Exploring Stress-Induced Neural Circuit Remodeling through Data-Driven Analysis and Artificial Neural Network Simulation

Feng Lin *[a]

[a] *Institute of blood transfusion, Chinese academy of medical science and Peking union medical college, No.26 Huacai Road, Chengdu, 610052, People's Republic of China*

*E-mail: linfeng@ibt.pumc.edu.cn

ABSTRACT: Chronic stress induces neural circuit remodeling, yet the computational principles linking these transitions to behavioral rigidity remain elusive. Performing a multi-level re-analysis of in vivo GCaMP8s recordings from BLA-DMS and CeA-DMS projections, we first quantified stress-induced shifts in stochastic state space using Kullback-Leibler (KL) divergence. To characterize the underlying recovery physics, we developed a phenomenological dynamical model, which revealed that the circuit's response to perturbations is governed by a robustly overdamped backbone, stress significantly biases the response toward increased damping with a concomitant reduction in autonomous restorative precision. To further explore the mechanistic origin of these shifts, we constructed an artificial neural network (ANN) as a proof-of-principle framework. By abstracting the pathways into functional sub-networks, the model successfully reproduces the experimental observation. Our results suggest that stress-induced remodeling emerges as a global shift in system gain. Furthermore, robustness tests reveal a functional–structural decoupling, where stable dynamical attractors persist despite substantial parameter degeneracy in coupling weights (CV $\approx$ 40%). This framework provides a quantitative computational perspective on how altered dynamical regimes emerge from a reconfigured, stable landscape.

KEYWORDS: Amygdala-Striatal Circuits; Chronic Stress; Artificial Neural Network Simulation; Cognitive Neuroscience; Data-Driven Modeling

1 Introduction

Establishing a precise mapping between psychological states and their underlying neural substrates remains a central tenet of cognitive neuroscience [1, chap. 1, p. 1-17]. Psychological processes arise from the dynamics of specific neural circuits, and external behavioral patterns therefore serve as observable readouts of internal neural computations[2]. By quantitatively analyzing the neural signals associated with these behaviors, it becomes possible to reverse-engineer the brain's operational logic and computational strategies across various psychological states [3]. Furthermore, this mind-brain-behavior relationship is inherently dynamic: environmental factors and constraints often drive shifts in psychological states, triggering a modulation of neural circuits that results in adaptive behavioral changes.

This dynamic mapping is particularly evident in the context of stress, a potent environmental constraint that precipitates significant shifts in both psychological states and behavioral strategies [4]. Chronic stress induces a transition from flexible goal-directed behavior to rigid habitual responding [5]. While classical models of habit formation emphasize the functional transition from the goal-directed dorsomedial striatum (DMS) to the habit-mediating dorsolateral striatum (DLS) [6], recent evidence indicates that chronic stress reconfigures the functional balance of amygdala inputs within the DMS circuit [5]. Specifically, Giovanniello et al. (2025) demonstrated that chronic stress differentially modulates two amygdala projections to the DMS. Reward-related activity in the BLA-DMS pathway, which supports action–outcome learning, is suppressed by stress, whereas the CeA-DMS pathway becomes progressively engaged during learning and mediates the formation of inflexible habits.

In the experimental paradigm of Giovanniello et al. (2025), mice were first assigned to control or stress groups (8-12 mice per group). The stress group underwent 14 days of chronic mild unpredictable stress, consisting of pseudorandom exposure to several mild stressors, including footshock, restraint, white noise, tilted cages, damp bedding, and continuous illumination. After the stress period, mice were trained in an operant lever-press task to obtain food rewards across multiple reinforcement schedules. Following training, outcome-based behavioral tests were performed to assess the balance between goal-directed and habitual control. Throughout the experiment, GCaMP8s fiber photometry was used to record calcium signals in the BLA-DMS and CeA-DMS projections.

Beyond experimental approaches, computational modeling has become a tool for investigating neuronal dynamics and calcium signaling processes. Mathematical and machine learning models provide a framework for capturing nonlinear neural interactions and predicting neuronal responses under varying conditions[7, 8].

Particularly, recent advancements have demonstrated that artificial neural networks can effectively resolve the complexities of local calcium signaling, offering a robust platform for integrating discrete experimental observations into continuous dynamical frameworks[7]. These approaches have increasingly contributed to the mechanistic understanding of complex neural systems [9].

From the perspective of dynamical systems theory, neural circuits must maintain a delicate balance between stability and flexibility. Stable dynamics enable reliable behavioral output, whereas flexible dynamics allow rapid adaptation to environmental changes [10, 11]. Theoretical studies in complex systems suggest that adaptive computation may emerge near the boundary between ordered and chaotic regimes, often referred to as the "edge of chaos"[12]. In this regime, systems retain sufficient stability to avoid runaway dynamics while maintaining the sensitivity required for flexible responses. Chronic stress may perturb this balance by reshaping the dynamical structure of neural circuits, potentially shifting them toward more rigid attractor states that favor habitual behavior. Although chaotic dynamics have been proposed as a potential source of adaptability in neural systems, the limited duration of the current recordings prevents reliable characterization of global chaos indicators. Therefore, in this study we focus on local dynamical stability as a measurable proxy for circuit adaptability.

In parallel, neural circuits can also be viewed as adaptive learning systems in which synaptic weights are continuously updated in response to environmental statistics. This view is closely related to the principles underlying artificial neural networks (ANNs), where adaptive learning emerges through iterative adjustments of internal parameters [13, 14, 15, 16]. From this perspective, the experimental design of Giovanniello et al. [5] can be interpreted as exposing mice to a particular statistical environment. In other words, the chronic stress paradigm effectively functions as a specific "training set": the mice are not merely exposed to stressors, but their neural circuits are repeatedly shaped by the statistical structure of a high-pressure environment, leading to long-term reconfiguration of circuit dynamics. This process aligns with emerging mechanistic insights suggesting that systemic stressors can leave persistent "memory traces" within neural stochastic landscapes, fundamentally altering the trajectory of neurodynamic recovery [9].

Despite demonstrating a quantitative dynamical association between chronic stress, pathway signal alterations, and behavioral rigidity, the precise computational mechanisms–specifically how these signals dynamically reconfigure neural circuits at the network level–remain unclear. Traditional methods, such as direct signal observation and behavioral analysis, do not capture the full complexity of this reorganization, thus limiting our understanding of its underlying logic. To address this gap, we analyze the publicly available GCaMP8s calcium fluorescence time-series dataset reported by Giovanniello et al. (2025). Using these data, we develop a multi-level analytical framework combining statistical distribution analysis, dynamical systems modeling, and artificial neural network simulations. This approach allows us to examine how chronic stress reshapes the dynamical organization of amygdala–striatal circuits and how such changes may give rise to distributional hysteresis in neural activity.

The analytical pipeline of this study follows a hierarchical modeling strategy. First, statistical analysis based on KL divergence quantifies distributional deviations in neural activity. Second, a dynamical systems model is employed to identify the mechanistic origin of these deviations in terms of altered stability and coupling. Finally, an artificial neural network toy model is constructed to test whether such dynamical signatures can emerge from computational trade-offs between precision tracking and robustness buffering.

2 Method

2.1 Data Source and Event Classification

We performed a secondary analysis on GCaMP8s fiber photometry recordings (sampling rate: 20 Hz) originally reported by Giovanniello et al. [5]. Signals were recorded from BLA–DMS and CeA–DMS projections in mice under two experimental conditions: Control and Chronic Stress (14 days). For each behavioral event, we extracted ± 10 s time windows to characterize peri-event neural activity.

For analytical purposes, behavioral events were categorized into two functional modalities according to whether the neural responses were associated with learned instrumental contingencies or non-contingent stimuli:

1. Non-Learned Modality (Non-Contingent Stimuli): These events occur independently of the mouse's actions. The primary event in this category is the aversive footshock stimulus delivered outside the instrumental task context. Because it is not contingent on ongoing behavior, this stimulus provides an externally imposed perturbation for examining neural responses under stress conditions.

2. Learned Modality (Contingent Events): This modality includes task-related events during the operant lever-press paradigm in which mice learned to associate lever pressing with food rewards. Task difficulty was manipulated using Random Ratio (RR) schedules, in which reward probability gradually decreased from 1.0 (FR1) to 0.1 (RR10). These schedules introduce increasing outcome uncertainty during learning. In addition, unpredicted rewards were occasionally delivered during the task. Although these rewards were not contingent on the mouse's immediate action, they occurred within a context where mice had already learned strong food–reward associations.

In the present study, footshock events from the non-learned modality were used as the primary dataset for dynamical analysis. Because these stimuli occur independently of instrumental behavior, they provide a relatively controlled perturbation for examining stress-related changes in neural activity. Data from the learned modality were subsequently analyzed as complementary evidence to evaluate whether the patterns captured by the proposed model are also consistent with neural dynamics observed during goal-directed task performance.

2.2 Data analysis

2.2.1 Statistical Perspective: Measuring Distributional Hysteresis via KL Divergence

To investigate latent state changes following perturbations, we quantified distributional hysteresis, defined as the persistent divergence of a signal's probability density function (PDF) from its baseline distribution, even when first-order statistics, such as the mean, have largely returned to pre-stimulus levels. The baseline distribution was estimated from the pre-stimulus period pooled across trials. We used the Kullback–Leibler (KL) divergence (D_{KL}) to measure the magnitude of this distributional shift over time. KL divergence was chosen because it provides a sensitive measure of distributional deformation beyond first-order statistics and has been widely used to quantify deviations between stochastic states.

In the present study, we use the term "distributional hysteresis" in an operational rather than a strictly physical sense, avoiding the implication of a strict thermodynamic analogy. In classical physics, hysteresis refers to path-dependent state transitions that form characteristic loops in the input–output relationship. Here, we extend the concept to the statistical domain: when the distribution of neural activity fails to return to its pre-stimulus baseline after the external perturbation has ceased, the system retains a residual distributional displacement. We refer to this persistent deviation as distributional hysteresis.

While D_{KL} is strictly an informational metric, we adopt the principle that persistent shifts in neural stochastic dynamics indicate incomplete relaxation toward the prior baseline statistical regime, which we interpret as an informational indicator of perturbation-induced circuit state changes. In the context of neural circuits, a sustained D_{KL} indicates that the system has not relaxed to its prior baseline statistical regime but has instead transitioned to a statistically distinct stochastic sub-state. We use this "distributional hysteresis" to characterize the persistence of stress-induced statistical deviations in circuit activity, representing a persistent informational deviation in the circuit's dynamical state space relative to its baseline distribution.

We reconstructed the PDFs ($p_{baseline}$ and p_t) of GCaMP8s fluorescence signals using Gaussian Kernel Density Estimation (KDE). To capture the dynamic evolution of the pathway's state, we utilized a sliding

window approach (window size = 1 s, step = 0.05s). PDFs were estimated using samples pooled across trials within each window to ensure sufficient sampling density for KDE estimation. The D_{KL} at each time step t is defined as:

$$D_{KL}(p_t||p_{baseline}) = \int_{-\infty}^{\infty} p_t(x) \ln \left(\frac{p_t(x)}{p_{baseline}(x)} \right) dx \tag{1}$$

where x represents the normalized neural signal intensity. While the fundamental shifts in the distribution's structure are reflected in shifts of lower-order moments such as mean displacement (ΔMean) and variance scaling (ΔStd) (see Supplementary Materials Table D8 for empirical values), the D_{KL} serves as an integrated metric that captures the total magnitude of these statistical deformations across the entire probability density landscape. The total accumulated information (I) shift over the recovery period is integrated as:

$$I = \sum_{t \in T_{recovery}} D_{KL}(p_t||p_{baseline}) \tag{2}$$

This integrated metric, which we term "distributional hysteresis", summarizes the cumulative magnitude of distributional deviations over the recovery period. Conceptually, this quantity reflects the persistence of stochastic state displacement relative to baseline activity. A resilient system dissipates this deviation and returns to its pre-stimulus statistical configuration, whereas a sustained I indicates that the circuit remains in a statistically distinct dynamical regime following perturbation.

To ensure robustness against sampling variability and parameter selection, two validation procedures were implemented:

- Automated Bandwidth Selection: We employed Scott's Rule to determine the KDE bandwidth, reducing subjective bias in PDF reconstruction. Scott's Rule adaptively determines the optimal bandwidth based on the data's variance and sample size, ensuring that the PDF faithfully represents the underlying stochastic structure without over-smoothing or introducing noise-driven artifacts. This selection is consistent with the smoothed inference framework proposed by McNealis et al. 2026 [17], which suggests that data-driven bandwidths optimize the mean-squared error and ensure the robustness of uncertainty quantification in finite-population contexts.

