LOOPED REASONING FINISHES EARLIER THAN YOU THINK

Anonymous authors
Paper under double-blind review

ABSTRACT

Looped reasoning with dynamic depth can be viewed as relaxation toward solution attractors in a learned dynamical landscape. In practice, however, latent trajectories often oscillate, cycle, or drift rather than cleanly converge to stable fixed points. How these non-stationary dynamics contribute to reasoning and constrain convergence remains unclear. We show that slowly decaying oscillatory modes can impose their convergence rate on coupled latent dynamics, creating a bottleneck. Yet damping these modes indiscriminately can disrupt the reasoning computation they support. This tension motivates an architectural approach: introducing rotational degrees of freedom along which symmetry-generated rotational motion leaves the reasoning-relevant quotient dynamics unchanged by construction. We propose **QRR** (**Q**uotient **R**ecurrent **R**easoner), a gauge-equivariant recurrent architecture with an exact per-token phase-rotation symmetry. Motion along these symmetry directions leaves the quotient representation unchanged and is dynamically decoupled from reasoning-relevant dynamics. Empirically, QRR achieves stable and faster convergence to solution roots in quotient space, without forming a dominant slow oscillatory bottleneck, while matching or surpassing recent recurrent reasoning models on Sudoku, Maze, and Mini-ARC.

1 INTRODUCTION

Looped reasoners trade architectural depth for recurrent computation (Dehghani et al., 2019; Saunshi et al., 2025; Fan et al., 2025). A shared neural block is applied repeatedly to refine an internal state, allowing additional computation without increasing parameter count (Bansal et al., 2022; Geiping et al., 2025; Wang et al., 2025; Jolicoeur-Martineau, 2025). Recent works make the dynamical interpretation explicit, viewing reasoning as the evolution of a latent state converging toward solution-aligned equilibria (Blayney et al., 2026; Kim, 2026). Fixed-Point Reasoning Model (FPRM) uses convergence of the recurrent state to a fixed point as an intrinsic stopping criterion, Attractor Models define the prediction itself as the equilibrium of a learned refinement process, and Equilibrium Reasoners (EqR) explicitly learns task-conditioned attractors whose stable fixed points correspond to valid solutions, using multiple initializations to explore their attraction basins (Movahedi et al., 2026; Fein-Ashley & Rashidinejad, 2026; Huang et al., 2026). In this view, reasoning is not just a sequence of refinements but a recurrent trajectory through the learned dynamical landscape.

However, viewing recurrent computation as convergence to a stable attractor does not capture the full range of dynamics it can exhibit. Trajectories can instead exhibit orbits, drifts, cyclic limiting behavior, and other structured forms of non-stationarity, rather than simply relaxing toward a single fixed point (Geiping et al., 2025; Pappone et al., 2025; Blayney et al., 2026; Kim, 2026; Yang et al., 2026). These studies reveal a wide range of latent geometries in recurrent models, yet how these different forms of oscillatory motion affect the underlying latent computation remains unexplored.

The key distinction is therefore not the geometry of the motion itself, but whether and how it influences the recurrent computation. Around a solution root $z^\star$, where the recurrence admits a fixed linearization with the Jacobian $J = Df(z^\star)$, we establish a local coupling mechanism by separating an oscillatory spectral coordinate from the complementary latent dynamics and quantifying how the former drives the latter. When the oscillatory mode decays more slowly, this coupling transfers its rate to the complementary dynamics and delays settlement. We then test this mechanism causally in the original nonlinear recurrences of FPRM and HRM (Wang et al., 2025) in Section 3.2. Across distinct oscillatory geometries, an initial perturbation of the identified mode produces latent responses and

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.

settling-time shifts that closely match the predictions of the local coupling model. This agreement directly links the convergence delay near solution roots to the coupled slow oscillatory mode.

Oscillatory appearance alone, however, does not reveal which motion can be safely removed. In FPRM, suppressing the alternating mode before semantic stabilization degrades task performance, showing that slow oscillatory motion can still support reasoning (Appendix A.4.4). Thus, rather than identifying and damping particular oscillations, we instead redesign what counts as convergence by introducing a controlled representational redundancy: motion along rotational symmetry directions is irrelevant to the underlying quotient state by construction. We propose **QRR**, the **Q**uotient **R**ecurrent **R**easoner, which imposes an exact per-token rotational symmetry on the recurrent state, inducing a quotient representation invariant to this motion. Gauge-equivariant updates preserve this equivalence, allowing symmetry-generated motion to persist in the raw trajectory without driving or limiting the convergence of quotient dynamics.

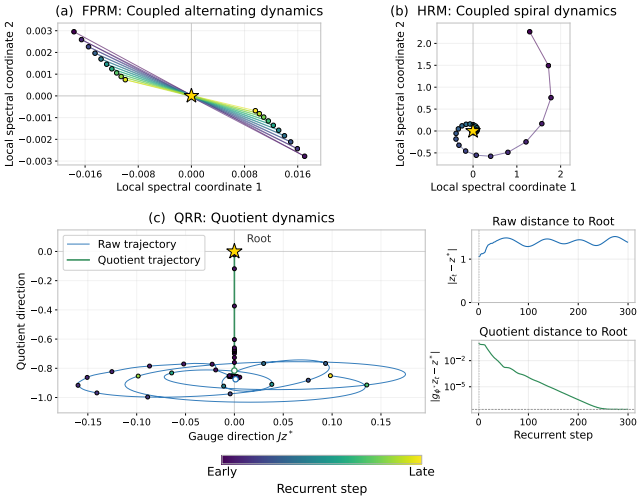


Figure 1: **Oscillatory geometry and quotient convergence.** FPRM and HRM exhibit distinct alternating and spiral dynamics near solution roots. In QRR, raw motion persists while quotient trajectory contracts.

As in Figure 1, FPRM and HRM exhibit distinct oscillatory geometries whereas QRR retains raw phase motion while its quotient trajectory converges. Oscillatory modes outside the gauge subspace remain free to participate in reasoning. Empirically, QRR also retains oscillatory modes in quotient space but does not develop the isolated dominant slow mode as in FPRM and HRM (Section 4.4).

2 Background

2.1 Looped reasoning models

Looped reasoning models increase effective computational depth by repeatedly applying shared neural blocks to a latent state, allowing additional test-time computation without increasing parameter count. Recent approaches can be broadly grouped by how they organize this recurrent computation. General recurrent-depth models iterate a shared Transformer block and scale reasoning by increasing the number of latent recurrences at inference (Geiping et al., 2025). Puzzle-oriented recursive reasoners instead structure the recurrence explicitly. HRM couples high- and low-level recurrent modules operating at different frequencies (Wang et al., 2025), whereas Tiny Reasoning Model (TRM) simplifies this hierarchy to a single small network that repeatedly refines a latent reasoning state and the current answer (Jolicoeur-Martineau, 2025). A complementary line organizes recurrent computation around equilibria or attractors: FPRM uses convergence to a fixed point as an adaptive halting criterion (Movahedi et al., 2026), EqR learns task-conditioned attractors whose stable equilibria correspond to solutions (Huang et al., 2026), and Attractor Models explicitly solve a learned fixed-point refinement with implicit differentiation (Fein-Ashley & Rashidinejad, 2026).

2.2 Gauge symmetry and quotient dynamics

Gauge symmetry formalizes a local representational redundancy where changing the local reference frame changes the coordinates used to represent a state, but not the underlying object itself (Cohen et al., 2019a; Bronstein et al., 2021; Gerken et al., 2023). If a map T is equivariant under the action of a group G, gauge-related states $z \sim g \cdot z$ evolve consistently and can be identified as a single equivalence class $[z]$ in the quotient space $\mathcal{Z}/G$ (Kondor & Trivedi, 2018; Cohen et al., 2019b). The quotient collapses symmetry-related representations into a single state, so motion along the symmetry orbit is ignored while all remaining state changes are preserved. This distinction also changes the notion of convergence. An equivariant system may continue moving along a symmetry orbit without changing its quotient state. Such motion is captured by a relative equilibrium, $T(z^\star) = g^\star \cdot z^\star$,

which is non-stationary in the original coordinates but becomes a fixed point after quotienting, $[T(z^\star)] = [z^\star]$ (Krupa, 1990; Marsden & Ratiu, 1999). In neural architectures, gauge equivariance provides a principled way to implement this invariance to local reference frames (Weiler et al., 2021; Wiersma et al., 2020). Gauge-equivariant convolutions extend conventional equivariance from global transformations to local gauge choices on manifolds (Cohen et al., 2019a), and Gauge Equivariant Mesh CNNs combine anisotropic convolution with parallel transport to communicate features in different local frames (Haan et al., 2021). The same principle has also been extended to attention, where the Gauge Equivariant Transformer constructs gauge-equivariant self-attention by comparing features across local coordinate systems after transporting them to a common frame (He et al., 2021).

2.3 Oscillatory and non-stationary dynamics in recurrent reasoning

Recent depth-recurrent language models exhibit structured latent trajectories that are substantially richer than monotone convergence to a fixed point, including multi-dimensional orbits and approximately one-dimensional drifts, particularly on tokens requiring extended computation (Geiping et al., 2025; Pappone et al., 2025; Kim, 2026). Other mechanistic studies find cyclic fixed-point trajectories, where recurrent layers repeatedly trace a stable cycle in latent space, as well as persistent non-stationarity even after successive state differences become small (Huang et al., 2026; Blayney et al., 2026). Oscillatory behavior also appears in recurrent reasoners based on solution equilibria. FPRM, for example, reports oscillation around solution fixed points and introduces damping to stabilize these trajectories (Movahedi et al., 2026). More broadly, looped Transformers with the same architecture, objective, and task performance can settle into different dynamical phases, including bounded non-stationary motion generated by complex-conjugate Jacobian modes (Kim, 2026).

What remains unclear is the dynamical role of this motion. An observed orbit or alternating trajectory may support recurrent computation, or instead reflect a slowly decaying mode that unnecessarily prolongs latent settlement after the remaining dynamics already converged.

3 Rate-Limiting Slow Oscillatory Modes in Recurrent Reasoners

Motivated by oscillatory structures observed in recurrent reasoners, we study whether an identified oscillatory mode can become rate-limiting for latent convergence. Around solution roots, the nonlinear recurrence admits a stationary first-order linearization, allowing us to isolate an oscillatory spectral component from its complementary latent dynamics.

We use a Sylvester decomposition as an analytic attribution tool to separate the latent state dynamics into its intrinsic evolution and the response induced by coupling with the identified spectral coordinate. This provides a direct way to quantify how the identified coordinate influences the remaining latent dynamics and their convergence behavior. The decomposition itself is general to any selected mode, particular oscillatory structure enters through local Jacobian block identified from the trained model.

3.1 Local spectral decomposition and rate dominance

We analyze how an oscillatory component of the local dynamics affects the remaining latent state near a solution root. Let $z_{t+1} = f(z_t)$ denote the trained nonlinear recurrence, and let $z^\star$ be a solution root satisfying $f(z^\star) = z^\star$. For a nearby trajectory state, write the root-centered displacement $\delta_t = z_t - z^\star$. Since $z^\star$ is stationary, the same local Jacobian governs subsequent steps as

$$\delta_{t+1} = J\delta_t + O(\|\delta_t\|^2), \qquad J = Df(z^\star). \tag{1}$$

This stationary linearization allows us to identify and track a fixed spectral component over recurrence. Consider an oscillatory spectral component of J, and let $U \in \mathbb{R}^{D\times m}$ be an orthonormal basis for the corresponding left-invariant subspace of J, and let S denote the restriction of J to this subspace. Equivalently, $J^\top U = US^\top$, with $S = U^\top JU$. Hence, defining a coordinate $u_t := U^\top \delta_t$, the local dynamics satisfy

$$u_{t+1} = U^\top J\delta_t + O(\|\delta_t\|^2) = Su_t + O(\|\delta_t\|^2). \tag{2}$$

Thus u_t is an autonomous spectral coordinate to first order, its next state update depends only on itself through the reduced Jacobian block S. We represent the remaining displacement by the orthogonal residual e_t,

$$u_t = U^\top \delta_t, \qquad e_t = Q\delta_t, \qquad Q = I - UU^\top, \tag{3}$$

so that $\delta_t = Uu_t + e_t$ and $U^\top e_t = 0$. Projecting the linearized dynamics onto these two coordinates, u_t and e_t, gives (see full derivation in Appendix A.2)

$$u_{t+1} = Su_t, \qquad e_{t+1} = Ae_t + Bu_t, \qquad A := QJQ, \quad B := QJU. \tag{4}$$

Here A governs the intrinsic evolution of the complementary state, whereas B measures how the isolated spectral coordinate drives that complement. The observed e_t therefore mixes two effects: its intrinsic transient under A and the response transmitted from u_t coupled by B. We separate these two distinct contributions to evolve e_t using the following Sylvester decomposition.

