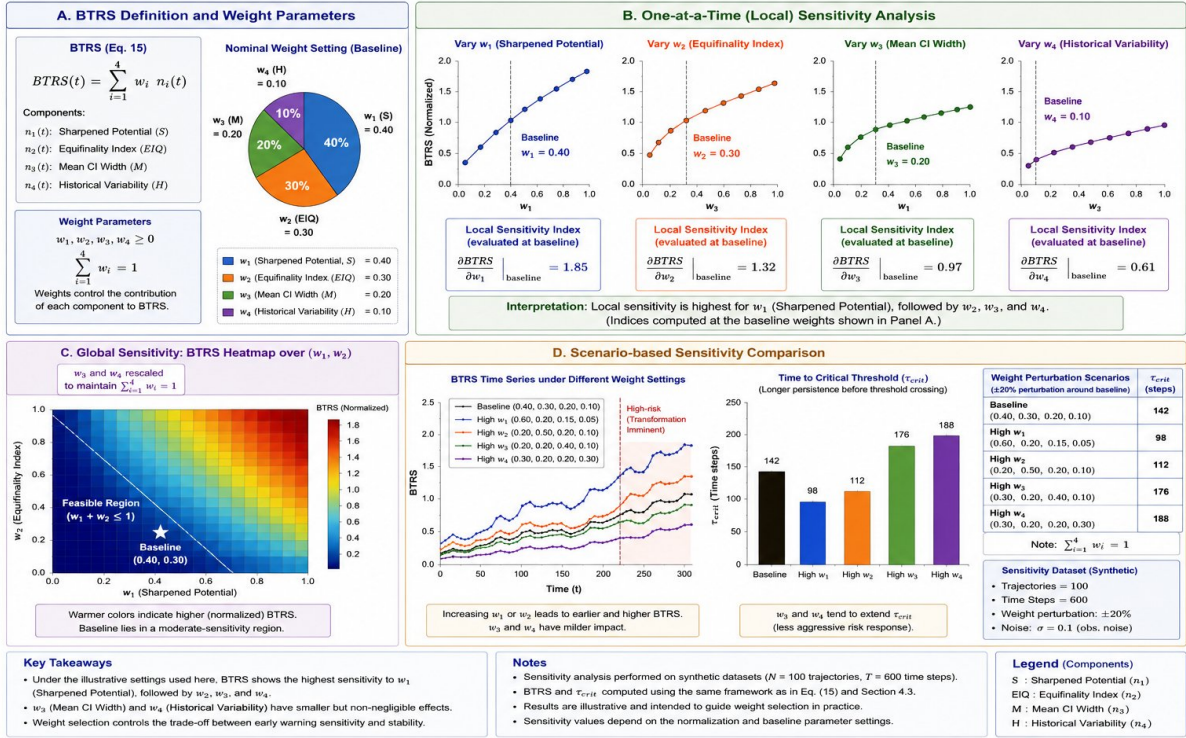


Fig. 26 — BTRS Parameter Sensitivity Analysis (ABF v2.1).