- Subject-level Bootstrap: We reported 95% confidence intervals (CIs) derived from 1000 bootstrap iterations. By resampling at the subject level, we ensure the observed hysteresis is a robust group-level phenomenon rather than an artifact of individual outliers.

Crucially, this statistical pipeline remains independent of our dynamical modeling (which operates directly on raw time-series data), thereby providing a cross-validated perspective on stress-induced pathway remodeling.

While the KL divergence analysis quantifies distributional shifts in the stochastic signal space, the dynamical modeling aims to identify the mechanistic origin of these shifts in terms of altered stability and pathway coupling.

2.2.2 Dynamical Perspective

We adopted a second-order phenomenological dynamical model. Although calcium signals are often approximated by first-order relaxation processes, projections of high-dimensional neural population dynamics onto low-dimensional observables can exhibit effective inertial dynamics such as delayed recovery or overshoot. A second-order phenomenological expansion therefore provides a minimal representation capable of capturing both restoring forces and viscous damping. Cross terms are retained to allow interactions

between state magnitude and recovery velocity. Such inertial terms naturally arise when high-dimensional neural population dynamics are projected onto a low-dimensional observable, producing effective second-order dynamics with memory-like effects. In addition, calcium indicators introduce intrinsic temporal filtering, which can effectively generate higher-order temporal dynamics in the observed signals.

Single-Variable Dynamical Modeling and Potential Landscape Analysis

To characterize the intrinsic stability of neural signals during recovery, we modeled the BLA–DMS and CeA–DMS pathways as autonomous dynamical systems. For each pathway, a phase space was constructed using the signal amplitude (x_{BLA} or x_{CeA}) and its first derivative ($\dot{x}_{BLA}$ or $\dot{x}_{CeA}$), representing the instantaneous state and rate of change of neural activity.

To ensure robustness against individual variability, we implemented a subject-level bootstrap procedure. For each of $n = 1000$ bootstrap iterations, subjects were resampled with replacement and the governing differential equation was estimated using Ridge Regression. From these fitted models, we derived the distribution of Jacobian eigenvalues ($E_i, i = 1, 2$), where the real part $\mathrm{Re}(E_i), i = 1, 2$ quantifies the intrinsic stability (recovery rate) of the pathway. Stability estimates were reported as 95% confidence intervals (CIs) derived from the bootstrap distribution. Further details on the regression formulations and parameter estimation can be found in Supplementary Section A.

To transcend the limitations of local linear stability, we reconstructed the effective quasi-potential landscapes $U(x)$ and phase portraits to provide a global characterization of the circuit's attractor geometry. Because the system exhibits a single stable equilibrium, the curvature of the effective potential near the fixed point provides a geometric interpretation of the local stability quantified by the Jacobian eigenvalues. This formulation is conceptually related to classical attractor-based descriptions of neural dynamics, where neural states evolve within an energy landscape toward stable fixed points[18]. By integrating the positional restoring-force terms in the empirical models, we mapped the effective potential landscape for each pathway, where the curvature near equilibrium reflects the strength of the homeostatic restoring force.

Attractors were identified as stable fixed points of the reconstructed dynamical system satisfying $\dot{x} = 0$ and $\ddot{x} = 0$. This global view suggests how chronic stress reshapes the potential landscape, flattening the effective potential landscape and altering the local curvature around the attractor.

Given the consistency of these dynamical features, we focused our global analysis on representative perturbations—footshock (non-learned) and unpredicted reward (learned)—which were sufficient to establish the link between local stability degradation and global attractor remodeling.

Finally, to bridge the statistical and dynamical scales, we performed a correlation analysis between the steady-state KL divergence and the principal eigenvalues ($Re(E_1)$). Steady state was defined as the KL divergence between the pre-perturbation baseline distribution and the neural population distribution during the final 5 seconds of the recovery period. This cross-layer validation suggests that distributional hysteresis is a statistical manifestation of the underlying dynamical decay. Because trial-to-trial neural variability effectively samples the local attractor basin, changes in attractor geometry translate into systematic shifts in the observed population distribution, thereby linking dynamical stability to KL divergence.

Latent Coupling and Interaction Modeling

To examine whether chronic stress alters interactions between pathways, we constructed a latent dynamical coupling model linking BLA–DMS and CeA–DMS activity. Although recordings from the two pathways were obtained in separate experimental sessions, we assume that the underlying circuit interaction rules remain statistically invariant across mice within the same experimental condition. This assumption allows population-level pooling to recover shared dynamical structure, a common strategy in neural population modeling [19]. This approach estimates cohort-level interaction rules rather than

moment-to-moment causal interactions between simultaneously recorded neurons.

To reconstruct these robust interaction rules, we extended the phase space to four dimensions: $\{x_{BLA}, \dot{x}_{BLA}, x_{CeA}, \dot{x}_{CeA}\}$. We employed a multivariate Ridge Regression–chosen for its ability to penalize non-robust terms and handle multi-collinearity–to model the acceleration of a target pathway:

$$\ddot{x}_{target} = \mathbf{A} \cdot \mathbf{X}_{intrinsic} + \mathbf{B} \cdot \mathbf{X}_{influence} + \text{higher-order terms} \tag{3}$$

where $\mathbf{X}_{intrinsic}$ and $\mathbf{X}_{influence}$ represent the state vectors $\{x, \dot{x}\}$ of the target and driving pathways, respectively. In this second-order framework, the coefficients map onto classical physical components: the position term (x) defines the restorative stiffness (gradient of the potential landscape), while the velocity term ($\dot{x}$) dictates the viscous damping (recovery rate). The explicit polynomial form of the regression model is provided in the Supplementary Materials Section B. The coupling terms should be interpreted as effective dynamical influences within the model, rather than direct synaptic or causal interactions between brain regions.

Parameter reliability was evaluated using a subject-level bootstrap procedure ($n = 1000$). Directional coupling was considered significant when the 95% confidence interval of the corresponding influence coefficient excluded zero.

Model Validation and Cross-Context Generalization

To investigate whether chronic stress leads to functional impairments in neural circuits, we categorized experimental perturbations into two modalities: (1) Non-Learned Modality (Unpredicted Footshock) probes the circuit's intrinsic homeostatic constraints and recovery capacity from high-arousal, un-modeled perturbations; (2) Learned Modality (Unpredicted Reward and Operant Actions) examines the circuit's adaptive stability within a learned operant context. Using Jacobian stability analysis, we aimed to compare the dynamical differences between the two modalities to determine whether chronic stress induces destructive changes in neural circuitry, particularly by examining the stability and interactions between the BLA-DMS and CeA-DMS pathways.

2.3 Neural Network Architecture: a Minimal Toy Model

Inspired by experimental observations (Nature, [5]), we constructed a simplified artificial neural network model to examine whether distinct computational objectives assigned to two pathways can reproduce the dynamical signatures observed in the empirical data.

We generated time-series data where a high-frequency periodic baseline with Gaussian noise served as the homeostatic equilibrium. Environmental stressors were operationalized as stochastic impulse perturbations (spikes). The Control group was trained exclusively on the noise-buffered baseline, whereas the Stress group was exposed to a regime characterized by frequent stochastic perturbations. To simulate experimental probes (e.g., reward delivery or foot-shocks), we utilized a test signal comprising the standard homeostatic background interrupted by a single, discrete spike (as shown in Figure 13).

The experimental observation that BLA-DMS activity dominates in the Control group while CeA-DMS activity prevails in the Stress group suggests a context-dependent functional bias under environmental pressure. We hypothesize that this pattern is consistent with a potential shift in computational priorities: the BLA-DMS may emphasize predictive precision in stable environments, whereas the CeA-DMS may prioritize robust buffering and persistence under volatility.

To mechanistically investigate these dynamics (see Section 3.1), we formulated the Unified Coupled System (UCS) within a simplified artificial neural network (ANN) framework. It is important to note that the UCS model is not intended to simulate the biological processes of BLA and CeA directly, but rather serves as a proof-of-principle model designed to explore the hypothesized trade-offs between precision-driven tracking and robust buffering. In this framework, the "BLA-DMS" and "CeA-DMS" pathways

represent abstract sub-networks that capture the functional roles of these brain regions within the model, rather than actual biological circuits.

Specifically, asymmetric loss functions were introduced as modeling hypotheses to reflect different functional objectives inferred from experimental observations. These loss functions simulate distinct functional constraints on the sub-networks representing the BLA-DMS and CeA-DMS pathways, with the relative weighting and transition dynamics between these sub-networks learned during training. The design implements a division of labor between the sub-networks, with a Meta-Controller dynamically gating their contributions based on the environmental context.

The UCS model should be viewed as a minimal computational framework to explore functional asymmetry between the pathways, rather than a biologically realistic representation of neural circuits. It is designed to test whether asymmetric objectives can reproduce the observed dynamical behaviors, with the understanding that it is a simplified, abstract representation rather than a literal simulation of biological processes.

The UCS is formulated as a dual-pathway adaptive controller. In this framework, we define the fundamental state of each pathway by the relation:

$$O_t = P_t + \delta_t \tag{4}$$

where O_t is the observable output sent to the DMS, P_t is the predicted calibration (baseline), and δ_t is the latent deviation (accumulated internal strain). The system's global objective is to minimize the total loss $L_{Total} = L_{BLA} + L_{CeA}$ relative to the homeostatic reference (T), which corresponds to the aforementioned periodic baseline.

The internal dynamics of the BLA-DMS sub-network (x_{BLA}) and the CeA-DMS sub-network (x_{CeA}) are governed by an adaptive LSTM cell state, where the gating signals are modulated by the meta-learned parameter set Φ (see Figure 1). The two sub-networks are therefore trained under distinct objective functions, creating different optimization pressures during learning (for detailed values please refer to Supplementary Materials C):

The BLA-like sub-network as a Precision Tracker: The BLA-DMS sub-network is constrained by a fidelity-centric objective (L_{BLA}), aiming to align O_{BLA} with T under a hard physical capacity barrier (θ_{BLA}):

$$L_{BLA} = W_{MSE} \cdot |O_{BLA} - T|^2 + W_{Sat} \cdot \text{Mean}(\text{Penalty}_{BLA}) \tag{5}$$

where W_{MSE} and W_{Sat} are weights for tracking precision and saturation penalty. The term $\text{Penalty}_{BLA} = \text{ReLU}(|\delta_{BLA}| - 0.2)$ simulates a "brittle" system: while highly accurate in stable regimes, the BLA-DMS sub-network becomes functionally paralyzed when its latent strain δ_{BLA} exceeds its metabolic ceiling (capacity constraint).

The CeA-like sub-network as a Robustness Buffer: CeA-DMS sub-network (x_{CeA}) is optimized for structural stability and informational residue (persistent latent state) under volatility. Its objective function (L_{CeA}) is engineered to reward the buffering of perturbations and the maintenance of latent strain δ_{CeA}, even when external stressors subside:

$$L_{CeA} = W_{Buf} \cdot L_{Buffer} + W_{Cal} \cdot L_{Calib} + W_{Pers} \cdot L_{Persist} \tag{6}$$

The interaction of these terms defines the "viscous" dynamics of the habitual state: Dynamic Buffering (L_{Buffer}):

$$L_{Buffer} = -\mathbb{E}[|O_{BLA} - O_{CeA}| \cdot \mathbf{M}_{pert}] \tag{7}$$

This term encourages the CeA-DMS sub-network to diverge from the BLA-DMS sub-network during perturbations ($\mathbf{M}_{pert}$). In this context, CeA-DMS sub-network contributes to modulating the overall system

dynamics by introducing a stabilizing (damping-like) effect on the response of the BLA-DMS sub-network. Homeostatic Anchoring (L_{Calib})is defined as MSE(O_{CeA}, T) during non-perturbation periods ($\mathbf{M}_{non-pert}$). This term ensures the system remains tethered to the baseline in stress-free environments, preventing unrestricted drift. The differentiation depending on $\mathbf{M}$ that denotes the binary masks identifying perturbation segments ensures that while the BLA-DMS sub-network tracks homeostatic precision, the CeA-DMS sub-network preserves the "memory" of environmental pressure. Informational Persistence ($L_{Persist}$):

$$L_{Persist} = -\frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} |\delta_{CeA}| \tag{8}$$

where $\mathcal{T}$ is the discrete temporal domain. By assigning a negative weight ($W_{Pers} = -2.0$), the model rewards the maintenance of a non-zero δ_{CeA}. This design choice ensures that the CeA-DMS sub-network does not return to a zero-mean state post-perturbation, thereby providing a direct computational implementation of the "distributional hysteresis" observed in our biological GCaMP8s data. This allows the UCS to bridge the gap between abstract informational metrics and mechanistic neural dynamics.

