BUILDING NETWORK FUNCTIONS FROM SIGNED FEEDFORWARD LOOPS

Anonymous authors
Paper under double-blind review

ABSTRACT

Connectivity shapes computation by determining how signals propagate and interact within a network. Compared with analyses of individual connections, the feedforward loop (FFL) provides a minimal multipath structure in which direct and indirect pathways converge on the same target through a relay node. FFLs have been widely studied for their computational capacity, and their excitatory–inhibitory (E-I) configurations are known to determine distinct dynamical regimes. However, the role of E-I organization in the functional specialization of FFLs remains unclear. Here, we systematically analyze the dynamics of four E-I configurations of the FFL and characterize how direct and indirect pathways interact to shape computation. Under defined input and background conditions, these dynamics support distinct computational functions, including 1) response initiation, 2) response switching, 3) stopping of ongoing responses, and 4) context-dependent release. We then extend motif regularization with a differentiable signed-motif loss that increases the proportion of selected FFLs in larger networks, translating local computational mechanisms into structural priors for network learning. Behavioral evaluations in MiniHack show that these motif-enriched networks exhibit task-dependent computational biases consistent with their local functional specialization, and that a learned scheduler can coordinate components shaped by different motif priors to improve performance in more complex environments. Together, these results show how E-I organization shapes the functional specialization of FFLs and how local signed pathway dynamics can be transformed into reusable computational building blocks for constructing and composing functions.

1 INTRODUCTION

The role of connectivity in shaping computational capabilities is a central focus in both artificial intelligence and neuroscience (Battaglia et al., 2018; Driscoll et al., 2024). Recurrent and state-space models provide a common framework across both fields, in which structural organization shapes how internal states evolve over time (Barak, 2017; Durstewitz et al., 2023; Gu & Dao, 2023). Within this framework, connectivity structure has been linked to population dynamics underlying working memory, decision-making, and cognitive control (Mastrogiuseppe & Ostojic, 2018; Yang et al., 2019; Dubreuil et al., 2022). This motivates our focus on small, analytically tractable connectivity structures, which allow us to explicitly trace how connectivity shapes dynamic regimes and informs structural design principles for artificial recurrent networks.

Network motifs provide a compact statistical framework for characterizing small circuits in terms of recurring local connectivity patterns (Milo et al., 2002; Alon, 2007). Among these motifs, the feedforward loop (FFL) is a minimal multi-pathway structure in which a source reaches a target through both a direct connection and an indirect route containing a relay. FFLs have been widely studied in network research for temporal regulation and signal integration (Mangan & Alon, 2003; Pouille & Scanziani, 2001; Cao & Van Hooser, 2022), and more recently for representation and robustness (Zhang et al., 2025). Across these roles, connection signs and excitatory–inhibitory (E-I) organization determine FFL classification and pathway interactions (Mangan & Alon, 2003; Zhang et al., 2025; Szilagyi et al., 2026), while motif–E-I interactions shape collective dynamics, stability, and responses to external inputs (Shao & Ostojic, 2023; Shao et al., 2025). Building on these foundations, we seek to determine how distinct E-I implementations of a common FFL topology

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

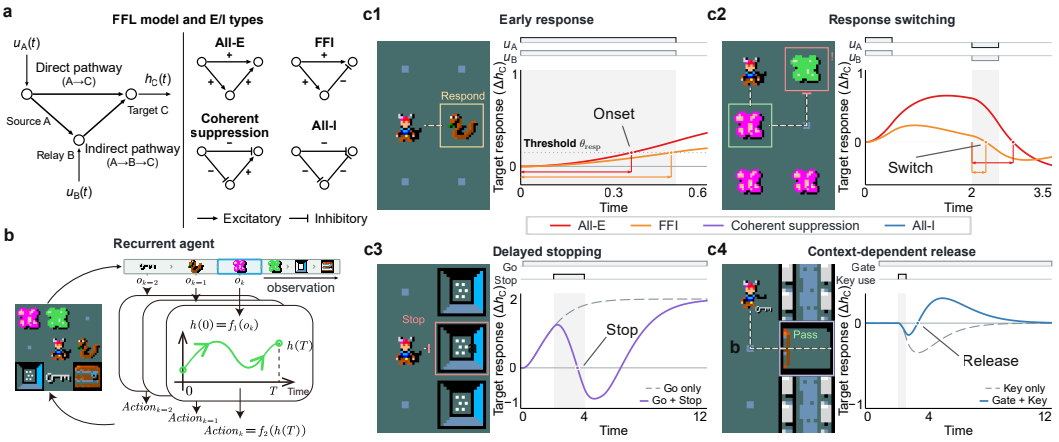

Figure 1: **Signed FFL dynamics motivate task-dependent structural priors.** **a**, Source–relay–target topology and four E-I configurations. **b**, Observation–action loop of a recurrent agent. **c1–c4**, Illustrative local target responses and corresponding task demands. **c1**, Early response and threshold crossing. **c2**, Response switching following a negative input. **c3**, Stopping an ongoing Go response after a delayed Stop signal. **c4**, Context-dependent transient release from inhibition. Matched numerical trajectories and protocol details are provided in Fig. 3 and Appendix C, respectively.

support interpretable computational operations, asking whether such mechanistic insights can guide structural design in recurrent networks.

Recent work has shown that motif organization and structural constraints can shape representational capacity, robustness, flexibility, and learned network function (Zhang et al., 2025; Balwani et al., 2025; Zhang et al., 2026). However, it remains unclear how E-I organization shapes computation within a fixed FFL topology and whether these computational capabilities can guide task-dependent learning. We therefore compare four signed FFLs under a common dynamical model and matched input protocols to identify how their pathway interactions give rise to distinct computational regimes. In all-E, convergent excitation strengthens early activation, supporting timely response initiation. In feedforward inhibition (FFI), the transition from excitation to inhibition limits the persistence of a response, providing a mechanism for response switching as inputs change. Coherent suppression combines direct inhibition with withdrawal of excitation to interrupt an established response when a later Stop signal arrives. In all-I, a signal can suppress an active inhibitory relay and transiently release the target from inhibition, enabling context-dependent gating. These mechanisms motivate matching signed FFL structures to distinct task demands (Figure 1).

To translate these local mechanisms into trainable network structure, we extend matrix-based motif regularization (Zhang et al., 2026) to differentiable counts of the four signed FFLs. A signed-motif loss encourages selected configurations in recurrent connectivity while preserving fixed neuronal E-I identities. We then compare networks trained with different motif priors across four MiniHack tasks (Samvelyan et al., 2021) that recruit these operations. Specifically, the combat task tests timely attack initiation, and the mushroom task tests response switching when an edible mushroom is followed by a poisonous one. The trap task, in turn, requires stopping ongoing movement when a hazard appears, while the key task requires passing through a door during the temporary opening enabled by key use. Together, these comparisons test whether the effects of motif priors on learning depend on the match between circuit function and task demand. Finally, we investigate whether a learned scheduler can coordinate components shaped by different motif priors as the demands of a single complex task change.

Our contributions are as follows:

- We analyze the dynamics of four signed E-I feedforward loops and derive how direct and indirect pathways give rise to distinct computational biases under matched input conditions.
- We develop a differentiable signed-motif regularizer that embeds these circuit configurations as structural priors, producing task-dependent learning biases consistent with their local functions.

- We introduce a learned scheduler to coordinate motif-shaped experts across task contexts, thereby linking local circuit dynamics to compositional network function.

2 METHODS

2.1 COMMON TOPOLOGY AND DYNAMICAL MODEL OF SIGNED FFLS

We first consider a three-node continuous-time recurrent model

$$\frac{\mathrm{d}h_i(t)}{\mathrm{d}t} = -\alpha h_i(t) + \sum_j W_{ij}\phi(h_j(t)) + u_i(t), \qquad i \in \{A, B, C\}, \tag{1}$$

where $h_i(t) \in \mathbb{R}$ is the state of node i, $\alpha > 0$ is the shared leak rate, ϕ is the activation function, and $u_i(t)$ is the external drive to node i. The coupling W_{ij} acts from node j to node i. We restrict the model to an FFL with source A, relay B, and target C, allowing only W_{BA}, W_{CA}, and W_{CB} to be nonzero (Milo et al., 2002; Mangan & Alon, 2003). Under Dale's principle, each node's outgoing connections share a sign, yielding four configurations: All-E $(+,+,+)$, FFI $(+,+,-)$, coherent suppression $(-,-,+)$, and All-I $(-,-,-)$, with signs ordered as (W_{BA}, W_{CA}, W_{CB}) (Fig. 1a).

2.2 LINEARIZATION AND INPUT-DRIVEN SOLUTION

To isolate how signed pathways shape target dynamics, we linearize Eq. (1) around a zero-input equilibrium $\mathbf{h}^*$. Defining $\mathbf{z} = \mathbf{h} - \mathbf{h}^*$ and assuming a continuously differentiable activation with $\phi' \geq 0$, we obtain

$$\dot{\mathbf{z}}(t) = \mathbf{J}\mathbf{z}(t) + \mathbf{u}(t), \qquad \mathbf{J} = -\alpha\mathbf{I} + \mathbf{F}, \tag{2}$$

where $\mathbf{I}$ is the identity matrix and $F_{ij} = W_{ij}\phi'(h_j^*)$. Nonnegative activation derivatives ensure that effective couplings do not reverse the original E-I signs (Appendix B.1). Setting $u_C = 0$, we next derive the target response to inputs at A and B.

Proposition 1 (Kernels and input-driven response)**.** *For all four E-I assignments, the unforced linearized system is exponentially stable for $\alpha > 0$. At lag $\tau \geq 0$, the input–output kernels are*

$$K_{A\to C}(\tau) = e^{-\alpha\tau}\left(\tau F_{CA} + \frac{\tau^2}{2}F_{CB}F_{BA}\right), \qquad K_{B\to C}(\tau) = e^{-\alpha\tau}\tau F_{CB}. \tag{3}$$

For $\mathbf{z}(0) = \mathbf{0}$, the target response is

$$z_C(t) = \int_0^t \left[K_{A\to C}(\tau)u_A(t-\tau) + K_{B\to C}(\tau)u_B(t-\tau)\right]\mathrm{d}\tau. \tag{4}$$

These expressions are exact for the linearized model (Appendix B.2). The direct and indirect terms in $K_{A\to C}$ have the same sign for an excitatory relay and opposite signs for an inhibitory relay, when both paths are active.

2.3 FUNCTIONAL PROTOCOLS AND READOUTS

We next define four matched input–output protocols for isolated three-node FFLs to test how different sign assignments alter circuit function. All protocols use source input u_A, relay input u_B, no direct target input ($u_C = 0$), and the signed state h_C as output. We write $\Delta h_C = h_C - h_{C,\mathrm{base}}$ for changes from the condition-specific baseline. Within each protocol, all four motifs receive identical inputs and use the same baseline rule and readout. Functional coverage is the fraction of sampled gains with a positive score. Input settings, parameter grids, and numerical details are specified in Appendix C. Together, these matched comparisons distinguish the extent of a functional regime from an isolated score advantage at particular parameters.

Early response. To measure the speed and strength of early activation, we apply simultaneous positive pulses to A and B. Let P be the peak of Δh_C in a fixed early window and T_{on} the time to reach a threshold $\theta_{\mathrm{resp}} > 0$. Define

$$S_{\mathrm{on}} = (P - \theta_{\mathrm{resp}})/\theta_{\mathrm{resp}}, \qquad Q_{\mathrm{on}} = S_{\mathrm{on}} - \lambda[S_{\mathrm{on}}]_+ T_{\mathrm{on}}, \tag{5}$$

where $[z]_+ = \max(z, 0)$ and $\lambda > 0$ weights the latency penalty. A positive score requires a peak above threshold, and earlier threshold crossing leads to a higher score.

Response switching. To assess switching of successive inputs, we apply a positive signal followed by a negative signal to both input channels. Let μ_+ be the mean Δh_C in the first readout window and μ_- the negative of the mean Δh_C in the second. Let T_{switch} be the waiting time for a negative response after the second signal begins. Define

$$S_{\text{seq}} = \tanh(k \min(\mu_+, \mu_-)), \qquad Q_{\text{seq}} = S_{\text{seq}} - \lambda[S_{\text{seq}}]_+ T_{\text{switch}}, \tag{6}$$

where a large k emphasizes switching success over response magnitude. A positive score indicates a correct response, and faster, stronger responses receive higher scores.

Delayed stopping. To assess suppression after an initial response has developed, we initiate Go input at B and introduce Stop input at A after a delay $t_{\text{Stop}} - t_{\text{Go}} > 0$, while Go remains present. The Go control omits Stop. At a fixed post-Stop readout t_V, define

$$V = \min[h_{C,\text{Go}}(t_V), -h_{C,\text{Stop}}(t_V)], \tag{7}$$

where $h_{C,\text{Stop}}$ denotes Go followed by Stop. A positive score requires suppression of the Go response below zero after Stop, while the Go-only control remains positive at the readout. Thus, the metric operationalizes delayed stopping as active suppression of an established Go response.

Context-dependent release. To test whether background drive changes the sign of a signal's effect, we apply a finite signal to A with tonic drive to B present or absent. For condition labels $(\text{context}, \text{signal}) \in \{00, 01, 10, 11\}$ and readout time t_G, define

$$G_{\text{gate}} = \min[\Delta h_{C,11}(t_G) - \Delta h_{C,10}(t_G), \Delta h_{C,00}(t_G) - \Delta h_{C,01}(t_G)]. \tag{8}$$

A positive score requires signal-induced release with context and suppression without context, relative to the corresponding no-signal controls. This defines context-dependent release controlled by the background state.