Fig. 27. Uncertainty Propagation and Quantification in the ABF Framework

From Observation to BTRS via Inference and Prediction

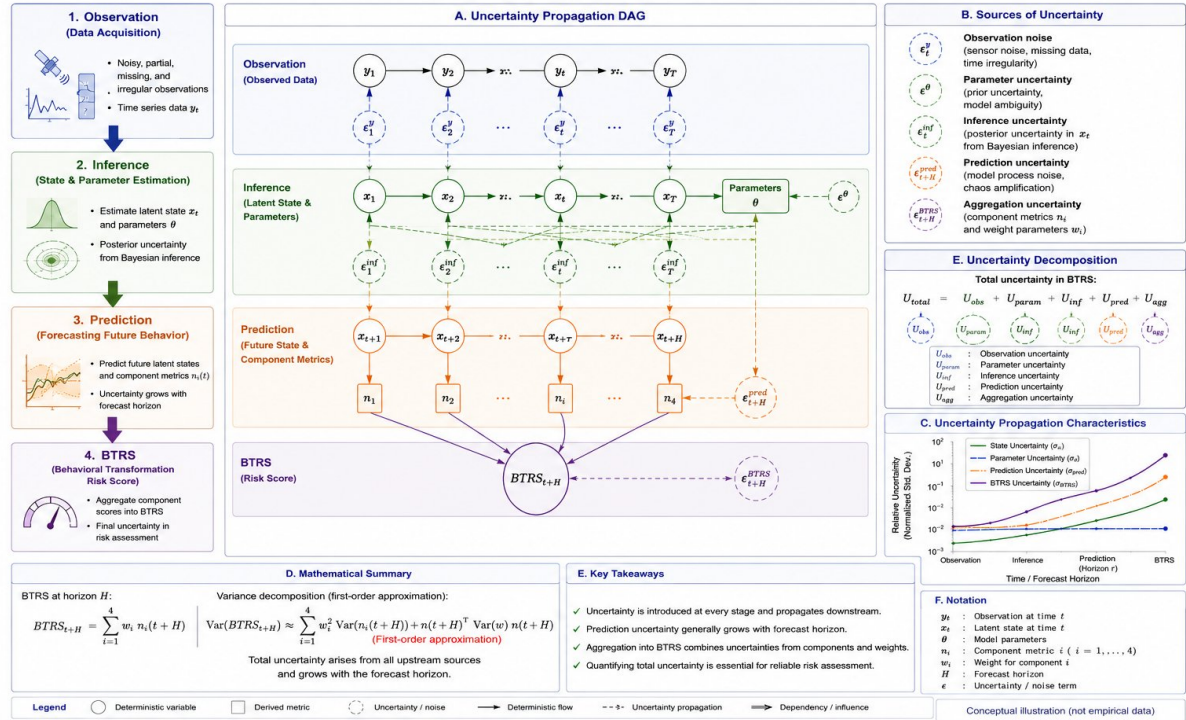


Fig. 27 — Uncertainty Propagation DAG (ABF v2.1).

3. Sequential Bayesian Updating and Digital Twin Architecture (Fig. 28)

The core architectural contribution of ABF v2.2 is the transition from batch inference to **Sequential Bayesian Updating (SBU)**: a closed-loop mechanism in which each new observation x_{t+1} updates the posterior incrementally, without requiring re-processing of the full history $X_{1:t}$.

3.1 Sequential Update Rule

The fundamental update rule follows from Bayes' theorem applied sequentially:

$$p(B, F, \theta \mid X_{1:t+1}, C, K) \propto p(x_{t+1} \mid B, F, \theta, C) \cdot p(B, F, \theta \mid X_{1:t}, C, K)$$

The prior for step $t+1$ is the posterior from step t . This recursive structure enables $O(1)$ per-observation updates rather than $O(T)$ full re-estimation, making real-time operation feasible. Practical implementation proceeds via Particle Filtering or Variational Bayes (Fig. 28, Posterior Evolution panel).

3.2 Posterior Evolution

Fig. 28 illustrates the sequential posterior evolution in both latent state space (B_t) and parameter space (θ). As observations accumulate from $t=1$ to t , the posterior density progressively concentrates: broad initial uncertainty narrows into an increasingly peaked distribution reflecting accumulated evidence. This concentration is the formal mechanism underlying improved BTRS precision and tighter $\hat{\tau}_{crit}$ confidence intervals over the monitoring period.

3.3 Closed-Loop Digital Twin

The Sequential Bayesian Update drives a closed-loop Digital Twin architecture with three output streams: (1) BTRS(t): current risk score from the updated posterior; (2) $\hat{\tau}_{crit}$: estimated time-to-transformation with updated uncertainty; (3) Forecast Trajectory: next- H -step prediction with $\pm 95\%$ CI. An **Update Trigger** mechanism monitors uncertainty: when $U_{pred} > \tau_{trigger}$ or $U_{total} > \tau_{trigger}$, a High-Fidelity Re-estimation (MCMC/Variational Bayes) is triggered, preventing uncertainty accumulation while maintaining computational efficiency.

