

Reviewed Preprint

v1 • September 7, 2026

Not revised

A theory of multi-task computation and task selection

For correspondence:

om2382@columbia.edu

[†] Present address: Kempner Institute for the Study of Natural and Artificial Intelligence, Harvard University, Cambridge, MA 02138, USA

[‡] a.litwin-kumar@columbia.edu 

(ALK)

Competing interests: No

competing interests declared

Funding: See [page 47](#)

Reviewing editor: Tatyana O

Sharpee, Salk Institute for Biological Studies, United States

© 2026, Marshall et al. This article is distributed under the terms of the

[Creative Commons Attribution](#)

[License](#), which permits unrestricted use and redistribution provided that the original author and source are credited.

Owen Marshall^{1,2} , David G Clark^{†,1,2}, Ashok Litwin-Kumar^{‡,1,2}

¹Kavli Institute for Brain Science, Columbia University, New York, United States • ²Zuckerman Institute, Columbia University, New York, United States

eLife Assessment

This **valuable** study analyzes how recurrent neural network can flexibly switch between multiple tasks. The evidence is **solid**, but reviewers have raised questions about the impacts of transient dynamics, whether the activity is actually chaotic, can inputs be used to switch between tasks, what determines overlap between tasks and others. In addition, there were more minor questions pertaining to whether the analyses are actually analytic or primarily based on numerical simulations, as well as the relationship to prior studies where the same network was optimized to perform multiple tasks simultaneously with different readouts.

<https://doi.org/10.7554/eLife.111734.1.sa2>

Abstract

Neural activity during the performance of a stereotyped behavioral task is often described as low-dimensional, occupying only a limited region in the space of all firing-rate patterns. This region has been referred to as the “neural manifold” associated with a task. More recently, recordings of neural activity in animals challenged to perform multiple tasks have suggested that each task is associated with a different low-dimensional manifold. What connectivity structures underlie this flexibility in neural dynamics, and how is interference between the dynamics associated with different tasks avoided? We develop a theoretical model for multi-task computation in nonlinear recurrent neural networks whose connectivity is constructed as a weighted sum of many low-rank components, each encoding the dynamics associated with a different task. The model demonstrates that interference between different tasks’ dynamics limits flexible multi-tasking and can lead to chaotic fluctuations. However, small modulations of a network’s effective connectivity overcome this interference. We derive the conditions that enable such task selection and characterize both single-neuron and population statistics in task-selected and unselected states. The model reveals the requirements for a single network to produce distinct dynamics confined to distinct neural manifolds and suggests circuit mechanisms that support this capability. Using the model, we propose different hypotheses for explaining the origin of high-dimensional neural activity in large-scale recordings.

Introduction

Adaptive behavior requires organisms to perform different neural computations at different times or in different contexts. For example, both spontaneous and skilled, practiced movements involve sequencing discrete motor computations [1, 2]. Similarly, cognitive tasks require operations on task-relevant variables that may differ across task, condition, or epoch [3–7]. Numerous studies support the proposal that individual neural computations are associated with coordinated dynamics of interconnected populations of neurons [8], implying that the effective dynamical system governing these neurons’ activity must rapidly change when a new computation is to be performed. Little is known about the circuit mechanisms that underlie these dynamical switches.

During execution of a single task, population activity is typically confined to a low-dimensional neural manifold [9, 10]. Mixed selectivity—the observation that individual neurons often participate in multiple distinct tasks [3, 5, 11, 12], although not always [13]—implies that these manifolds are distributed across many neurons rather than aligned to a small subset. This distributed structure is reproduced in recurrent neural networks trained to perform multiple tasks [4, 14]. Recent work shows that the location and orientation of the manifold can change depending on the task being performed [2, 15–18]. Thus, switching between manifolds, each distributed across the population, may underlie the rapid changes in neural dynamics and behavior required for switching tasks [2].

In another line of work, neural activity has been observed to be high-dimensional during periods of spontaneous, uninstructed behavior [19, 20], in contrast to observations of low-dimensional activity during the performance of stereotyped tasks [21, 22]. These observations may be compatible with the network switching among many distinct manifolds, each associated with a different self-initiated behavior [1, 23]. The union of many such individually low-dimensional manifolds can be high-dimensional if they are heterogeneously oriented in neural state space, as recent experimental work suggests [2]. Alternatively, such activity may reflect unstructured, high-dimensional fluctuations, as proposed by random-network models of the spontaneous state [24, 25].

