ChatGPT

Cross-Domain Modeling Techniques for Behavioral AI: Finance, Evolutionary Biology & Astrophysics

Introduction

Models such as Markov chains, Hidden Markov Models (HMMs), state-space models, stochastic differential equations, reinforcement learning and graph/agent-based models are widely used to model complex systems. These techniques were developed for diverse domains – from predicting stock-market regimes to identifying evolutionary patterns and modelling orbital dynamics. This report investigates how these same models can be repurposed to understand and predict psychological tipping points in behavioral AI. By exploring the cross-domain literature, we identify the strengths of each technique and show how they can be stacked into a layered system for behavioural prediction.

Why Cross-Domain Methodology Matters

- **Transferable theory** – The mathematical structures underlying Markov chains, Kalman filters or RL algorithms are domain-agnostic. A state-transition probability in a stock market regime model is mathematically equivalent to a state transition in a mental state model. Therefore the knowledge gained from finance or astrophysics can inform mental-health modelling.
- **Shared challenges** – Complex systems in finance, biology and astrophysics are high-dimensional, non-linear and noisy. Techniques developed to handle non-stationary financial time series or wandering astrophysical signals may help detect subtle psychological shifts before they become crises.
- **Evidence of cross-application** – Hidden Markov Models originally developed for speech recognition were later used to map genome sequences and to detect gravitational-wave frequency drift. Reinforcement learning developed for robotics is now used for stock-trading and astronomical scheduling. These examples show that modelling tools are portable when there is a similar underlying problem structure.

The sections below summarise how each modelling technique has been used in the three domains (finance, evolutionary biology, astrophysics) and discuss how they can be re-imagined for behavioural AI. A final section proposes a layered architecture combining these models.

Hidden Markov Models and Regime-Switching

Finance

In financial markets, returns often oscillate between bullish (upward), calm and bearish (downward) regimes. A study of Middle-Eastern and North African (MENA) markets used a Hidden Markov Model to relate a **financial stress index** to stock-market dynamics. The authors estimated transition probabilities between regimes and found that the bullish state often persists longer than other states; the probability of transitioning from bullish to calm was low, and mean returns indicated that bearish and calm regimes may be more attractive for risk-averse investors [1]. Another study trained a Hidden Markov Model on historical closing prices to predict future prices and evaluated performance using error metrics and directional prediction accuracy [2]. These works show that HMMs can reveal latent market conditions and predict regime shifts.

Evolutionary Biology

In genomics, researchers used custom HMM profiles to analyse **prokaryotic genome space**. By scanning ~11 000 species with sensitive HMMER profiles, they traced the distribution of NADH-quinone oxidoreductase subunits (Complex I) and found that about **51 %** of species possessed a complete Complex I, indicating multiple evolutionary variants [3]. HMMs thus provide a probabilistic way to annotate genes and infer hidden evolutionary states.

Astrophysics

Continuous gravitational-wave signals from spinning neutron stars may experience **spin wandering** due to stochastic variations in rotational frequency. A 2025 study compared several search algorithms – including a **hidden Markov model solved by the Viterbi algorithm** – for detecting these signals. The HMM approach models the gravitational-wave frequency as a stochastically wandering hidden variable and recovers the signal by tracking the most probable sequence of frequencies [4]. The study demonstrates that HMMs can efficiently handle stochastic frequency drift in astrophysical data.

Implications for Behavioral AI

Mental health data often exhibit sudden shifts (e.g., relapse vs. remission). HMMs can model latent psychological states (stress, calm, manic) and transitions between them. Techniques for estimating transition matrices in finance or astrophysics can be applied to **behavioural regime-switching**. For example, a high probability of remaining in a manic state could signal a need for more intensive intervention, analogous to the persistence of a bullish market.

State-Space Models, Kalman Filters & Stochastic Differential Equations

Finance

Regime-switching is often combined with **state-space models** to handle noisy observations and unobserved factors. State-space models use a latent state vector that evolves over time and generate observable data through a measurement equation. Kalman filters and their extensions (e.g., particle filters) provide optimal estimates when noise is Gaussian. Financial economists employ these filters to extract hidden volatility or macroeconomic factors from asset prices.

Evolutionary Biology

Ecological systems are dynamic and subject to noise. Multivariate state-space models have been used to study fish populations and ecological time series. Stochastic differential equations (SDEs) offer another continuous-time framework; the drift and diffusion terms capture deterministic trends and random fluctuations. For example, Ornstein–Uhlenbeck models describe how trait values revert to an optimum while experiencing random shocks.