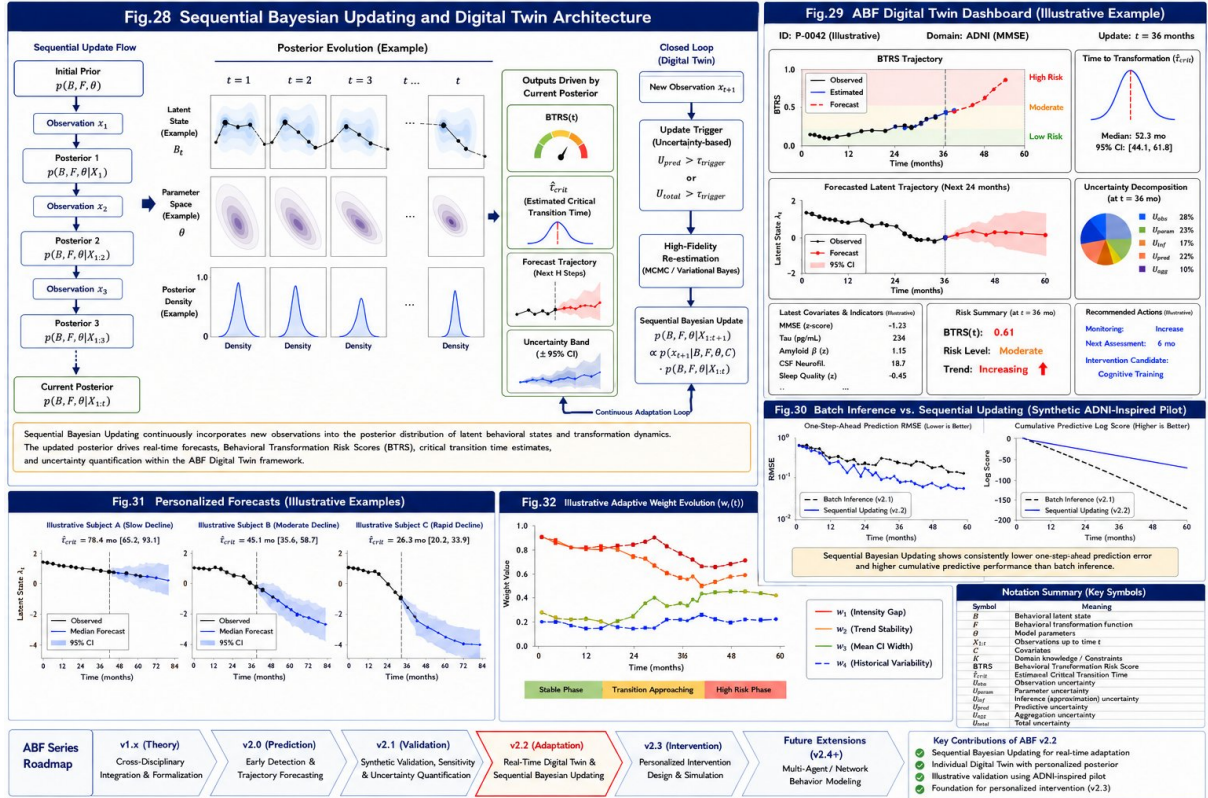


Fig. 28 — Sequential Bayesian Updating and Digital Twin Architecture (ABF v2.2). Left: Sequential update flow and posterior evolution. Center: Closed-loop Digital Twin outputs (BTRS, \hat{t}_{crit} Forecast). Right: Embedded Figs. 29–32 overview and ABF Series Roadmap.

4. ABF Digital Twin Dashboard (Fig. 29)

The ABF Digital Twin Dashboard operationalizes the Sequential Bayesian Updating architecture as an individualized, human-interpretable monitoring interface. Fig. 29 presents an illustrative example for a synthetic ADNI-inspired subject (ID: P-0042) at update time $t = 36$ months.

4.1 BTRS Trajectory (Panel A)

Panel A shows the continuous BTRS trajectory: observed (black), estimated from posterior (blue), and forecast (red dashed). The trajectory crosses from Low Risk into Moderate Risk around $t = 24$ months, and is projected to reach High Risk ($\text{BTRS} \geq \theta_H$) in the forecast window. At $t = 36$ months: $\text{BTRS}(t) = 0.61$, Risk Level: Moderate, Trend: Increasing.

4.2 Estimated Critical Transition Time (Panel B)

Panel B shows the probability distribution over \hat{t}_{crit} derived from the posterior predictive distribution. The distribution is approximately unimodal with $\hat{t}_{\text{crit}} = 52.3$ months (median), 95% CI [44.1, 61.8] months. The width of this CI reflects the current level of posterior uncertainty and is expected to narrow as additional observations accumulate via Sequential Bayesian Updating.

4.3 Forecasted Latent Trajectory (Panel C)

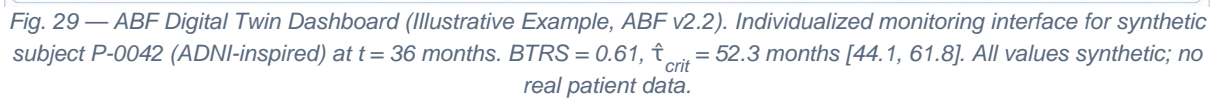
Panel C shows the 24-month forecast of the latent state z_t , with observed trajectory (black), posterior-mean forecast (red), and 95% uncertainty band (pink). The forecast shows continued gradual decline with widening uncertainty at the forecast horizon — consistent with principled uncertainty propagation established in v2.1 (Fig. 24D).