This reward structure fundamentally transforms the system's potential landscape. Unlike "elastic" systems that gravitate toward zero deviation, the CeA-DMS sub-network develops a "plastic" memory of stress. This formalizes the pathological rigidity observed in our GCaMP8s data: the CeA-DMS sub-network remains locked in an active "stress mode" even after the homeostatic reference (T) has stabilized. Consistent with the observed dominance of CeA-DMS sub-network in stressed subjects, L_{CeA} is specifically engineered to reward informational persistence, allowing the pathway to maintain an active latent state even as external perturbations subside.

The interaction between pathways is governed by a hierarchical control structure: a Parameter Predictor sets macroscopic interaction scales based on long-term environmental statistics, while a Meta-Controller provides microscopic adaptation to instantaneous uncertainty.

The coupling of these pathways is governed by a hierarchical control structure: a Parameter Predictor sets macroscopic interaction scales based on environmental statistics, while a Meta-Controller provides microscopic adaptation to instantaneous uncertainty. Schematic construction of the neural network is shown in Figure 1.

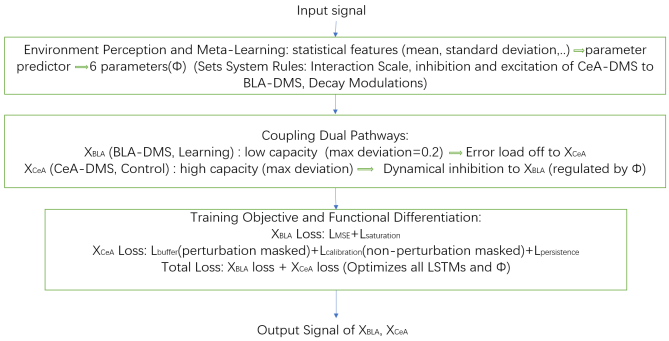

Figure 1: Schematic of the adaptive controller Unified Coupled System (UCS) architecture. The model formalizes the transition between pathways through three functional layers: (1) Macro-Perception Layer: A parameter predictor (Φ) that senses environmental volatility to dynamically adjust system coupling rules. (2) Dual-Pathway Coupling Layer: Parallel Adaptive LSTMs representing the BLA-DMS sub-network (x_{BLA}, low-capacity) and CeA-DMS sub-network (x_{CeA}, high-capacity) circuits. (3) Optimization Layer: Functional differentiation driven by Asymmetric Loss functions, where L_{BLA} prioritizes tracking precision and L_{CeA} optimizes for robustness and informational persistence. The core mechanism is adaptive learning, enabling adaptive responses and pathway specialization through dynamic parameter modulation.

in Figure 1, there is target asymmetry (Functional Division). This is a design where the BLA-DMS sub-network and CeA-DMS sub-network are optimized for different functional objectives, thereby formalizing the functional division between the precision-focused BLA-DMS sub-network and the stability-focused CeA-DMS sub-network.

As shown in Figure 1, this leads to a principled topological transition: as the BLA-DMS sub-network reaches its saturation limit under high-entropy stress, the system spontaneously shifts the computational load to the high-capacity CeA-DMS sub-network. This transition is quantified by the Max Deviation, defined as the peak internal strain $\max|\delta_t|$ during a perturbation. The systematic reallocation of deviation from the BLA-DMS sub-network to the CeA-DMS sub-network provides a computational metric for the shift from flexible, goal-directed tracking to a rigid, stress-locked state. We refer to the two units as "BLA-DMS sub-network" and "CeA-DMS sub-network" purely for interpretability, based on their correspondence to the empirical signals. This naming does not imply anatomical or mechanistic equivalence.

To explore stress-induced circuit remodeling, we trained independent UCS instances on distinct environmental conditions (Control vs. Stress). Initially, both subnetworks were trained using the same MSE loss function (see Supplementary Materials C.2). However, this approach led to results inconsistent with experimental observations. This discrepancy highlights that the UCS design depends on the interaction mechanisms between the two pathways, which are crucial for modeling their distinct roles in precision and robustness.

3 Results

3.1 Data Analysis

3.1.1 Non-Learned Modality: Pathological State-Locking under Aversive Stress

Statistical Perspective

We first analyzed the temporal evolution of KL divergence following acute footshock to quantify the circuit's homeostatic recovery (Figure 2). In the Control group, both the BLA-DMS and CeA-DMS pathways exhibited a transient surge in informational distance followed by a rapid decay, with probability distributions relaxing toward baseline within approximately 4 seconds. In contrast, chronic stress induced pathway-specific distributional hysteresis. The BLA-DMS pathway of stressed mice exhibited a sustained elevation in KL divergence, failing to return to baseline and plateauing at a high level.

In the stressed group, from $t = 6$ s onward, the mean KL divergence of the BLA-DMS and CeA-DMS pathways became distinctly separated, indicating that the two pathways responded differently to the footshock. In contrast, in the Control group, the temporal dynamics of the two pathways converged after the perturbation, suggesting a more coordinated response. This divergence in the stressed group suggests a shift in the balance of functional roles between the pathways, which may have implications for the altered dynamics of the stressed circuit.

To validate the reliability of these informational shifts, we performed a subject-level bootstrap procedure ($n = 1000$ iterations) to calculate the cumulative remodeling cost (Table 1). Despite individual variability, the Mean Cumulative D_{KL} for the Stressed BLA-DMS pathway was approximately double that of the Control group (51663 vs. 26979), with a substantial separation in the 95% CI lower bounds (25083 for Stress vs. 6137 for Control). These results confirm that the observed circuit-wide redistribution is a robust group-level phenomenon, independent of sampling artifacts or individual outliers.

To bridge the temporal dynamics with steady-state outcomes, we compared the time-resolved D_{KL} with the pre- and post-perturbation moments (Supplementary Materials Table D8). While low-order moments (mean and variance) exhibited relatively stable shifts, the informational divergence (KL divergence) showed a markedly more pronounced and persistent reconfiguration. This suggests that the impact of stress is more

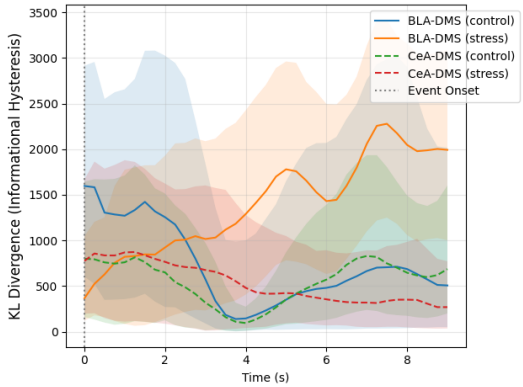

Figure 2: Time-resolved Kullback-Leibler (KL) divergence relative to the pre-stimulus baseline ($t < 0$) following a single footshock ($t = 0$, gray dotted line). Solid and dashed lines represent the mean KL divergence for BLA-DMS and CeA-DMS pathways, respectively; shaded areas denote the 95% confidence intervals (CIs) derived from subject-level bootstrapping ($n = 1000$ iterations). In control mice, both pathways exhibit efficient relaxation, with D_{KL} returning to baseline within 4s. In stressed mice, the BLA-DMS pathway shows prominent distributional hysteresis, with sustained elevation in KL divergence from $t > 6$ s. Conversely, the stressed CeA-DMS pathway shows a blunted response, reflecting a loss of informational flexibility.

Table 1: Statistical Robustness of Integrated KL Divergence Following Acute Footshock

Pathway	Group	Mean Cumulative D_{KL}	95% CI
BLA-DMS	Control	26979.6135	[6137.2289, 76874.4732]
BLA-DMS	Stress	51663.6477	[25083.7458, 84200.4346]
CeA-DMS	Control	20525.2011	[8192.5872, 37799.9575]
CeA-DMS	Stress	19665.9537	[6240.9008, 37132.0447]

discernible in the high-order stochastic structure of the ensemble activity than in its aggregate statistical averages.

Dynamical Perspective

To bridge the observed statistical shifts with their underlying neural kinetics, we analyzed the stability regimes of the two pathways using Jacobian eigenvalues. Table 2 reports the mean coefficients and their corresponding 95% confidence intervals (CIs) derived from 1000 bootstrap iterations.

Interestingly, the quadratic coefficient (x_{BLA}^2) for the Stress group's BLA-DMS pathway becomes consistently positive and statistically significant across all behavioral modalities–including footshock (Table 2), operant pressing, and reward delivery (see Supplementary Tables D5, D13 and D14). This positive coefficient in the phenomenological model suggests that the system inherently resists divergence, maintaining homeostatic stability even under perturbations. This reconfiguration may be linked to the observed decline in behavioral flexibility in stressed mice, indicating a trade-off between stability and adaptability in the stressed BLA-DMS pathway.

To ensure our findings represent the entire cohort rather than individual outliers, we derived the probability density of the principal eigenvalues ($Re(E_1)$) via a subject-level bootstrap-regression pipeline (Figure 3). These principal eigenvalues ($Re(E_1)$) were derived by linearizing the discovered governing equations at their respective steady states, capturing the local stability of the system's vector field.

Table 2: Population-level Statistics of Governing Dynamical Equations for Footshock Response

		Control		Stress	
Pathway	Term	Mean Coefficients	95% CI	Mean Coefficients	95% CI
BLA-DMS	Intercept	-1.9133	[-11.0969, 6.5834]	-11.7957	[-25.3569, -3.9355] *
	x_{BLA}	-12.7695	[-52.9504, 3.9237]	-8.0333	[-16.1302, -1.5321] *
	$\dot{x}_{BLA}$	-27.5925	[-30.3323, -23.1249] *	-26.5726	[-29.2188, -22.9836] *
	x^2_{BLA}	-9.9306	[-50.8948, 19.6649]	18.2556	[2.2446, 49.5706] *
	$x_{BLA}\dot{x}_{BLA}$	-0.381	[-6.3485, 5.3484]	2.9982	[-1.5469, 8.5104]
	$\ddot{x}_{BLA}$	0.2087	[-0.5831, 1.3130]	-0.069	[-1.0170, 0.9458]
CeA-DMS	Intercept	-6.2783	[-12.0011, -1.0509] *	-4.167	[-10.3045, 6.9254]
	x_{CeA}	-8.3031	[-19.1907, -3.5590] *	-7.9044	[-18.7739, 2.7631]
	$\dot{x}_{CeA}$	-28.7741	[-31.5479, -25.3304] *	-27.4448	[-30.5188, -23.8047] *
	x^2_{CeA}	3.558	[-20.5504, 14.2200]	1.6879	[-16.1770, 14.3044]
	$x_{CeA}\dot{x}_{CeA}$	3.2412	[0.0413, 6.9261] *	1.2205	[-2.9535, 4.8927]
	$\ddot{x}_{CeA}$	0.5948	[-0.5281, 1.6150]	1.1538	[-0.5994, 2.8350]

* Indicates statistical significance (95% CI does not overlap with zero).

In the BLA-DMS pathway, the Control group exhibits a distribution centered near zero, suggesting a marginally stable state. The Stress group shows a significant leftward shift in its $Re(E_1)$ distribution, with the 95% confidence interval (CI) strictly localized in the negative domain (Table 3). This shift suggests that chronic stress has reconfigured the BLA-DMS pathway into a pathological attractor–a state that is dynamically stable but topologically distinct from the original baseline (e.g., the significant x^2 term in Table 2).

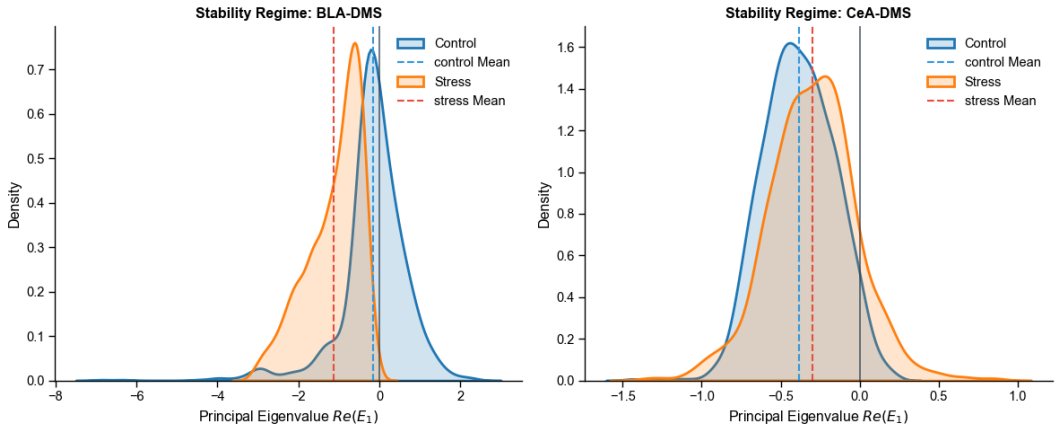

Figure 3: Population-level stability analysis of neural trajectories. The probability density functions (PDFs) of the principal Jacobian eigenvalues ($Re(E_1)$) were estimated via subject-level bootstrapping ($n = 1000$ iterations) following acute footshock. (Left) In the BLA-DMS pathway, the stress group (orange) exhibits a more negative mean principal eigenvalue compared to the control group (blue), indicating an altered damping profile. (Right) In the CeA-DMS pathway, the stress group shows a rightward shift in the distribution, signifying a decrease in recovery efficiency. Dashed lines represent the mean of the bootstrapped distributions. The consistently negative real parts of the eigenvalues across both pathways confirm the global asymptotic stability of the underlying dynamical regimes.