Lemma 3.1 (Exact Sylvester decomposition of a coupled spectral coordinate)**.** *Consider the linear dynamics in Eq. 4, with $A \in \mathbb{R}^{d\times d}$, $S \in \mathbb{R}^{m\times m}$, and $B \in \mathbb{R}^{d\times m}$, and assume $\sigma(A) \cap \sigma(S) = \emptyset$. Then there exists a unique matrix $P \in \mathbb{R}^{d\times m}$ satisfying $PS - AP = B$. The complementary state admits the exact decomposition*

$$e_t = A^t(e_0 - Pu_0) + PS^t u_0. \tag{5}$$

Lemma 3.1 separates the complementary state e_t into two subsystems. The first term, $A^t(e_0 - Pu_0)$, is its intrinsic transient, whereas $PS^t u_0$ is the response transmitted from the coupled spectral coordinate u_t. The lemma is not specific to oscillatory dynamics, it provides a general framework for attributing the effect of one autonomous spectral component on the rest of a coupled linear system.

Slow oscillatory specialization. We now specialize this decomposition to the reduced block $S = U^\top JU$ observed in the actual recurrent reasoners, where the isolated mode is oscillatory and decays more slowly than the complementary dynamics. When the measured S has a complex-conjugate pair with modulus r and frequency ω, its real two-dimensional dynamics admit

$$\sigma(S) = \{re^{\pm i\omega}\}, \qquad S = M(rQ_\omega)M^{-1}, \qquad 0 < \omega < \pi, \tag{6}$$

where Q_ω is the planar rotation through ω and M is real and invertible. When $\rho(A) < r$, the intrinsic transient in Eq. 5 decays faster than the coupled response $PS^t u_0$. For a trajectory $e = (e_t)_{t\geq 0}$, define

$$\tau_\delta(e) = \inf\{t \in \mathbb{N}_0 : \|e_k\|_2 \leq \delta \text{ for all } k \geq t\}, \qquad \mathcal{E}(e) = \inf_{c\in\mathbb{R}^d} \limsup_{t\to\infty} \|e_t - c\|_2. \tag{7}$$

Theorem 3.2 (Rate inheritance from a slow oscillatory mode)**.** *Under the setting of Lemma 3.1, suppose S satisfies Eq. 6, with $\rho(A) < r \leq 1$, $B \neq 0$, and $u_0 \neq 0$. If $r < 1$, then $e_t \to 0$,*

$$\limsup_{t\to\infty} \|e_t\|_2^{1/t} = r > \rho(A), \tag{8}$$

and

$$\tau_\delta(e) = \frac{\log(1/\delta)}{-\log r} + O(1), \qquad \delta \downarrow 0. \tag{9}$$

If $r = 1$, then e_t remains bounded but does not converge to any constant, with

$$\mathcal{E}(e) \geq \frac{\|u_0\|_2 \|B\|_F}{\sqrt{2}\,\kappa_2(M)(1 + \|A\|_2)} > 0. \tag{10}$$

Theorem 3.2 shows that a slower coupled mode transfers its rate to the complementary dynamics, delaying settlement when $r < 1$ and preventing stationary convergence when $r = 1$. It covers general two-dimensional spiral modes, while $M = I$ recovers a scaled rotation and $S = -r$ gives the alternating flip case of Corollary A.1. All proofs are provided in Appendix A.1.

3.2 Causal verification in nonlinear recurrent dynamics

We now use this framework to ask whether the slow oscillatory modes identified in trained recurrent reasoners actually exert the predicted rate-limiting effect on their nonlinear trajectories. For each solution root, we identify a leading oscillatory eigenspace of the local Jacobian, construct its left-invariant basis U, and obtain S, A, and B matching the decomposition in Section 3.1. We focus on cases satisfying the rate-dominance condition $\rho(A) < r \leq 1$ (Appendix A.4.3).

To test the contribution of this mode in the original nonlinear dynamics, we take a state from the native trajectory once it has entered the local linear regime around the solution root $z^\star$. Writing

$$\delta_0 := z_{t_0} - z^\star = e_0 + Uu_0, \qquad e_0 = Q\delta_0, \qquad u_0 = U^\top \delta_0, \tag{11}$$

separates the displacement into the identified oscillatory coordinate u_0 and the remaining local state e_0. We then vary only the amplitude of u_0 by adjusting α,

$$z_{t_0}^{(\alpha)} = z^\star + e_0 + \alpha Uu_0, \qquad \alpha \in \left\{1, \tfrac{1}{2}, 0\right\}, \tag{12}$$

while keeping e_0 exactly fixed. Here $\alpha = 1$ is the native state, $\alpha = 0.5$ attenuates the oscillatory coordinate, and $\alpha = 0$ removes it at the intervention point. This is the only modification and all trajectories subsequently evolve using the original trained nonlinear recurrence f. Because all arms share the same complementary state e_0, the linearized rollout gives $\widehat{e}_k^{(\alpha)} = Q\delta_k = QJ^k(e_0 + \alpha U u_0)$ (Eq. 3), so subtracting the $\alpha = 0$ arm isolates the response to the oscillatory coordinate. Using the decomposition in Eq. 5, this prediction can be further written equivalently as

$$\widehat{\Delta e}_k^{(\alpha)} = \widehat{e}_k^{(\alpha)} - \widehat{e}_k^{(0)} = \alpha QJ^k U u_0 = \alpha \left(PS^k - A^k P\right) u_0. \tag{13}$$

This response prediction is determined entirely by the local Jacobian and the native displacement at t_0, with no quantities fit to the subsequent nonlinear rollout.

Alternating oscillatory mode in FPRM. We apply the intervention protocol to the rate-dominant oscillatory mode identified in FPRM. As illustrated in Figure 1, this mode appears as an approximately one-dimensional alternating flip, with $S = -r$ and $\rho(A) < r \leq 1$. The corresponding response prediction is (derivations in Appendix A.1)

$$\widehat{\Delta e}_k^{(\alpha)} = \alpha \left[P(-r)^k - A^k P\right] u_0 = \alpha P(-r)^k u_0 + o(r^k). \tag{14}$$

Thus, changing α scales the leading response while leaving its asymptotic decay rate unchanged: $\alpha = \frac{1}{2}$ halves its leading amplitude, whereas $\alpha = 0$ removes the r^k contribution. We compare this closed-form local prediction directly with the response measured from the original nonlinear FPRM recurrence. For $\alpha = 0$, the linearized counterfactual reduces to $\widehat{e}_k^{(0)} = A^k e_0$, allowing us to additionally test whether removing the oscillatory coordinate accelerates root-relative settlement.

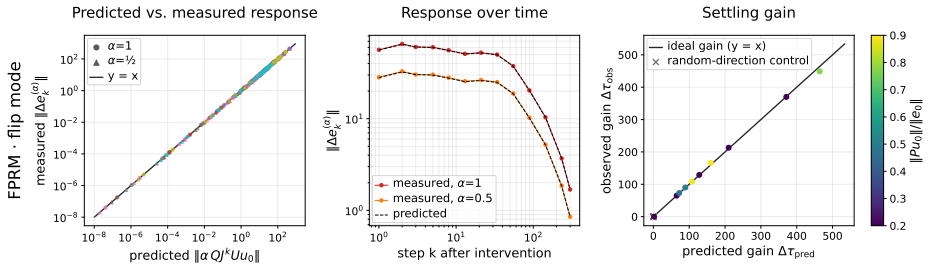


Figure 2: **Local oscillatory-coupling theory predicts the response and settling cost in the original FPRM recurrence.** We control only the Jacobian-identified oscillatory coordinate once with $\alpha \in \{0, 0.5, 1\}$ and then resume unmodified recurrent updates. HRM results in Appendix A.4.1.

Figure 2 tests the oscillatory-coupling mechanism at three levels: (i) the response induced in the complementary state, (ii) its evolution over recurrent time, and (iii) the resulting cost in settling time.

The left panel measures how a perturbation of the oscillatory coordinate subsequently propagates into the complementary latent dynamics over k recurrent steps. For each intervention strength α and rollout step k, the horizontal axis shows the prediction from the root Jacobian, while the vertical axis shows the corresponding change measured from the original nonlinear rollout:

$$\underbrace{\|\Delta e_k^{(\alpha)}\|_2 = \left\| Q\left(z_{t_0+k}^{(\alpha)} - z_{t_0+k}^{(0)}\right)\right\|_2}_{\text{actual nonlinear rollout}} \quad \text{vs.} \quad \underbrace{\|\widehat{\Delta e}_k^{(\alpha)}\|_2 = \|\alpha QJ^k U u_0\|_2}_{\text{local linear prediction}}. \tag{15}$$

Agreement with the identity line shows that the response induced in the complementary state is quantitatively predicted by the local Jacobian.

The middle panel tests whether this prediction remains accurate throughout the subsequent nonlinear rollout. Consistent with Eq. 14, the measured response follows the predicted alternating temporal envelope, including the proportional reduction under $\alpha = 0.5$.

The right panel tests whether the same local mode produces the predicted delay in convergence. Using the root-relative settling criterion $\tau_\delta(e) = \inf\{k : \|e_j\| \leq \delta \text{ for all } j \geq k\}$ from Eq. 7, we define $\hat{\tau}^{(\alpha)} := \tau_\delta(\hat{e}^{(\alpha)})$ and $\tau^{(\alpha)} := \tau_\delta(e^{(\alpha)})$ using the same sustained-window implementation. The corresponding gains from removing the oscillatory coordinate are $\Delta\tau_{\text{pred}} = \hat{\tau}^{(1)} - \hat{\tau}^{(0)}$ and $\Delta\tau_{\text{obs}} = \tau^{(1)} - \tau^{(0)}$. The identity-line agreement therefore confirms that the local prediction quantitatively matches the actual settling-time gain from removing the identified oscillatory component.

Together, the three columns test the causal chain $u_0 \rightarrow \Delta e_k \rightarrow \Delta\tau_\delta$ in the original FPRM. The root Jacobian predicts how the oscillatory coordinate drives the complementary latent state, and removing that coordinate accelerates settlement while preserving the decoded output. Extended analysis for distinct two-dimensional spiral mode in HRM is provided in Appendix A.4.1, where we obtain the same prediction–measurement agreement.

4 Quotient Dynamics Convergence

Our analysis in Section 3 shows that slowly decaying oscillatory modes can limit convergence speed in a locally linearized regime near a solution root. This does not, however, justify suppressing oscillatory motion throughout the entire reasoning trajectory. Rotational or alternating dynamics may still participate in the reasoning computation, and removing them indiscriminately can disrupt reasoning. Indeed, persistently suppressing an alternating trajectory component in FPRM before output stabilization degrades task performance (see Appendix A.4.4). Rather than attempting to directly correct the oscillatory bottlenecks identified above, QRR takes a complementary architectural approach by introducing an exact token-local gauge symmetry.

Importantly, QRR does not assume any correspondence between these gauge directions and the slow modes identified in other reasoners. QRR learns its own recurrent trajectory. The symmetry instead introduces token-local rotational degrees of freedom whose motion is separated from the task-relevant quotient dynamics. How this symmetry changes the learned dynamics is therefore an empirical question: QRR still retains oscillatory modes in quotient space, but does not develop the isolated dominant slow mode observed in FPRM or HRM.