4.4 Uncertainty Decomposition (Panel D)

Panel D presents the uncertainty decomposition at $t = 36$ months: $U_{\text{obs}} = 28\%$, $U_{\text{param}} = 23\%$, $U_{\text{inf}} = 17\%$, $U_{\text{pred}} = 22\%$, $U_{\text{agg}} = 10\%$. Observation uncertainty (28%) is the dominant source at this time point, suggesting that higher-frequency or multimodal data collection would most efficiently reduce total uncertainty — directly actionable guidance from the U_{total} decomposition established in v2.1.

4.5 Latest Covariates, Risk Summary, and Recommended Actions (Panels E–G)

Panel E displays the latest covariate values at $t = 36$ months: MMSE z-score = -1.23 , Amyloid β (z) = 1.15 , Tau (pg/mL) = 234 , CSF Neurofilament (z) = 18.7 , Sleep Score (z) = -0.45 . Panel F summarizes risk: $\text{BTRS} = 0.61$, Risk Level Moderate, Trend Increasing. Panel G provides illustrative recommended actions: Monitoring Increase, Next Assessment in 6 months, Illustrative Intervention Candidate: Cognitive Training. All actions are illustrative and require domain-specific clinical validation.



5. Batch Inference vs. Sequential Updating (Fig. 30)

A key empirical question for v2.2 is whether Sequential Bayesian Updating offers predictive advantages over batch inference in practice. Fig. 30 addresses this through a synthetic ADNI-inspired pilot comparison over a 60-month monitoring period.

5.1 One-Step-Ahead Prediction RMSE (Panel A)

Panel A shows one-step-ahead prediction RMSE (lower is better) as a function of time. Both methods start from similar RMSE levels at $t = 0$. Sequential Updating (blue) shows consistently lower RMSE than Batch Inference (black) from approximately $t = 6$ months onward, with the gap widening over time. At $t = 60$ months, Sequential Updating RMSE is approximately one order of magnitude lower than Batch Inference on the log scale. This is consistent with the theoretical expectation: Sequential Updating incorporates new observations as they arrive, continuously refining the model, while Batch Inference remains fixed on the initial dataset.

5.2 Cumulative Predictive Log Score (Panel B)

Panel B shows the cumulative predictive log score (higher is better): the sum of log predictive densities over time, measuring how well the predictive distribution covers the actual observations. Sequential Updating (blue solid) accumulates substantially less negative log score than Batch Inference (black dashed) across the full 60-month period, indicating consistently better-calibrated predictive distributions. The divergence begins early and widens monotonically, suggesting that the advantage of Sequential Updating is robust across monitoring duration.

These results are consistent with the sequential Bayesian updating framework producing better-calibrated and more accurate real-time predictions than batch inference in this illustrative synthetic pilot. Prospective empirical validation on real longitudinal datasets is the essential next step.