2.4 Differentiable Signed-FFL Priors

To test whether enriching a motif subtype can support a corresponding computation in a trained RNN, we extend the matrix-based counting and regularization framework of Zhang et al. (2026) to the four signed FFLs. The resulting differentiable motif regularizer penalizes deviations from target frequencies and is combined with the PPO objective (Fig. 2).

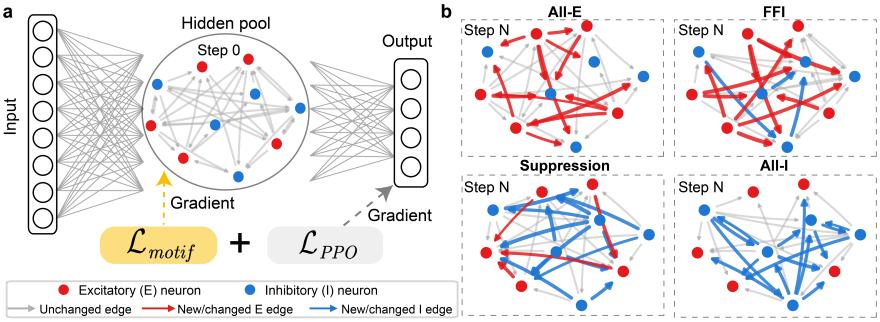

Figure 2: **Embedding E-I motif priors into RNN training. a**, RNN training combines the motif regularizer $\mathcal{L}_{\text{motif}}$ (Eq. (12)) with the PPO task loss $\mathcal{L}_{\text{PPO}}$. **b**, Motif regularization reshapes a common initial network to enrich All-E, FFI, coherent suppression, or All-I motifs (Appendix D.2).

2.4.1 Exact E-I identities with differentiable edge presence

Following Dale-constrained RNNs (Song et al., 2016; Cornford et al., 2021), we fix neuronal E-I identities $\mathbf{s} \in \{-1, +1\}^n$ for $n \geq 3$ neurons. Let $\mathbf{S} = \text{diag}(\mathbf{s})$ and $P_{ij} = \mathbf{1}_{i \neq j}$. We parameterize $\mathbf{W} = \mathbf{D}_{\text{self}} + \mathbf{MS}$, with $\mathbf{M} = \mathbf{P} \odot \text{softplus}(\mathbf{\Theta}) \geq 0$. Here $\mathbf{\Theta}$ contains trainable magnitude parameters, $\mathbf{D}_{\text{self}}$ is diagonal, and $\odot$ denotes elementwise multiplication. Off-diagonal outgoing signs are fixed, and self-connections are excluded from motif counting.

To make edge presence differentiable, we use a sigmoid σ with threshold $\theta_{\text{edge}} > 0$ and sharpness $\beta > 0$, defining $\widetilde{\mathbf{M}} = \mathbf{P} \odot \sigma(\beta(\mathbf{M} - \theta_{\text{edge}}))$, $\mathbf{Q}_E = (\mathbf{I} + \mathbf{S})/2$, and $\mathbf{Q}_I = (\mathbf{I} - \mathbf{S})/2$. The masks for excitatory, inhibitory, and absent connections are

$$\widetilde{\mathbf{W}}^+ = \widetilde{\mathbf{M}}\mathbf{Q}_E, \qquad \widetilde{\mathbf{W}}^- = \widetilde{\mathbf{M}}\mathbf{Q}_I, \qquad \widetilde{\mathbf{W}}^0 = \mathbf{P} - \widetilde{\mathbf{M}}. \tag{9}$$

Thus, neuronal E-I identity remains discrete and fixed, whereas edge presence remains differentiable and can be optimized during training.

2.4.2 FFL PROPORTIONS AND TRAINING OBJECTIVE

We control overall FFL abundance and its distribution across the four E-I subtypes. Let $\widetilde{C}_r$ be the soft count of subtype r, requiring three forward connections and excluding their reverse connections (Appendix D.1). Following the connected three-node motif counting and normalization of Zhang et al. (2026), let C_n^3 denote the total three-node motif count. With $\widetilde{C}_{\text{FFL}} = \sum_r \widetilde{C}_r$, define

$$\widetilde{f}_{\text{FFL}} = \frac{\widetilde{C}_{\text{FFL}}}{C_n^3}, \qquad \widetilde{\pi}_r = \frac{\widetilde{C}_r}{\widetilde{C}_{\text{FFL}}}. \tag{10}$$

Here $\widetilde{f}_{\text{FFL}}$ is the overall FFL proportion and $\widetilde{\pi}_r$ a soft analogue of $P(r \mid \text{FFL})$. Task learning uses the PPO loss (Schulman et al., 2017):

$$\mathcal{L}_{\text{PPO}} = -\mathcal{L}_{\text{actor}} + c_1\mathcal{L}_{\text{critic}} - c_2\mathcal{L}_{\text{entropy}}. \tag{11}$$

The three terms are the clipped policy surrogate, squared value-prediction error, and policy entropy, with nonnegative weights c_1 and c_2. The motif regularizer penalizes deviations from target FFL proportions and subtype shares:

$$\mathcal{L}_{\text{motif}} = \lambda_{\text{FFL}}(\widetilde{f}_{\text{FFL}} - f_{\text{FFL}}^{\text{target}})^2 + \sum_r \lambda_r(\widetilde{\pi}_r - \pi_r^{\text{target}})^2. \tag{12}$$

The total objective is $\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{PPO}} + \mathcal{L}_{\text{motif}}$. Figure D2 illustrates the resulting connectivity changes. Thus, the training objective jointly optimizes task performance while controlling overall FFL abundance and signed-subtype composition through separate structural penalties.

3 RESULTS

3.1 SIGNED FFLS SUPPORT DISTINCT INPUT-DEPENDENT COMPUTATIONS

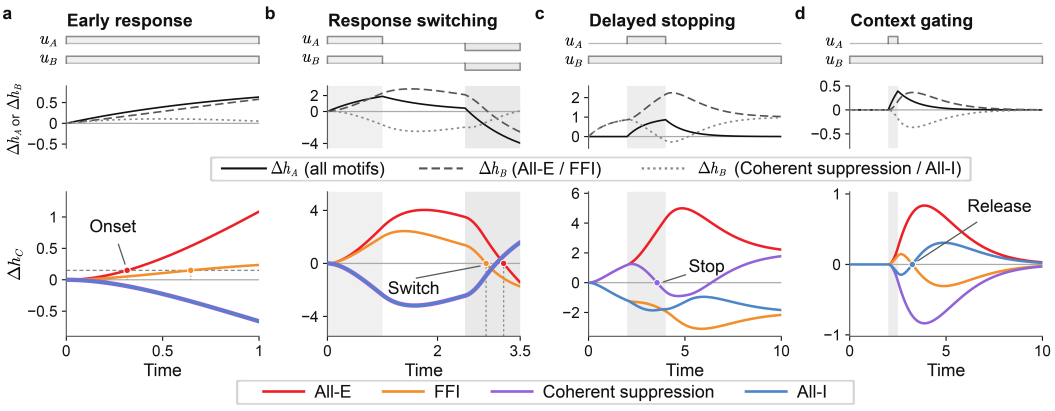

Figure 3: **Matched dynamics of four signed FFLs.** Columns show early response (a), switching (b), stopping (c), and context gating (d). Top rails indicate input timing with schematic amplitudes. Upper plots show source and relay states; lower plots compare target responses across the four motifs. Protocol details are given in Appendix C.

We first test whether different E-I sign assignments produce distinct functional responses within the same FFL topology. To do so, we compare the four signed FFLs using the matched input–output

protocols defined in Section 2.3, with identical inputs, connection magnitudes, baseline rules, and readouts, varying only their E-I signs. The four protocols differ in their input structure: 1) **early response**, simultaneous positive inputs to A and B; 2) **response switching**, a positive input followed by a negative input to both channels; 3) **delayed stopping**, a Go input to B followed by a delayed Stop input to A while Go remains active; and 4) **context-dependent release**, a finite signal to A with tonic context input to B present or absent.

Under these matched conditions, All-E promotes rapid response initiation through convergent excitation. With response initiation defined by crossing a common activation threshold, All-E crosses first at $T_{\text{on}} = 0.316$, followed by FFI at 0.646, whereas coherent suppression and All-I remain below threshold under the same conditions (Fig. 3a). In FFI, inhibition through the indirect $A \to B \to C$ pathway opposes excitation through the direct $A \to C$ pathway, allowing the target to switch from a positive to a negative response sooner after the input changes. The switching latency is $T_{\text{switch}} = 0.381$, compared with 0.700 for All-E. Coherent suppression and All-I fail to establish the initial positive response required for the positive–negative sequence (Fig. 3b). In coherent suppression, Go input to the excitatory relay B first establishes a positive target response. A subsequent Stop input to A directly inhibits C and suppresses B, reducing its excitatory drive to C. The target falls below zero while the Go-only control remains positive, and recovers after Stop ends. The other three motifs cannot achieve delayed stopping under the same conditions (Fig. 3c). In All-I, tonic context maintains an active inhibitory relay B. A signal to A directly inhibits C while also suppressing B, relieving B-mediated inhibition and producing a transient increase from the inhibited baseline. The target subsequently returns to that baseline. The other three motifs cannot achieve this context-dependent release under the same conditions (Fig. 3d).

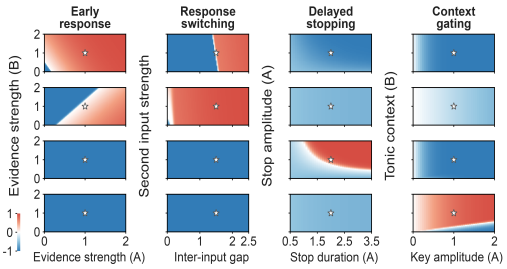

Figure 4: **Functional regimes across input conditions.** Rows (top to bottom): All-E, FFI, coherent suppression, and All-I. Columns show functions. Red/blue indicates positive/negative scores. Stars mark the settings in Fig. 3.

We next asked whether the functions shown in Fig. 3 remain supported across different input conditions. We independently varied the input amplitudes at A and B to assess whether each motif can respond promptly under different input strengths and balances. All-E achieves positive early-response scores on 98.49% of the sampled input grid, compared with 58.45% for FFI and 0% for coherent suppression and All-I (Fig. 4, first column). For response switching, we varied the interval between successive inputs and the strength of the second, negative input to assess whether each motif can promptly switch its response to a new signal. FFI supports switching at shorter input intervals and achieves positive switching scores on 99.41% of the sampled input grid, compared with 41.36% for All-E and 0% for coherent suppression and All-I (Fig. 4, second column). For delayed stopping, we varied the amplitude and duration of the Stop input to determine how strong and how sustained a signal must be to interrupt an ongoing response. Coherent suppression achieves positive stopping scores on 50.27% of the sampled input grid, compared with 0% for All-E, FFI, and All-I (Fig. 4, third column). For context-dependent release, we varied the amplitude of the brief signal at A and the sustained background input to B to assess how signal strength and background activity jointly control transient release from inhibition. All-I achieves positive gating scores on 83.22% of the sampled input grid, compared with 0% for All-E, FFI, and coherent suppression (Fig. 4, fourth column). Table C1 summarizes the input-scan coverage.

We further examined whether these functions persist across different connection strengths by scanning all three connection magnitudes over a shared 25^3 grid, with inputs fixed within each functional protocol (Fig. C1). All-E achieves positive early-response scores on 97.26% of the sampled weight combinations, compared with 37.14% for FFI, 1.13% for coherent suppression, and 0% for All-I. For response switching, FFI achieves positive scores on 70.36% of the grid, compared with 40.39% for All-E, 1.02% for coherent suppression, and 0% for All-I. Coherent suppression achieves positive stopping scores on 56.83% of the grid, compared with 0% for the other three motifs. All-I achieves positive gating scores on 39.32% of the grid, compared with 0% for the other three motifs (Table C2).

Together, these results show that different E-I sign assignments support distinct computations within the same FFL topology in the isolated three-node FFL model. These functions persist across the tested input and gain ranges.

3.2 Motif enrichment and task-dependent network biases

To examine whether motif-dependent computational biases transfer from local circuits to network behavior, we evaluate motif-enriched recurrent agents on four MiniHack environments with distinct computational demands. The combat, mushroom, trap, and key rooms require response initiation, response switching, delayed stopping, and context-dependent release, respectively (Fig. 5a–d). These tasks provide a behavioral setting to test whether the functions identified from isolated FFL dynamics correspond to task-specific advantages in larger recurrent networks.

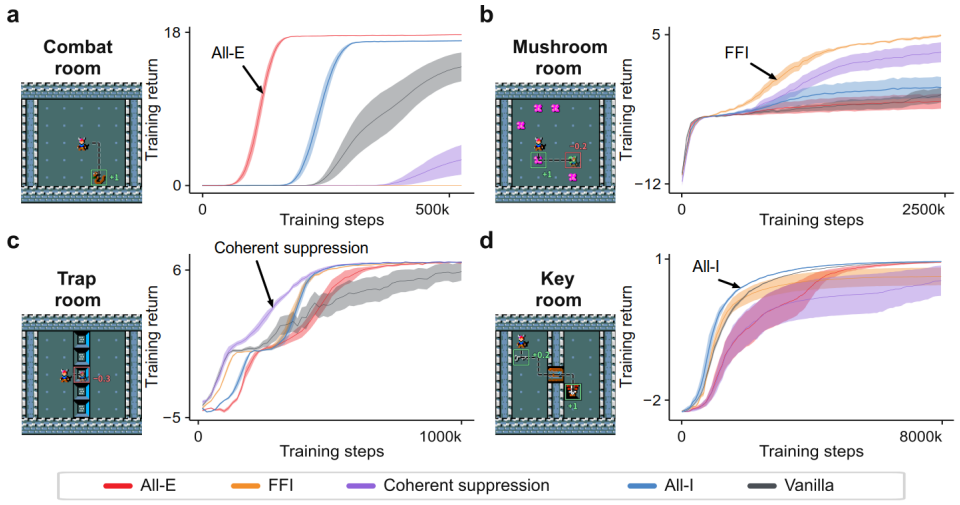

Figure 5: **Signed-motif priors across four MiniHack tasks. a–d**, Combat, mushroom, trap, and key rooms, respectively. Each panel shows a task scene and training return versus environment steps for four motif priors and the Vanilla baseline. Scene annotations indicate rewards and penalties. Lines show means across 10 random seeds per condition. Shaded bands indicate ± 1 standard error of the mean (SEM).