4.1 Token-local gauge symmetry

Let the recurrent state consist of S token states $h = (h_1, \ldots, h_S)$, with $h_s \in \mathbb{C}^d$, together with momentum states $m_s \in \mathbb{C}^d$. In our implementation, each real-valued token state $y_s \in \mathbb{R}^{2d}$ is reinterpreted as $h_s \in \mathbb{C}^d$ through the fixed pairing $h_{s,c} = y_{s,2c-1} + i\, y_{s,2c}$, $c = 1, \ldots, d$. The input representation x is gauge-neutral. We introduce the group $G = U(1)^S$, with one phase ϕ_s for each token. For $g_\phi = (e^{i\phi_1}, \ldots, e^{i\phi_S}) \in G$, the group action is

$$g_\phi \cdot (h, m) = \big((e^{i\phi_s} h_s)_{s=1}^S,\, (e^{i\phi_s} m_s)_{s=1}^S\big), \qquad g_\phi \cdot x = x. \tag{16}$$

Thus, all complex channels within a token share a common phase, while different tokens may rotate independently. Equivalently, in real coordinates, let $\mathcal{J}$ denote the block-diagonal infinitesimal generator of the phase rotation, so that the gauge tangent associated with token s is generated by $\mathcal{J} y_s$ (or, equivalently, by $i h_s$ in complex coordinates). We require the recurrent map T to be G-equivariant while the prediction readout $\mathcal{R}$ is G-invariant,

$$T(g \cdot z; x) = g \cdot T(z; x), \qquad \mathcal{R}(g \cdot z; x) = \mathcal{R}(z; x), \qquad z = (h, m). \tag{17}$$

Hence, states that differ only by token-local phase are identified as the same task-relevant state, and motion along the corresponding gauge directions is invisible to the prediction readout.

4.2 Quotient dynamics

The symmetry identifies all states on the same group orbit, $[z] = G \cdot z = \{g \cdot z : g \in G\}$, yielding the quotient state space $\mathcal{Q} = \mathcal{Z}/G$. Because the recurrent map is G-equivariant, it induces a well-defined map on the quotient,

$$\overline{T}([z]; x) = [T(z; x)]. \tag{18}$$

The quotient removes only symmetry-generated degrees of freedom, while preserving all remaining latent variation. Accordingly, a state may satisfy $T(z^\star; x) = g^\star \cdot z^\star$ for some $g^\star \in G$, making $z^\star$ a *relative equilibrium* where the unreduced trajectory may continue along the group orbit, while $\overline{T}([z^\star]; x) = [z^\star]$ is an ordinary fixed point in the quotient. Moreover, exact equivariance ensures that a perturbation along the gauge orbit remains gauge-tangent after the recurrent update, and therefore produces no change in the quotient state.

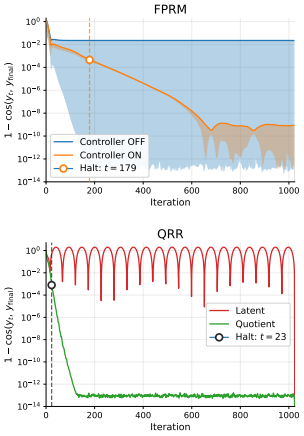


Figure 3: Convergence and halting in quotient space.

At a relative equilibrium, the quotient analogue of Eq. 4 takes the form $e_{t+1} = A_Q e_t + B_G u_t$ with $B_G = 0$, so this symmetry-generated motion cannot transmit its timescale into the quotient dynamics (see Proposition A.2).

Quotient-based halting. The quotient view changes what it means for recurrent reasoning to have converged. At a relative equilibrium, the state may continue moving along its gauge orbit while its quotient state remains fixed. Thus, requiring raw-state stationarity can mistake symmetry-generated phase motion for unfinished computation even after the quotient state has converged. Figure 3 illustrates this distinction between raw-state and quotient convergence. In FPRM, residual oscillatory motion prevents the raw latent state from becoming stationary until it is damped by the step-size controller, making its fixed-point halting criterion sensitive to this motion. In QRR, by contrast, the raw state may continue moving along its gauge orbit after the quotient trajectory has converged to numerical precision, allowing convergence to be detected substantially earlier. We therefore measure convergence modulo the gauge,

$$r_G(z) = \inf_{g \in G} \frac{\|T(z;x) - g \cdot z\|}{\|z\| + \epsilon}. \tag{19}$$

This residual vanishes when the current state and its next update differ only by a gauge transformation, while remaining sensitive to motion in the quotient. At inference, QRR halts once r_G remains below a calibrated threshold for a prescribed persistence window.

4.3 Gauge-equivariant recurrent block

QRR adopts the recurrent Transformer block of FPRM as its backbone, while replacing its internal operations with gauge-equivariant counterparts. Under the token-local symmetry in Equations 16–17, the latent state and momentum transform covariantly with the local token phase, while attention weights and prediction readouts remain gauge-invariant. Figure 4 provides an overview of the overall QRR architectural workflow. Full architectural details are in Appendix A.3.

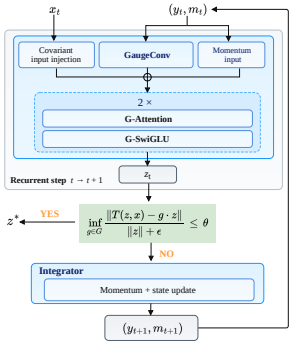


Figure 4: Overview of the QRR recurrent block.

Relative transport. Because each token has its own local phase, features from different tokens must be aligned before they can be mixed. We associate each token with a covariant phase $\gamma_s \in U(1)$, satisfying $\gamma_s \mapsto e^{i\phi_s}\gamma_s$, and define the relative transport $U_{s\leftarrow t} = \gamma_s \gamma_t^*$. Under a pure gauge rotation, $U_{s\leftarrow t} \mapsto e^{i(\phi_s - \phi_t)} U_{s\leftarrow t}$, providing the relative phase required to transport features from token t into the local frame of token s. We use this transport for both token mixing and self-attention.

Gauge-equivariant attention. Queries, keys, and values are produced by complex-linear maps. Using the relative transport, we define the attention scores as

$$a_{st} = \frac{1}{\sqrt{d}} \operatorname{Re}\left[U_{s\leftarrow t}\, q_s^\dagger k_t\right], \tag{20}$$

which are gauge-invariant. Values are likewise transported into the local frame of token s before aggregation, $o_s = W_O \sum_t \operatorname{softmax}_t(a_{st})\, U_{s\leftarrow t} v_t$, so that o_s transforms covariantly with token s.

Gauge-covariant input injection and feed-forward layers. Because the input x_s is gauge-neutral, it cannot be added directly to the recurrent state. We construct a covariant phase carrier

$$D_s = \frac{\langle W_d x_s, h_s \rangle}{|\langle W_d x_s, h_s \rangle| + \epsilon}, \qquad h_s \leftarrow \alpha_0 h_s + \beta_0 D_s x_s, \tag{21}$$

where $D_s \mapsto e^{i\phi_s} D_s$, so that the injected input transforms in the local frame of token s. For the feed-forward block, we use an invariant gate with an equivariant value branch,

$$\text{G-SwiGLU}(h) = W_{\text{down}}\left[\text{SiLU}(|W_g h|) \odot W_v h\right]. \tag{22}$$

Momentum-carrying recurrent dynamics. We use a momentum-carrying second-order recurrence to evolve the latent state. At each recurrent step, the gauge-equivariant trunk produces a state proposal $z_t = \Phi_\theta(y_t; x, m_t)$, defining the residual $F_t = z_t - y_t$. We update the momentum as

$$p_t = \beta m_t + (1-\beta) F_t, \qquad \beta = \sigma(b_{\text{H}}), \tag{23}$$

where $b_{\text{H}} \in \mathbb{R}$ is a learned global scalar. The recurrent state is then updated by

$$y_{t+1} = y_t + p_t, \qquad m_{t+1} = p_t. \tag{24}$$

The carried momentum is also fed back into the recurrent trunk at the next step.

4.4 Spectral Dominance and Convergence in Reasoning Dynamics

We next examine how these dynamical constraints are reflected in the learned spectra of QRR.

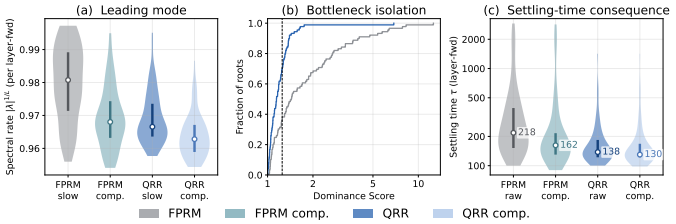


Figure 5: Spectral analysis and settling near solution roots for FPRM and QRR. Results for HRM in Appendix A.4.2

We characterize the local spectral structure of FPRM and QRR around certified solution roots that decode to the ground truth. FPRM roots are ordinary fixed points of the latent dynamics, whereas QRR roots are relative equilibria, fixed points in quotient space. We linearize FPRM in latent space and QRR in quotient space, and report spectral rates as $|\lambda|^{1/L}$, where L is the number of recurrent layer-forwards (lfwd) per native step. Following Section 3, we decompose each system into its leading spectral mode and complementary dynamics. For FPRM, the leading mode is slow oscillatory block identified in Section 3.2. For QRR, both are derived from quotient space. Figure 5 (a) shows QRR operates in a faster spectral regime with less separation between its leading and remaining dynamics. Figure 5 (b) quantifies this separation using the dominance score $D = \log r_{\text{comp}} / \log r_{\text{slow}}$, $D = 1$ indicates matched timescales, while larger D indicates a more isolated slow leading mode. FPRM frequently exhibits strong spectral dominance, whereas QRR remains much closer to $D = 1$ despite retaining some weaker oscillatory modes. Panel (c) shows the corresponding time settlement. QRR settles earlier than FPRM, while factoring out the leading-mode response accelerates FPRM ($218 \rightarrow 162$) but the effect remains only modest in QRR ($138 \rightarrow 130$). Thus, QRR operates in a faster convergence regime with weaker dependence on a single slow leading mode.

5 Experiments

Datasets and baselines We evaluate on four reasoning benchmarks covering constraint satisfaction, path finding, and abstract transformation. **Sudoku-Extreme** consists of challenging partially filled 9×9 Sudoku puzzles with unique solutions; following the standard HRM protocol, we train on 1,000 base puzzles with $1{,}000\times$ augmentation and evaluate using exact sequence accuracy on the full test set (Wang et al., 2025). **Maze-Hard** is a 30×30 shortest-path benchmark with 1,000 training and 1,000 test mazes. We follow the FPRM setup with no augmentation (Movahedi et al., 2026). For **Mini-ARC**, a compact 5×5 ARC variant for abstract rule induction (Kim et al., 2022; Huang et al., 2026), we use the released tasks with 100 data augmentations per task. We additionally evaluate on **Maze-Unique**, following the EqR construction. Unlike Maze-Hard, where multiple optimal paths may correspond to a single supervised target, Maze-Unique uses perfect tree-structured mazes with a unique start–goal solution (Kim et al., 2022). This provides a cleaner convergence analysis, avoiding the ambiguous solution geometry that can hinder formation of a single stable basin (Huang et al., 2026). We compare against HRM, TRM, EqR, and FPRM, spanning hierarchical recurrence, recursive refinement, learned attractors, and fixed-point adaptive halting.

5.1 Results

Table 1 summarizes both reasoning performance and computational cost. In Table 1 (a), we prioritize matched evaluation from an official checkpoint when available, otherwise use the official result, and reproduce the baseline under its native setup. HRM, TRM, and EqR use their own fixed inference budgets, whereas FPRM and QRR halt adaptively according to fixed-point residuals evaluated in the latent and quotient space, respectively. For the latter, we report halt-or-cap performance with caps of 35k recurrent steps on Sudoku and 3k on the Maze tasks. QRR achieves the strongest performance on Sudoku, Maze-Unique, and Mini-ARC while remaining competitive on Maze-Hard. Table 1 (b) reports corresponding runtime cost in single A100 GPU on Sudoku-Extreme. QRR takes approximately $1.5\times$ higher LFWD latency and $1.3\times$ higher memory than FPRM, reflecting the additional cost of its complex-valued gauge-equivariant operations and explicit phase transport. Nevertheless, faster convergence outweighs this overhead. The median end-to-end rollout time decreases from 5.4 s to 1.9 s, and the p95 from 168.1 s to 60.0 s ($2.8\times$ faster). Thus, despite its higher per-LFWD cost, QRR completes full reasoning rollout faster than FPRM.