The deterministic parameters of the "average" model (Table 4) align closely with these bootstrapped mean distributions, providing a mechanistic explanation for the KL divergence results. The total recovery speed, quantified by the L^2-norm of the eigenvalues ($||E||_2$), remains consistently high across most pathways

Table 3: Population-Level Statistics of Principal Jacobian Eigenvalues for Footshock

Pathway	Group	Mean Principal $Re(E_1)$	95% CI
BLA-DMS	Control	-0.1585	[-2.5636, 1.2795]
	Stress	-1.1217	[-2.6978, -0.2651]*
CeA-DMS	Control	-0.3855	[-0.7820, 0.0215]
	Stress	-0.3000	[-0.8698, 0.2210]

* Indicates statistical significance (95% CI does not overlap with zero).

(> 26), except for the highly flexible Control BLA-DMS. This suggests that while the Stress BLA-DMS pathway recovers "quickly" in a kinetic sense, it is stably trapped in an aberrant sub-state, manifested as the persistent "distributional hysteresis" in our informational analysis. Conversely, the Stress CeA-DMS pathway maintains negative eigenvalues and a relatively high recovery speed, indicating a more rigid, less adaptable response compared to controls. This aligns with its role as a stabilizing controller, but the observed diminished informational flexibility suggests that, in chronic stress, it may prioritize stability over adaptability. The CeA-DMS pathway's role is more consistent with facilitating rigid, compulsive-like behaviors, rather than strictly habitual ones. Additionally, while the DMS is implicated in goal-directed behaviors, its connections to the prefrontal cortex, rather than the BLA, are more central to such processes. These connections between CeA-DMS and behavioral rigidity, however, remain an important point of investigation.

Table 4: Deterministic Stability Parameters and Recovery Kinetics Post-footshock

Group	Pathway	Steady State		Eigenvalues		Total Recovery Speed
		x_s	$\dot{x}_s$	E_1	E_2	
Control	BLA-DMS	-0.397	0	1.60E-04	-0.27275	0.273
	CeA-DMS	-0.3924	0	-0.4044	-29.469	29.472
Stress	BLA-DMS	-0.6465	0	-0.837	-27.2812	27.294
	CeA-DMS	-0.1597	0	-0.2615	-26.8716	26.873

For notation simplicity in the table, x_s and $\dot{x}_s$ denote the steady-state amplitude and its derivative, respectively, for the corresponding target pathway (either BLA-DMS or CeA-DMS) identified by the dynamical model.

Our coupling analysis suggests a stress-induced breakdown in the inter-pathway homeostatic regulation (Table 5). In Control mice, we identified a robust restorative mechanism in the CeA $\rightarrow$ BLA direction: the CeA-DMS exerts a significant positive drive on BLA-DMS dynamics ($x_{CeA} = 80.88$, 95% CI: [23.88, 144.23]), acting as an extrinsic stabilization force. In the Stress group, however, this regulatory term collapses and loses statistical significance (Figure 4). This "decoupling" from CeA-mediated control, combined with the emergence of a rigid, non-linear self-interaction term in the CeA itself ($x^2_{CeA} = 29.96$, Table 6), suggests that chronic stress leads to altered dynamics in the amygdala-striatal axis. While the exact nature of these changes remains to be fully characterized, the data imply that stress disrupts the normal regulatory balance between these pathways, potentially contributing to maladaptive states.

This erosion of inter-pathway compensatory rules provides a quantitative basis for the distributional hysteresis observed in stressed mice. Unlike the "state-locking" triggered by the innate Footshock (Non-Learned Modality), trajectories in the Reward context (Learned Modality) remained stable (Supplementary Materials Table D6). This suggests that while stress systematically dismantles innate homeostatic con-

Table 5: Coefficients of Latent Dynamical Coupling (Influence Direction: CeA-DMS to BLA-DMS)

	Control		Stress	
Variables	Mean Coefficients	95% CI	Mean Coefficients	95% CI
x_{BLA}	-128.9280	[-239.7098, -33.8852]*	-8.7873	[-71.4567, 35.6582]
$\dot{x}_{BLA}$	-32.5210	[-37.6582, -26.5185]*	-29.2282	[-35.7247, -23.9503]*
x_{BLA}^2	-22.6382	[-164.0459, 105.3996]	56.2701	[-7.3992, 116.3720]
$x_{BLA}\dot{x}_{BLA}$	-9.5272	[-29.8239, 7.8918]	1.0535	[-9.3516, 9.0369]
$\dot{x}_{BLA}^2$	-0.2753	[-1.4229, 0.7945]	-0.2047	[-1.4371, 0.9042]
x_{CeA}	80.8810	[23.8786, 144.2312]*	8.2858	[-23.6643, 44.6718]
$\dot{x}_{CeA}$	1.8124	[-1.6390, 5.1101]	1.4626	[-1.1804, 4.5300]
$x_{BLA}x_{CeA}$	46.7142	[-53.7033, 187.4726]	-25.8531	[-85.4224, 40.9728]
$x_{BLA}\dot{x}_{CeA}$	4.7782	[-1.3883, 15.1915]	3.4656	[-2.2191, 11.0586]
$\dot{x}_{BLA}x_{CeA}$	8.7838	[-5.3136, 28.8629]	1.3723	[-4.9063, 8.9948]
$\dot{x}_{BLA}\dot{x}_{CeA}$	0.4897	[-0.7424, 1.9411]	0.8116	[-0.5674, 2.3390]

* Indicates statistical significance (95% CI does not overlap with zero).

Table 6: Coefficients of Latent Dynamical Coupling (Influence Direction: BLA-DMS to CeA-DMS)

	Control		Stress	
Variables	Mean Coefficients	95% CI	Mean Coefficients	95% CI
x_{CeA}	-51.4172	[-113.5551, -8.0192]*	-65.9611	[-187.4206, -15.7276]*
$\dot{x}_{CeA}$	-29.3639	[-32.8705, -25.6075]*	-27.1137	[-33.9473, -20.8905]*
x_{CeA}^2	-21.9424	[-117.9379, 35.3150]	29.9629	[1.4243, 84.2127]*
$x_{CeA}\dot{x}_{CeA}$	-2.2302	[-17.9178, 12.2457]	0.2663	[-5.9641, 5.6891]
$\dot{x}_{CeA}^2$	0.1302	[-1.0966, 1.3994]	0.3835	[-0.8955, 1.7972]
x_{BLA}	54.8296	[11.2364, 123.4697]*	58.1055	[6.7276, 162.0060]*
$\dot{x}_{BLA}$	0.5796	[-2.1818, 3.6602]	0.8940	[-2.5930, 5.7940]
$x_{CeA}x_{BLA}$	26.4998	[-44.4399, 107.2743]	-44.9540	[-141.7844, 2.6415]
$x_{CeA}\dot{x}_{BLA}$	1.6595	[-2.3736, 7.8194]	3.4341	[-3.0802, 12.2634]
$\dot{x}_{CeA}x_{BLA}$	5.7149	[-12.6906, 24.2480]	4.6674	[-5.0549, 14.1652]
$\dot{x}_{CeA}\dot{x}_{BLA}$	0.2680	[-0.8885, 1.4286]	0.4331	[-1.0463, 1.9704]

* Indicates statistical significance (95% CI does not overlap with zero).

straints, the formation of an internal model during operant learning may provide a dynamical "buffer" against total circuit instability, a hypothesis we explore through longitudinal analysis in Supplementary Materials.

3.1.2 Learned Modality: Functional Resilience and Cross-Context Validation

Statistical Perspective

During the Contingency Acquisition phase, we monitored the informational hysteresis as mice transitioned from deterministic (FR1) to stochastic (RR10) reinforcement. In control mice, the BLA-DMS pathway showed stable dynamical behavior under the RR10 reinforcement schedule, as indicated by KL divergence remaining near baseline. This suggests that the pathway effectively stabilized the operant act over time, with minimal informational shifts, reflecting a well-learned and consistent motor program

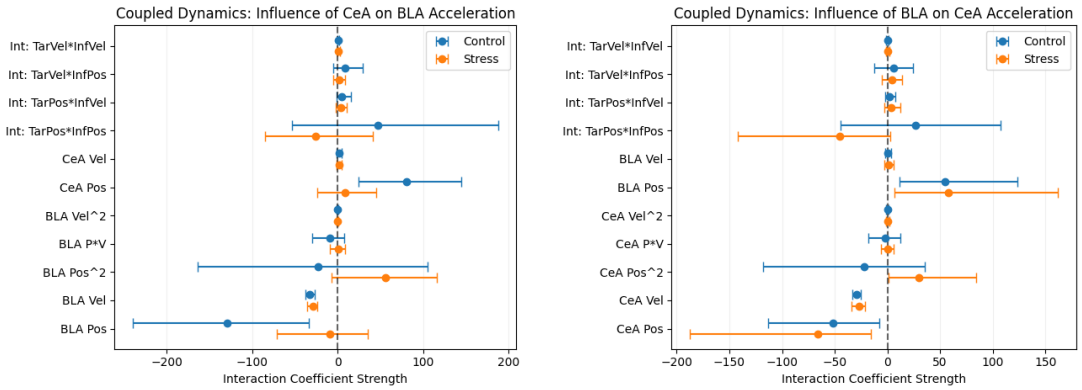

(a) Influence direction: CeA-DMS to BLA-DMS (b) Influence direction: BLA-DMS to CeA-DMS

Figure 4: Forest plots illustrate the interaction coefficients derived from the latent dynamical model across two primary influence directions. Variable Mapping: To maintain a generalizable computational framework, variables are labeled as *Target* (Tar) and *Influence* (Inf). "Tar" denotes the pathway being modeled (dependent variable), while "Inf" denotes the pathway exerting an external drive (independent variable). Specifically, *Pos* (Position) corresponds to the signal amplitude (x), and *Vel* (Velocity) corresponds to its first derivative ($\dot{x}$). Nonlinear Modulation Term: Terms involving products (e.g., $TarPos \times InfPos$) represent non-linear coupling effects. Significance: Each data point represents the mean coefficient from subject-level bootstrapping ($n = 1000$), with error bars indicating 95% confidence intervals (CIs). A coupling rule is considered statistically robust if its CI excludes the zero-axis (dashed vertical line), signifying a consistent directional effect across the cohort.

Terminology Note: We adopted a standardized nomenclature (Target and Influence) in the graphical representation to emphasize the interaction direction between pathways, independent of specific anatomical labels. The mapping between these graphical terms and the anatomical variables (x_{BLA}, x_{CeA}) used in the regression tables is detailed in the Main Text.

(Figure 5).

However, the stressed BLA-DMS pathway (orange line) showed a progressive failure to restore its informational distribution after lever-press. While early acquisition (FR1) resulted in a surge in KL divergence, the RR10 session revealed a sustained increase in KL divergence, suggesting that the stressed pathway continues to process a well-learned behavior with significant informational cost, comparable to the handling of a novel challenge. This failure to return the distribution to baseline under uncertainty, as evidenced by the larger and more sustained rise in KL divergence in the stressed group compared to controls, relates to the dynamical signature of stress-induced cognitive rigidity.

The temporal evolution of the 95% confidence intervals (shaded areas) further illuminates the divergent learning trajectories between groups. In the initial acquisition phase (FR1), control mice exhibited a relatively broad CI around 4s post-press, reflecting the inherent behavioral variability as individuals first consolidate the action-outcome association. As uncertainty increased (RR2 and RR5), the CI in the control group remained wide, signifying a period of learning bifurcation: while some individuals successfully adapted their internal models to the sparse reinforcement, others remained in an exploratory state (see Figure 5).