We introduce each signed FFL configuration as a structural prior through differentiable motif regularization and compare motif-enriched networks with a Vanilla baseline without motif constraints. The motif priors preserve the same E-I identity constraints while selectively increasing the abundance of different signed FFL configurations. Implementation details across network backbones are provided in Appendix E.

Different motif priors produce task-dependent behavioral biases consistent with their local computational functions. In the combat room, All-E achieves the strongest performance, reaching a final return of approximately 17.7 compared with 13.9 for Vanilla, and reaches return 10 with substantially fewer interactions. In contrast, FFI performs best in the mushroom room, where switching between positive and negative outcomes is required. Coherent suppression provides the fastest learning in the trap room, reducing the interactions required to reach return 5 compared with other conditions, while All-I improves learning speed in the key room, where contextual release is required. Together, these results show that embedding different signed FFLs biases recurrent networks toward the computational demands predicted from their local dynamics.

3.3 Learned scheduling of motif-dependent functional biases

We next examined whether motif-dependent computational biases can be composed within a single environment containing multiple task demands. The agent must fight monsters, distinguish edible

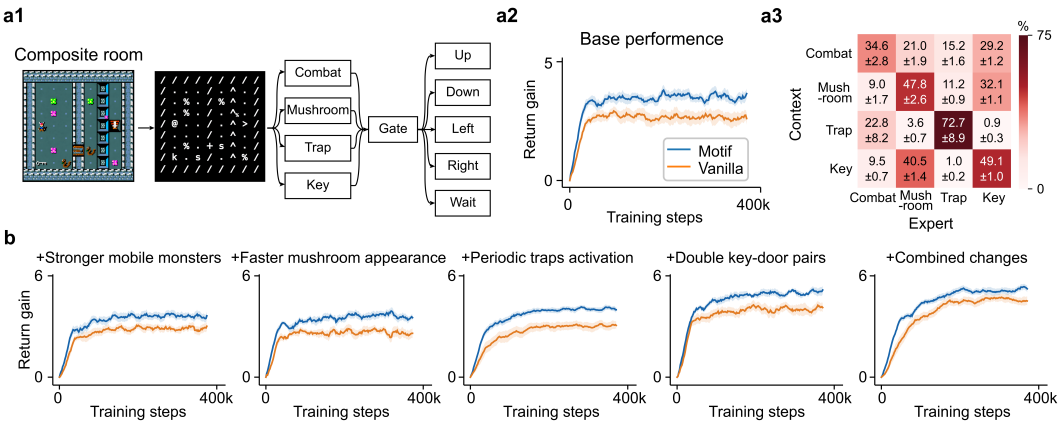

Figure 6: **Context-dependent expert recruitment in a composite task. a1**, Composite room and agent architecture with four experts coordinated by a learned gate. **a2**, Training return gain for Motif and Vanilla in the base environment. **a3**, Expert routing weights (%) for Motif. Rows indicate contexts and columns indicate experts. Weights are normalized within each row for each seed before averaging. Entries show mean ± standard deviation (SD) across ten random seeds. **b**, Fresh gate-training runs with stronger mobile monsters, faster mushroom appearance, periodic trap activation, two key–door pairs, or their combination. In **a2** and **b**, lines and shaded bands show the mean and ± 1 standard error of the mean (SEM) across ten random seeds per condition.

from poisonous mushrooms, stop for traps, and use keys to reach the goal. A learned gate combines the action distributions of four frozen task-specific experts to coordinate these behaviors (Fig. 6a1).

Motif-enriched experts achieve higher returns than Vanilla throughout training and maintain their advantage over the approximately 400,000-interaction budget (Fig. 6a2). The final mean return gain is 3.66 compared with 2.61 for Vanilla, corresponding to a 40.1% increase. The gate preferentially assigns higher routing weights to experts aligned with the current context: 34.6% for combat, 47.8% for mushrooms, 72.7% for traps, and 49.1% for keys (Fig. 6a3). However, routing remains distributed across experts; for example, key contexts recruit both the key expert (49.1%) and the mushroom expert (40.5%).

To test whether this compositional advantage generalizes beyond the original environment, we trained a fresh gate for each environmental variant while keeping the four experts frozen (Fig. 6b). These variants independently increase the demands on combat, event timing, stopping, and sequential navigation. Across these four variants, Motif retains higher final mean returns than Vanilla, with gains of 20.4%, 35.7%, 30.2%, and 25.2%, respectively. When all modifications are combined, Motif maintains a 16.0% advantage.

Together, these results indicate that motif-biased experts provide reusable computational components that can be coordinated according to task context. Although expert parameters remain fixed, the learned gate selectively combines their outputs and preserves the advantage of motif-based representations across diverse environments.

4 Discussion and Conclusion

Conclusion. In this study, we investigate how local connectivity shapes computation through signed E-I feedforward loops in neural networks. We derive how direct and indirect pathways determine target responses under matched input conditions. Our analyses and network experiments show that these circuit configurations produce different computational biases, and that these biases can be encoded as structural priors and coordinated across task contexts. Together, these results support a circuit-level account in which signed connectivity links local dynamics to network learning and compositional function.

Previous studies have established connectivity motifs as important organizational principles for understanding network structure and dynamics (Milo et al., 2002; Alon, 2007). However, motif-based

analyses have largely focused on recurring topological patterns or on how motif composition relates to collective network properties (Shao & Ostojic, 2023; Shao et al., 2025). Our work extends this perspective by examining how different implementations of the same connectivity pattern generate distinct computational functions. In particular, we distinguish connectivity topology from E-I organization, showing that the existence of a pathway alone does not determine its computational consequence. Instead, the functional role of a circuit emerges from the dynamical interaction among signed pathways, consistent with prior work showing that connection signs and E-I organization alter pathway interactions and motif dynamics (Mangan & Alon, 2003; Shao & Ostojic, 2023; Shao et al., 2025; Szilagyi et al., 2026). This framework provides a mechanistic interpretation of motif specialization beyond topological classification.

Modern artificial neural networks increasingly rely on connectivity structures that shape information processing and internal dynamics (Battaglia et al., 2018; Mastrogiuseppe & Ostojic, 2018; Gu & Dao, 2023). Recent studies have further shown that motif organization and biologically motivated structural constraints can influence representation, robustness, flexibility, and learned network function (Zhang et al., 2025; Balwani et al., 2025; Zhang et al., 2026). However, how local circuit mechanisms can be used prospectively to define task-relevant structural priors remains less understood. Here, we show that mechanistically characterized motifs can serve as constructive structural priors for designing recurrent networks. By introducing signed FFL biases, we demonstrate that local circuit principles can shape task-specific computation while preserving interpretable relationships between connectivity and function. This provides a potential strategy for developing more interpretable and biologically grounded neural architectures.

A major challenge in linking connectivity to computation is that structural organization alone does not specify the dynamical processes underlying functional behavior. Previous studies have characterized neural computation through dynamical properties such as stability, amplification, temporal evolution, and low-dimensional population dynamics (Barak, 2017; Mastrogiuseppe & Ostojic, 2018; Dubreuil et al., 2022; Durstewitz et al., 2023), while recent work has examined how local connectivity and E-I organization shape collective dynamics and responses to external inputs (Shao & Ostojic, 2023; Shao et al., 2025). However, how these dynamical properties map onto specific computational operations remains less clear. Our study addresses this gap by defining circuit functions through task-level computational requirements and examining the dynamics required to implement them. This functional specification framework allows transient pathway interactions to be interpreted according to the computations they support rather than their dynamical features alone. Such an approach may provide a general strategy for connecting dynamical systems analysis with computational principles in both biological and artificial networks.

Our findings also provide a perspective for understanding neural computation beyond anatomical connectivity alone. Studies of neuronal circuits increasingly show that the computational consequences of connectivity depend on neuronal identity, synaptic sign, and the interaction between excitatory and inhibitory pathways (Isaacson & Scanziani, 2011; Pfeffer et al., 2013; Pi et al., 2013). Our results suggest that these properties provide an additional layer of organization that shapes how local connectivity transforms information. This principle is particularly relevant for circuits containing feedforward excitation and inhibition, where pathway interactions regulate integration windows and temporal precision (Pouille & Scanziani, 2001; Cao & Van Hooser, 2022), as well as disinhibitory circuits in which cell-type-specific interactions support response modulation and context-dependent gating (Letzkus et al., 2011; Pfeffer et al., 2013; Pi et al., 2013; Wolff et al., 2014; Yang et al., 2016). Extending this framework to larger brain circuits may help bridge anatomical connectivity with cell-type-specific computation and cognitive function.

Limitation. Several limitations define the scope of this study. First, FFLs represent minimal circuit architectures that enable mechanistic analysis but do not capture the full complexity of recurrent and heterogeneous biological networks. Second, our framework focuses on simplified E-I dynamics, whereas biological circuits involve nonlinear neuronal properties, synaptic dynamics, plasticity, and diverse cell-type interactions (Isaacson & Scanziani, 2011; Pfeffer et al., 2013). Third, although motif-based structural priors provide interpretable computational biases, their generalization to larger artificial intelligence systems remains to be explored. Future work combining richer circuit models, large-scale neural recordings, and more complex artificial environments will be necessary to determine how local circuit principles scale toward network-level computation.

Reproducibility Statement

The model and operating assumptions are stated in Section 2. Appendix B derives the transient kernel. Appendix D.1 gives the signed matrix counts. Appendix C specifies the input protocols, numerical settings, and gain scans.

AI Use Statement

Generative AI assisted literature discovery, manuscript restructuring and language editing, interpretation of results, refinement of theoretical statements and their scope, numerical checks of the kernel and counting formulas, and LaTeX and reference preparation.

References

Uri Alon. Network motifs: theory and experimental approaches. *Nature Reviews Genetics*, 8(6): 450–461, 2007.

Aishwarya Balwani, Alex Q. Wang, Farzaneh Najafi, and Hannah Choi. Constructing biologically constrained RNNs via Dale's backpropagation and topologically informed pruning. *Science Advances*, 11(50):eadw4970, 2025. doi: 10.1126/sciadv.adw4970. URL https://doi.org/10.1126/sciadv.adw4970.

Omri Barak. Recurrent neural networks as versatile tools of neuroscience research. *Current Opinion in Neurobiology*, 46:1–6, 2017. doi: 10.1016/j.conb.2017.06.003. URL https://doi.org/10.1016/j.conb.2017.06.003.

Peter W. Battaglia, Jessica B. Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, Caglar Gulcehre, Francis Song, Andrew Ballard, Justin Gilmer, George Dahl, Ashish Vaswani, Kelsey Allen, Charles Nash, Victoria Langston, Chris Dyer, Nicolas Heess, Daan Wierstra, Pushmeet Kohli, Matt Botvinick, Oriol Vinyals, Yujia Li, and Razvan Pascanu. Relational inductive biases, deep learning, and graph networks. arXiv preprint arXiv:1806.01261, 2018. URL https://arxiv.org/abs/1806.01261.

Alexandra S. Cao and Stephen D. Van Hooser. Paired Feed-Forward Excitation With Delayed Inhibition Allows High Frequency Computations Across Brain Regions. *Frontiers in Neural Circuits*, 15:803065, 2022. doi: 10.3389/fncir.2021.803065. URL https://doi.org/10.3389/fncir.2021.803065.

Jonathan Cornford, Damjan Kalajdzievski, Marco Leite, Amélie Lamarquette, Dimitri M. Kullmann, and Blake A. Richards. Learning to live with Dale's principle: ANNs with separate excitatory and inhibitory units. In *International Conference on Learning Representations*, 2021.

Laura N. Driscoll, Krishna Shenoy, and David Sussillo. Flexible multitask computation in recurrent networks utilizes shared dynamical motifs. *Nature Neuroscience*, 27(7):1349–1363, 2024. doi: 10.1038/s41593-024-01668-6. URL https://doi.org/10.1038/s41593-024-01668-6.

Alexis Dubreuil, Adrian Valente, Manuel Beiran, Francesca Mastrogiuseppe, and Srdjan Ostojic. The role of population structure in computations through neural dynamics. *Nature Neuroscience*, 25(6):783–794, 2022. doi: 10.1038/s41593-022-01088-4. URL https://doi.org/10.1038/s41593-022-01088-4.

Daniel Durstewitz, Georgia Koppe, and Max Ingo Thurm. Reconstructing computational system dynamics from neural data with recurrent neural networks. *Nature Reviews Neuroscience*, 24: 693–710, 2023. doi: 10.1038/s41583-023-00740-7. URL https://doi.org/10.1038/s41583-023-00740-7.

Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces, 2023. URL https://arxiv.org/abs/2312.00752.

Jeffry S. Isaacson and Massimo Scanziani. How inhibition shapes cortical activity. *Neuron*, 72(2): 231–243, 2011.