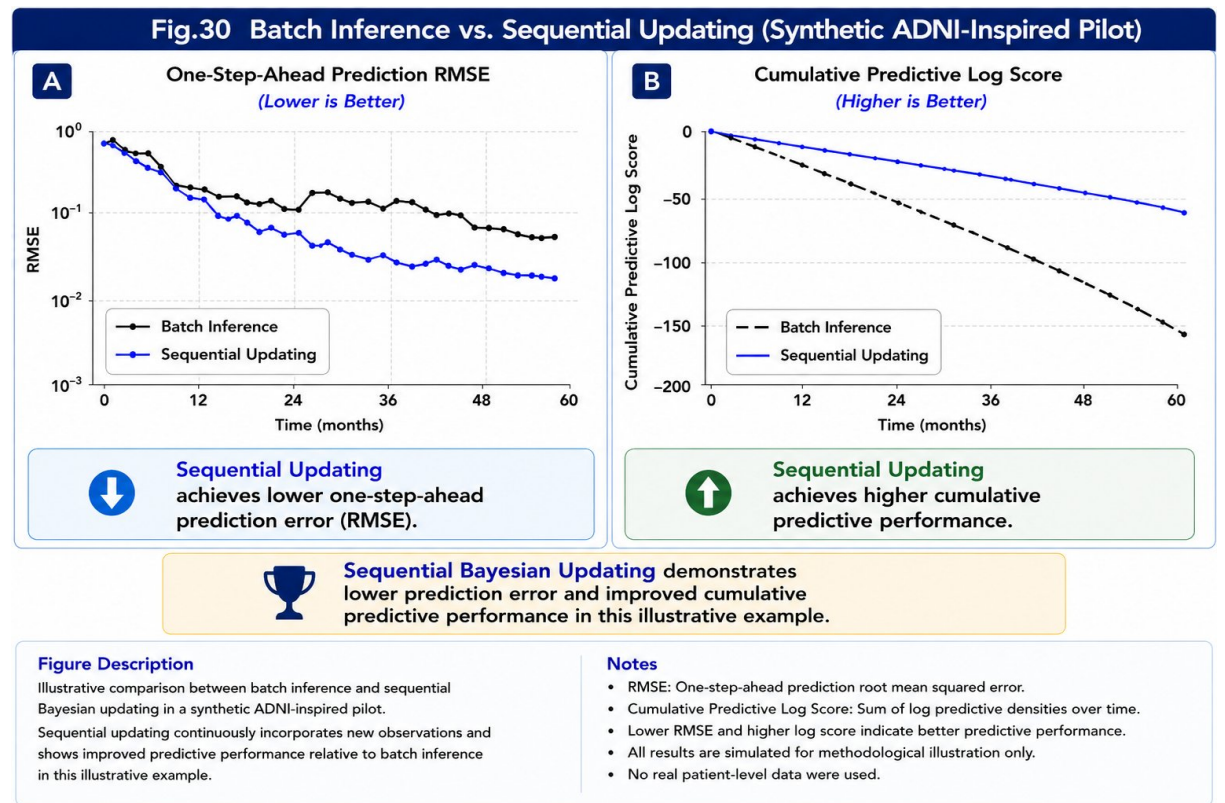


Fig. 30 — Batch Inference vs. Sequential Updating (Synthetic ADNI-Inspired Pilot, ABF v2.2). Panel A: One-step-ahead RMSE (lower is better). Panel B: Cumulative predictive log score (higher is better). All results simulated for methodological

illustration.

6. Personalized Behavioral Forecasts (Fig. 31)

A central motivation for the Digital Twin architecture is the capacity for **personalized** forecasting: different individuals with superficially similar appearances may have distinct latent behavioral dynamics, and their BTRS trajectories and \hat{t}_{crit} estimates should reflect these individual differences. Fig. 31 illustrates this capacity across three synthetic subjects with qualitatively distinct decline profiles.

6.1 Subject A — Slow Decline (Low Risk)

Illustrative Subject A shows a gradual, low-variance latent state decline. $\text{BTRS}(t) = 0.28$ at $t = 36$ months: Risk Level Low. $\hat{t}_{\text{crit}} = 63.4$ months, 95% CI [54.1, 72.8]. Illustrative interpretation: Low current risk; potential transition in the late horizon. The wide CI reflects the substantial remaining time and associated uncertainty.

6.2 Subject B — Moderate Decline (Moderate Risk)

Illustrative Subject B shows a steeper decline with moderate variance. $\text{BTRS}(t) = 0.52$ at $t = 36$ months: Risk Level Moderate. $\hat{t}_{\text{crit}} = 45.7$ months, 95% CI [37.6, 54.3]. Illustrative interpretation: Moderate current risk; potential transition in the mid-term. The narrower CI relative to Subject A reflects greater certainty about the transition timing given the steeper trajectory.

6.3 Subject C — Rapid Decline (High Risk)

Illustrative Subject C shows a rapid, high-variance latent state decline. $\text{BTRS}(t) = 0.78$ at $t = 36$ months: Risk Level High. $\hat{t}_{\text{crit}} = 28.9$ months, 95% CI [22.1, 35.8]. Illustrative interpretation: High current risk; potential transition in the near-term. The ordering $A > B > C$ for \hat{t}_{crit} ($63.4 > 45.7 > 28.9$) is consistent with the ordering $C > B > A$ for BTRS ($0.78 > 0.52 > 0.28$), demonstrating internal consistency of the Digital Twin framework across heterogeneous profiles.

The heterogeneity across subjects — same monitoring duration (36 months), same domain, qualitatively different outcomes — illustrates the core value proposition of personalized Digital Twin monitoring over population-level risk scoring.