Table 2 (a) evaluates root convergence on Sudoku-Extreme using Anderson acceleration applied to each trained recurrent map. Full solver, certification, and compute-normalization details are provided in Appendix A.5.2. *Correct root found* is the fraction of examples for which the solver finds a root decoding to the ground-truth solution, while τ_{root} is the first sustained entry of the native trajectory into a neighborhood of that root, measured in recurrent layer-forward equivalents. For QRR, distance

Table 1: **Reasoning performance and runtime.** (a) exact accuracy on Sudoku-Extreme, Maze-Hard, and Maze-Unique, and task-level Pass@1 on Mini-ARC. Models with adaptive halting (FPRM and QRR) report halt-or-cap results. Superscripts denote: [1] evaluation matched from an official checkpoint, [2] official report, and [3] full reproduction following the native setup. (b) Comparison of per-layer-forward (LFWD) computational cost and end-to-end rollout time (s).

(a) Reasoning performance

Model	Params	Sudoku-Extreme	Maze		Mini-ARC
	(M)	Exact ↑	Hard Exact ↑	Unique Exact ↑	Pass@1 ↑
HRM	27	55.4[1]	74.5[2]	51.9[3]	56.4[3]
TRM	7	75.7[1]	85.3[2]	58.0[3]	40.9[3]
EqR	7	89.9[1]	–	88.9[2]	41.6[3]
FPRM	7	91.5[1]	**86.4**[1]	–	–
FPRM (Rep.)	7	83.7[3]	85.3[3]	94.7[3]	57.1[3]
QRR	**7**	**97.9**	85.8	**97.0**	**62.4**

(b) Runtime comparison

Setting	QRR	FPRM
Per LFWD ($B = 512$)		
Runtime (ms)	34.7	22.3
Peak memory (GB)	2.88	2.14
End-to-end rollout		
Coverage (%)	98.0	100.0
Halt step, p50 / p95	27 / 865	122 / 3,769
Runtime (s), p50 / p95	1.9 / 60.0	5.4 / 168.1

Table 2: **Convergence and adaptive halting.** (a) Root convergence on Sudoku-Extreme. (b) Adaptive halting of FPRM and QRR across tasks. Coverage is the fraction halted within finite cap.

(a) Root convergence on Sudoku-Extreme

Model	τ_{root} median [IQR] ↓	Correct root found (%) ↑
HRM	1,032 [744, 1,272]	55.2
EqR	672 [546, 966]	86.1
FPRM	234 [152, 396]	90.4
QRR	**122 [108, 148]**	89.6

(b) Adaptive halting (cap Sudoku 35k, Maze-Unique 3k)

Task	Model	Coverage (%) ↑	τ_{halt} median [IQR] ↓	Acc.@halt (%) ↑	False halt (%) ↓
Sudoku-Extreme	FPRM	100.0	122 [39, 295]	91.1	8.9
	QRR	98.0	27 [23, 63]	99.9	0.1
Maze-Unique	FPRM	96.4	36 [27, 76]	98.2	1.8
	QRR	98.0	42 [32, 61]	98.5	1.5

is measured in quotient space after removing token-local gauge motion. QRR reaches solution roots significantly earlier than the baselines while maintaining a root-finding rate comparable to FPRM. Table 2 (b) evaluates deployed adaptive halting under finite rollout caps. *Coverage* is the fraction halted within the cap, while τ_{halt} reports the median halting step and IQR over halted examples. Among these, *Acc.@halt* is the fraction with an exactly correct prediction, and *False halt* the fraction halted incorrectly. QRR halts earlier with fewer false halts on Sudoku, while achieving comparable halting depth with higher coverage and accuracy on Maze-Unique.

Ablations

We ablate the main structural components of QRR. Restricting the token-local $U(1)^S$ symmetry to a single global $U(1)$ preserves moderate task accuracy but never triggers halting, indicating that a shared global phase is too restrictive to account for the residual token-wise phase motion that persists after the decoded answer settles. Removing momentum and using first-order dynamics disrupts both exact accuracy and coverage, while leaving the faster halting time and quotient halting precision of trajectories that do settle largely intact. Finally, replacing the G-equivariant trunk with FPRM trunk while applying momentum sharply degrades accuracy and leads to unreliable convergence. All examples eventually halt, but only after longer trajectories and with a high false-halt.

Variant	Exact (%) ↑	Coverage (%) ↑	False halt (%) ↓	τ_{halt} [IQR] ↓
Training				
QRR	97.9	98.0	0.1	27 [23, 63]
global $U(1)$	91.3	0.0	–	CAP (never fires)
first order w/o momentum	87.8	79.6	0.1	28 [22, 44]
FPRM trunk + momentum	77.0	100	23.0	12,650 [12,522, 12,871]

Table 3: **QRR ablations on Sudoku-Extreme.**

6 Conclusion

We studied how non-stationary latent dynamics constrain convergence in recurrent reasoners. Our analysis shows that a slowly decaying oscillatory mode can transmit its rate to coupled latent dynamics and become the bottleneck to settlement, while direct suppression of such motion can also disrupt the reasoning computation. This motivates QRR, which does not eliminate oscillations based on their dynamical appearance, but instead introduces an exact token-local gauge symmetry under which symmetry-generated motion leaves the quotient state unchanged by construction. Exact gauge equivariance prevents this motion from forcing the quotient dynamics, allowing convergence and halting to be defined modulo the symmetry. Across various reasoning tasks, QRR achieves strong reasoning performance while reaching solution roots much earlier and supporting reliable adaptive halting. These results suggest that convergence in recurrent reasoning is better understood in terms of task-relevant quotient space than raw latent space.

AI USE STATEMENT

We used generative AI tools to assist with writing and editing, literature retrieval and discovery, and research ideation and execution. All AI-assisted content was reviewed by the authors, who take full responsibility for the final manuscript.

ETHICS STATEMENT

This work studies computational reasoning using synthetic and publicly available benchmark tasks. It does not involve human participants, sensitive personal data, or deployment in high-stakes settings. We are not aware of ethical concerns specific to this work beyond the general considerations associated with machine learning research.

REPRODUCIBILITY STATEMENT

We provide complete proofs of the theoretical results in Appendix A.1– A.2, detailed architectural and training specifications in Appendix A.3 and A.5, and additional experimental analyses in Appendix A.4. We also provide anonymized source code, configurations, and evaluation scripts as supplementary material to reproduce the reported experiments.

REFERENCES

Arpit Bansal, Avi Schwarzschild, Eitan Borgnia, Zeyad Emam, Furong Huang, Micah Goldblum, and Tom Goldstein. End-to-end algorithm synthesis with recurrent networks: Extrapolation without overthinking. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 20232–20242. Curran Associates, Inc., 2022. doi: 10.52202/068431-1471. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/7f70331dbe58ad59d83941dfa7d975aa-Paper-Conference.pdf.

Hugh Blayney, Álvaro Arroyo, Johan Obando-Ceron, Pablo Samuel Castro, Aaron Courville, Michael M. Bronstein, and Xiaowen Dong. A mechanistic analysis of looped reasoning language models, 2026.

Michael M. Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021.

Taco Cohen, Maurice Weiler, Berkay Kicanaoglu, and Max Welling. Gauge equivariant convolutional networks and the icosahedral CNN. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 1321–1330. PMLR, 2019a. URL https://proceedings.mlr.press/v97/cohen19d.html.

Taco S. Cohen, Mario Geiger, and Maurice Weiler. A general theory of equivariant cnns on homogeneous spaces. *arXiv preprint arXiv:1811.02017*, 2019b.

Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Lukasz Kaiser. Universal transformers. In *International Conference on Learning Representations*, 2019.

Ying Fan, Yilun Du, Kannan Ramchandran, and Kangwook Lee. Looped transformers for length generalization. In *International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=2edigk8yoU.

Jacob Fein-Ashley and Paria Rashidinejad. Solve the loop: Attractor models for language and reasoning, 2026. URL https://arxiv.org/abs/2605.12466.

Jonas Geiping, Sean McLeish, Neel Jain, John Kirchenbauer, Siddharth Singh, Brian R. Bartoldson, Bhavya Kailkhura, Abhinav Bhatele, and Tom Goldstein. Scaling up test-time compute with latent reasoning: A recurrent depth approach, 2025. URL https://arxiv.org/abs/2502.05171.

Jan E. Gerken, Jimmy Aronsson, Oscar Carlsson, Hampus Linander, Fredrik Ohlsson, Christoffer Petersson, and Daniel Persson. Geometric deep learning and equivariant neural networks. *Artificial Intelligence Review*, 56:14605–14662, 2023. doi: 10.1007/s10462-023-10502-7.

Pim De Haan, Maurice Weiler, Taco Cohen, and Max Welling. Gauge equivariant mesh cnns: Anisotropic convolutions on geometric graphs. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=Jnspzp-oIZE.

Lingshen He, Yiming Dong, Yisen Wang, Dacheng Tao, and Zhouchen Lin. Gauge equivariant transformer. In *Advances in Neural Information Processing Systems*, volume 34, pp. 27331–27343, 2021.

Benhao Huang, Zhengyang Geng, and Zico Kolter. Equilibrium reasoners: Learning attractors enables scalable reasoning, 2026. URL https://arxiv.org/abs/2605.21488.

Alexia Jolicoeur-Martineau. Less is more: Recursive reasoning with tiny networks, 2025. URL https://arxiv.org/abs/2510.04871.

Gunn Kim. Dynamical phase selection controls compute scaling in looped transformers, 2026. URL https://arxiv.org/abs/2608.26556.

Subin Kim, Prin Phunyaphibarn, Donghyun Ahn, and Sundong Kim. Playgrounds for Abstraction and Reasoning. In *NeurIPS Workshop on neuro Causal and Symbolic AI*, 2022.

Risi Kondor and Shubhendu Trivedi. On the generalization of equivariance and convolution in neural networks to the action of compact groups. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 2747–2755. PMLR, 2018.

Martin Krupa. Bifurcations of relative equilibria. *SIAM Journal on Mathematical Analysis*, 21(6): 1453–1486, 1990. doi: 10.1137/0521081.

Jerrold E. Marsden and Tudor S. Ratiu. *Introduction to Mechanics and Symmetry: A Basic Exposition of Classical Mechanical Systems*, volume 17 of *Texts in Applied Mathematics*. Springer, 1999.

Sajad Movahedi, Vera Milovanović, Shlomo Libo Feigin, Alexander Theus, Thomas Hofmann, Valentina Boeva, T. Konstantin Rusch, and Antonio Orvieto. Fixed-point reasoners: Stable and adaptive deep looped transformers, 2026. URL https://arxiv.org/abs/2606.18206.

Francesco Pappone, Donato Crisostomi, and Emanuele Rodolà. Two-scale latent dynamics for recurrent-depth transformers, 2025. URL https://arxiv.org/abs/2509.23314.

Nikunj Saunshi, Nishanth Dikkala, Zhiyuan Li, Sanjiv Kumar, and Sashank J. Reddi. Reasoning with latent thoughts: On the power of looped transformers. In *International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=din0lGfZFd.

Guan Wang, Jin Li, Yuhao Sun, Xing Chen, Changling Liu, Yue Wu, Meng Lu, Sen Song, and Yasin Abbasi Yadkori. Hierarchical reasoning model, 2025. URL https://arxiv.org/abs/2506.21734.

Maurice Weiler, Patrick Forré, Erik Verlinde, and Max Welling. Coordinate independent convolutional networks: Isometry and gauge equivariant convolutions on riemannian manifolds. *arXiv preprint arXiv:2106.06020*, 2021.

Ruben Wiersma, Elmar Eisemann, and Klaus Hildebrandt. Cnns on surfaces using rotation-equivariant features. *ACM Transactions on Graphics*, 39(4), 2020. doi: 10.1145/3386569.3392437.