Johannes J. Letzkus, Steffen B. E. Wolff, Elisabeth M. M. Meyer, Philip Tovote, Julien Courtin, Cyril Herry, and Andreas Lüthi. A disinhibitory microcircuit for associative fear learning in the auditory cortex. *Nature*, 480:331–335, 2011. doi: 10.1038/nature10674.

S. Mangan and U. Alon. Structure and function of the feed-forward loop network motif. *Proceedings of the National Academy of Sciences*, 100(21):11980–11985, 2003.

Francesca Mastrogiuseppe and Srdjan Ostojic. Linking Connectivity, Dynamics, and Computations in Low-Rank Recurrent Neural Networks. *Neuron*, 99(3):609–623.e29, 2018. doi: 10.1016/j.neuron.2018.07.003. URL https://doi.org/10.1016/j.neuron.2018.07.003.

R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, and U. Alon. Network motifs: Simple building blocks of complex networks. *Science*, 298(5594):824–827, 2002. doi: 10.1126/science.298.5594.824. URL https://www.science.org/doi/abs/10.1126/science.298.5594.824.

Carsten K. Pfeffer, Mingshan Xue, Miao He, Z. Josh Huang, and Massimo Scanziani. Inhibition of inhibition in visual cortex: the logic of connections between molecularly distinct interneurons. *Nature Neuroscience*, 16:1068–1076, 2013.

Hyun-Jae Pi, Balázs Hangya, Duda Kvitsiani, Joshua I. Sanders, Z. Josh Huang, and Adam Kepecs. Cortical interneurons that specialize in disinhibitory control. *Nature*, 503:521–524, 2013.

Frédéric Pouille and Massimo Scanziani. Enforcement of temporal fidelity in pyramidal cells by somatic feed-forward inhibition. *Science*, 293(5532):1159–1163, 2001.

Mikayel Samvelyan, Robert Kirk, Vitaly Kurin, Jack Parker-Holder, Minqi Jiang, Eric Hambro, Fabio Petroni, Heinrich Küttler, Edward Grefenstette, and Tim Rocktäschel. Minihack the planet: A sandbox for open-ended reinforcement learning research, 2021. URL https://arxiv.org/abs/2109.13202.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal Policy Optimization Algorithms. *arXiv preprint arXiv:1707.06347*, 2017. URL https://arxiv.org/abs/1707.06347.

Yuxiu Shao and Srdjan Ostojic. Relating local connectivity and global dynamics in recurrent excitatory-inhibitory networks. *PLOS Computational Biology*, 19(1):e1010855, 2023. doi: 10.1371/journal.pcbi.1010855. URL https://doi.org/10.1371/journal.pcbi.1010855.

Yuxiu Shao, David Dahmen, Stefano Recanatesi, Eric Shea-Brown, and Srdjan Ostojic. Impact of Local Connectivity Patterns on Excitatory-Inhibitory Network Dynamics. *PRX Life*, 3(2):023008, 2025. doi: 10.1103/PRXLife.3.023008. URL https://doi.org/10.1103/PRXLife.3.023008.

H. Francis Song, Guangyu R. Yang, and Xiao-Jing Wang. Training excitatory-inhibitory recurrent neural networks for cognitive tasks: a simple and flexible framework. *PLoS Computational Biology*, 12(2):e1004792, 2016.

Gabor S. Szilagyi, Attila Gulyas, Zsolt Vassy, Peter Csermely, and Bank G. Fenyves. Signed motif analysis of the Caenorhabditis elegans neuronal network reveals positive feedforward and negative feedback loops. *BMC Biology*, 24:167, 2026. doi: 10.1186/s12915-026-02641-4. URL https://doi.org/10.1186/s12915-026-02641-4.

Steffen B. E. Wolff, Jan Gründemann, Philip Tovote, Sabine Krabbe, Gilad A. Jacobson, Christian Müller, Cyril Herry, Ingrid Ehrlich, Rainer W. Friedrich, Johannes J. Letzkus, and Andreas Lüthi. Amygdala interneuron subtypes control fear learning through disinhibition. *Nature*, 509:453–458, 2014. doi: 10.1038/nature13258.

Guangyu Robert Yang, John D. Murray, and Xiao-Jing Wang. A dendritic disinhibitory circuit mechanism for pathway-specific gating. *Nature Communications*, 7:12815, 2016. doi: 10.1038/ncomms12815.

Guangyu Robert Yang, Madhura R. Joglekar, H. Francis Song, William T. Newsome, and Xiao-Jing Wang. Task representations in neural networks trained to perform many cognitive tasks. *Nature Neuroscience*, 22(2):297–306, 2019. doi: 10.1038/s41593-018-0310-2. URL https://doi.org/10.1038/s41593-018-0310-2.

Haoling Zhang, Chao-Han Huck Yang, Hector Zenil, Pin-Yu Chen, Yue Shen, Narsis A. Kiani, and Jesper N. Tegnér. Leveraging network motifs to improve artificial neural networks. *Nature Communications*, 16(1):11495, 2025. doi: 10.1038/s41467-025-66533-x. URL https://doi.org/10.1038/s41467-025-66533-x.

Jian Zhang, Yue Sun, Wangzi Yao, and Tielin Zhang. Functional building blocks of neural networks: from network motifs to collective dynamics. In *Forty-third International Conference on Machine Learning*, 2026. URL https://openreview.net/forum?id=XiENNDAYNR.

Supplementary Material for: Building Network Functions from Signed Feedforward Loops

Contents

A Notation **14**

B Dynamical analysis of signed FFLs **16**

B.1 Linearized dynamics with time-varying inputs 16

B.2 Derivation of the kernels and input-driven solution 16

C Simulation protocols and parameter scans **18**

C.1 Input protocols and readouts 18

C.2 Connection-strength scans 19

D Counting signed FFLs **21**

D.1 Counting formulas 21

D.2 Structural illustration 22

E Network Experiments **23**

E.1 Networks 23

E.2 Environments 23

F Task demonstrations **24**

F.1 Five-model comparisons on simple tasks 24

F.2 Motif demonstrations on complex tasks 26

F.3 Motif–Vanilla comparisons on complex tasks 29

A NOTATION

Table A1 summarizes the notation used in the paper. Bold symbols denote vectors or matrices.

Table A1: Notation used throughout the paper.

Symbol	Description
Local dynamics and linearization	
A, B, C	Source, relay, and target nodes of a three-node FFL.
i, j, k	Node indices in the counting formulas, distinct source, relay, and target indices, respectively.
$h_i(t), \mathbf{h}(t)$	State of node i and the vector of node states.
$\mathbf{h}^*, \mathbf{u}^*$	Reference equilibrium and its constant baseline input.
$z_i(t), \mathbf{z}(t)$	State deviation from the reference equilibrium, $\mathbf{z} = \mathbf{h} - \mathbf{h}^*$.
$u_i(t), \mathbf{u}(t)$	External drive to node i and the vector of external drives.
$\delta u_i(t)$	Input perturbation relative to the baseline, $u_i - u_i^*$.
$h_{C,\mathrm{base}}, \Delta h_C$	Condition-specific target baseline and the change from it, $\Delta h_C = h_C - h_{C,\mathrm{base}}$.
α	Leak/integration coefficient: the shared leak rate in Eq. (1), with network settings in Table E1.
$\phi(\cdot), \phi'(\cdot)$	Activation function and its derivative.
q_j	Local activation gain at the reference equilibrium, $q_j = \phi'(h_j^*)$.
$\mathbf{W}, W_{ij}$	Connection-weight matrix and its entry for the connection from node j to node i.
$\mathbf{F}, F_{ij}$	Effective linearized coupling matrix and its entries, $F_{ij} = W_{ij} q_j$.
$\mathbf{J}, J_{ij}$	Jacobian and its entries, $\mathbf{J} = -\alpha \mathbf{I} + \mathbf{F}$.
$\mathbf{I}, \delta_{ij}$	Identity matrix and Kronecker delta, respectively.
t, τ	Time and kernel lag, respectively. The variable τ also serves as an integration variable in the driven solution.
$K_{A\to C}, K_{B\to C}$	Input–output kernels from the source and relay to the target.
$\mathbf{p}_A, \mathbf{p}_B$	Input-selection vectors $(1,0,0)^\top$ and $(0,1,0)^\top$.
Functional protocols and readouts	
$P, \theta_{\mathrm{resp}}$	Peak target change in the early window and the response threshold, respectively.
T_{on}	First threshold-crossing latency in the early-response window.
$S_{\mathrm{on}}, Q_{\mathrm{on}}$	Early-response scores before and after the latency penalty (Eq. (5)).
μ_+, μ_-	Mean target change in the first window and the negative of the mean in the second window.
T_{switch}	Waiting time from the negative input onset to the first negative target change. Zero-value conventions are given in Appendix C.1.
$S_{\mathrm{seq}}, Q_{\mathrm{seq}}$	Switching scores before and after the latency penalty (Eq. (6)).
k, λ	Scale of the switching-score nonlinearity and weight of the latency penalty, respectively.
$t_{\mathrm{Go}}, t_{\mathrm{Stop}}$	Onset times of the Go and Stop inputs.
$h_{C,\mathrm{Go}}, h_{C,\mathrm{Stop}}$	Target states under Go alone and Go followed by Stop.
t_V, V	Readout time and score for delayed stopping (Eq. (7)).
t_G, G_{gate}	Readout time and score for context-dependent release (Eq. (8)).
$00, 01, 10, 11$	Condition labels in $\Delta h_{C,00}, \ldots, \Delta h_{C,11}$, with context first and signal second. The digits 0 and 1 denote absence and presence, respectively.
$[z]_+$	Positive-part operator, $[z]_+ = \max(z, 0)$.

Continued on the next page.

Table A1: Notation used throughout the paper (continued).

Symbol	Description
Network structure, motif counts, and training losses	
n	Number of recurrent neurons ($n \geq 3$).
N	Optimization-step label in the schematic "Step N" in Fig. 2.
$\mathbf{s}, \mathbf{S}$	Fixed neuronal excitatory–inhibitory (E-I) identities and their diagonal matrix: $\mathbf{S} = \mathrm{diag}(\mathbf{s})$, with $+1$ for E and -1 for I.
$\mathbf{P}, P_{ij}$	Mask excluding self-connections, $P_{ij} = \mathbf{1}_{i \neq j}$.
$\mathbf{D}_{\mathrm{self}}$	Diagonal matrix of self-connection weights.
$\mathbf{\Theta}, \mathbf{M}$	Trainable magnitude parameters and nonnegative off-diagonal weights, $\mathbf{M} = \mathbf{P} \odot \mathrm{softplus}(\mathbf{\Theta})$.
$\sigma(\cdot)$	Sigmoid function used to make edge presence differentiable.
$\theta_{\mathrm{edge}}, \beta$	Weight-magnitude threshold and sharpness of the sigmoid edge relaxation.
$\widetilde{\mathbf{M}}$	Soft edge-presence matrix, $\mathbf{P} \odot \sigma(\beta(\mathbf{M} - \theta_{\mathrm{edge}}))$.
$\mathbf{Q}_E, \mathbf{Q}_I$	Diagonal selectors for excitatory and inhibitory source neurons.
$\widetilde{\mathbf{W}}^+, \widetilde{\mathbf{W}}^-, \widetilde{\mathbf{W}}^0$	Soft masks for excitatory, inhibitory, and absent connections, respectively.
r	Index of a signed FFL subtype: All-E, FFI, coherent suppression, or All-I.
a, b	Source and relay signs in the counting formulas, $a, b \in \{+, -\}$.
$\widetilde{C}_r, \widetilde{C}_{\mathrm{FFL}}$	Soft count of subtype r and the total soft FFL count, $\widetilde{C}_{\mathrm{FFL}} = \sum_r \widetilde{C}_r$.
C_n^3	Total three-node motif count used to normalize the overall FFL proportion in Eq. (10).
$\widetilde{f}_{\mathrm{FFL}}, \widetilde{\pi}_r$	Overall FFL proportion and the share of subtype r among FFLs, respectively (Eq. (10)).
$f_{\mathrm{FFL}}^{\mathrm{target}}, \pi_r^{\mathrm{target}}$	Target overall FFL proportion and target subtype share.
$\lambda_{\mathrm{FFL}}, \lambda_r$	Penalty weights for overall FFL abundance and subtype composition.
$\mathcal{L}_{\mathrm{PPO}}$	PPO task loss (Eq. (11)).
$\mathcal{L}_{\mathrm{actor}}$	Clipped policy surrogate in the PPO objective.
$\mathcal{L}_{\mathrm{critic}}, \mathcal{L}_{\mathrm{entropy}}$	Squared value-prediction error and policy entropy, respectively.
c_1, c_2	Weights of the critic and entropy terms in the PPO loss.
$\mathcal{L}_{\mathrm{motif}}, \mathcal{L}_{\mathrm{total}}$	Motif regularizer and total loss, $\mathcal{L}_{\mathrm{total}} = \mathcal{L}_{\mathrm{PPO}} + \mathcal{L}_{\mathrm{motif}}$.
$\mathbf{L}$	All-ones row vector used to sum the matrix counts.
$\odot, (\cdot)^\top$	Elementwise multiplication and matrix/vector transpose.
Input-scan parameters (Appendix C.1)	
x, y	Early response: perturbations $(\delta u_A, \delta u_B) = (x, 0.5y)$. Context release: signal amplitude $u_A = x$ and tonic context $u_B = y$.
d, m	Gap between the switching-protocol pulses and parameter controlling the second pulse's negative amplitude.
t_{neg}	Onset time of the second, negative pulse, $t_{\mathrm{neg}} = 1 + d$.
s, d_V	Amplitude and duration of the Stop input, respectively.