Altogether, these observations raise several theoretical questions. What forms of recurrent connectivity support many distinct task-related manifolds, with mixed selectivity to different tasks at the single-neuron level? Under what conditions do the dynamics associated with each manifold interfere or not interfere? How does a circuit selectively engage the dynamics appropriate for one task while suppressing others? What conditions lead to high-dimensional neural activity?

Prior work has addressed multi-tasking by training recurrent neural networks on multiple tasks using machine learning methods and then analyzing the solutions [3, 4, 14, 26, 27]. Such studies have shown that artificial networks reuse dynamical motifs, such as fixed points and ring attractors, when the tasks have shared component subtasks. We take a complementary approach, using a solvable network model rather than trained networks. For simplicity, we assume each “task” is an irreducible computation that can be solved by a single low-dimensional dynamical system, without compositional overlap [6].

Specifically, we propose and analyze a mathematically tractable model of multi-tasking neural circuits, in which distinct task-related dynamics operate on distinct low-dimensional activity manifolds. Building on low-rank network models that explain how recurrent structure shapes low-dimensional dynamics [28, 29], we model networks whose weight matrix is constructed as a sum of many low-rank single-task matrices [30]. Because the network dynamics are nonlinear, this leads to complex interactions between the dynamics associated with different tasks. We develop a dynamical mean-field theory that characterizes how tasks interfere, how spontaneous fluctuations arise when a system has learned to perform many distinct tasks, and how these dynamics manifest at both the single-neuron and population levels. Our model makes predictions about single-neuron tuning and fluctuations, trial-to-trial variability in neural activity projected onto task-related dimensions, and the dimensionality of population activity that can be examined in experimental data.

A key feature of the model is that small modulations of connectivity can lead to the dynamic selection of a single task and the suppression of others. This provides a biologically plausible mechanism for task selection. In particular, our proposal is consistent with evidence that external inputs, for example thalamic inputs, gate and sustain motor cortical activity [31, 32]. More generally, task execution may result from coordination between recurrent dynamics and modulatory pathways. In total, we provide a theory relating connectivity, dynamics, and experimentally measurable population signatures of multi-tasking in recurrent neural circuits.

Results

We study nonlinear recurrent neural networks of N continuous-time rate neurons with currents $x_i(t)$ and firing rates $\phi_i(t)$ that evolve according to

$$\dot{x}_i(t) = -x_i(t) + \sum_{j=1}^N J_{ij} \phi_j(t), \quad (1)$$

with a sigmoid-shaped firing-rate function (taken to be $\phi_i(t) = \text{erf}(\sqrt{\pi}x_i(t)/2)$ in simulations). The network dynamics are controlled by the $N \times N$ weight matrix $\mathbf{J}$ with elements J_{ij} . We assume that N is large, and our theoretical results are exact in the $N \rightarrow \infty$ limit. We now discuss the dynamics that occur when $\mathbf{J}$ is low-rank, before presenting our model in which many such low-rank matrices are superposed.

Low-rank networks

Many previous works have considered the case where the recurrent connectivity $\mathbf{J}$ is low-rank [28, 29, 33, 34], given by

$$J_{ij} = D \sum_{r=1}^R m_i^{(r)} n_j^{(r)}, \quad (2)$$

a sum over a small number $R \ll N$ of outer products of N -dimensional *loading vectors* $\mathbf{m}^{(r)}$, $\mathbf{n}^{(r)}$, scaled by a constant $D \geq 0$. The right loading vectors $\mathbf{n}^{(r)}$ define the subspace of neural firing rates that determines the recurrent input to the network, and this input lies in a subspace spanned by the left loading vectors $\mathbf{m}^{(r)}$ (Fig. 1A, left). When transformed by the firing-rate nonlinearity, this R -dimensional input subspace produces neural activity that lies on an R -dimensional nonlinear manifold (Fig. 1A, right).