Astrophysics

Orbit determination is a classical state-space problem. A 2025 preprint proposed modelling orbital dynamics as a **stochastic system** and solving it with **sequential filters** (e.g., Kalman and particle filters). The authors argued that linearised models assuming Gaussian errors are insufficient for highly nonlinear orbital motion. By adopting stochastic numerical integrators and sequential filters, they simultaneously estimated the process noise and achieved accurate tracking of low-Earth orbit objects in simulated scenarios [5] . This demonstrates that state-space approaches combined with stochastic integrators can handle non-linear, noisy dynamics.

Implications for Behavioral AI

Human behaviours—mood, attention or relapse risk—are dynamic and noisy. State-space models can describe these processes as latent variables evolving over time, with ecological momentary assessments (EMA) serving as observations. Kalman or particle filters used to track orbits or financial volatility can be used to estimate a person's underlying state. SDEs (e.g., Ornstein–Uhlenbeck processes) can model the tendency to return to an equilibrium mood while capturing random shocks (stressful events). Identifying **process noise** and **measurement noise** will enable better forecasts of psychological tipping points.

Reinforcement Learning & Sequence Models

Finance

Reinforcement learning (RL) has become popular in algorithmic trading. A 2025 PeerJ Computer Science article proposed an **actor–critic RL model** combined with an adaptive data window to trade commodities like gold, Euro and crude oil. The RL agent uses a neural network to estimate policy and value functions; the proposed approach achieved average loss reductions of **0.03 %** (Euro), **0.25 %** (gold) and **0.13 %** (crude oil) in the test phase [6] . The authors note that RL is attractive because it allows agents to learn directly from interactions with the market and adapt to constant fluctuations; prior studies show RL **surpasses traditional machine-learning approaches** by optimizing policies to maximize profits while minimizing losses in real time [7] .

Evolutionary Biology

Evolutionary processes involve adaptation driven by rewards (fitness). Evolutionary game theory and RL share conceptual similarities: an agent learns to maximise a fitness function through trial and error. RL algorithms have been used to model foraging behaviour, predator–prey dynamics and even protein folding. They capture the notion that organisms adapt strategies based on environmental feedback, akin to RL agents.

Astrophysics

Astronomical observatories must allocate observation time optimally. The QuarkNet project described a reinforcement-learning program that **optimizes observation schedules** for celestial objects by receiving the states of objects and allocating time based on RL policies. Neural networks built using TensorFlow learn from data, with reward functions defining the objectives (e.g., maximizing useful observations). The project aims to incorporate RL into **surveying strategies for large astronomical projects** [8]. This shows that RL can handle scheduling problems in high-stakes environments like astrophysics.

Sequence Models

Long-term dependencies often span across time and modalities. Sequence models such as long short-term memory (LSTM) networks and transformers have been used to predict human decisions and astrophysical time series. These models learn complex patterns without assuming Markovian dynamics. They can be integrated with RL (e.g., deep RL) for improved policy learning.

Implications for Behavioral AI

RL provides a framework for adaptive interventions: an agent interacts with a person's digital environment and chooses interventions based on reward signals (improvements in mood, reduced risk). Sequence models can process long histories of mood data, detect precursors to tipping points and feed this information into RL agents. For example, the actor–critic architecture used for trading could be adapted to schedule therapy sessions or deliver behavioural nudges.

Graph & Agent-Based Models

Graph neural networks (GNNs) and agent-based models (ABMs) capture interactions among entities. In finance, ABMs represent heterogenous traders whose local rules can produce market crashes or bubbles. In evolutionary biology, ABMs simulate the interactions of individuals and their environment to explain emergent phenomena. In astrophysics, networks represent relationships among celestial bodies or sensors.

GNNs generalise neural networks to graph-structured data and allow information to propagate across nodes. They can integrate multimodal data (e.g., EEG and speech) for mental-health detection. ABMs are particularly powerful for modelling social behaviours where individuals influence each other; such models capture heterogeneity and emergent patterns not accessible to equation-based models.

Information-Theoretic Approaches

Information theory underlies statistical modelling by quantifying **uncertainty** and **mutual information** between variables. In cognitive neuroscience, mutual information between tasks and neural activation helps decompose sensorimotor, contextual and episodic components [9]. In astrophysics and finance, information theory informs trading strategies and detection thresholds. Integrating information-theoretic measures into behavioural AI can help quantify complexity, detect early warning signals (rising entropy) and optimise information flow across model layers.

Layered Behavioral AI System

We propose a **layered system** (diagram below) that stacks the modelling approaches used in finance, evolutionary biology and astrophysics into a unified behavioural AI framework:

1. **Domain Knowledge Layer** – Use domain-specific insights (finance, biology, astrophysics) to inspire modelling strategies. For example, regime-switching in markets suggests looking for hidden mood states in humans.
2. **Modelling Layer** – Apply models appropriate for the behavioural data: Hidden Markov models for latent states, state-space models and stochastic differential equations for continuous processes, RL & sequence models for adaptive decision-making, and graph & agent-based models for social interactions.
3. **Information Layer** – Use information-theoretic measures to quantify uncertainty, detect tipping points and guide model selection.