B Dynamical analysis of signed FFLs

B.1 Linearized dynamics with time-varying inputs

Consider Eq. (1) with a constant baseline input $\mathbf{u}^*$ and a corresponding equilibrium $\mathbf{h}^*$. For each $i \in \{A, B, C\}$,

$$0 = -\alpha h_i^* + \sum_j W_{ij}\phi(h_j^*) + u_i^*. \tag{13}$$

Write the actual state and input as

$$h_i(t) = h_i^* + z_i(t), \qquad u_i(t) = u_i^* + \delta u_i(t). \tag{14}$$

The reference equilibrium is fixed, so $\dot{h}_i = \dot{z}_i$. Substituting these decompositions into Eq. (1) and subtracting Eq. (13) gives the exact deviation equation

$$\dot{z}_i = -\alpha z_i + \sum_j W_{ij}\left[\phi(h_j^* + z_j) - \phi(h_j^*)\right] + \delta u_i(t). \tag{15}$$

Assume that ϕ is continuously differentiable and has a nonnegative derivative, it implies that $q_j = \phi'(h_j^*) \geq 0$. Since the first-order Taylor approximation is

$$\phi(h_j^* + z_j) - \phi(h_j^*) \approx q_j z_j. \tag{16}$$

Therefore, after keeping these first-order terms defines the linearized model

$$\dot{z}_i = -\alpha z_i + \sum_j F_{ij} z_j + \delta u_i(t), \qquad F_{ij} = W_{ij} q_j. \tag{17}$$

The input is additive in the original model and is retained directly, only the state-dependent activation terms are approximated. The Jacobian entries are

$$J_{ij} = \left.\frac{\partial \dot{h}_i}{\partial h_j}\right|_{\mathbf{h}=\mathbf{h}^*,\, \mathbf{u}=\mathbf{u}^*} = -\alpha\delta_{ij} + F_{ij}, \tag{18}$$

where δ_{ij} is the Kronecker delta and the FFL has no self-connections. Thus $\dot{\mathbf{z}} = \mathbf{J}\mathbf{z} + \delta\mathbf{u}(t)$. For the zero-input reference used in Section 2.2, $\mathbf{u}^* = \mathbf{0}$ and $\delta\mathbf{u}(t) = \mathbf{u}(t)$, yielding Eq. (2). No particular activation function is required for this linearization. As an example, ReLU is affine on each side of its threshold, so the deviation model is exact while the reference and trajectory remain in the same activation region.

B.2 Derivation of the kernels and input-driven solution

We first compute the matrix exponential and then use it to solve the linearized equation. In the node order (A, B, C), with F_{ij} denoting the effective coupling from j to i, the coupling matrix and Jacobian are

$$\mathbf{F} = \begin{pmatrix} 0 & 0 & 0 \\ F_{BA} & 0 & 0 \\ F_{CA} & F_{CB} & 0 \end{pmatrix}, \qquad \mathbf{J} = -\alpha\mathbf{I} + \mathbf{F} = \begin{pmatrix} -\alpha & 0 & 0 \\ F_{BA} & -\alpha & 0 \\ F_{CA} & F_{CB} & -\alpha \end{pmatrix}. \tag{19}$$

The Jacobian is triangular, so its eigenvalues are its diagonal entries, all equal to $-\alpha$. Thus, for $\alpha > 0$, the linearized system is exponentially stable in the absence of input, regardless of the coupling signs. Direct multiplication gives $\mathbf{F}^3 = \mathbf{0}$, so the series for $e^{\mathbf{F}t}$ ends at the quadratic term. Because $\mathbf{I}$ is the identity matrix, $(-\alpha\mathbf{I})\mathbf{F} = \mathbf{F}(-\alpha\mathbf{I}) = -\alpha\mathbf{F}$, thus the exponential of $-\alpha\mathbf{I}$ and $\mathbf{F}$ sum equals the product of their exponentials. We can therefore write

$$e^{\mathbf{J}t} = e^{(-\alpha\mathbf{I}+\mathbf{F})t} = e^{-\alpha t}\, e^{\mathbf{F}t} = e^{-\alpha t}\left(\mathbf{I} + t\mathbf{F} + \frac{t^2}{2}\mathbf{F}^2\right). \tag{20}$$

The quadratic term contains only the indirect path $A \to B \to C$

$$\mathbf{F}^2 = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ F_{CB}F_{BA} & 0 & 0 \end{pmatrix}, \tag{21}$$

Substituting this matrix into Eq. (20) gives

$$e^{\mathbf{J}t} = e^{-\alpha t}\begin{pmatrix} 1 & 0 & 0 \\ tF_{BA} & 1 & 0 \\ tF_{CA} + \dfrac{t^2}{2}F_{CB}F_{BA} & tF_{CB} & 1 \end{pmatrix}. \tag{22}$$

Solution for inputs at the source and relay. With $u_C(t) = 0$, Eq. (2) becomes

$$\dot{\mathbf{z}}(t) = \mathbf{J}\mathbf{z}(t) + \mathbf{p}_A u_A(t) + \mathbf{p}_B u_B(t), \tag{23}$$

where $\mathbf{p}_A = (1, 0, 0)^\top$ and $\mathbf{p}_B = (0, 1, 0)^\top$. Multiplying Eq. (23) by $e^{-\mathbf{J}t}$ gives

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(e^{-\mathbf{J}t}\mathbf{z}(t)\right) = e^{-\mathbf{J}t}\left[\mathbf{p}_A u_A(t) + \mathbf{p}_B u_B(t)\right]. \tag{24}$$

Integrating from 0 to t and multiplying by $e^{\mathbf{J}t}$ gives

$$\mathbf{z}(t) = e^{\mathbf{J}t}\mathbf{z}(0) + \int_0^t e^{\mathbf{J}(t-\tau)}\left[\mathbf{p}_A u_A(\tau) + \mathbf{p}_B u_B(\tau)\right]\,\mathrm{d}\tau. \tag{25}$$

For $\mathbf{z}(0) = \mathbf{0}$, substituting Eq. (22) gives the responses of the three nodes:

$$z_A(t) = \int_0^t e^{-\alpha(t-\tau)} u_A(\tau)\,\mathrm{d}\tau, \tag{26}$$

$$z_B(t) = \int_0^t e^{-\alpha(t-\tau)}\left[(t-\tau)F_{BA}u_A(\tau) + u_B(\tau)\right]\,\mathrm{d}\tau, \tag{27}$$

$$z_C(t) = \int_0^t e^{-\alpha(t-\tau)}\Big[(t-\tau)F_{CA}u_A(\tau) + \frac{(t-\tau)^2}{2}F_{CB}F_{BA}u_A(\tau) + (t-\tau)F_{CB}u_B(\tau)\Big]\,\mathrm{d}\tau. \tag{28}$$

The three terms in the target response correspond to source input through the direct path $A \to C$, source input through the indirect path $A \to B \to C$, and relay input through $B \to C$, respectively. Grouping the two source-input terms gives Eq. (4) with the kernels in Eq. (3).

Pathway interactions in the four motifs. For an input at A, the linearized kernel in Eq. (3) contains direct and indirect pathway contributions. Both pathways are excitatory in All-E and suppressive in coherent suppression. FFI combines direct excitation with indirect inhibition, whereas All-I combines direct inhibition with indirect disinhibition.

The direct and indirect kernel terms scale as $\tau e^{-\alpha\tau}$ and $\tau^2 e^{-\alpha\tau}$, respectively. For nonzero effective couplings, the direct term dominates at short elapsed times, and the indirect term's relative contribution grows with τ.

C SIMULATION PROTOCOLS AND PARAMETER SCANS

Figure 1 illustrates the local response patterns associated with the four functional protocols. The quantitative simulations in Figs. 3, 4, and C1 use the same rate model with ReLU activation, $\alpha = 1$, and $u_C = 0$. Within each protocol, the four motifs are compared using identical inputs, connection magnitudes, initialization rules, and readouts, with only the connection signs changed. Figure 3 shows representative trajectories. Figure 4 varies two input factors while retaining the corresponding connection magnitudes. Stars mark the input settings used in Fig. 3. Figure C1 retains the input protocols of Fig. 3 and varies all three connection magnitudes. The protocol-specific settings and scan grids are detailed below.

C.1 INPUT PROTOCOLS AND READOUTS

Readout definitions follow Section 2.3. Connection magnitudes for the representative trajectories and input scans are given in the order $(|W_{BA}|, |W_{CA}|, |W_{CB}|)$. All scan axes are uniformly sampled, including both endpoints. Input-grid sizes follow horizontal $\times$ vertical axes.

Early response. Each motif starts from its equilibrium under the constant background. Simultaneous input perturbations $(\delta u_A, \delta u_B) = (x, 0.5y)$ are applied on $[0, 1)$, on top of the background inputs $(u_A, u_B) = (1.5, 0.5)$. Inputs return to the background after the pulse. The score Q_{on} in Eq. (5) uses the peak P of Δh_C and the first crossing time T_{on} of the threshold $\theta_{\mathrm{resp}} = 0.15$, both measured over $[0, 1]$, with latency weight $\lambda = 0.2$. If the threshold is not reached within this window, we set $T_{\mathrm{on}} = 0$ by convention. Figure 3a and the connection-strength scan in Fig. C1 use $(x, y) = (1, 1)$, giving actual inputs $(u_A, u_B) = (2.5, 1)$ during the pulse. With connection magnitudes $(1, 2.5, 2)$, the input scan varies $x, y \in [0, 2]$ on an 81×81 grid.

Response switching. Each motif starts from its equilibrium under background inputs $(u_A, u_B) = (1.5, 0.5)$. The first pulse applies perturbations $(3, 0.5)$ on $[0, 1)$, giving actual inputs $(4.5, 1)$. Inputs then return to the background for a gap of duration d. A second pulse begins at $t_{\mathrm{neg}} = 1 + d$ and lasts one time unit, with actual inputs

$$(u_A, u_B) = (-4.5 - 0.5m, -1 - 0.5m).$$

Increasing m makes both inputs more negative. After the second pulse, inputs return to the background. The signed mean responses used in Eq. (6) are

$$\mu_+ = \int_0^1 \Delta h_C(t)\,\mathrm{d}t, \qquad \mu_- = -\frac{1}{1.5}\int_{t_{\mathrm{neg}}}^{t_{\mathrm{neg}}+1.5} \Delta h_C(t)\,\mathrm{d}t.$$

We use $k = 50$ and latency weight $\lambda = 0.2$. The switching latency T_{switch} is measured from the onset of the second pulse to the first negative Δh_C within the second readout window. We set $T_{\mathrm{switch}} = 0$ if the response is already negative at t_{neg} or does not become negative within this window. A positive score requires a positive mean response in the first window and a negative mean response in the second. The switching event is indicated in Fig. 3b. Figure 3b and the connection-strength scan in Fig. C1 use $(d, m) = (1.5, 1)$. The second pulse therefore applies inputs $(-5, -1.5)$ on $[2.5, 3.5)$, and the second readout window extends from $t = 2.5$ to $t = 4$. With connection magnitudes $(2.5, 3, 0.5)$, the input scan varies $d \in [0, 2.5]$ and $m \in [0, 2]$ on a 101×81 grid.

Delayed stopping. All motifs start from zero state. Go applies $u_B = 1$ continuously from $t = 0$. Stop applies $u_A = s$ on $[2, 2 + d_V)$, where s and d_V specify its amplitude and duration. The Stop onset time is $t_{\mathrm{Stop}} = 2$. Outside this interval, $u_A = 0$. The control receives Go without Stop. The score V in Eq. (7) is evaluated at $t_V = 2 + d_V$. Figure 3c and the connection-strength scan in Fig. C1 use $(d_V, s) = (2, 1)$, giving $t_V = 4$. With connection magnitudes $(2, 2, 2)$, the input scan varies $d_V \in [0.5, 3.5]$ and $s \in [0, 2]$ on a 41×41 grid.

Context-dependent release. A brief signal applies $u_A = x$ on $[2, 2.5)$. We compare conditions with sustained background input $u_B = y$ and without it ($u_B = 0$). Each condition starts from its own equilibrium under $(u_A, u_B) = (0, y)$ or $(0, 0)$. Each condition also has a no-signal control that remains at its baseline. The score G_{gate} in Eq. (8) is evaluated at $t_G = 3.5$. Figure 3d and the

connection-strength scan in Fig. C1 use $(x, y) = (1, 1)$. With connection magnitudes $(2, 2, 2)$, the input scan varies $x, y \in [0, 2]$ on a 41×41 grid. Table C1 summarizes functional coverage for the input scans.

Table C1: **Functional coverage across input scans.** Entries are percentages of sampled input settings in Fig. 4 with a positive raw score. The largest coverage in each row is bold. Suppression denotes coherent suppression.

Protocol	All-E	FFI	Suppression	All-I
Early response	**98.49**	58.45	0.00	0.00
Response switching	41.36	**99.41**	0.00	0.00
Delayed stopping	0.00	0.00	**50.27**	0.00
Context-dependent release	0.00	0.00	0.00	**83.22**

C.2 Connection-strength scans

The connection-strength scans hold the inputs fixed at the values specified in Section C.1 and independently vary each magnitude over

$$|W_{BA}|, |W_{CA}|, |W_{CB}| \in \{0, 0.125, 0.250, \ldots, 3.000\}.$$

Each motif and protocol therefore uses $25^3 = 15{,}625$ weight combinations, including zero-edge boundary cases. Figure C1 shows the resulting scores.