Xiao-Wen Yang, Ziyu Han, Xi-Hua Zhang, Wen-Da Wei, Jie-Jing Shao, Lan-Zhe Guo, and Yu-Feng Li. Stabilizing recurrent dynamics for test-time scalable latent reasoning in looped language models, 2026. URL https://arxiv.org/abs/2605.26733.

A Appendix

A.1 Theoretical Proofs

Lemma 3.1 (Exact Sylvester decomposition of a coupled spectral coordinate)**.** *Consider the linear dynamics in Eq. 4, with $A \in \mathbb{R}^{d\times d}$, $S \in \mathbb{R}^{m\times m}$, and $B \in \mathbb{R}^{d\times m}$, and assume $\sigma(A) \cap \sigma(S) = \emptyset$. Then there exists a unique matrix $P \in \mathbb{R}^{d\times m}$ satisfying $PS - AP = B$. The complementary state admits the exact decomposition*

$$e_t = A^t(e_0 - Pu_0) + PS^t u_0. \tag{5}$$

Proof. Define the Sylvester operator

$$\mathcal{L}(X) := XS - AX, \qquad \mathcal{L} : \mathbb{R}^{d\times m} \to \mathbb{R}^{d\times m}.$$

Using vectorization,

$$\mathrm{vec}(\mathcal{L}(X)) = \left(S^\top \otimes I_d - I_m \otimes A\right) \mathrm{vec}(X). \tag{25}$$

The eigenvalues of $S^\top \otimes I_d - I_m \otimes A$ are

$$\mu - \lambda, \qquad \mu \in \sigma(S), \quad \lambda \in \sigma(A).$$

Because $\sigma(A) \cap \sigma(S) = \emptyset$, none of these eigenvalues is zero. Hence the matrix in Eq. 25 is invertible, so $\mathcal{L}$ is bijective. Therefore, for every B, there exists a unique $P \in \mathbb{R}^{d\times m}$ satisfying

$$PS - AP = B.$$

Now define the reparameterized auxiliary variable using P,

$$w_t := e_t - Pu_t. \tag{26}$$

Using $e_{t+1} = Ae_t + Bu_t$, $u_{t+1} = Su_t$, and $PS - AP = B$, we obtain

$$\begin{aligned} w_{t+1} &= e_{t+1} - Pu_{t+1} && (27)\\ &= Ae_t + Bu_t - PSu_t && (28)\\ &= Ae_t - APu_t && (29)\\ &= A(e_t - Pu_t) = Aw_t. && (30) \end{aligned}$$

Thus

$$w_t = A^t w_0 = A^t(e_0 - Pu_0).$$

Since $u_t = S^t u_0$, substituting back into $e_t = w_t + Pu_t$ gives

$$e_t = A^t(e_0 - Pu_0) + PS^t u_0,$$

which proves Eq. 5. □

Theorem 3.2 (Rate inheritance from a slow oscillatory mode)**.** *Under the setting of Lemma 3.1, suppose S satisfies Eq. 6, with $\rho(A) < r \le 1$, $B \neq 0$, and $u_0 \neq 0$. If $r < 1$, then $e_t \to 0$,*

$$\limsup_{t\to\infty} \|e_t\|_2^{1/t} = r > \rho(A), \tag{8}$$

and

$$\tau_\delta(e) = \frac{\log(1/\delta)}{-\log r} + O(1), \qquad \delta \downarrow 0. \tag{9}$$

If $r = 1$, then e_t remains bounded but does not converge to any constant, with

$$\mathcal{E}(e) \ge \frac{\|u_0\|_2 \|B\|_F}{\sqrt{2}\,\kappa_2(M)(1 + \|A\|_2)} > 0. \tag{10}$$

Proof. Since

$$S = M(rQ_\omega)M^{-1}$$

and every eigenvalue of S has modulus $r > \rho(A)$, we have

$$\sigma(A) \cap \sigma(S) = \emptyset.$$

Hence Lemma 3.1 gives a unique solution P of $PS - AP = B$ and the exact decomposition

$$e_t = A^t(e_0 - Pu_0) + PS^t u_0. \tag{31}$$

Since $B \neq 0$, the Sylvester equation implies $P \neq 0$. Define

$$q := M^{-1}u_0, \qquad C := PM, \qquad a := \|q\|_2.$$

Invertibility of M implies $a > 0$ and $C \neq 0$. Moreover,

$$S^t = r^t M Q_\omega^t M^{-1},$$

so

$$PS^t u_0 = r^t C Q_\omega^t q.$$

This change of oscillator coordinates leaves the complementary state e_t and its Euclidean norm unchanged. Choose ν such that

$$\rho(A) < \nu < r.$$

There exists $C_\nu < \infty$ satisfying

$$\|A^t\|_2 \leq C_\nu \nu^t, \qquad t \geq 0.$$

Writing

$$\xi_t := A^t(e_0 - Pu_0), \qquad p_t := CQ_\omega^t q,$$

we obtain

$$e_t = r^t p_t + \xi_t, \qquad \|\xi_t\|_2 = O(\nu^t) = o(r^t), \tag{32}$$

where ξ_t takes the A-transient part and $r^t p_t$ corresponds to the oscillatory response. We first establish a uniform two-step lower bound on the oscillatory response. This accounts for possible cancellation at individual steps. Let

$$v_t := \frac{Q_\omega^t q}{a} \in \mathbb{R}^2, \qquad V_t := (v_t \quad Q_\omega v_t) \in \mathbb{R}^{2\times 2}.$$

Because Q_ω is orthogonal,

$$\|v_t\|_2 = \|Q_\omega v_t\|_2 = 1, \qquad \langle v_t, Q_\omega v_t\rangle = \cos\omega.$$

Thus

$$V_t^\top V_t = \begin{pmatrix} 1 & \cos\omega \\ \cos\omega & 1 \end{pmatrix}$$

has eigenvalues $1 \pm |\cos\omega|$. Moreover, because v_t and $Q_\omega v_t$ are non-collinear for $0 < \omega < \pi$, the Gram matrix $V_t^\top V_t$ is positive definite. The same eigenvalues occur in $V_t V_t^\top$, giving

$$(1 - |\cos\omega|)I_2 \preceq V_t V_t^\top \preceq (1 + |\cos\omega|)I_2.$$

Since

$$p_t = aCv_t, \qquad p_{t+1} = aCQ_\omega v_t,$$

we have

$$\|p_t\|_2^2 + \|p_{t+1}\|_2^2 = a^2\|CV_t\|_F^2 = a^2 \operatorname{tr}(CV_tV_t^\top C^\top).$$

Consequently,

$$a^2(1 - |\cos\omega|)\|C\|_F^2 \leq \|p_t\|_2^2 + \|p_{t+1}\|_2^2 \leq a^2(1 + |\cos\omega|)\|C\|_F^2. \tag{33}$$

The lower bound is strictly positive because $0 < \omega < \pi$, $a > 0$, and $C \neq 0$. The results show that p_t can be temporarily collapsed close to 0 but p_t and p_{t+1} cannot become arbitrarily small at the same time.

Extending this result to $e_t = r^t p_t + \xi_t$, define the two-step envelope

$$W_t := \left(\|e_t\|_2^2 + r^{-2}\|e_{t+1}\|_2^2\right)^{1/2}. \tag{34}$$

Eq. 32 gives

$$\begin{pmatrix} e_t \\ r^{-1}e_{t+1} \end{pmatrix} = r^t \begin{pmatrix} p_t \\ p_{t+1} \end{pmatrix} + \begin{pmatrix} \xi_t \\ r^{-1}\xi_{t+1} \end{pmatrix},$$

where

$$\left\| \begin{pmatrix} \xi_t \\ r^{-1}\xi_{t+1} \end{pmatrix} \right\|_2 = O(\nu^t) = o(r^t).$$

The triangle and reverse triangle inequalities therefore imply

$$\left| W_t - r^t \left(\|p_t\|_2^2 + \|p_{t+1}\|_2^2 \right)^{1/2} \right| = O(\nu^t).$$

Combining this with Eq. 33, there exist constants $c_-, c_+ > 0$ and t_0 such that

$$c_- r^t \le W_t \le c_+ r^t, \qquad t \ge t_0. \tag{35}$$

Now suppose $r < 1$. Since

$$\|p_t\|_2 \le \|C\|_2 a,$$

Eq. 32 implies $e_t \to 0$ and

$$\limsup_{t\to\infty} \|e_t\|_2^{1/t} \le r.$$

If this limsup were strictly smaller than r, there would exist $0 < s < r$ such that

$$\|e_t\|_2 = O(s^t).$$

It would follow from Eq. 34 that

$$W_t = O(s^t),$$

contradicting the lower bound $W_t \ge c_- r^t$ in Eq. 35. Hence

$$\limsup_{t\to\infty} \|e_t\|_2^{1/t} = r,$$

which proves Eq. 8. We next derive the settling-time asymptotic. Since $e_t \to 0$, $\tau_\delta(e)$ is finite for every $\delta > 0$. The positive lower bound in Eq. 35 rules out an eventually zero trajectory,

$$\tau_\delta(e) \longrightarrow \infty \qquad \text{as } \delta \downarrow 0.$$

Thus, for sufficiently small δ, the relevant settling times are at least t_0. If $t \ge t_0$ satisfies

$$c_+ r^t \le \delta,$$

then, for every $k \ge t$,

$$\|e_k\|_2 \le W_k \le c_+ r^k \le c_+ r^t \le \delta.$$

Therefore, for sufficiently small δ,

$$\tau_\delta(e) \le \left\lceil \frac{\log(c_+/\delta)}{-\log r} \right\rceil \le 1 + \frac{\log(c_+/\delta)}{-\log r} = \frac{\log(1/\delta)}{-\log r} + O(1). \tag{36}$$

Conversely for the lower bound, if $\|e_k\|_2 \le \delta$ for every $k \ge t$, then

$$W_t \le \delta\sqrt{1 + r^{-2}}.$$

At $t = \tau_\delta(e) \ge t_0$, the lower bound in Eq. 35 therefore gives

$$c_- r^{\tau_\delta(e)} \le \delta\sqrt{1 + r^{-2}},$$

and hence

$$\tau_\delta(e) \ge \frac{\log\big(c_-/(\delta\sqrt{1 + r^{-2}})\big)}{-\log r} = \frac{\log(1/\delta)}{-\log r} + O(1). \tag{37}$$

Combining Eq. 36 and Eq. 37 yields

$$\tau_\delta(e) = \frac{\log(1/\delta)}{-\log r} + O(1), \qquad \delta \downarrow 0,$$

which proves Eq. 9. This shows that e_t inherits the slow rate of the oscillatory mode with $r < 1$ due to the coupling.