This layered architecture is illustrated in Figure 1 below. The arrow directions indicate how knowledge from each domain feeds into specific modelling approaches; the models then propagate information down to the information-theory layer, which assesses uncertainty and guides decision-making.

Cross-Domain Summary Table

Domain	Modelling technique	Summary of key findings
Finance	Hidden Markov models; state-space models; reinforcement learning	HMMs uncover latent bullish/calm/bearish regimes and show persistence of regimes [1]. RL trading agents using actor-critic networks reduce average losses in currency and commodity markets [6] and adapt to market fluctuations better than traditional ML [7].
Evolutionary Biology	Hidden Markov models; SDEs; agent-based models	Custom HMM profiles map the distribution of Complex I subunits across ~11,000 species, revealing that ~51 % possess complete complexes [3]. Agent-based models simulate interactions and heterogeneity to explain emergent behaviours in ecological systems.
Astrophysics	Hidden Markov models; stochastic state-space models; reinforcement learning	HMM-Viterbi algorithms track stochastic frequency variations in continuous gravitational-wave signals [4]. Sequential filters using stochastic numerical integrators accurately track orbital dynamics and process noise [5]. RL is used to optimise observation schedules for astronomical surveys [8].

Discussion: Re-purposing Models for Psychological Tipping Points

1. **Similarity of Dynamics** – Financial markets, biological populations and astrophysical systems exhibit **regime shifts**, **non-linear dynamics** and **noise**. Mental-health trajectories similarly alternate

between stable periods and sudden transitions (relapse or improvement). Models successful in those domains can thus inform psychological predictions.
2. **Latent States & Regimes** – HMMs used to detect market stress or gravitational-wave frequencies can be applied to detect latent psychological states from wearable data or self-reports. Transition probabilities estimated in financial studies may guide interventions (e.g., if the probability of transitioning from moderate stress to high stress is high, proactively deliver support).
3. **Stochastic Trajectories** – State-space models and SDEs developed for orbit determination or trait evolution can model continuous mood trajectories. Kalman/particle filters can produce real-time estimates of underlying mental states and forecast tipping points.
4. **Adaptive Decisions** – RL agents that trade stocks or schedule telescopes can be repurposed as **therapeutic intervention schedulers**, learning to deliver support at the right moment and adapt to individual responses. Sequence models can capture long-term dependencies in behaviour data and inform RL policies.
5. **Networks & Emergence** – Graph and agent-based models used in social networks, ecology and finance can simulate how peer influence or social support shapes mental health. Incorporating heterogeneity and emergent effects is crucial for designing community-level interventions.
6. **Information Theory** – Measuring entropy and mutual information across behavioural features can detect increasing complexity or unpredictability preceding a tipping point. Information-theory metrics can also guide which model layer (HMM, SDE, RL) is most informative for a given individual.

Conclusion

Modelling techniques originally developed for finance, evolutionary biology and astrophysics offer powerful tools for behavioural AI. Hidden Markov models reveal latent regimes; state-space models, Kalman filters and stochastic differential equations capture continuous dynamics; reinforcement-learning agents make adaptive decisions; graph and agent-based models account for social interactions; and information-theory measures quantify uncertainty. These techniques can be **stacked** into a layered system where domain knowledge informs model selection and information-theory metrics regulate decision-making. Repurposing cross-domain models will help predict **psychological tipping points**, optimize interventions and deepen our understanding of human behaviour.

---

[1] The Predictive Power of Financial Stress on the Financial Markets Dynamics: Hidden Markov Model | Journal of Economics and Finance
https://link.springer.com/article/10.1007/s12197-022-09600-z

[2] [2310.03775] Hidden Markov Models for Stock Market Prediction
https://arxiv.org/abs/2310.03775

[3] Hidden Markov Model-Based Prokaryotic Genome Space Mining Reveals the Widespread Pervasiveness of Complex I and Its Potential Evolutionary Scheme - PMC
https://pmc.ncbi.nlm.nih.gov/articles/PMC12361115/

[4] 2504.08163
https://arxiv.org/pdf/2504.08163

[5] 2504.07515
https://arxiv.org/pdf/2504.07515

[6] [7] Stock market trading via actor-critic reinforcement learning and adaptable data structure - PMC
https://pmc.ncbi.nlm.nih.gov/articles/PMC11888913/

[8] Reinforcement Learning in Astrophysics
https://quarknet.fnal.gov/fnal-uc/quarknet-summer-research/QNET2018/Nicholas-Anne.shtml

[9] Information Theory and Cognition: A Review - PMC
https://pmc.ncbi.nlm.nih.gov/articles/PMC7513233/