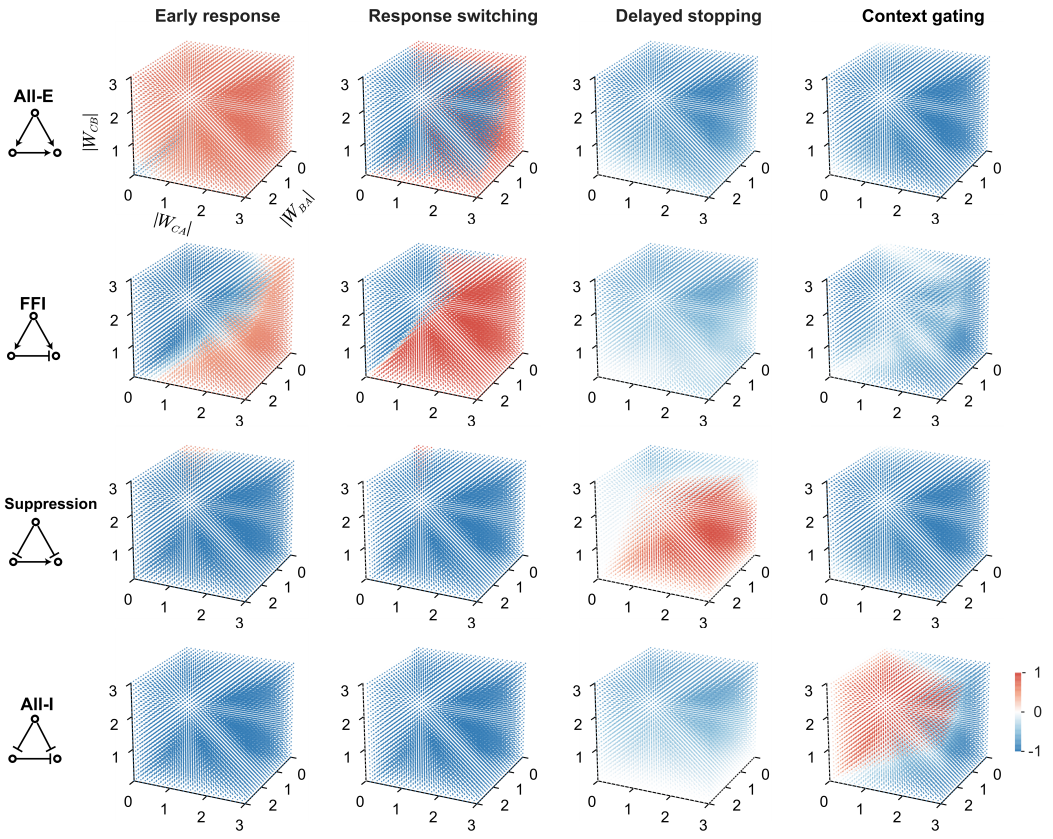

Figure C1: **Functional scores across connection strengths.** Rows compare motifs and columns show the four protocols. Each panel contains the same 25^3 weight combinations in $[0, 3]^3$. Red/blue indicates positive/negative scores, with zero in white. Inputs and readouts are specified in Section C.1.

Functional coverage is the percentage of sampled settings with a strictly positive raw score. Table C2 reports coverage for the connection-strength scans.

Table C2: **Functional coverage across connection-strength scans.** Entries are percentages of sampled weight combinations in Fig. C1 with a positive raw score. The largest coverage in each row is bold. Suppression denotes coherent suppression.

Protocol	All-E	FFI	Suppression	All-I
Early response	**97.26**	37.14	1.13	0.00
Response switching	40.39	**70.36**	1.02	0.00
Delayed stopping	0.00	0.00	**56.83**	0.00
Context-dependent release	0.00	0.00	0.00	**39.32**

D COUNTING SIGNED FFLS

D.1 COUNTING FORMULAS

We use W_{ij} for the connection from node j to node i, with the soft masks defined in Eq. (9). Let i, j, k denote distinct source, relay, and target nodes. For subtype r, let $a, b \in \{+, -\}$ denote the source and relay signs. The three forward edges have signs (a, a, b), giving the induced soft count

$$\widetilde{C}_r = \sum_{\substack{i,j,k \\ \text{distinct}}} (\widetilde{W}^a_{ji}\widetilde{W}^0_{ij})(\widetilde{W}^a_{ki}\widetilde{W}^0_{ik})(\widetilde{W}^b_{kj}\widetilde{W}^0_{jk}). \tag{29}$$

Each parenthesized factor pairs a required forward connection with the absence of its reverse. The choices $(a, b) = (+,+), (+,-), (-,+), (-,-)$ give All-E, FFI, coherent suppression, and All-I, respectively.

Summing over the shared source index gives the matrix expression

$$\widetilde{C}_r = \mathbf{L}\Big[\Big(\widetilde{\mathbf{W}}^a \odot (\widetilde{\mathbf{W}}^0)^\top\Big)\Big(\widetilde{\mathbf{W}}^0 \odot (\widetilde{\mathbf{W}}^a)^\top\Big) \odot \Big(\widetilde{\mathbf{W}}^0 \odot (\widetilde{\mathbf{W}}^b)^\top\Big)\Big]\mathbf{L}^\top. \tag{30}$$

Here $\mathbf{L}$ is the all-ones row vector and $\odot$ denotes elementwise multiplication. The masks have zero diagonals, which eliminates repeated-node terms. In a hard induced FFL, the source, relay, and target roles determine a unique ordering, so no division for repeated counting is needed. Figure D1 illustrates the FFI case. The four subtype counts sum to the total FFL count. The corresponding proportions are defined in Eq. (10).

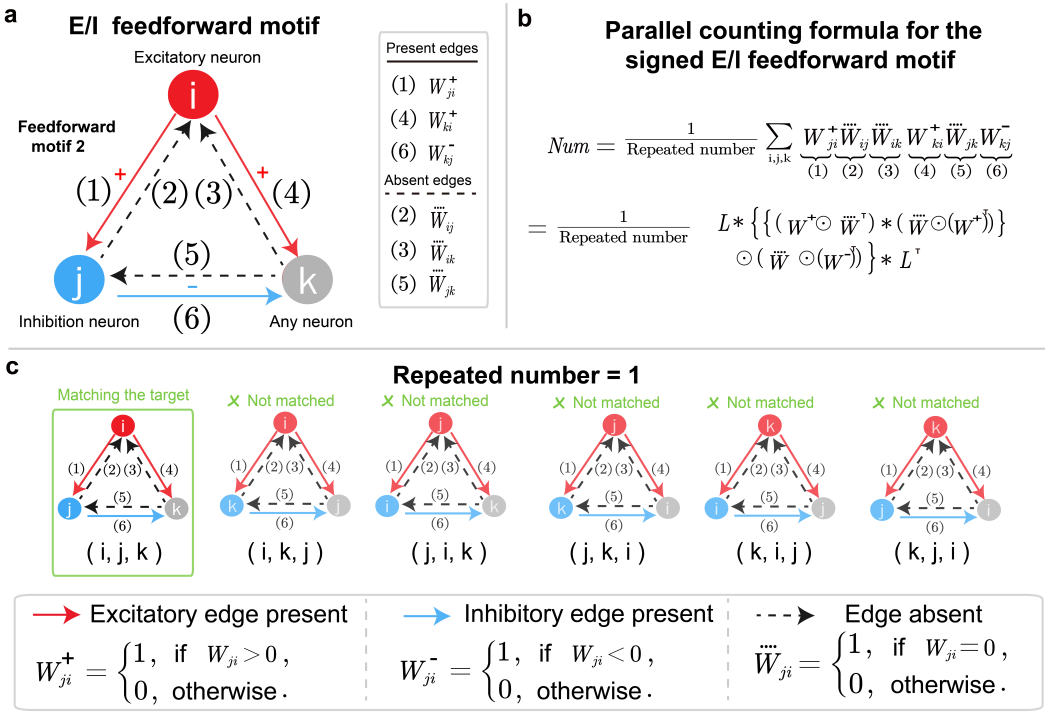

Figure D1: **Matrix counting illustrated for FFI. a**, Three forward edges have signs $+, +, -$, and the reverse edges are absent. **b**, The six edge indicators are combined through matrix operations. **c**, Only one node ordering matches the required roles, giving a repeated-count factor of one. The hard indicators shown here are replaced by the soft masks in Eq. (9) for differentiable training.

D.2 STRUCTURAL ILLUSTRATION

To demonstrate that motif-frequency constraints can be realized through weight optimization, we trained four copies of the same initial network under constraints favoring All-E, FFI, coherent suppression, or All-I motifs, respectively. Figure D2 shows the evolution of connectivity within the same subset of neurons across optimization steps. Starting from identical recurrent weights, the networks progressively develop distinct connection patterns as their respective target motifs become enriched. These results demonstrate that the differentiable motif regularizer can guide changes in network connectivity toward specified motif compositions, providing a practical means of incorporating motif-based structural priors into network training.

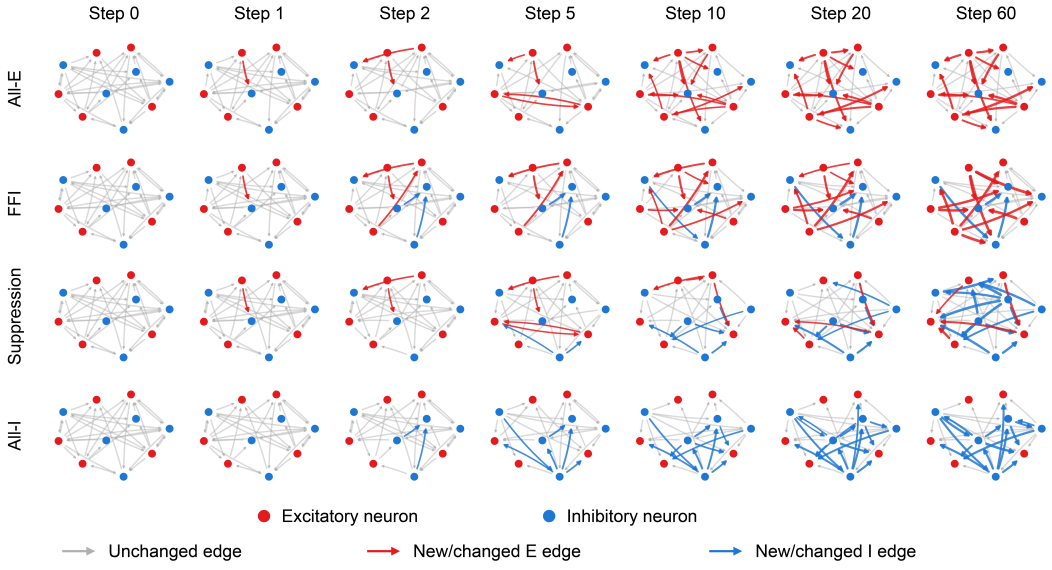

Figure D2: **Connectivity changes during motif-loss optimization.** Rows target All-E, FFI, coherent suppression, and All-I. Columns show steps 0, 1, 2, 5, 10, 20, and 60. Red and blue nodes are excitatory and inhibitory. Gray edges are unchanged from initialization. Red and blue edges mark new or changed excitatory and inhibitory connections, respectively.

E NETWORK EXPERIMENTS

E.1 NETWORKS

All comparisons use five conditions: four signed-FFL priors plus a motif-free baseline (Vanilla), with matched architectures and 10 seeds per condition. Detailed hyperparameters are listed in the table below.

Table E1: Hyperparameters of the four single-task policies and the gating network.

	Combat	**Mushroom**	**Trap**	**Key**	**Gating**
Network					
Recurrent units	64	64	64	64	64
Integration coefficient α	0.2	0.2	0.2	0.2	0.2
Input features	57	57	105	105	461
Actions	2	5	5	5	5
Optimizer	Adam	Adam	Adam	Adam	Adam
PPO					
Parallel environments	32	64	256	256	32
Rollout length	32	64	64	64	64
Epochs	4	4	4	4	4
Minibatches	4	4	4	4	4
Learning rate	3×10^{-4}	3×10^{-4}	3×10^{-4}	3×10^{-4}	3×10^{-4}
Entropy coefficient	0.01	0.01	0.01	0.01	0.01
Clip	0.2	0.2	0.2	0.2	0.2
Seeds	10	10	10	10	10

E.2 ENVIRONMENTS

Each of the four rooms isolates a single demand. In the combat room, the agent must approach a monster and initiate an attack when an encounter occurs, testing **timely initiation**. In the mushroom room, the agent must eat edible mushrooms but suppress the same action whenever the mushroom encountered is poisonous. Since toxicity can only be told apart up close, the response must follow the object rather than the scene. The trap room interrupts forward movement with a trigger and requires waiting a fixed duration before movement resumes, testing a **self-clearing temporary veto**. The key room places a hidden key before a locked door, testing a **permission that must be maintained from pickup until passage**. All four rooms provide shaped rewards for the target behavior, a small cost per action, and a fixed cap on actions per episode. The agent observes a local object crop, its previous action, and a task cue.

The composite environment places all four demands within a single episode, so role switching occurs inside the episode rather than across rooms. Geometry and movement rules are provided by MiniHack. A wrapper adds a monster guarding the key and the area behind the door, mushrooms that appear and vanish in batches, a warning column that requires waiting on first entry, and a key–door–goal chain on a 7×7 interior region. Hit points are limited, so the agent must clear or avoid adjacent monsters while completing this chain. The variants respectively change the monster's durability and movement speed, the temporal structure of the mushrooms, the periodicity of the trap, or add a second key–door pair. The merged variant applies all of these changes simultaneously and grants higher initial hit points. Reward values, layout constants, and the gating objective are as specified in the released configuration files.

F TASK DEMONSTRATIONS

The demonstrations below cover four simple tasks and six complex tasks. Selected trajectories illustrate behavior rather than aggregate performance.