At the highest level of uncertainty (RR10), the control group's confidence intervals (CI) significantly contracted, indicating more consistent neural responses. In contrast, the stressed BLA-DMS pathway exhibited a significant and sustained increase in KL divergence, suggesting that chronic stress led to

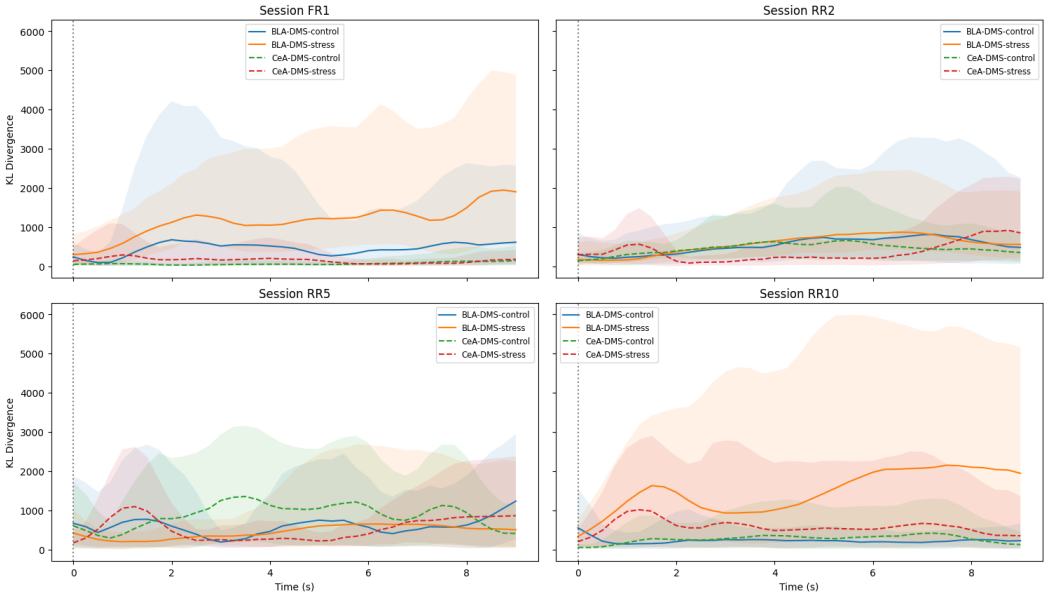

Figure 5: Evolution of informational hysteresis during operant act execution. Time-resolved KL divergence of BLA-DMS (solid lines) and CeA-DMS (dashed lines) pathways aligned to the onset of lever pressing across four training stages (FR1, RR2, RR5, and RR10). (Top-Left) In the initial acquisition phase (FR1), chronic stress (orange/red) triggers a premature and elevated informational load in the BLA-DMS compared to controls (blue), suggesting higher initial circuit friction during goal-directed initiation. (Bottom-Right) Late-stage uncertainty (RR10): As the reward probability drops to 0.1, control circuits maintain a "quiet" and efficient informational distribution. Conversely, the stressed BLA-DMS exhibits a significant late-phase surge. Shaded areas represent 95% CIs derived from subject-level bootstrapping ($n = 1000$). Note the relatively constrained scale of D_{KL} compared to reward-driven responses, reflecting the stability of the informational distribution during motor execution.

greater variability and inconsistency in neural responses. While control mice's neural responses converged to a stable pattern, the stressed group exhibited more erratic and less predictable neural activity, reflecting a loss of consistency and a failure to adapt to the reward structure effectively.

The narrowing of confidence intervals in control mice at RR10 marks the transition from active exploration to refined decision-making, a process absent in the chronically stressed brain.

Following the operant execution, the delivery of reinforcement (Reward-onset) triggered a significant stochastic reconfiguration within the amygdala-striatal axis. Mirroring the inefficiency observed during the preceding motor act, the feedback phase revealed a notable functional shift in the CeA-DMS pathway (Figure 6).

In the initial acquisition stage (FR1), chronic stress induced a "suppressed" informational response, characterized by reduced KL divergence trajectories in the CeA-DMS (red dashed line) relative to controls. However, as the operant contingency solidified into high-uncertainty regimes (RR5, RR10), this suppression transitioned into an elevated responsiveness. In these later sessions, the stressed CeA-DMS exhibited significantly higher and more persistent informational hysteresis peaks than the control group. This transition suggests that while control circuits successfully restore the informational distribution as a reward becomes unpredictable, the stressed circuit fails to restore the distribution, with greater informational divergence. Similarly, the state of the BLA-DMS pathway in the stress group further confirms that chronic stress prevents the circuit from achieving the pre-perturbation baseline distribution, as evidenced by persistent deviations in informational distribution.

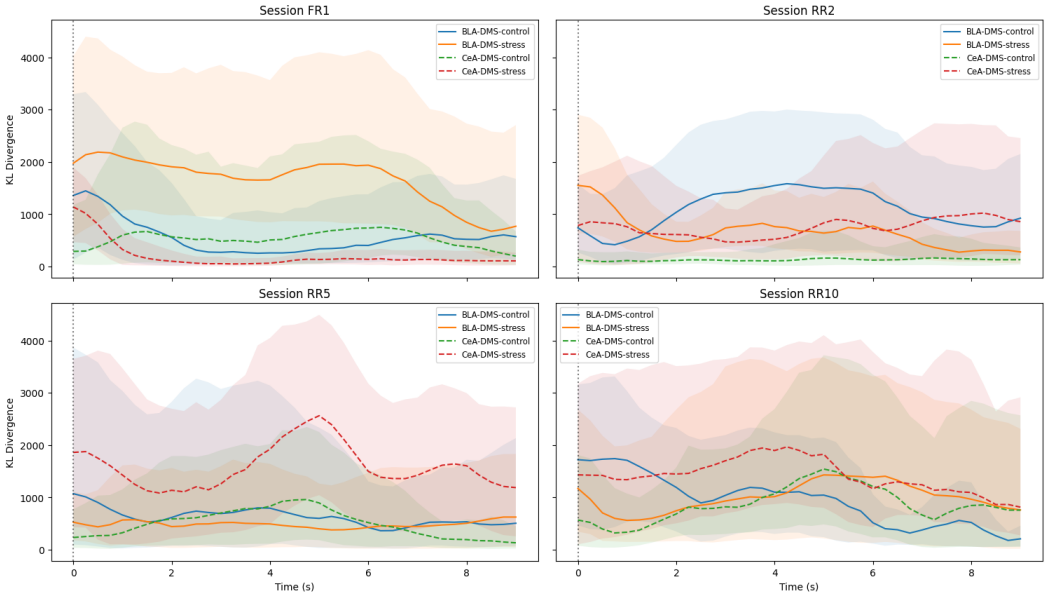

Figure 6: Evolution of informational hysteresis following lever-press reinforcement. Time-resolved KL divergence trajectories aligned to the delivery of food reinforcement following lever pressing across progressive learning sessions. In the early session (FR1), chronic stress induces a suppressed informational response in the CeA-DMS pathway (red dashed line), characterized by reduced KL divergence compared to controls. However, as the reward ratio becomes lower (RR5-RR10), the stressed CeA-DMS pathway exhibits significantly higher and more persistent KL divergence peaks compared to controls (green dashed line), indicating an increased informational divergence.

Dynamical Perspective

Beyond informational metrics, we employed both univariate and bivariate dynamical modeling to characterize the evolving state-space of the pathways. Due to the extensive number of estimated parameters, the full coefficient tables for all training sessions and reinforcement regimes are provided in the Supplementary Materials (Tables D13–D14 and Tables D18-D19).

A pivotal finding from our univariate modeling is the universal stability of the circuit during operant tasks. As visualized in the distributions of the principal eigenvalue's real part ($Re(E_1)$, Figures 7 and 8), both the BLA-DMS and CeA-DMS pathways remain strictly within the negative semi-plane ($Re(E_1) < 0$) across all training sessions, from deterministic FR1 to high-uncertainty RR10. Detailed information of 95% CIs of $Re(E_1)$ could be referred in Supplementary Materials D.

This pervasive asymptotic stability stands in contrast to the transient dynamical instabilities occasionally triggered by acute, unpredicted footshocks in the control group. It suggests that the operant learning process itself acts as a fundamental top-down constraint, reshaping the biological neural network's attractor landscape into a globally dissipative regime to ensure reliable behavioral execution.

While the macro-stability ($Re(E_1) < 0$) is preserved in both groups, the underlying parameter configurations reveal a stress-induced remodeling of the regulatory rules. Through systemic screening of coupling coefficients across all experimental phases (Tables 5–6 and Supplementary Tables D9-D10 and D18-D19), several statistically robust patterns emerge.

Across all sessions (FR1–RR10) and both reward/pressing phases, the consistently significant and negative coefficients of velocity terms ($\dot{x}_{BLA}, \dot{x}_{CeA}$) constitute the dominant dissipative diagonal elements of the system's Jacobian matrix (ranging from -24 to -32; 95% CI strictly below zero). This robust negative feedback represents strong dissipative damping. While full stability formally depends on the

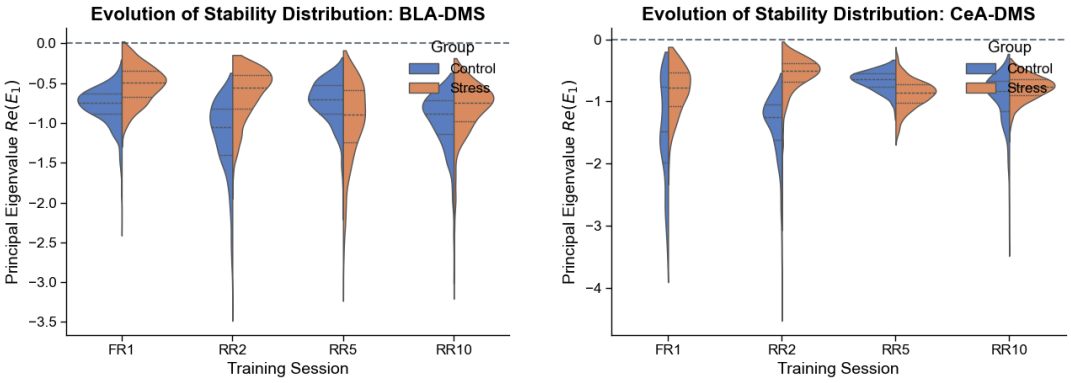

(a) Violin plot of principal eigenvalue's real part distribution of BLA-DMS pathway

(b) Violin plot of principal eigenvalue's real part distribution of CeA-DMS pathway

Figure 7: Violin plots and range trajectories represent the distribution of the principal eigenvalue's real part ($Re(E_1)$) derived from bootstrapped ridge regression ($n = 1000$) for bout-initiating pressing. Across all reinforcement regimes (FR1 to RR10), the $Re(E_1)$ distributions are strictly confined to the negative semi-plane ($Re(E_1) < 0$), indicating that the BLA-DMS and CeA-DMS pathways maintain global asymptotic stability.

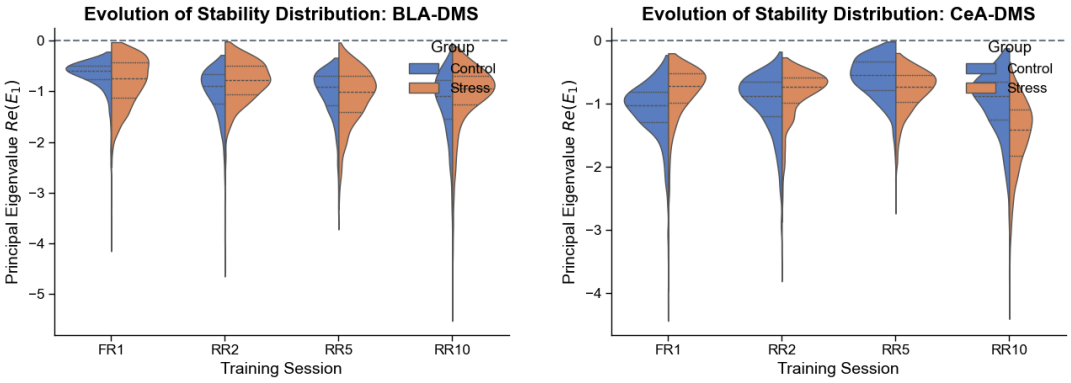

(a) Violin plot of principal eigenvalue's real part distribution of BLA-DMS pathway

(b) Violin plot of principal eigenvalue's real part distribution of CeA-DMS pathway

Figure 8: Violin plots and range trajectories represent the distribution of the principal eigenvalue's real part ($Re(E_1)$) derived from bootstrapped ridge regression ($n = 1000$) for reward. Across all reinforcement regimes (FR1 to RR10), the $Re(E_1)$ distributions are strictly confined to the negative semi-plane ($Re(E_1) < 0$), indicating that the BLA-DMS and CeA-DMS pathways maintain global asymptotic stability.

eigenvalues of the coupled Jacobian matrix, these large damping terms strongly bias the system toward an overdamped regime, effectively suppressing sustained oscillatory dynamics (Supplementary Tables D18-D19).