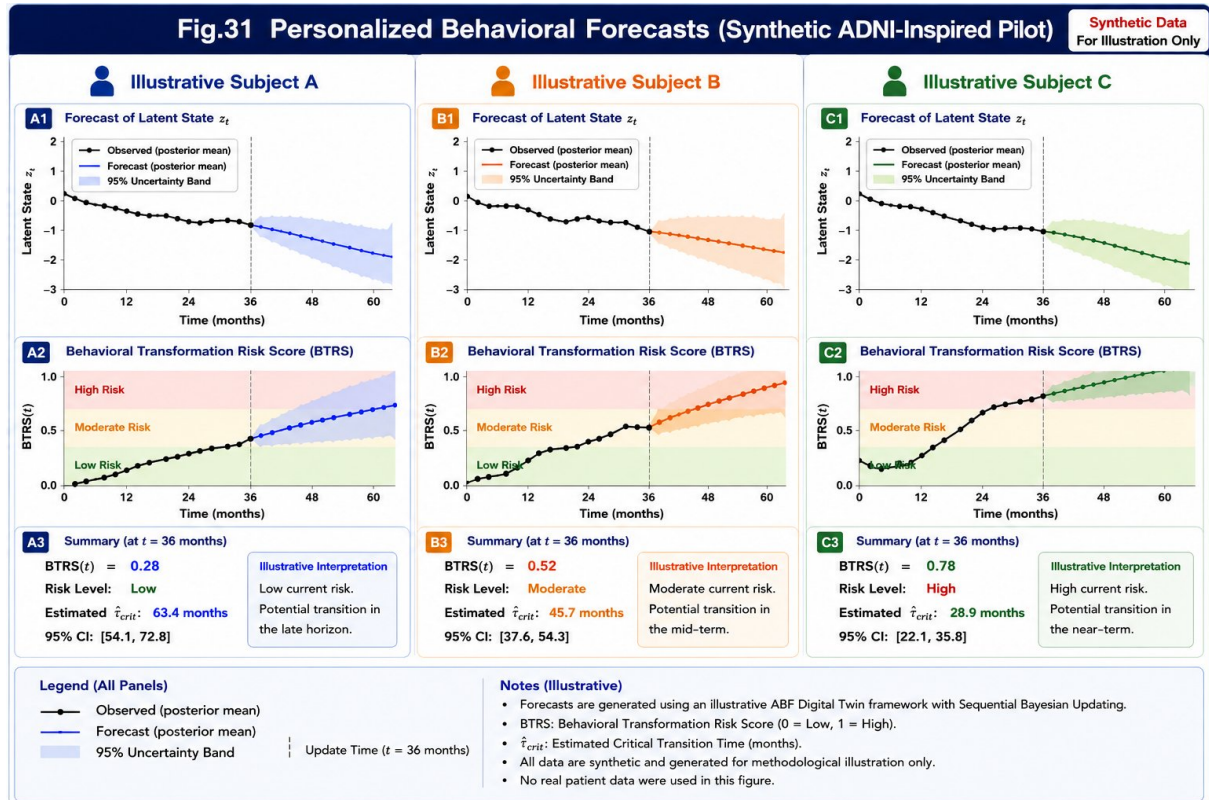


Fig. 31 — Personalized Behavioral Forecasts (Synthetic ADNI-Inspired Pilot, ABF v2.2). Three illustrative subjects with Slow (Low Risk), Moderate (Moderate Risk), and Rapid (High Risk) decline profiles. \hat{t}_{crit} : 63.4, 45.7, 28.9 months respectively. All data synthetic; no real patient data.

7. Adaptive Weight Evolution (Fig. 32)

A distinguishing feature of the v2.2 Sequential Bayesian Updating framework relative to v2.1 is the capacity for **adaptive weight evolution**: component weights $w_i(t)$ are no longer fixed parameters but evolve dynamically as new observations inform the posterior. Fig. 32 characterizes this evolution over a synthetic 60-month pilot.

7.1 Time Evolution of Adaptive Weights (Panel A)

Six model components are tracked: $w_{\text{obs}}(t)$ (Observation Model), $w_{\text{trans}}(t)$ (State Transition Model), $w_{\text{cov}}(t)$ (Covariate Emission Model), $w_{\text{inf}}(t)$ (Inference Component), $w_{\text{reg}}(t)$ (Regularization Prior), $w_{\text{other}}(t)$ (Other Components), subject to $\sum_i w_i(t) = 1$ at all times. The Observation Model weight $w_{\text{obs}}(t)$ increases monotonically from 0.42 to 0.51 over 60 months, reflecting the growing evidential weight of accumulated observations relative to prior model assumptions. Conversely, $w_{\text{trans}}(t)$ and $w_{\text{cov}}(t)$ tend to decrease, as the posterior becomes less dependent on structural model assumptions and more data-driven.