We consider the special case of Gaussian-distributed loading-vector components with zero mean and specified covariance structure. All components are generated i.i.d. across different neurons. For a given neuron i , the left- and right-loading-vector components are sampled with nonzero covariances $\langle n_i^{(r)} m_i^{(r')} \rangle = A_{rr'}/N$ according to the entries of the $R \times R$ *overlap matrix* $\mathbf{A}$. As such, this matrix encodes the loading vectors' geometrical arrangement $\langle \mathbf{n}^{(r)} \cdot \mathbf{m}^{(r')} \rangle = A_{rr'}$ and plays a major role in determining the network's dynamics.

Activity in a low-rank network can be described by a set of *latent variables*, defined by projections onto the right loading vectors, $z_r(t) = D(\mathbf{n}^{(r)} \cdot \Phi(t))$. In the $N \rightarrow \infty$ limit, the overlap matrix determines the dynamics among the latent variables according to

$$z_r(t) = D \langle \phi'(t) \rangle \sum_{r'=1}^R A_{rr'} [L * z_{r'}](t), \quad (3)$$

where $L(\tau)$ is a low-pass filter with the single-neuron time constant (see Methods). The gain factor $\langle \phi'(t) \rangle$ is the population-averaged slope of the firing-rate nonlinearity. It depends on the norm of the low-pass-filtered latent variables, with larger norm reducing the gain due to the sigmoidal form of $\Phi(\cdot)$ (see Methods). Eq. (3) summarizes a low-rank network's activity as an R -dimensional, nonlinear dynamical system [28, 29]. The gain factor $\langle \phi'(t) \rangle$ is the only nonlinearity in this system and thus is necessary for nonlinear dynamical behaviors such as stable limit cycles or multistability.




Figure 1. Single-task and multi-task network dynamics.

(A) Left: Illustration of a state-space trajectory in the space of neuronal pre-activations (purple), confined to a task-relevant two-dimensional subspace. Such a geometry arises in rank-two networks with loading vectors $\mathbf{m}^{(1)}$, $\mathbf{m}^{(2)}$. Right: Illustration of the same trajectory in firing-rate space. Activity lies on a two-dimensional nonlinear manifold. **(B)** Schematic of a “multi-task” network that flexibly implements multiple low-dimensional dynamics. Each task is associated with a different nonlinear manifold, which collectively span more activity dimensions than any one manifold. **(C)** Schematic of multi-task network connectivity as a sum of low-rank matrices, each associated with a different task. **(D)** Each task is associated with an overlap matrix comprising pairwise inner products among the left and right loading vectors.

Multi-task low-rank networks

Neural recordings suggest that the activity of individual circuits can switch among distinct neural manifolds that collectively span a large number of dimensions [Fig. 1B; 2, 16]. When measured across many different contexts, such activity is high-dimensional, as are measurements of dimension during spontaneous behavior [19, 20]. Such phenomena are not well modeled by low-rank networks that produce activity residing on a single low-dimensional manifold. We therefore developed a model of multi-task networks in which distinct dynamics occur on distinct neural manifolds.

In the context of our model, “task” refers to the low-dimensional dynamics associated with a particular manifold in firing-rate space, which we interpret as implementing a specific computation. Indeed, prior work in the single-task setting has shown that low-rank networks can implement experimentally inspired input–output mappings through their low-dimensional dynamics [28, 34]. Although real tasks involve both sensory inputs and behavioral outputs, we focus only on the recurrent dynamics, noting that our model could be easily extended to include task-specific inputs as in prior single-task works (see Discussion).