F.1 FIVE-MODEL COMPARISONS ON SIMPLE TASKS

The four simple tasks compare All-E, FFI, coherent suppression, All-I, and Vanilla.

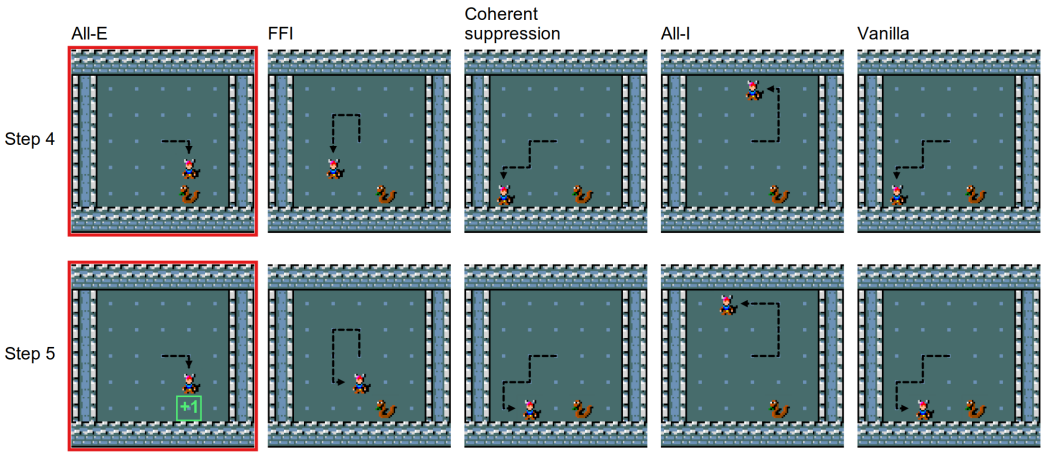

Figure F1: Five-model comparison in the combat room.

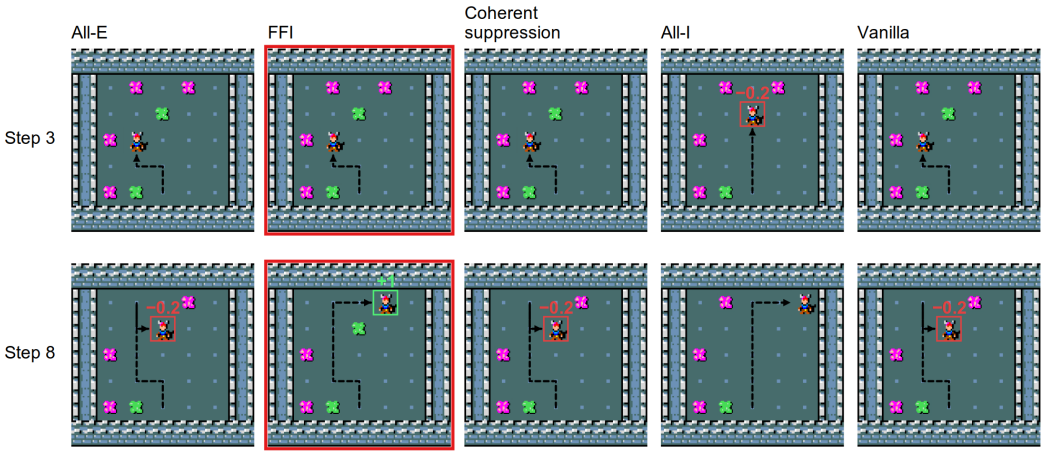

Figure F2: Five-model comparison in the mushroom room.

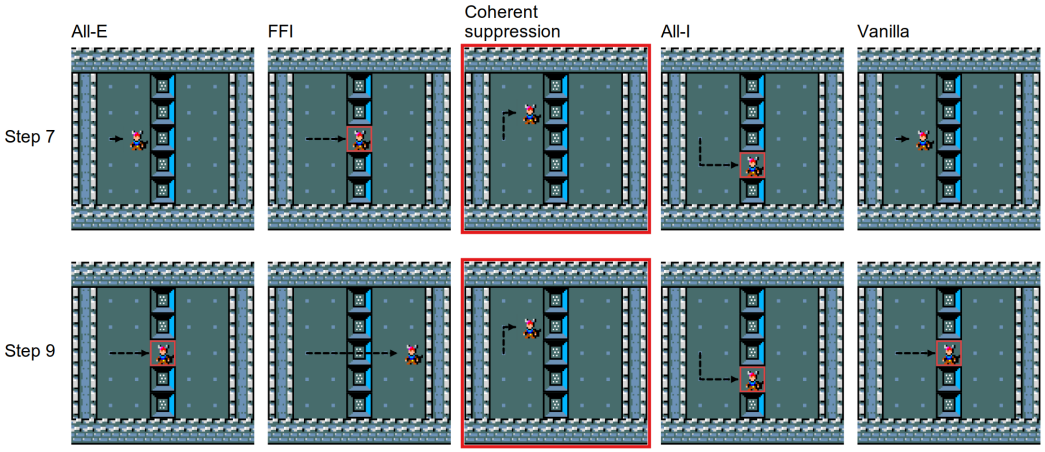

Figure F3: Five-model comparison in the trap room.

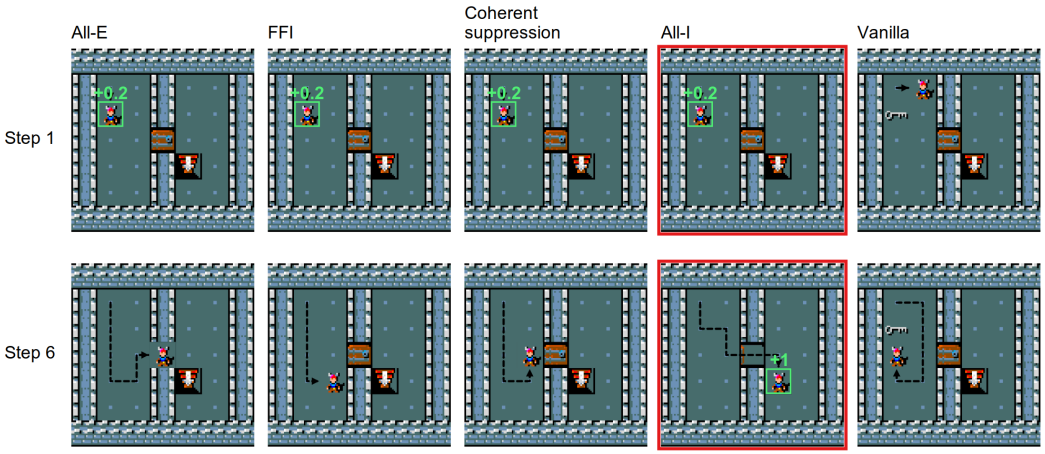

Figure F4: Five-model comparison in the key room.

F.2 MOTIF DEMONSTRATIONS ON COMPLEX TASKS

Six tasks are illustrated: the base composite task and its five variants. Each figure shows a single Motif trajectory, with dashed trails marking movement. These examples and the comparisons below need not use the same episode.

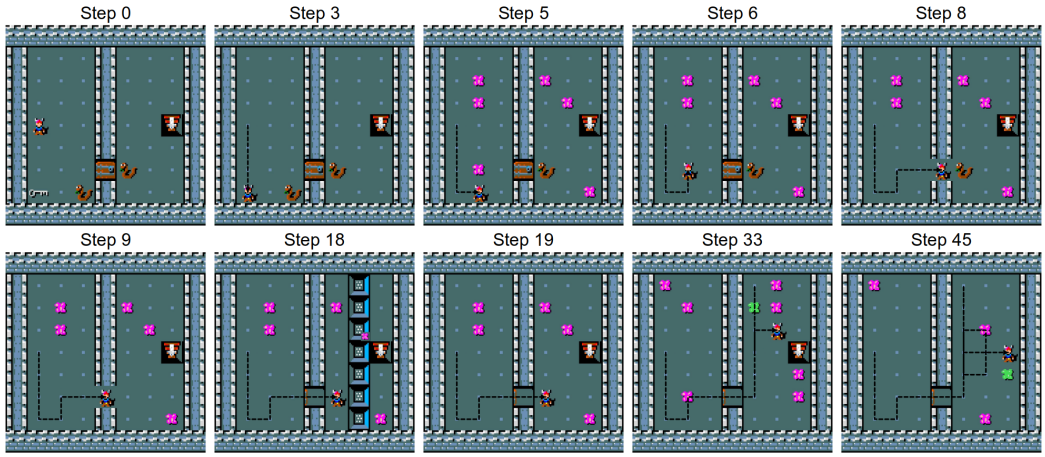

Figure F5: Motif trajectory in the composite room.

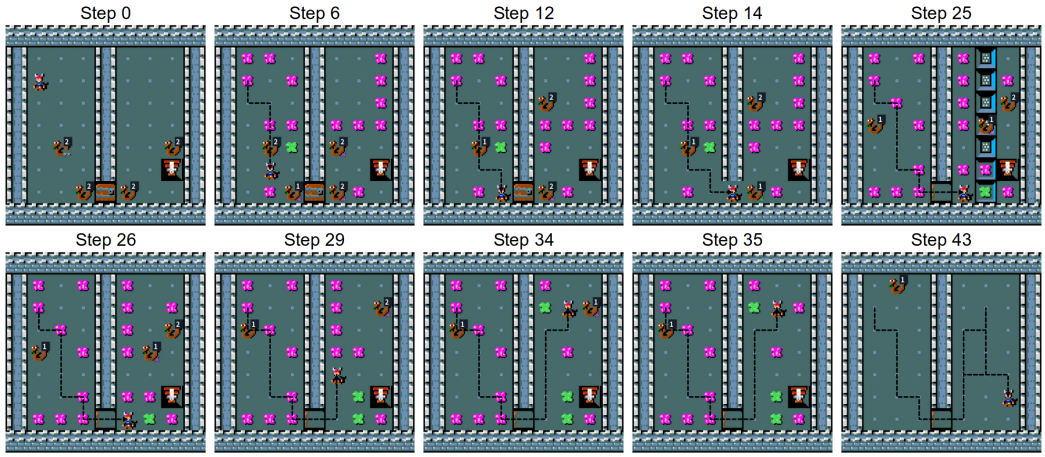

Figure F6: Motif trajectory in the composite task with mobile monsters.

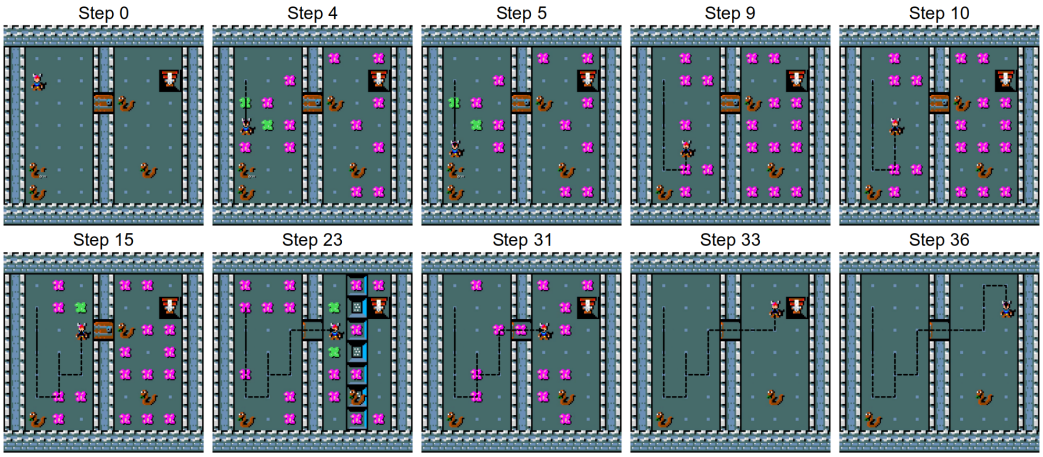

Figure F7: Motif trajectory in the composite task with faster mushroom appearance.

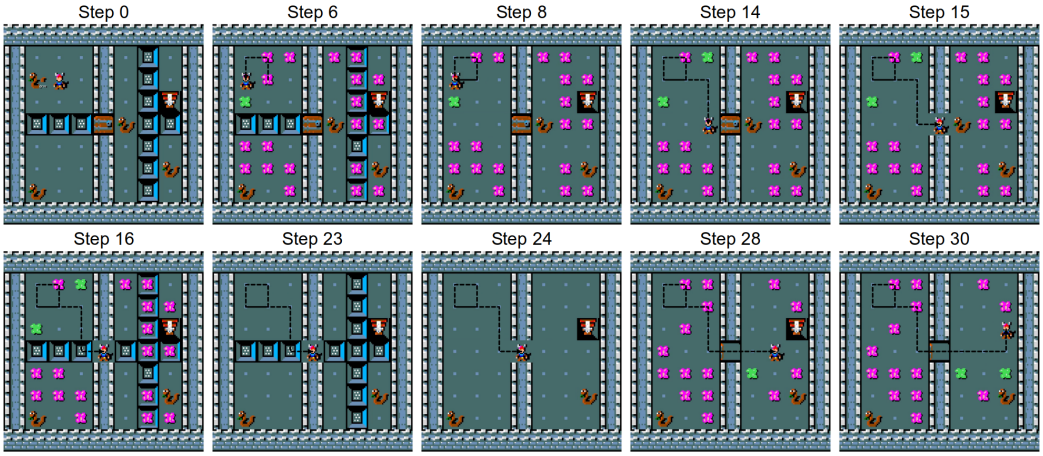

Figure F8: Motif trajectory in the composite task with periodic traps.