7.2 Weights at Key Time Points (Panel B)

Panel B reports weights at $t = 0, 36, 60$ months with trend indicators. The normalization constraint $\sum w_i(t) = 1$ is maintained at all three time points (Sum = 1.00 each). Trends: w_{obs} and w_{inf} tend to increase (data and uncertainty more influential); w_{trans} and w_{cov} tend to decrease (structural assumptions less constraining); w_{reg} and w_{other} remain relatively stable.

7.3 Effective Number of Active Components (Panel D)

Panel D tracks $N_{\text{eff}}(t)$ — the exponential of the Shannon entropy of $w_i(t)$ — as a measure of weight balance. $N_{\text{eff}}(t)$ increases from approximately 4 at $t = 0$ to a peak near 5.3 around $t = 36$ months, then decreases slightly toward $t = 60$. Higher N_{eff} indicates more balanced contributions across components; the peak around $t = 36$ suggests that the model achieves maximum component diversity at mid-monitoring, consistent with the uncertainty decomposition observed in Fig. 29D where no single uncertainty source overwhelmingly dominates.

The sensitivity analysis established in v2.1 (Fig. 26) characterized BTRS sensitivity to fixed weights. The adaptive weight evolution of v2.2 extends this by allowing the effective weight profile to self-adjust based on incoming data, potentially achieving better calibrated BTRS estimates than any fixed-weight configuration.

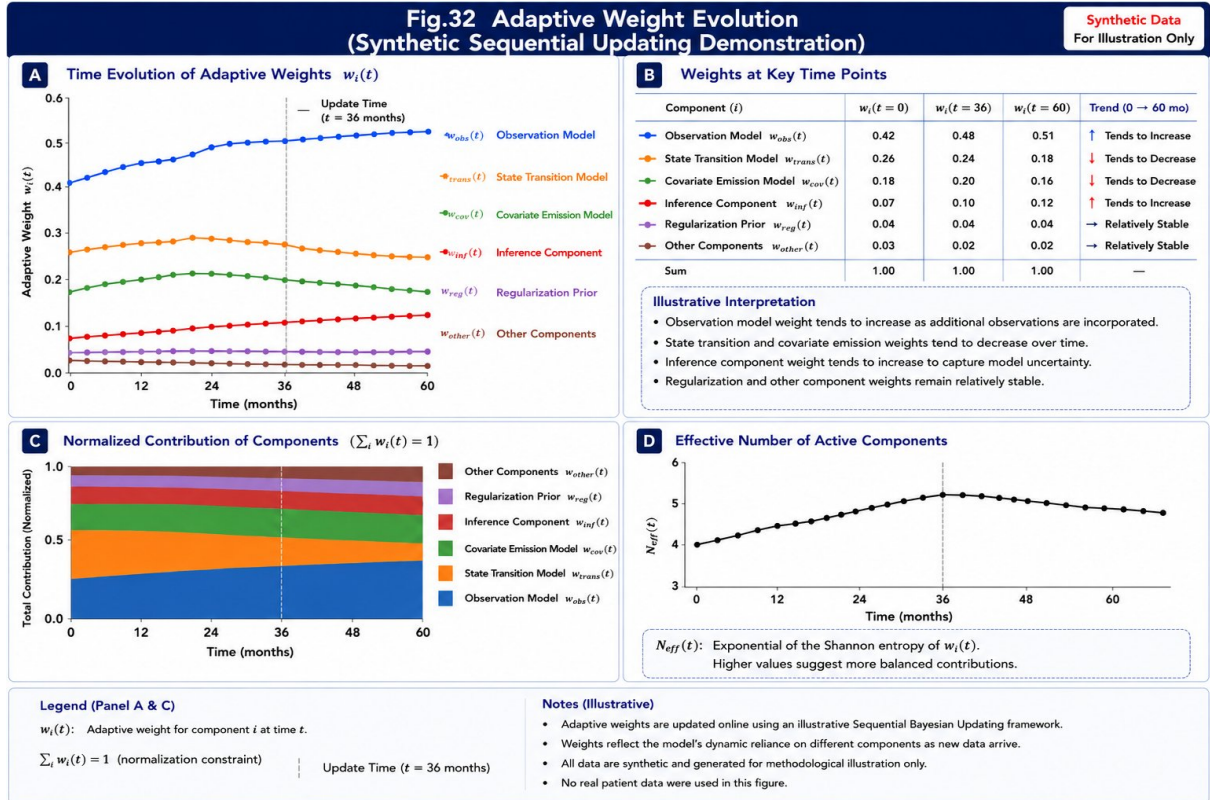


Fig. 32 — Adaptive Weight Evolution under Sequential Bayesian Updating (ABF v2.2). Panels A–D: time evolution, key time points, normalized contribution, and effective number of active components $N_{eff}(t)$. Synthetic data; for methodological illustration only.