In control mice, both the BLA-DMS and CeA-DMS pathways exhibit consistent negative self-feedback in their positional terms (x_{BLA}, x_{CeA}) across most learning sessions (Figures 9(a)-10(c)). These negative coefficients act as linear restoring forces, defining a steep attractor basin. This suggests that under baseline conditions, each pathway possesses an autonomous ability to return to its steady state, operating with a high degree of independent positional control.

In contrast, the stressed circuit shows an attenuation of this autonomous linear coordination. The self-feedback positional terms (x_{BLA} or x_{CeA}) frequently drift toward non-significance (CI crossing zero, Figures 9-10). This reduction in local restorative precision suggests that the control over stability in the stressed circuit has migrated from individual pathway restoration to collective dissipative dynamics.

Notably, the significant positive x_{BLA}^2 terms observed in our univariate analysis for the stress group become non-significance in this multivariate coupling model. This indicates that the nonlinear stabilizing effect previously attributed to the BLA-DMS pathway's univariate dynamics is, in the coupled framework, accounted for by inter-pathway interactions. While individual cross-pathway terms may not always reach statistical significance, their inclusion absorbs the variance previously assigned to the self-feedback nonlinearities.

Reviewing the structure of both the univariate and coupling models, we find that the coefficients of $\dot{x}_{BLA}$ and $\dot{x}_{CeA}$ are consistently negative and statistically significant in both frameworks. This cross-model consistency confirms that the system's fundamental physical structure is defined by second-order overdamped dynamics. The disappearance of x (restoring) significance specifically in the stressed coupled model suggests that stability control has shifted to a distributed manifold of interaction terms. These terms likely represent higher-order dynamics of the complex system; while they may not emerge as primary linear drivers, they collectively sustain an alternative stable state, distinct from the original baseline equilibrium.

Ultimately, chronic stress does not necessarily damage individual circuit components (variables), but rather distorts the dynamical relationships (coupling) between them, favoring robust damping over flexible restoration.

3.1.3 Stress-induced Attractor Remodeling and Informational Hysteresis

To characterize the mechanistic origin of the persistent informational footprint, we examined the global attractor geometry of the BLA–DMS and CeA–DMS pathways. The reconstructed potential landscapes $U(x)$ revealed that chronic stress fundamentally reconfigures the circuit's steady-state energy profile (Figure 11).

Following footshock (a non-learned aversive modality), the BLA–DMS pathway in control subjects exhibited a restorative potential with a minimum centered near the pre-perturbation baseline ($x \approx 0$) (Figure 11(a)). In contrast, chronic stress induced a statistically significant displacement of the attractor, shifting the potential minimum toward a biased, negative state. This observation was further corroborated by phase portraits, where the vector fields in the stressed group converged toward a displaced equilibrium. This suggests a transition from baseline-oriented recovery to a reconfigured steady state, where the circuit's dynamical balance is maintained at a distinct, offset coordinate.

Interestingly, this stress-induced remodeling in the BLA–DMS pathway was modality-specific. While the circuit underwent displacement following footshock, it maintained comparable dynamical stability across groups during the processing of unpredicted reward (learned modality) (Figure 11(c); detailed metrics in Supplementary Materials D). This suggests that chronic stress selectively prioritizes the structural reorganization of threat-processing attractors while leaving reward-related dynamical templates relatively intact in the BLA-DMS pathway.

To bridge the statistical and dynamical scales, we correlated the steady-state KL divergence with the principal Jacobian eigenvalues ($Re(E_1)$) (Table 7; Figure 12). In the footshock condition, the stressed BLA–DMS circuit exhibited a significant negative correlation ($-0.215, p < 10^{-11}$). This relationship reflects attractor displacement: stronger local stability around the displaced equilibrium increases the persistence of deviation from the baseline distribution. Conversely, for unpredicted reward, the CeA–DMS pathway showed a robust positive correlation (0.327 for stress, $p < 10^{-25}$), indicating that reward-related information integrity is directly tethered to the circuit's restoring force.

These results demonstrate that stress-induced informational persistence is not merely a transient delay but a structural consequence of attractor remodeling. Depending on the circuit and stimulus valence,

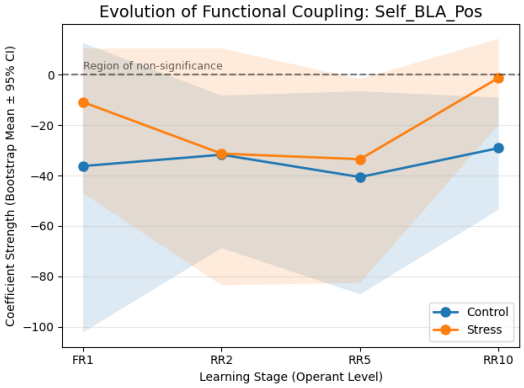

(a) Influence direction: $CeA \rightarrow BLA$

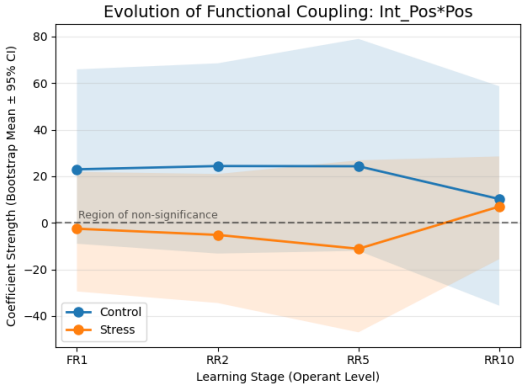

(b) Influence direction: $CeA \rightarrow BLA$

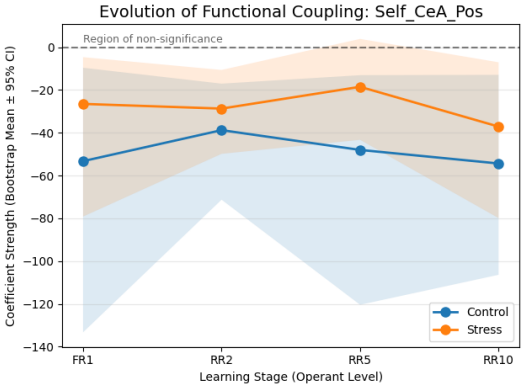

(c) Influence direction: $BLA \rightarrow CeA$

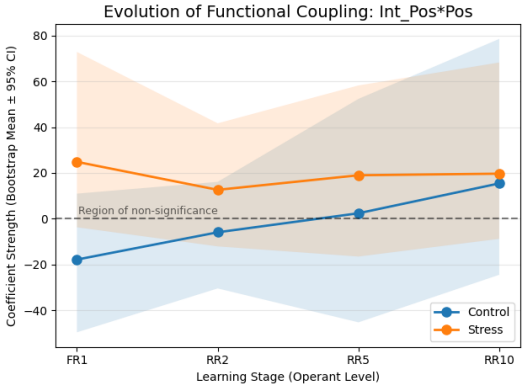

(d) Influence direction: $BLA \rightarrow CeA$

Figure 9: Evolution of key coupling coefficients during bout-initiating pressing. Parameter trajectories ($n = 1000$ bootstraps) show functional circuit reorganization from FR1 to RR10. (9(a), 9(c)) Self-feedback positional terms (x_{CeA} and x_{BLA}). Control mice (blue) maintain significantly negative coefficients (deterministic restoring forces), while stressed mice (orange) show a drift toward non-significance (dynamical erosion). (9(b), 9(d)) Nonlinear modulation term ($x_{CeA}x_{BLA}$). Most terms remain non-significant across groups, indicating modular stability. Shaded areas represent 95% CIs.

Table 7: Correlation Coefficients Between Steady-state KL Divergence and Real Part of Principle Eigenvalue

Pathway		Footshock		Unpredicted Reward	
		Correlation Coefficients	P-value	Correlation Coefficients	P-value
BLA-DMS	Control	0.212	1.177e-11*	-0.08	1.19e-02*
	Stress	-0.215	6.820e-12*	0.016	6.24E-01
CeA-DMS	Control	-0.357	1.832e-31*	0.49	1.738e-61*
	Stress	-0.092	3.548e-03*	0.327	2.647e-26*

* Indicates statistical significance.

this manifests as either the formation of a displaced pathological attractor (BLA-DMS, footshock) or the degradation of homeostatic restoring forces (CeA-DMS, reward).

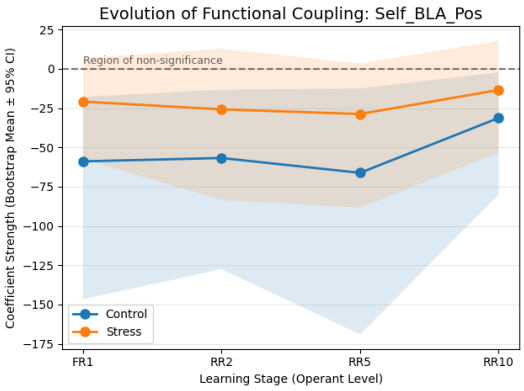

(a) Influence direction: $CeA \rightarrow BLA$

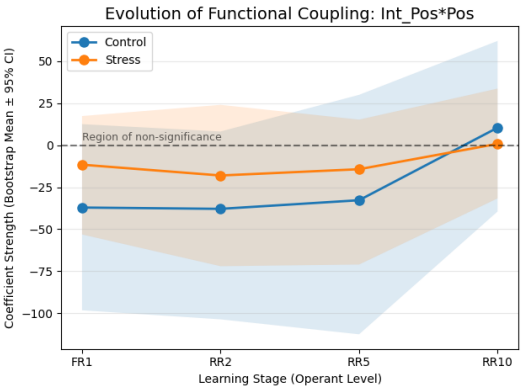

(b) Influence direction: $CeA \rightarrow BLA$

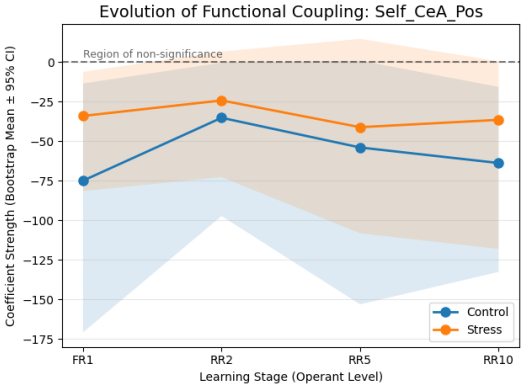

(c) Influence direction: $BLA \rightarrow CeA$

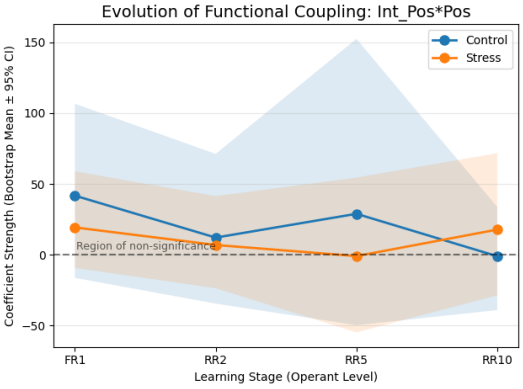

(d) Influence direction: $BLA \rightarrow CeA$

Figure 10: Dynamical coupling coefficients during reward processing. (10(a), 10(c)) Self-positional gains (x_{CeA} and x_{BLA}). Controls maintain significant negative set-points; stressed mice exhibit "slack" in gain control with high-variance, non-significant intervals. (10(b), 10(d)) Cross-pathway nonlinear modulation term. Consistent non-significance suggests stability is maintained via internal damping rather than altered inter-pathway coupling. Shaded areas represent 95% CIs.

Remarks: It is noteworthy that while the subject-level bootstrap procedure exhibits inherent stochastic fluctuations across iterations—reflecting biological heterogeneity and individual variability in stress susceptibility—the emergent statistical patterns remain globally consistent. The fundamental dynamical bifurcation between footshock-induced decoupling and the functional resilience observed in the Learned Modality is robustly preserved across multiple resampling validations. This cross-context stability provides a reliable empirical constraint for the parametric configuration of our subsequent generative neural network model.