Now it remains to consider $r = 1$. Eq. 32 becomes

$$e_t = p_t + \xi_t, \qquad p_t = CQ_\omega^t q, \qquad \xi_t \to 0. \tag{38}$$

The sequence (e_t) is bounded because (p_t) is bounded and ξ_t decays geometrically. Write

$$q = a(\cos\phi, \sin\phi)^\top.$$

Since $0 < \omega < \pi$, both $e^{i\omega}$ and $e^{2i\omega}$ differ from 1. For $\ell \in \{1, 2\}$, the geometric-sum formula gives

$$\frac{1}{T}\sum_{t=0}^{T-1} e^{i\ell t\omega} = \frac{1 - e^{i\ell T\omega}}{T(1 - e^{i\ell\omega})} \longrightarrow 0.$$

Applying the $\ell = 1$ identity to

$$Q_\omega^t q = a\begin{pmatrix}\cos(\phi + t\omega)\\ \sin(\phi + t\omega)\end{pmatrix}$$

yields zero-mean for the rotating vector $Q_\omega^t q$

$$\frac{1}{T}\sum_{t=0}^{T-1} Q_\omega^t q \longrightarrow 0. \tag{39}$$

For the second moment, the double-angle identities give

$$Q_\omega^t qq^\top (Q_\omega^t)^\top = \frac{a^2}{2} I_2 + \frac{a^2}{2}\begin{pmatrix}\cos(2\phi + 2t\omega) & \sin(2\phi + 2t\omega)\\ \sin(2\phi + 2t\omega) & -\cos(2\phi + 2t\omega)\end{pmatrix}.$$

The $\ell = 2$ identity therefore implies

$$\frac{1}{T}\sum_{t=0}^{T-1} Q_\omega^t qq^\top (Q_\omega^t)^\top \longrightarrow \frac{a^2}{2} I_2. \tag{40}$$

These limits require no irrationality assumption on $\omega/(2\pi)$. Since $p_t = CQ_\omega^t q$, we obtain

$$\frac{1}{T}\sum_{t=0}^{T-1} p_t = C\left(\frac{1}{T}\sum_{t=0}^{T-1} Q_\omega^t q\right) \longrightarrow 0,$$

and

$$\frac{1}{T}\sum_{t=0}^{T-1} \|p_t\|_2^2 = \operatorname{tr}\left[C\left(\frac{1}{T}\sum_{t=0}^{T-1} Q_\omega^t qq^\top (Q_\omega^t)^\top\right) C^\top\right] \longrightarrow \frac{a^2}{2}\|C\|_F^2.$$

Thus

$$\frac{1}{T}\sum_{t=0}^{T-1} p_t \longrightarrow 0, \qquad \frac{1}{T}\sum_{t=0}^{T-1} \|p_t\|_2^2 \longrightarrow \frac{a^2}{2}\|C\|_F^2. \tag{41}$$

Fix any $c \in \mathbb{R}^d$. The squared distance satisfies

$$\|e_t - c\|_2^2 = \|p_t - c\|_2^2 + 2\langle p_t - c, \xi_t\rangle + \|\xi_t\|_2^2.$$

Because ξ_t decays geometrically, the last two terms vanish as

$$\frac{1}{T}\sum_{t=0}^{T-1} \|\xi_t\|_2 \longrightarrow 0, \qquad \frac{1}{T}\sum_{t=0}^{T-1} \|\xi_t\|_2^2 \longrightarrow 0.$$

Moreover,

$$\left|\frac{1}{T}\sum_{t=0}^{T-1} \langle p_t - c, \xi_t\rangle\right| \le (\|C\|_2 a + \|c\|_2)\,\frac{1}{T}\sum_{t=0}^{T-1} \|\xi_t\|_2 \longrightarrow 0.$$

Using

$$\|p_t - c\|_2^2 = \|p_t\|_2^2 - 2\langle p_t, c\rangle + \|c\|_2^2,$$

Eq. 41 therefore yields

$$\lim_{T\to\infty}\frac{1}{T}\sum_{t=0}^{T-1}\|e_t-c\|_2^2=\frac{a^2}{2}\|C\|_F^2+\|c\|_2^2. \tag{42}$$

For every c, the limsup is at least the limiting Cesàro average,

$$\limsup_{t\to\infty}\|e_t-c\|_2^2\geq\lim_{T\to\infty}\frac{1}{T}\sum_{t=0}^{T-1}\|e_t-c\|_2^2.$$

Taking square roots and then the infimum over c gives

$$\mathcal{E}(e)\geq\frac{a}{\sqrt{2}}\|C\|_F=\frac{\|q\|_2\|C\|_F}{\sqrt{2}}=\frac{\|M^{-1}u_0\|_2\|PM\|_F}{\sqrt{2}}. \tag{43}$$

Finally, we express this bound in the original coordinates. Multiplying $PS-AP=B$ on the right by M, and using $SM=MQ_\omega$ when $r=1$, gives

$$BM=PSM-APM=CQ_\omega-AC.$$

Since Q_ω is orthogonal,

$$\|CQ_\omega\|_F=\|C\|_F.$$

Consequently,

$$\begin{aligned}\|BM\|_F&\leq\|CQ_\omega\|_F+\|AC\|_F\\&\leq\|C\|_F+\|A\|_2\|C\|_F\\&=(1+\|A\|_2)\|C\|_F.\end{aligned}$$

Substituting into Eq. 43 gives

$$\mathcal{E}(e)\geq\frac{\|M^{-1}u_0\|_2\|BM\|_F}{\sqrt{2}(1+\|A\|_2)}.$$

The identities $u_0=M(M^{-1}u_0)$ and $B=(BM)M^{-1}$ imply

$$\|u_0\|_2\leq\|M\|_2\|M^{-1}u_0\|_2$$

and

$$\|B\|_F\leq\|BM\|_F\|M^{-1}\|_2.$$

Therefore,

$$\|M^{-1}u_0\|_2\|BM\|_F\geq\frac{\|u_0\|_2\|B\|_F}{\|M\|_2\|M^{-1}\|_2}=\frac{\|u_0\|_2\|B\|_F}{\kappa_2(M)}.$$

We conclude that

$$\mathcal{E}(e)\geq\frac{\|u_0\|_2\|B\|_F}{\sqrt{2}\,\kappa_2(M)(1+\|A\|_2)}>0,$$

which proves Eq. 10. Hence the complementary state remains bounded but cannot converge to any constant. □

Corollary A.1 (Alternating modes)**.** *Consider the alternating oscillatory mode*

$$e_{t+1}=Ae_t+bu_t,\qquad u_{t+1}=-ru_t, \tag{44}$$

where $e_t,b\in\mathbb{R}^d$, $u_t\in\mathbb{R}$, $A\in\mathbb{R}^{d\times d}$, $\rho(A)<r\leq 1$, and $bu_0\neq 0$. If $r<1$, the rate and settling-time conclusions from 3.2 hold. More precisely, we can further derive

$$\lim_{t\to\infty}\frac{\|e_t\|_2}{r^t}=\|(-rI-A)^{-1}bu_0\|_2>0. \tag{45}$$

If $r=1$, the trajectory approaches an alternating two-cycle, and

$$\mathcal{E}(e)=\|(I+A)^{-1}bu_0\|_2>0. \tag{46}$$

Proof. Since $\rho(A) < r$, the matrix $-rI - A$ is invertible. Define

$$v = (-rI - A)^{-1} b u_0. \tag{47}$$

Then $v \neq 0$ and $-rv - Av = bu_0$. Direct substitution into Eq. 44 gives

$$e_t = A^t(e_0 - v) + (-r)^t v. \tag{48}$$

The matrix-power bound implies $A^t(e_0 - v) = o(r^t)$. Therefore,

$$r^{-t} e_t = (-1)^t v + o(1), \tag{49}$$

which proves Eq. 45.

For $r < 1$, the norm is consequently bounded above and below by positive multiples of r^t for all sufficiently large t. The exponential-rate and settling-time conclusions follow as in Theorem 3.2.

For $r = 1$, Eq. 48 gives

$$e_{2t} \longrightarrow v, \qquad e_{2t+1} \longrightarrow -v. \tag{50}$$

Thus

$$\mathcal{E}(e) = \inf_{c \in \mathbb{R}^d} \max\{\|v - c\|_2, \| - v - c\|_2\} = \|v\|_2. \tag{51}$$

The last equality follows because any ball containing v and $-v$ has radius at least $\|v\|_2$, attained at $c = 0$. Since $v = -(I + A)^{-1} b u_0$ when $r = 1$, this proves Eq. 46. □

Perturbation response In the alternating case $S = -r$, the Sylvester equation gives

$$P(-r) - AP = b, \qquad P = (-rI - A)^{-1} b.$$

Hence the linearized complementary trajectory is

$$\widehat{e}_k^{(\alpha)} = A^k e_0 + \alpha\big(P(-r)^k - A^k P\big) u_0.$$

Subtracting the $\alpha = 0$ counterfactual isolates the response induced by the oscillatory coordinate,

$$\widehat{\Delta e}_k^{(\alpha)} := \widehat{e}_k^{(\alpha)} - \widehat{e}_k^{(0)} = \alpha\big(P(-r)^k - A^k P\big) u_0.$$

Since $\rho(A) < r$, we have $A^k P u_0 = o(r^k)$, and therefore

$$\widehat{\Delta e}_k^{(\alpha)} = \alpha P(-r)^k u_0 + o(r^k).$$

A.1.1 Structural non-forcing of gauge motion.

The quotient construction in Section 4.2 is stronger than readout invariance, exact equivariance prevents symmetry-generated motion from forcing dynamics transverse to the gauge orbit.

Proposition A.2 (Gauge-to-quotient non-forcing)**.** *Let T be differentiable and equivariant under the token-local $G = U(1)^S$ action. Define the gauge tangent space $\mathcal{V}_z := \mathrm{Tan}_z(G \cdot z)$, and let Π_z denote the orthogonal projector onto $\mathcal{V}_z^{\perp}$. Then*

$$\Pi_{T(z)} DT(z) \xi = 0, \qquad \xi \in \mathcal{V}_z. \tag{52}$$

At a relative equilibrium $T(z^\star) = g^\star \cdot z^\star$, define $\widehat{T} = (g^\star)^{-1} \cdot T$ and $\widehat{J} = D\widehat{T}(z^\star)$. If U_G has orthonormal columns spanning $\mathcal{V} * z^\star$, then

$$\widehat{J} U_G = U_G, \qquad B_G := \Pi * z^\star \widehat{J} U_G = 0. \tag{53}$$

Proof. Let $\xi \in \mathcal{V}_z$. By definition of the tangent space to the group orbit, there exists a differentiable curve $g(\phi) \in G$ with $g(0) = e$ such that

$$\xi = \frac{d}{d\phi} g(\phi) \cdot z \Big|_{\phi = 0}. \tag{54}$$

By equivariance,

$$T(g(\phi) \cdot z) = g(\phi) \cdot T(z). \tag{55}$$

Differentiating at $\phi = 0$ gives

$$DT(z)\xi = \frac{d}{d\phi} g(\phi) \cdot T(z)\Big|_{\phi=0} \in \mathcal{V}_{T(z)}. \tag{56}$$

Hence

$$DT(z)\mathcal{V}_z \subseteq \mathcal{V}_{T(z)}, \tag{57}$$

and therefore

$$\Pi_{T(z)} DT(z)\xi = 0, \qquad \xi \in \mathcal{V}_z. \tag{58}$$

Now let $z^\star$ satisfy $T(z^\star) = g^\star \cdot z^\star$, and define

$$\widehat{T} = (g^\star)^{-1} \cdot T. \tag{59}$$

Then $\widehat{T}(z^\star) = z^\star$. Since $G = U(1)^S$ is Abelian, $gg^\star = g^\star g$ for all $g, g^\star \in G$, and hence

$$\widehat{T}(g \cdot z) = (g^\star)^{-1} \cdot T(g \cdot z) \tag{60}$$

$$= (g^\star)^{-1} g \cdot T(z) \tag{61}$$

$$= g(g^\star)^{-1} \cdot T(z) = g \cdot \widehat{T}(z). \tag{62}$$

Thus $\widehat{T}$ is also G-equivariant.

Thus, for any $\xi \in \mathcal{V}_{z^\star}$ and any curve $g(\phi)$ generating ξ,

$$\widehat{J}\xi = \frac{d}{d\phi} \widehat{T}(g(\phi) \cdot z^\star)\Big|_{\phi=0} \tag{63}$$

$$= \frac{d}{d\phi} g(\phi) \cdot \widehat{T}(z^\star)\Big|_{\phi=0} \tag{64}$$

$$= \frac{d}{d\phi} g(\phi) \cdot z^\star\Big|_{\phi=0} = \xi. \tag{65}$$

Therefore $\widehat{J}$ acts as the identity on $\mathcal{V}_{z^\star}$. If U_G spans this space,

$$\widehat{J} U_G = U_G. \tag{66}$$

Thus, in the co-moving frame, while the raw state advances along the gauge orbit by $g^\star$, removing this common rotation leaves infinitesimal gauge perturbations unchanged to first order.