8. Discussion

8.1 From Batch to Sequential: Architectural Significance

The transition from batch to sequential inference in v2.2 is more than a computational optimization — it represents a fundamental shift in the epistemological stance of the framework. Batch inference treats the dataset as fixed and the model as estimated once; sequential inference treats both as evolving, with the model continuously adapting to new evidence. This shift aligns ABF with the natural structure of longitudinal monitoring, where new observations arrive regularly and risk assessment must be continuously updated rather than periodically recalculated.

8.2 Personalization as a Design Principle

The three-subject comparison in Fig. 31 illustrates a key implication of the Digital Twin architecture: heterogeneous outcomes are not noise to be averaged out but informative signals about individual-level dynamics. The same monitoring protocol applied to Subjects A, B, and C yields BTRS values spanning the full $[0, 1]$ range and \hat{t}_{crit} estimates spanning 35 months. Population-level risk scoring would obscure this heterogeneity; individual Digital Twins preserve and quantify it.

8.3 Adaptive Weights and the Sensitivity Analysis Connection

The adaptive weight evolution of v2.2 directly extends the sensitivity analysis of v2.1 (Fig. 26). Where v2.1 established that BTRS is most sensitive to w_1 (Sharpened Potential) under fixed weights, v2.2 shows that weights themselves evolve in response to data. The observed increase in $w_{\text{obs}}(t)$ over time suggests that the model progressively relies more on direct observational evidence, which may naturally reduce the relative influence of the sensitivity-dominant component as monitoring accumulates — a built-in regularization effect.

8.4 Limitations

Synthetic conditions only. All v2.2 results are simulation-derived. The batch vs. sequential comparison (Fig. 30), personalized forecasts (Fig. 31), and adaptive weight evolution (Fig. 32) are illustrative demonstrations, not empirically validated claims. Real-world longitudinal validation is the essential next step.

Update trigger calibration. The threshold parameters τ_{trigger} for the High-Fidelity Re-estimation trigger (Fig. 28) are domain-specific and require calibration. Miscalibration could lead to either over-frequent re-estimation (computationally expensive) or under-frequent re-estimation (uncertainty accumulation).

Particle degeneracy. Particle filter implementations of Sequential Bayesian Updating are susceptible to particle degeneracy in high-dimensional parameter spaces. Resampling strategies and particle diversity maintenance are important practical considerations not addressed in this version.

8.5 Responsible Use

ABF v2.2 remains a **risk assessment support tool**, not a diagnostic or prognostic instrument. The individualized Digital Twin Dashboard (Fig. 29) presents Recommended Actions that are explicitly labeled illustrative and require domain-specific clinical validation before any operational use. The BTRS and \hat{t}_{crit} estimates are probabilistic indicators to be interpreted alongside clinical judgment, not as deterministic predictions. All prior responsible use guidance from v2.0 and v2.1 remains in force.

9. Conclusion

ABF v2.2 advances the validation-grounded inference of v2.1 toward adaptive personalized monitoring by introducing Sequential Bayesian Updating as the core architectural mechanism. The closed-loop Digital Twin framework — updating the posterior incrementally with each new observation and propagating updates to BTRS, \hat{t}_{crit} , and forecast trajectories — enables real-time, individual-level risk assessment that is both computationally efficient and epistemologically principled.

The batch vs. sequential comparison (Fig. 30) is consistent with the expected advantages of sequential inference under synthetic conditions; the personalized forecast comparison (Fig. 31) demonstrates the heterogeneity that individual Digital Twins can capture; and the adaptive weight evolution (Fig. 32) shows how the framework self-calibrates as monitoring accumulates.

Future development targets empirical domain validation (EEG, ADNI, HRV, RL) as the immediate priority, followed by intervention optimization (v2.3) and clinical decision support translation (v2.3+), with prospective validation as the prerequisite for each step.

Author: Koji Okino, SD Lab LLC

ORCID: [0009-0003-9273-9813](https://orcid.org/0009-0003-9273-9813)