3.2 ANNs Model Simulation

The architecture of the Unified Coupled System (UCS) was not merely a black-box fit; rather, it was explicitly designed to synthesize the two key findings from our preceding empirical analyses:

Statistical Inspiration (The Persistent State): Our KL divergence analysis (Section 2.2) revealed a "distributional hysteresis"—a persistent information divergence after perturbation. This suggested that the circuit requires a latent "memory" or "cell state" to maintain a reconfigured stochastic profile. Consequently, we utilized Adaptive LSTMs as the core units of the UCS, as their gated cell states provide a

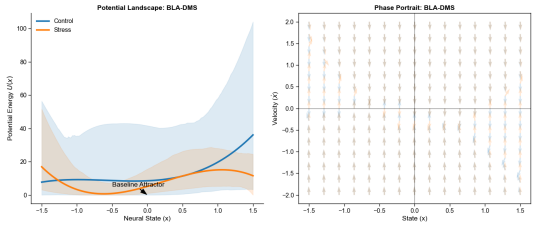

(a)

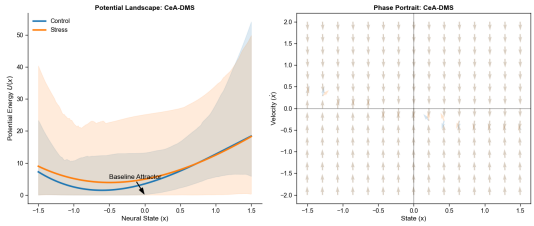

(b)

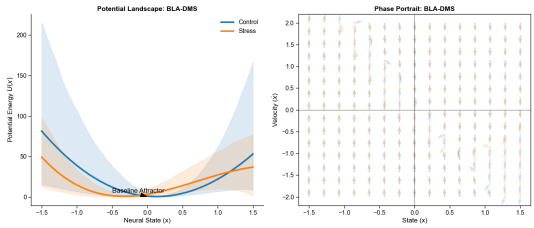

(c)

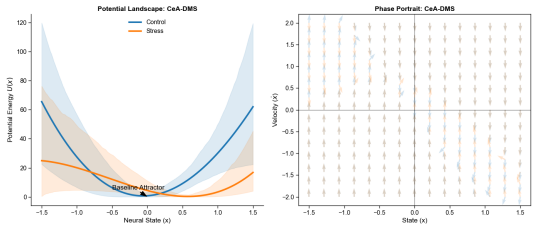

(d)

Figure 11: Global Attractor Geometry across Modalities. Potential landscapes and phase portraits of neural circuit dynamics. Blue: Control; Orange: Stress. Shaded areas: 95% bootstrap CIs. For each subgraph, left: Effective quasi-potential $U(x)$, right: phase portraits $(\dot{x}, x)$. 11(a) and 11(b) Non-learned Aversive Modality (Footshock): Stressed BLA–DMS (11(a)) exhibits a significant attractor displacement toward a biased state, while CeA–DMS (11(b)) maintains stability. 11(c)-11(d) Learned Reward Modality (Unpredicted Reward): Both BLA–DMS 11(c) and CeA–DMS 11(d) show qualitatively similar attractor geometries across groups, with noticeable landscape flattening in the stressed CeA–DMS.

natural computational substrate for this observed hysteresis.

Dynamical Inspiration (The Overdamped Set-point): Our dynamical modeling (Section 3.1) identified that the BLA-DMS sub-network is a high-precision but capacity-limited restorer, whereas the CeA-DMS sub-network acts as a high-gain, viscous damper. To formalize this, we translated these physical properties into Asymmetric Loss Functions (L_{BLA} and L_{CeA}). Specifically, L_{BLA} was constrained by a "brittle" saturation penalty to mimic the BLA-DMS sub-network's precision-limited dynamics, while L_{CeA} was

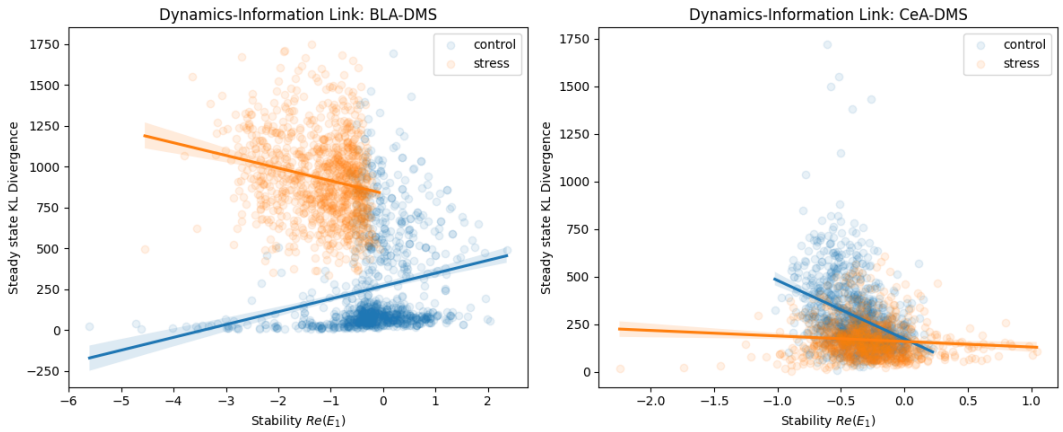

(a)

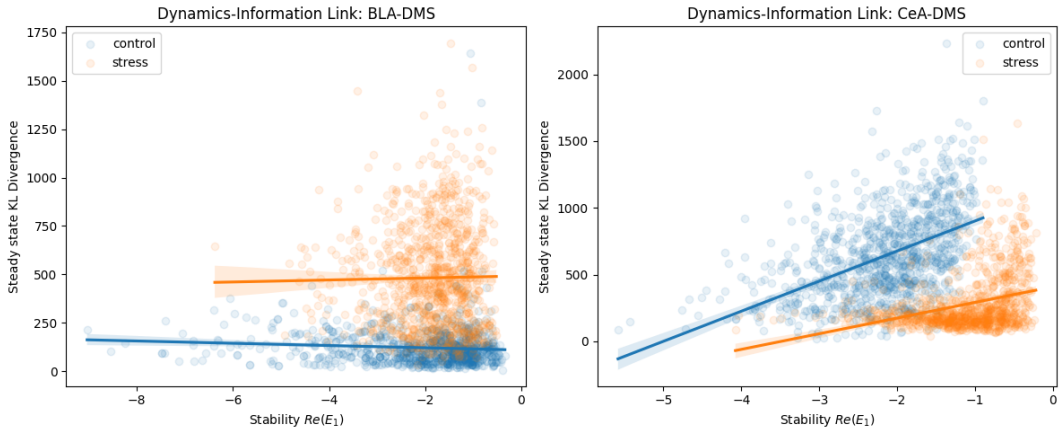

(b)

Figure 12: Correlation Between Dynamics and Information. Cross-layer validation of the stability-hysteresis link. Scatter plots: Relationship between the principal eigenvalue $Re(E_1)$ and steady-state KL divergence. In Fig. 12(a) the negative correlation in the stressed BLA–DMS (orange) highlights pathological state-locking at a displaced attractor. In Fig. 12(b) strong positive correlation in CeA-DMS underscores that informational integrity is tethered to the circuit's homeostatic restoring force ($Re(E_1)$). Correlation coefficients values and significance levels are detailed in Table 7

rewarded for informational persistence and buffering to mimic the CeA-DMS sub-network's damping-dominated regime.

By integrating these empirical constraints, the UCS functions as a mechanistic bridge: it tests whether the macroscopic shift from goal-directed to rigid behavior can emerge from the microscopic optimization of precision and robustness under stochastic environmental challenges. This allows us to move beyond descriptive modeling and provide a possible computational interpretation for the observed circuit remodeling.

To investigate the computational principles underlying stress-induced circuit remodeling, we trained independent instances of the Unified Coupled System (UCS). The models were tasked with adapting to distinct environmental statistics: a noise-buffered periodic baseline (Control) versus a regime characterized by frequent stochastic impulse perturbations (Stress).

The synthetic input sequences were designed to capture the core statistical features of the experimental environment (Figure 13): a periodic cosine component ($T = 2\pi/80$) representing predictable cycles,

Gaussian noise ($\sigma = 0.3$) for internal fluctuations, and discrete anomalous pulses (amplitude 1.5) as computational analogs of acute footshocks.

The stable convergence of the total loss function (L_{Total}) across training epochs (Figure 14) confirms that both models successfully internalized their respective environmental statistics. However, it was through this simulation that we identified a critical necessity: reproducing biological observations required the pathways to be optimized using asymmetric objectives (L_{BLA} vs. L_{CeA}). This suggests that functional divergence is not merely a consequence of connectivity but is driven by distinct optimization principles.

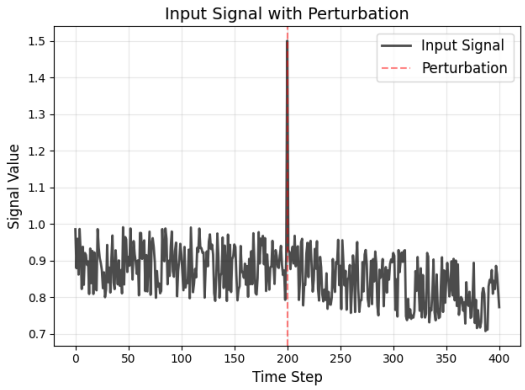

Figure 13: Simulation of test signals. The signal combines predictable periodicity with stochastic noise. The central spike mimics an acute high-impact perturbation (such as footshock), serving as a probe to test circuit stability.

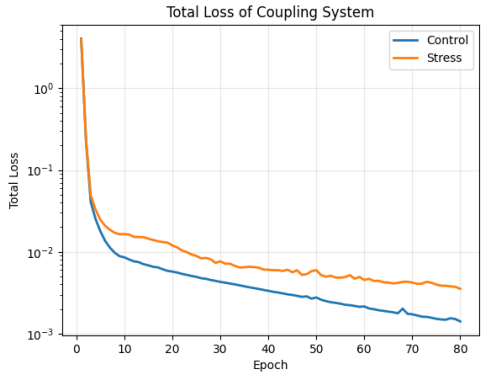

Figure 14: Training Convergence and Learning Stability. The trajectories of the total loss function for Control (blue) and Stress (orange) models demonstrate stable convergence across 80 epochs. The higher converged loss in the Stress group reflects the persistent informational overhead required to model a high-entropy environment.

To address the sensitivity of the model to initialization, we validated the architecture using 10 independent random seeds. The results demonstrate that the functional bifurcation is a robust emergent property of the asymmetric loss functions rather than a stochastic artifact.

We utilized the Max Deviation ($\text{Max}\delta = \max_t |\delta_t|$) as a quantitative metric for the peak computational load or "internal strain" experienced by each circuit axis (Figure 16).

- Mechanism of Brittle Precision: In the Control group, the BLA-DMS sub-network reaches a high Max Deviation (0.339 ± 0.07), which significantly exceeds its intrinsic capacity threshold ($\theta_{BLA} = 0.2$).

This structural overload manifests as a prolonged ecovery Time (3.00 ± 0.77 steps), indicating that precision-optimized systems are inherently vulnerable to high-impact perturbations.

- Mechanism of Functional Hardening: In the Stress group, the BLA-DMS sub-network is clamped to a low-load regime (Maxδ = 0.057), while the computational burden is shifted to the high-capacity CeA-DMS sub-network, which accommodates a significantly higher strain (Maxδ = 0.670). This allows for an instantaneous return to steady state (Recovery Time = 1.00 step).

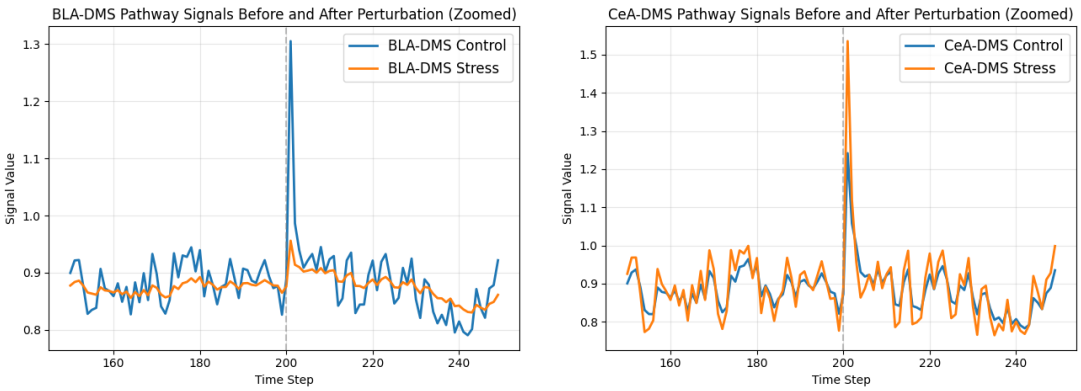

(a) Simulated signals of BLA-DMS sub-network (b) Simulated signals of CeA-DMS sub-network

Figure 15: Simulated Pathway Responses to Acute Perturbation. 15(a) The BLA-DMS sub-network shows high-fidelity tracking in controls but is suppressed in the stress group to prevent saturation. 15(b) The CeA-DMS sub-network exhibits amplified activation in the stress group, acting as a compensatory buffer to absorb the perturbation-induced strain.