In direct analogy with the coupling matrix $B = QJU$ in Section 3.1, define the coupling from the gauge coordinate into the transverse quotient dynamics as

$$B_G := \Pi_{z^\star} \widehat{J} U_G. \tag{67}$$

Since $\Pi_{z^\star} U_G = 0$ and $\widehat{J} U_G = U_G$,

$$B_G = \Pi_{z^\star} \widehat{J} U_G = \Pi_{z^\star} U_G = 0. \tag{68}$$

Finally, let W be an orthonormal basis of $\mathcal{V}_{z^\star}^\perp$ and write $\delta = We + U_G u$. Projecting the co-moving linearization onto the quotient coordinates gives

$$e^+ = W^\top \widehat{J}(We + U_G u) \tag{69}$$

$$= W^\top \widehat{J} W\, e + W^\top \widehat{J} U_G\, u \tag{70}$$

$$= A_Q e, \tag{71}$$

where $A_Q := W^\top \widehat{J} W$. Thus the quotient component evolves independently of the gauge coordinate. □

To connect this result to the coupled-mode analysis in Section 3, let W be an orthonormal basis of $\mathcal{V}_{z^\star}^\perp$ and write $\delta = We + U_G u$. Since $W^\top \widehat{J} U_G = 0$, the quotient component evolves locally as

$$e_{k+1} = A_Q e_k, \qquad A_Q := W^\top \widehat{J} W,$$

independently of the gauge coordinate u_k. Hence a symmetry-generated gauge rotation cannot transmit its timescale into the quotient dynamics. This statement applies only to the gauge subspace so other oscillatory modes transverse to the gauge orbit remain unconstrained and may continue to participate in the quotient dynamics.

A.2 Local theory and empirical prediction

Derivation of the local block dynamics.

Let $U \in \mathbb{R}^{D\times m}$ have orthonormal columns and satisfy

$$J^\top U = US^\top. \tag{72}$$

Transposing gives

$$U^\top J = SU^\top. \tag{73}$$

Define

$$Q = I - UU^\top. \tag{74}$$

Since $U^\top U = I$, Q is the orthogonal projector onto the complement of $\mathrm{span}(U)$, and

$$QU = 0, \qquad U^\top Q = 0. \tag{75}$$

For any root-centered displacement δ, define

$$u = U^\top \delta, \qquad e = Q\delta. \tag{76}$$

Then $\delta = e + Uu$, with $U^\top e = 0$. Under the linearized dynamics

$$\delta^+ = J\delta, \tag{77}$$

the oscillatory coordinate evolves as

$$\begin{aligned} u^+ &= U^\top \delta^+ && (78)\\ &= U^\top J(e + Uu) && (79)\\ &= SU^\top e + SU^\top Uu && (80)\\ &= Su. && (81) \end{aligned}$$

Thus the complementary coordinate does not feed back into u to first order.

Similarly,

$$\begin{aligned} e^+ &= Q\delta^+ && (82)\\ &= QJ(e + Uu) && (83)\\ &= QJQ\, e + QJU\, u, && (84) \end{aligned}$$

where we used $e = Qe$. Defining

$$A := QJQ, \qquad B := QJU, \tag{85}$$

gives

$$e^+ = Ae + Bu, \qquad u^+ = Su. \tag{86}$$

Hence the left-invariant oscillatory coordinate evolves autonomously, while it can still drive the complementary state through the coupling block B.

A.3 QRR Architecture

QRR follows the recurrent-block organization of FPRM, but replaces the recurrent trunk with the gauge-equivariant construction described in Section 4.3 and augments recurrence with momentum.

Input and recurrent state. Let x denote the embedded task input, which is kept fixed and re-injected at every recurrent step. The recurrent state consists of a latent state y_t and a momentum state m_t, both transforming covariantly under the token-local $G = U(1)^S$ action. In real coordinates, $y_t, m_t \in \mathbb{R}^{S\times 2d}$, with the fixed identification $\mathbb{R}^{2d} \simeq \mathbb{C}^d$ used throughout the gauge-equivariant trunk. For each input, we initialize $m_0 = 0$ and draw y_0 independently at random. Specifically, each four-dimensional real cell of y_0 is sampled coordinate-wise from a truncated normal distribution and normalized to unit norm.

Gauge-equivariant recurrent trunk. At recurrent step t, the trunk receives the current state, carried momentum, and input,

$$z_t = \Phi_\theta(y_t, m_t; x), \tag{87}$$

where Φ_θ is G-equivariant,

$$\Phi_\theta(g\cdot y, g\cdot m; x) = g\cdot \Phi_\theta(y, m; x). \tag{88}$$

The trunk retains the pre-norm Transformer structure of FPRM, while token mixing, input injection, attention, and feed-forward transformations are replaced by their gauge-equivariant counterparts. The input x is injected through a covariant carrier, while the carried momentum m_t is also fed back into the trunk. The resulting output z_t is a *state proposal*, rather than the next recurrent state directly.

Momentum-carrying update. We define the proposal residual

$$F_t = z_t - y_t. \tag{89}$$

The deployed recurrence applies momentum directly to this residual,

$$p_t = \beta m_t + (1-\beta)F_t, \qquad \beta = \sigma(b_{\mathrm{H}}), \tag{90}$$

where b_{H} is a learned global scalar. Ignoring an equivalent gauge rephasing described below, the update is

$$\widetilde{y}_{t+1} = y_t + p_t, \qquad \widetilde{m}_{t+1} = p_t. \tag{91}$$

Gauge rephasing. The implementation additionally permits a state-dependent gauge rotation

$$a_t = \frac{\langle F_t, \mathcal{J}y_t\rangle}{\|y_t\|^2 + \epsilon}, \qquad R_t = \exp(a_t\mathcal{J}), \tag{92}$$

with a_t detached from the computational graph. The stored recurrent state is then

$$y_{t+1} = R_t\widetilde{y}_{t+1}, \qquad m_{t+1} = R_t\widetilde{m}_{t+1}. \tag{93}$$

Since R_t is itself a gauge transformation, this operation doesn't change quotient state.

Readout and adaptive halting. Predictions are produced by a gauge-invariant readout

$$\widehat{y}_t = \mathcal{R}(y_t; x), \qquad \mathcal{R}(g\cdot y_t; x) = \mathcal{R}(y_t; x). \tag{94}$$

Consequently, prediction is insensitive to motion confined to the gauge orbit. Adaptive halting is based on the quotient residual r_G from Equation 19, computed from the joint state (y_t, m_t) after phase alignment. A trajectory halts once this residual remains below the task-specific threshold for the prescribed persistence window; otherwise recurrence continues until the rollout cap.

Recurrent training. The pair (y_t, m_t) is carried across recurrent steps and across successive deep-supervision segments for the same input, and is reset only when a new example is started. Within each segment, the model is unrolled for a task-specific number of recurrent steps, while gradients are propagated only through the final truncated-BPTT window. Earlier steps advance the recurrent state without storing gradients. After each deep-supervision segment, the prediction loss is applied and the carried state is detached before the next segment.

A.4 Additional Experimental Results

A.4.1 Causal verification on HRM spiral dynamics

We additionally verify the local oscillatory-coupling prediction on HRM, whose solution-root dynamics exhibit a qualitatively different oscillatory geometry from the alternating mode analyzed for FPRM in Section 3.1. Here, the identified slow spectral component is a two-dimensional real invariant block

$$S \in \mathbb{R}^{2\times 2}, \qquad \sigma(S) = \{re^{\pm i\omega}\}, \qquad 0 < \omega < \pi, \tag{95}$$

with $\rho(A) < r \le 1$. Thus, rather than alternating along a single direction, the slow coordinate rotates while decaying toward the solution root. We use exactly the same intervention protocol used for FPRM in main paper. The local linearization from Eq. 13 predicts

$$\widehat{\Delta e}_k^{(\alpha)} = \alpha Q J^k U u_0 = \alpha\left(PS^k - A^kP\right)u_0, \tag{96}$$

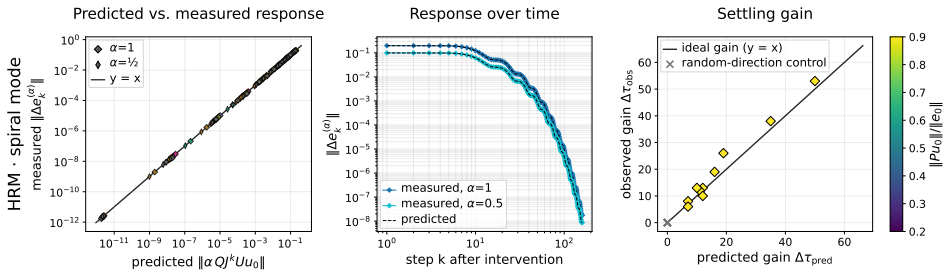


Figure 6: **The local oscillatory-coupling prediction extends to the two-dimensional spiral dynamics of HRM.** We modify only the Jacobian-identified oscillatory coordinate once with $\alpha \in \{0, 0.5, 1\}$ and subsequently resume the original nonlinear HRM recurrence.

where the full matrix power S^k retains the rotation of the two-dimensional spiral mode.

Figure 6 mirrors the three tests reported for FPRM in Figure 2. First, the response measured from the nonlinear rollout agrees with the Jacobian-based prediction

$$\left\| Q\left(z_{t_0+k}^{(\alpha)} - z_{t_0+k}^{(0)} \right) \right\|_2 \qquad \text{vs.} \qquad \left\| \alpha Q J^k U u_0 \right\|_2 . \tag{97}$$

Second, the temporal response follows the predicted rotating and decaying evolution, rather than the one-dimensional alternation observed in FPRM. Finally, removing the oscillatory component with $\alpha = 0$ reduces root-relative settling time by an amount consistent with the local prediction, using the same settling criterion and sustained-window implementation as in the main-text analysis. These results show that the intervention agreement in Figure 2 is not specific to the nearly one-dimensional flip geometry of FPRM. The same local decomposition predicts the nonlinear response and convergence cost for a qualitatively distinct two-dimensional spiral mode.

A.4.2 Spectral Analysis on HRM vs. QRR

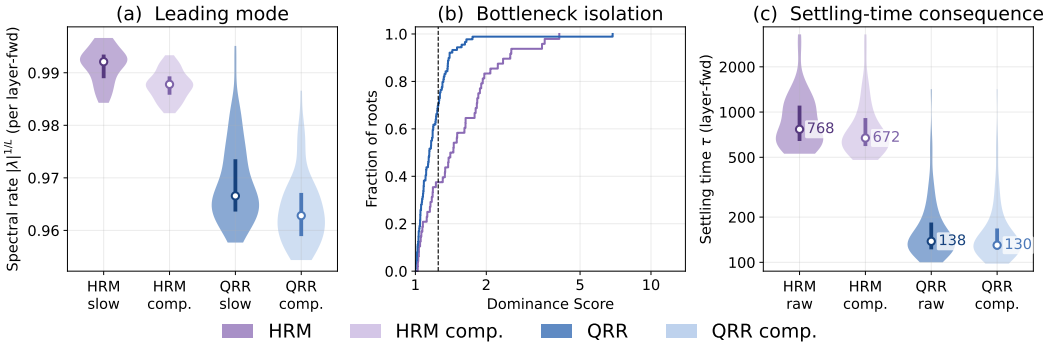


Figure 7: **Spectral organization and settling in HRM and QRR.** (a) Leading and complementary spectral rates near certified solution equilibria. QRR rates are measured in quotient space. (b) Distribution of the dominance score $D = \log r_{\text{comp}} / \log r_{\text{lead}}$. (c) Root-relative settling time before and after factoring out the leading-mode response.

We repeat the spectral analysis of Section 4.4 for HRM, whose local dynamics exhibit the spiral geometry studied in Appendix A.4.1. For HRM, we linearize the recurrent H-cycle map around certified solution fixed points. For QRR, we analyze the corresponding linearization in quotient space around relative equilibria. As in the main analysis, we decompose each system into its leading spectral mode and complementary dynamics and normalize spectral rates to recurrent layer-forward.

Figure 7 shows the same distinction at a different absolute timescale. Although the leading–complementary spectral gaps appear similar in magnitude, HRM operates much closer to unit spectral rate, producing a larger separation in decay timescales and a stronger dominance score than QRR. This is reflected in finite-time settlement. Factoring out the leading-mode response reduces

the median HRM settling time from 768 to 672 layer-forwards, whereas the corresponding change in the QRR quotient is only 138 to 130. Thus, QRR remains in a substantially faster regime and is less dependent on its leading mode for settlement, even relative to the qualitatively different spiral dynamics of HRM.