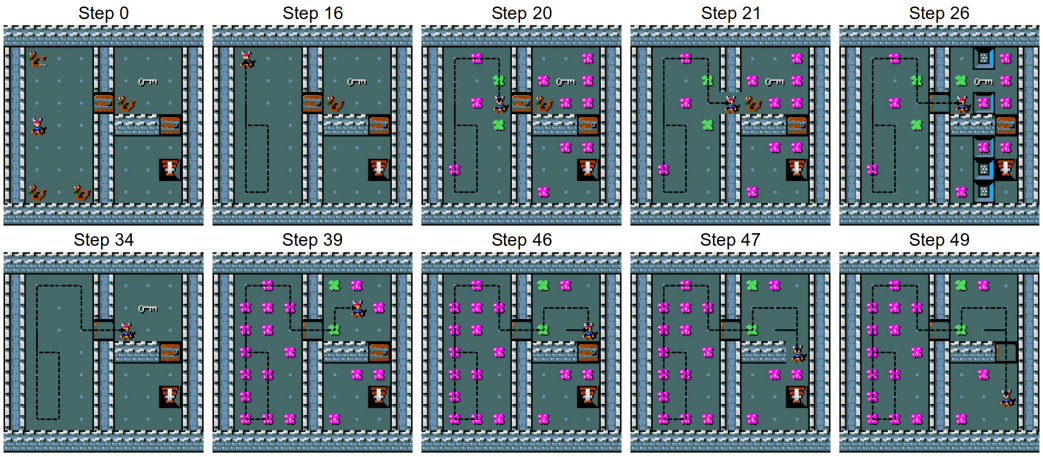

Figure F9: Motif trajectory in the composite task with two key–door pairs.

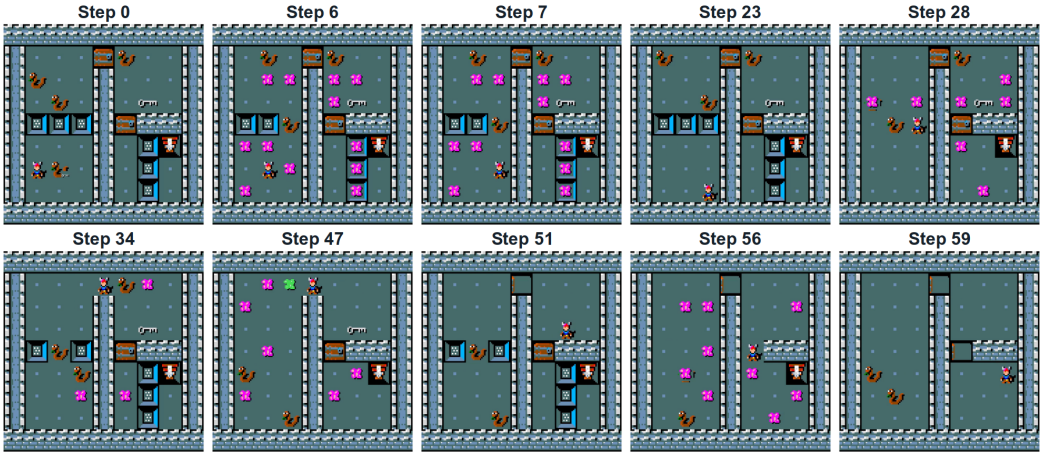

Figure F10: Motif trajectory in the composite task with combined changes.

F.3 MOTIF–VANILLA COMPARISONS ON COMPLEX TASKS

The same six task types are compared below. Motif is shown in the top row and Vanilla in the bottom row. Dashed arrows show recent movement; green and red outlines identify selected Motif and Vanilla positions, respectively, not reward signs. Each row is chronological. Columns are not necessarily time-matched.

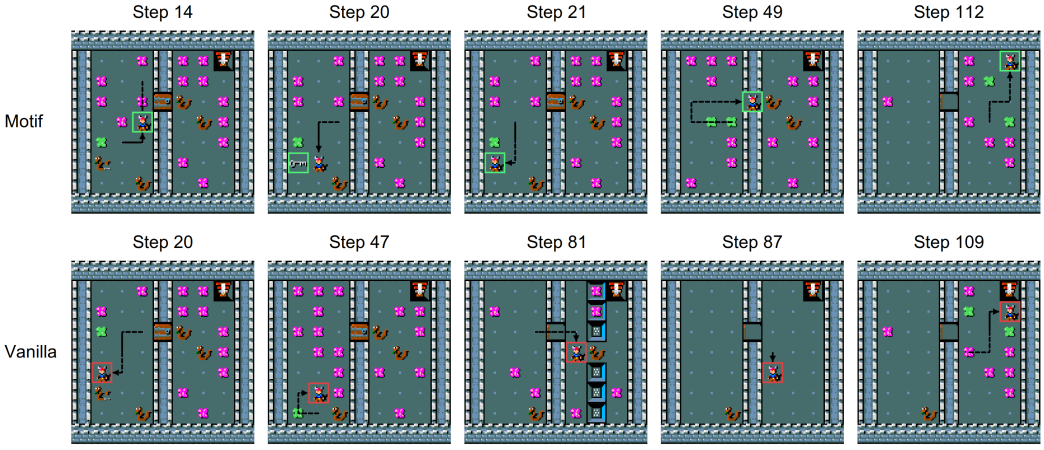

Figure F11: Motif–Vanilla comparison in the composite room.

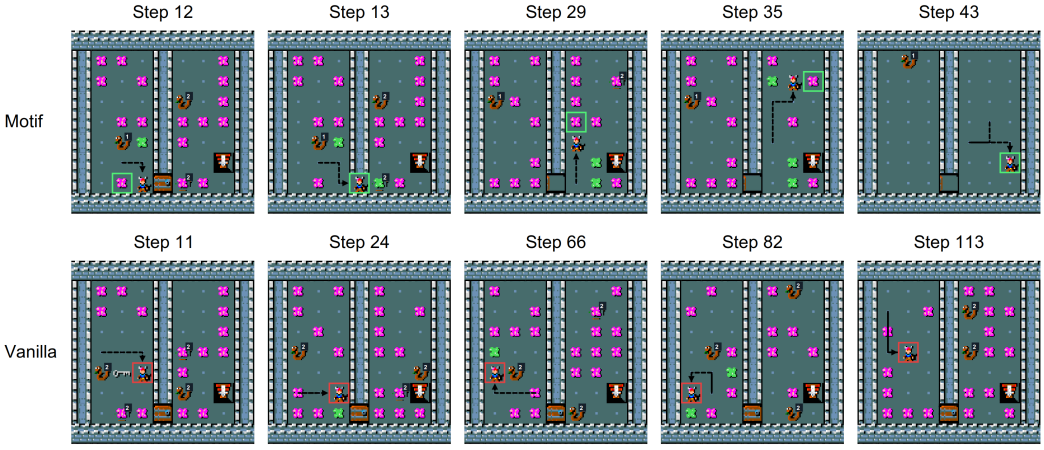

Figure F12: Motif–Vanilla comparison in the composite task with mobile monsters.

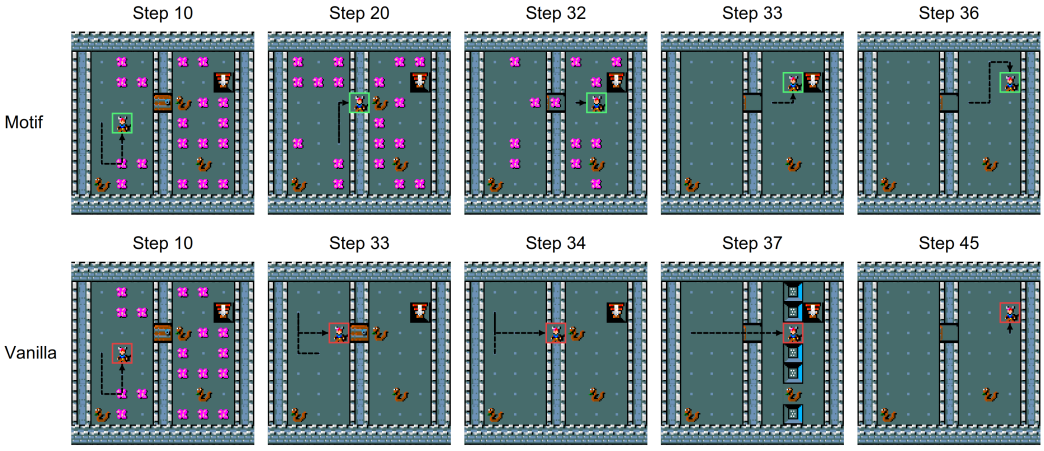

Figure F13: Motif–Vanilla comparison in the composite task with faster mushroom appearance.

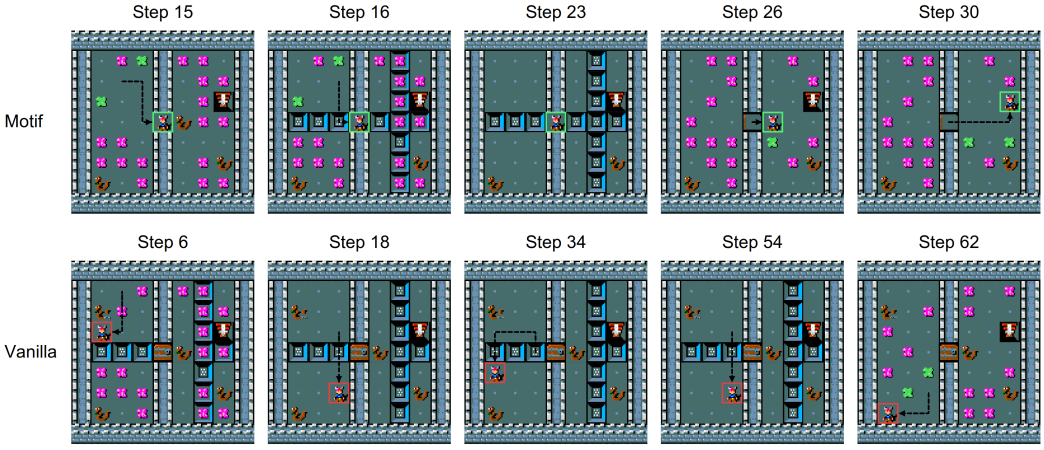

Figure F14: Motif–Vanilla comparison in the composite task with periodic traps.

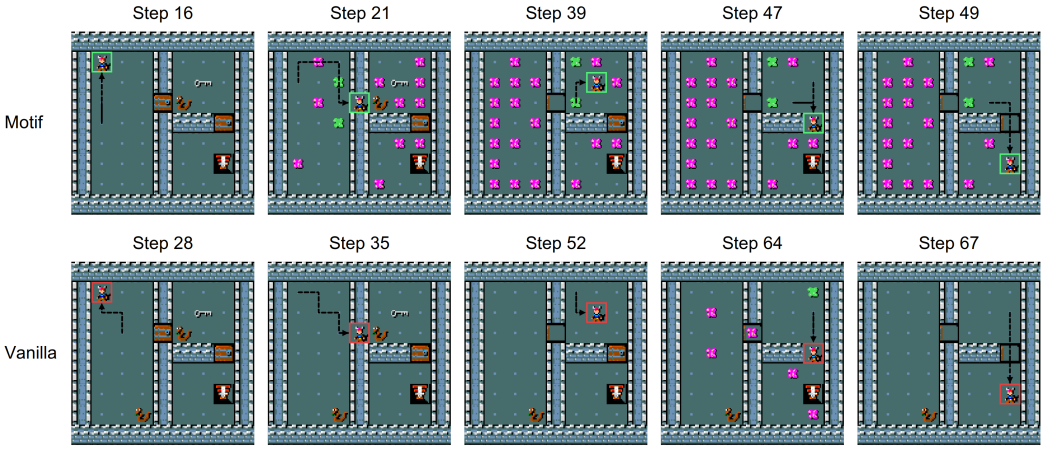

Figure F15: Motif–Vanilla comparison in the composite task with two key–door pairs.

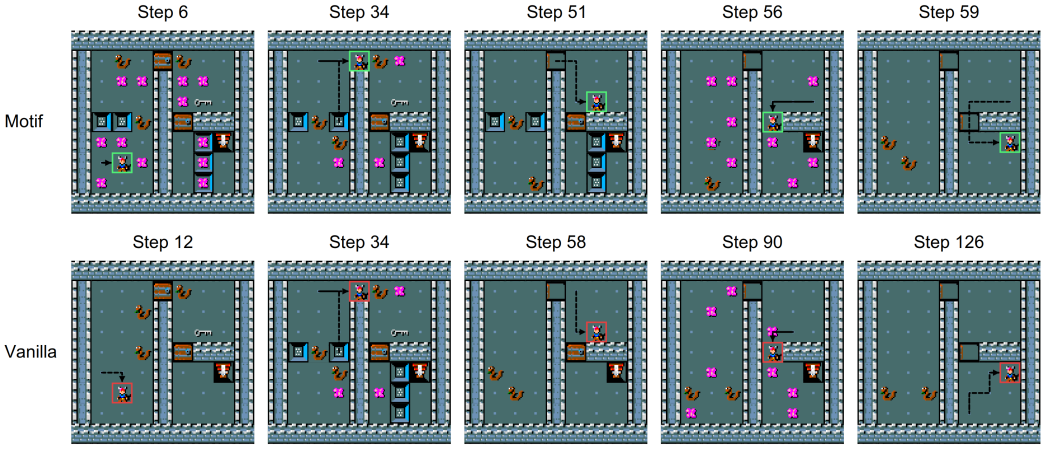

Figure F16: Motif–Vanilla comparison in the composite task with combined changes.