This reallocation of internal strain from the BLA-DMS sub-network to the CeA-DMS sub-network provides a computational analog consistent with the observed shift toward reduced behavioral flexibility under stress. The CeA-DMS sub-network in the toy model undergoes a "functional hardening" that mimics habitual dynamics by prioritizing robust damping over predictive fidelity (Fig. 16).

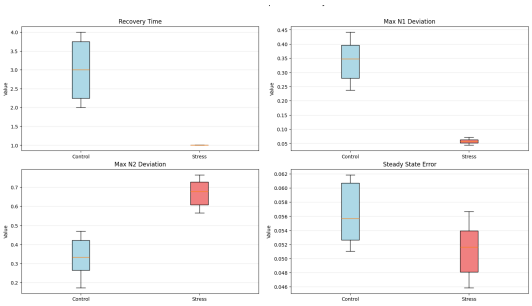

Figure 16: Robustness Metrics and Mechanistic Handover. Analysis across 10 independent random seeds suggests the stability of the UCS architecture. Top Left: Recovery time is significantly shorter in the Stress group, indicating a "hardened" equilibrium. Top Right/Bottom Left: Max Deviation (Maxδ) metrics illustrate the systematic shift of computational load from the capacity-limited BLA-DMS sub-network in controls to the robust CeA-DMS sub-network in stressed models. Bottom Right: Lower steady-state error in the stress group supports the hypothesis that CeA-DMS sub-network-mediated damping enhances global stability at the cost of informational fidelity.

A key emergent feature of our ANN simulations is the decoupling of structural variability from functional stability. Statistical analysis across 10 independent seeds revealed that while the sub-networks consistently enforced a rigid recovery trajectory under stress conditions (Recovery Time CV = 0), the underlying inhibitory coupling weights displayed significant stochasticity (CV $\approx$ 40%, Supplementary Table C3). This suggests that the transition to a "stress-hardened" state is a topological necessity of the optimization objective rather than a result of precise synaptic-like weight matching. Such a property aligns with the concept of computational degeneracy in neural circuits, where stable functional outputs are maintained despite high variability in the underlying structural components.

4 Conclusion and Discussion

The present study provides a concordant multi-level computational analysis of amygdala-striatal dynamics under chronic stress, integrating distributional, dynamical, and modeling perspectives. Across these levels, our results consistently indicate that stress reshapes the structure of circuit activity, leading to more constrained and less flexible state-space trajectories.

At the statistical level, we quantified changes in neural activity distributions using KL divergence. This measure captures the degree of deviation between baseline and post-perturbation states, and thus serves as an informational descriptor of how circuit activity reorganizes under stress. Importantly, we interpret these results strictly in a statistical sense, without attributing them to specific biological processes such as energy expenditure or synaptic plasticity. Within this framework, the observed persistence of distributional differences following perturbation–referred to here as distributional hysteresis–reflects a sustained alteration in the circuit's stochastic state space.

At the dynamical level, we modeled the evolution of neural activity using a second-order system. This formulation is motivated by the temporal filtering properties of calcium signals, which naturally give rise to higher-order dynamics in the observed data. Within this phenomenological framework, the identified coefficients describe effective restoring forces, damping, and nonlinear corrections that shape the trajectory of the system in state space. Across conditions, we observe that stress is associated with increased damping and more constrained trajectories. Crucially, the dominance of negative velocity coefficients ($\dot{x}$) ensures that the system remains in a robustly overdamped regime, precluding oscillatory dynamics and favoring rapid state stabilization. The analysis of coupling coefficients provides further insight. Notably, the nonlinear self-stabilization terms (e.g., x_{BLA}^2) that appear significant in univariate stress models lose their dominance when inter-pathway interactions are included. This redistribution of explanatory power suggests that the observed circuit rigidity is not merely an intrinsic property of a single projection, but an emergent consequence of distorted coupling between pathways. Rather than simple excitation or inhibition, these parameters reflect a state-dependent reconfiguration of the system's Jacobian, prioritizing global dissipation over autonomous node restoration.

The nonlinear terms in the model further suggest that the structure of the state space differs between conditions. In particular, stress is associated with stronger state-dependent corrections, which are consistent with a more restricted set of accessible trajectories. These effects should be understood as properties of the fitted dynamical system and not as direct indicators of structural or synaptic changes in the underlying circuit. Our analysis of coupling coefficients provides additional insight into how interactions between pathways shape circuit dynamics. Rather than interpreting these coefficients as direct excitatory or inhibitory interactions, we treat them as effective parameters that govern how variables influence one another within the dynamical system. In this sense, cross-pathway terms reflect state-dependent modulation of system stability and trajectory geometry, as formalized through the local Jacobian.

To further explore the computational implications of these dynamics, we introduced a minimal artificial neural network (ANN) model. This model was not designed to replicate biological circuitry, but to test whether asymmetric optimization constraints are sufficient to reproduce the observed dissociation in

dynamical behavior. Within this framework, we observe a redistribution of activity patterns across units under stress conditions. This result demonstrates that relatively simple computational asymmetries can give rise to distinct dynamical regimes. It should be emphasized that this model provides a proof-of-principle illustration and does not imply that biological circuits implement the same optimization strategy or that the observed behavior is uniquely determined by this architecture.

Several limitations should be noted. First, the models employed in this study are phenomenological and are not intended to capture detailed biophysical mechanisms. Second, while bootstrap and sensitivity analyses support the robustness of our findings, other sources of variability (such as inter-subject differences and sampling effects) may also influence the estimated parameters. Third, the ANN framework represents only one possible computational instantiation, and alternative models may equally account for the observed patterns.

In summary, chronic stress drives a systematic reorganization of amygdala–striatal dynamics, characterized by increased damping, a more restricted state-space structure, and persistent distributional shifts. Our findings suggest that stress-induced circuit remodeling may be understood as a transition toward a rigid, alternative stable equilibrium, providing a potential computational basis for the behavioral shift from goal-directed to habit-like patterns. By moving the analytical focus from individual component damage to the distortion of dynamical relationships, this study establishes a rigorous foundation for quantifying the loss of neural flexibility under chronic stress.

Declaration of Competing Interest

None.

Declaration of generative AI in scientific writing

Authors declare no using of generative AI.

Acknowledgements

We thank the authors of Giovanniello et al. (2025) [5] for sharing their datasets publicly.

This article is dedicated to the memory of my father, Weize Lin (born 1941-07-29 – passed away 2025-05-19). For a long time, I sought to understand his complex and challenging behaviors through the lens of psychological theory. However, at the beginning of 2025, I had a sudden epiphany: that these behavioral patterns, while appearing psychological on the surface, are essentially the outputs of a neural network trained and shaped by experience. This realization led me to develop a computational model to simulate these dynamics. Shortly thereafter, in February 2025, I encountered the relevant research published in Nature that provided the precise experimental datasets needed to validate this modeling approach. Although he passed away before the completion of this work, this study stands as a tribute to his profound influence on my thinking, representing the culmination of a journey from personal observation to scientific discovery during a time of great loss.

Funding Statement

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Author contributions

Contribution of the authors: Feng Lin analyzed the data and constructed the framework, developed the analytical methodology, implemented the codes and wrote the manuscript.

Data availability

The data used in this study were originally reported by Giovanniello et al. [5] and are publicly available in their supplementary materials. The datasets used for secondary analysis can be accessed through the publication and relevant repositories, subject to any access restrictions specified by the original authors.

Ethics Statement

This article is based on a secondary analysis of data originally reported by Giovanniello et al. [5]. As this is a secondary data analysis, no new ethical approval or animal research was conducted for this study. All experimental procedures involving animals were performed in accordance with the ethical standards established by the original study.

Statements

All authors have read and approved the final version of the manuscript. Feng Lin had access to the publicly available data used in this study and takes full responsibility for the integrity of the data analysis.

References

[1] Jamie Ward. *The student's guide to cognitive neuroscience*. Routledge, 4th edition, 2019. ISBN 9781351035187. doi: https://doi.org/10.4324/9781351035187.

[2] Christian A Klöckner. A comprehensive model of the psychology of environmental behaviour–a meta-analysis. *Global environmental change*, 23(5):1028–1038, 2013. doi: https://doi.org/10.1016/j.gloenvcha.2013.05.014.

[3] Hang-Yee Chan, Maarten AS Boksem, Vinod Venkatraman, Roeland C Dietvorst, Christin Scholz, Khoi Vo, Emily B Falk, and Ale Smidts. Neural signals of video advertisement liking: Insights into psychological processes and their temporal dynamics. *Journal of Marketing Research*, 61(5):891–913, 2024. doi: https://doi.org/10.1177/00222437231194319.

[4] Ben Hannigan, Deborah Edwards, and Philip Burnard. Stress and stress management in clinical psychology: Findings from a systematic review. *Journal of Mental Health*, 13(3):235–245, 2004. doi: https://doi.org/10.1080/09638230410001700871.

[5] Jacqueline R Giovanniello, Natalie Paredes, Anna Wiener, Kathia Ramírez-Armenta, Chukwuebuka Oragwam, Hanniel O Uwadia, L Yu Abigail, Kayla Lim, Jenna S Pimenta, Gabriela E Vilchez, et al. A dual-pathway architecture for stress to disrupt agency and promote habit. *Nature*, 2025. doi: https://doi.org/10.1038/s41586-024-08580-w.

[6] Christina M Gremel and Rui M Costa. Orbitofrontal and striatal circuits dynamically encode the shift between goal-directed and habitual actions. *Nature communications*, 4(1):2264, 2013. doi: https://doi.org/10.1038/ncomms3264.

[7] Dharmik Upadhyay and Hardik Joshi. Mathematical modeling of local calcium signaling in neurons using artificial neural networks. *Discrete and Continuous Dynamical Systems-S*, 18(5):1392–1415, 2025. doi: https://doi.org/10.3934/dcdss.2025010.

[8] Hardik Joshi and Brajesh Kumar Jha. 2d dynamic analysis of the disturbances in the calcium neuronal model and its implications in neurodegenerative disease, 2023.

[9] Hardik Joshi. Mechanistic insights of covid-19 dynamics by considering the influence of neurodegeneration and memory trace. *Physica Scripta*, 99(3):035254, 2024. doi: https://doi.org/10.1088/1402-4896/ad2ad0.

[10] KE Kürten and JW Clark. Chaos in neural systems. *Physics Letters A*, 114(7):413–418, 1986. doi: https://doi.org/10.1016/0375-9601(86)90729-2.

[11] Christine A Skarda and Walter J Freeman. How brains make chaos in order to make sense of the world. *Behavioral and brain sciences*, 10(2):161–173, 1987. doi: https://doi.org/10.1017/S0140525X00047336.

[12] Feng Lin. Hidden complexity of biological clocks that adapt to external control and internal regulation: a case study on body temperature simulation. *Biological Rhythm Research*, 55(5-6):334–353, 2024. doi: https://doi.org/10.1080/09291016.2024.2384635.

[13] Geoffrey E Hinton. How neural networks learn from experience. *Scientific American*, 267(3):144–151, 1992. doi: https://doi.org/10.1038/scientificamerican0992-144.

[14] Vincent A Magnotta, Dan Heckel, Nancy C Andreasen, Ted Cizadlo, Patricia Westmoreland Corson, James C Ehrhardt, and William TC Yuh. Measurement of brain structures with artificial neural networks: two-and three-dimensional applications. *Radiology*, 211(3):781–790, 1999.

[15] Bokai Cao, Xiangnan Kong, and Philip S Yu. A review of heterogeneous data mining for brain disorder identification. *Brain informatics*, 2(4):253–264, 2015. doi: https://doi.org/10.1007/s40708-015-0021-3.

[16] David Pascucci, Maria Rubega, and Gijs Plomp. Modeling time-varying brain networks with a self-tuning optimized kalman filter. *PLoS computational biology*, 16(8):e1007566, 2020. doi: https://doi.org/10.1371/journal.pcbi.1007566.

[17] Vanessa McNealis and Christian Léger. Smoothed pseudo-population bootstrap methods with applications to finite population quantiles. *Journal of Survey Statistics and Methodology*, page smaf036, 2026. doi: https://doi.org/10.1093/jssam/smaf036.

[18] John J Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, 79(8):2554–2558, 1982. doi: https://doi.org/10.1073/pnas.79.8.2554.

[19] Mark M Churchland, John P Cunningham, Matthew T Kaufman, Justin D Foster, Paul Nuyujukian, Stephen I Ryu, and Krishna V Shenoy. Neural population dynamics during reaching. *Nature*, 487(7405):51–56, 2012. doi: https://doi.org/10.1038/nature11129.