A.4.3 Spectral Properties of Leading Oscillatory Modes

For each certified solution root, we linearize the one-step recurrent map and identify the leading spectral mode of the local Jacobian J. We retain roots for which this mode is oscillatory, corresponding to either a complex-conjugate pair or a real negative flip mode. Let U span this leading block, with modulus r, and define its complement by $Q = I - UU^\top$ and $A = QJQ$. Because the selected block contains the leading eigenmode of J, $r \geq \rho(A)$ holds by construction. We therefore use $r - \rho(A)$ to measure the spectral separation between the leading oscillatory mode and the remaining dynamics.

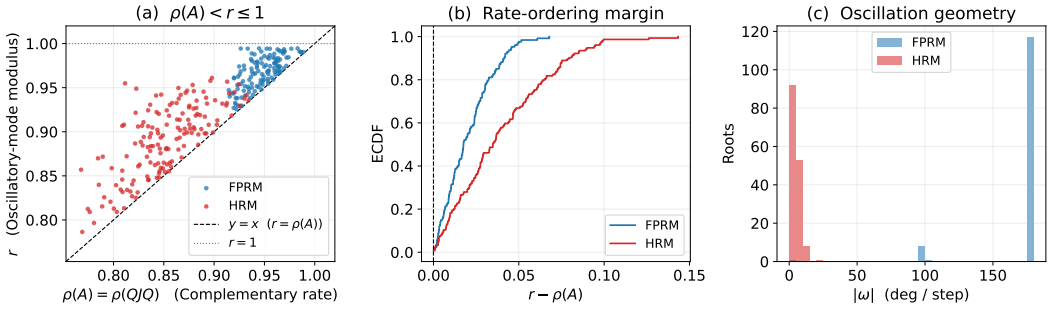


Figure 8: **Spectral separation of leading oscillatory modes.** (a) Leading-mode modulus r and complementary spectral radius $\rho(A)$. (b) Spectral margin $r - \rho(A)$. (c) Oscillation angle $|\omega|$ for FPRM and HRM.

Figure 8 shows a strict spectral gap for all selected roots. The median margin $r - \rho(A)$ is 0.0176 for FPRM and 0.035 for HRM, with minimum margins of 2.2×10^{-4} and 1.5×10^{-4}. The leading modes also exhibit distinct geometries. FPRM is concentrated near $|\omega| = \pi$ with median $179.7°$, whereas HRM exhibits small-angle spirals with median $3.9°$. Thus, the selected oscillatory modes are cleanly separated from the complementary spectrum across both flip and spiral regimes.

A.4.4 FPRM flip mode suppression

The local analysis in Section 3 shows that the slowly decaying flip mode can limit late-stage convergence in FPRM. This does not imply, however, that the same mode is dispensable throughout reasoning. We therefore test whether persistently suppressing the alternating component before the output has stabilized interferes with the computation itself.

For each trajectory, we estimate a one-dimensional alternating direction U from the preceding recurrent updates and, from an intervention time t_0 onward, modify every subsequent update d_t as

$$z_{t+1}^{(\alpha)} = z_t + \left[I - (1-\alpha)UU^\top\right] d_t. \tag{98}$$

We use $\alpha = \frac{1}{2}$ as the suppression intervention and compare against the unmodified rollout. Early interventions start before the native output-stabilization (decoded output stabilizes and remains equal for the remainder of the rollout), whereas late interventions start after stabilization. To distinguish suppression of the oscillatory mode from a generic reduction in update magnitude, we additionally suppress a non-oscillatory control direction while matching the removed energy. All comparisons are paired against the corresponding original rollout, and confidence intervals are obtained by bootstrapping over boards.

Early suppression decreases final exact accuracy by 5.9 percentage points, and delays output stabilization. In contrast, removing a comparable amount of update energy along a non-oscillatory direction produces no statistically distinguishable degradation, and suppressing the flip mode only after the output has stabilized has no measurable effect. These results indicate that the oscillatory motion of the same alternating form can still participate in the earlier reasoning computation. This motivates QRR to avoid suppressing oscillations based on their dynamical appearance alone, and instead quotient out only motion whose representational redundancy is guaranteed by symmetry.

Table 4: **Persistent suppression of the FPRM flip mode.** Suppressing the alternating component before output stabilization reduces final exact accuracy, whereas an energy-matched non-oscillatory intervention does not show a significant effect.

Intervention	Δ final accuracy	95% CI
Oscillatory suppression, early	-5.9 pp	$[-13.1, -0.6]$
Energy-matched control, early	-1.1 pp	$[-2.8, 0.0]$
Oscillatory suppression, late	0.0 pp	$[0.0, 0.0]$

A.4.5 Halting time sensitivity to task difficulty

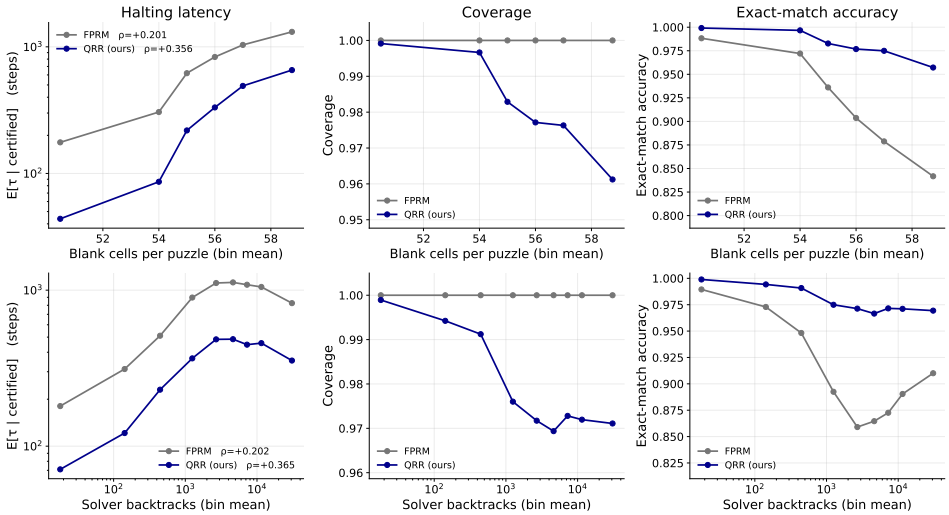


Figure 9: **Adaptive halting as a function of Sudoku difficulty.** Difficulty as the number of blank cells (top) and the backtrack count (bottom). Left: mean halting time among certified examples. Middle: fraction of examples certified within the rollout budget. Right: exact-match accuracy.

On 42,023 held-out Sudoku-Extreme boards, we evaluate QRR and FPRM under the same halting rules and 35,000-step cap as in Table 2. As shown in Figure 9, halting latency increases with difficulty for both models, but more strongly for QRR. The Spearman correlation is $\rho = 0.356$ versus 0.201 for blank count and 0.365 versus 0.202 for solver backtracks. QRR nevertheless certifies significantly earlier across the difficulty range while maintaining 98.0% overall coverage and consistently higher exact-match accuracy, indicating stronger difficulty-adaptive compute allocation with lower cost.

A.5 Experimental Setup

A.5.1 Reasoning Tasks

We evaluate on four reasoning benchmarks. **Sudoku-Extreme** consists of partially filled 9×9 Sudoku boards with unique solutions and is evaluated by exact-match accuracy of the completed grid (Wang et al., 2025). **Maze-Hard** contains 30×30 shortest-path problems and is evaluated by exact match to the supervised solution path (Wang et al., 2025). Because multiple optimal paths may exist in Maze-Hard, we additionally use **Maze-Unique**, constructed from perfect tree-structured mazes so that each start–goal pair has a unique solution (Kim et al., 2022; Huang et al., 2026). **Mini-ARC** is a compact 5×5 abstract rule-induction benchmark in which the model infers a transformation from demonstration pairs and applies it to a held-out input (Kim et al., 2022; Huang et al., 2026). We follow the standard augmentation-based evaluation, map predictions from each augmented view back to the original frame, aggregate identical candidate grids by voting, and report task-level Pass@1 from the highest-ranked prediction.

A.5.2 ROOT CONVERGENCE ANALYSIS

Here we provide detailed protocol for root convergence analysis used in this paper, which is used for the local Jacobian and spectral analyses in Sections 3, 4.4, and Table 2 (a). We firstly identify solution equilibria using Anderson acceleration with history $m = 6$, Tikhonov regularization 10^{-8}, and $\beta = 1$ in fp64. The solver runs for a fixed budget and returns the iterate with the smallest residual. For FPRM, HRM, and EqR, the solved maps are one FPRM recurrent step, two HRM H-cycles, and three EqR cycles, respectively. For QRR, Anderson acceleration is applied to the gauge-aligned map.

For each certified solution root, we run the original model from its native initialization and measure when the trajectory reaches the corresponding root. We define τ_{root} as the first step after which the root-relative distance remains below $\delta = 10^{-3}$ for a persistence window of approximately 96 recurrent layer-forwards. This corresponds to 48 native steps for FPRM and QRR and 4 trajectory steps for HRM. For EqR, whose native step contains 42 layer-forwards, we use two steps, corresponding to 84 layer-forwards. *Correct root found* denotes the fraction of evaluation examples for which the solver returns a certified solution equilibrium decoding to the ground truth.

To compare convergence across architectures with different computation per native recurrent step, we report convergence times in recurrent layer-forward equivalents (LFWD). An LFWD counts one forward evaluation of one recurrent network layer. If one native step of model m evaluates ℓ_m recurrent layers, then

$$\tau^{(m)}_{\mathrm{lfwd}} = \ell_m \, \tau^{(m)}_{\mathrm{native}},$$

where $\tau^{(m)}_{\mathrm{native}}$ is the number of native recurrent steps required to reach the certified solution equilibrium. Table 5 summarizes the conversion used for each model.

Table 5: Conversion from native recurrent steps to recurrent layer-forward equivalents. An LFWD counts one forward evaluation of one recurrent network layer.

Model	Layer-forward computation per native step	ℓ_m (LFWD/step)
QRR	$1 \times L_{\mathrm{layers}} = 1 \times 2$	2
FPRM	$1 \times L_{\mathrm{layers}} = 1 \times 2$	2
HRM	$H_{\mathrm{cyc}}\big(L_{\mathrm{cyc}} L_{\mathrm{layers}} + H_{\mathrm{layers}}\big) = 2(2 \cdot 4 + 4)$	24
EqR	$H_{\mathrm{cyc}}(L_{\mathrm{cyc}} + 1)L_{\mathrm{layers}} = 3(6 + 1)2$	42

A.5.3 HYPERPARAMETERS

We use the same core QRR architecture across all tasks and tune only the task-specific training, recurrent-compute, and inference settings summarized in Table 6. Architectural details shared across tasks are given in Section A.3.

Table 6: **Task-specific training and inference configurations for QRR.** The core architecture is shared across tasks. The reported result is from the best in-run evaluation checkpoint.

Setting	Sudoku-Extreme	Maze-Hard	Maze-Unique	Mini-ARC
Data and optimization				
Training data	1K $\times$ 1000 aug.	1K, no aug.	1K, no aug.	149 tasks, 100 aug.
Training epochs	60K	60K	6K	100K
Global batch size	768	768	384	768
Learning rate	3×10^{-4}	1×10^{-4}	2×10^{-4}	6×10^{-4}
Weight decay	10^{-3}	1.0	1.0	0.1
Gradient clipping	1.0	1.0	1.0	1.0
Compute setup				
Training GPUs	4$\times$ H100	4$\times$ B200	4$\times$ H100	1$\times$ A100
Parallelism	DDP $\times4$	DDP $\times4$	DDP $\times4$	single GPU
Per-GPU batch size	192	192	96	768
Recurrence and deep supervision				
Max DS segment	12	24	24	8
Recurrent steps / DS segment	6	6	6	6
TBPTT depth	6	6	6	6
Inference and halting				
Halting criterion	$r_G < 0.03$	$r_G < 0.075$	$r_G < 0.15$	$r_G < 0.08$
Persistence window W	1	1	1	2
Rollout cap	35,000	3,000	3,000	100
Evaluation	halt-or-cap	halt-or-cap	halt-or-cap	halt-or-cap + vote