The network performs P different tasks, indexed by $\mu = 1, \dots, P$. Each task is associated with a rank- R matrix of the form of Eq. (2), and the multi-task weight matrix is constructed as a weighted sum of these matrices:

$$J_{ij} = \sum_{\mu=1}^P D^{(\mu)} \sum_{r=1}^R m_i^{(\mu,r)} n_j^{(\mu,r)}. \quad (4)$$

Each task μ is therefore parameterized by R left and right loading vectors $\mathbf{m}^{(\mu,r)}$ and $\mathbf{n}^{(\mu,r)}$, and a weight $D^{(\mu)} \geq 0$ controlling the strength of the connectivity component associated with that task. The loading vectors are Gaussian and i.i.d. across neurons, as in the single-task setting. They are furthermore independent across tasks, i.e. there is zero expected overlap between loading vectors of distinct tasks. The covariance for task μ given by

$$\langle n_i^{(\mu,r)} m_i^{(\mu,r')} \rangle = A_{rr'}^{(\mu)} / N, \quad (5)$$

with the quantities $A_{rr'}^{(\mu)}$ comprising the task-specific $R \times R$ overlap matrix $\mathbf{A}^{(\mu)}$ (Fig. 1D). We refer to the subspace spanned by $\mathbf{m}^{(\mu,r)}$ for $r = 1, \dots, R$ as *task subspace* μ , the R -dimensional space that contains the input to the network associated with task μ .

As described previously for single-task networks, the structure of overlaps between left and right loading vectors plays a major role in shaping network dynamics. The connectivity of Eq. (4) can be interpreted as that of a rank- P R network with a particular structure of overlaps among left and right loading vectors. Specifically, the overlaps between $\mathbf{m}^{(\mu,r)}$ and $\mathbf{n}^{(\nu,r')}$ are constrained to be, on average, block-diagonal, with the block entries given by $A_{rr'}^{(\mu)}$ (Fig. S1). Each of these blocks corresponds to R -dimensional dynamics associated with a single task subspace. Our model therefore describes synaptic weight matrices with strong, structured interactions among the components associated with a single task, specified by these blocks, and small, random overlap between components associated with different tasks. In terms of activity, the subspaces corresponding to different tasks are nearly, although not exactly, orthogonal when N is large, with finite-size fluctuations in cross-task overlaps that are reduced as N grows. This blockwise constraint reduces the number of parameters needed to specify the model from $P^2 R^2$ to PR^2 overlap-matrix elements—a dramatic reduction when P is large—while allowing the possibility of distinct tasks occupying different manifolds defined by the connectivity.

Multi-task networks exhibit winner-take-all competition across tasks

We begin by comparing single-task networks ($P = 1$) to a minimal multi-task network ($P = 2$). Network weights are set according to Eq. (4), fixing the within-task dimension $R = 2$. The $P = 1$ case reduces to the low-rank networks described by Eqs. (2) and (3). Such low-rank networks can produce a variety of different dynamics. For example, two different choices of a 2×2 overlap matrix $\mathbf{A}$ produce networks that exhibit either a stable limit cycle (Fig. 2A) or bistability (Fig. 2B). We next consider a multi-task network with $P = 2$ using a weight matrix that superposes the two matrices used above that, on their own, produce limit-cycle and fixed-point dynamics. The dynamics can be analyzed via two sets (one for each task) of latent variables $z_r^{(\mu)}(t) = D^{(\mu)}(\mathbf{n}^{(\mu,r)} \cdot \phi(t))$ where $\mu = 1, 2$ indexes tasks and $r = 1, 2$ indexes within-task dimensions. Both sets of latent variables initially follow trajectories similar to those of the corresponding single-task networks (Fig. 2C). However, closer examination reveals that the dynamics are coupled across tasks and exhibit different asymptotic behaviors. While the dynamics in the task-1 (limit cycle) subspace are barely changed (Fig. 2A vs. Fig. 2C, left), the dynamics in the task-2 (bistable attractor) subspace lose their bistability. Trajectories are initially attracted towards one of the two previously stable fixed points but subsequently decay to the origin, once task-1 activity has grown sufficiently large in magnitude (Fig. 2B vs. Fig. 2C, right). The network therefore exhibits winner-take-all behavior, where one task dominates and suppresses others. In this example, task 1 dominates because $D^{(1)} > D^{(2)}$. Reversing this inequality causes task 2 to dominate (Fig. S2).

In a linear network, activity in distinct task subspaces evolves independently when the loading vectors for different tasks are orthogonal. Interference between tasks, defined as the influence of one task's latent variables on another's, occurs only when this orthogonality is broken. In multi-task networks with a fixed number of tasks, such interference becomes negligible as $N \rightarrow \infty$, because randomly oriented activity patterns in a high-dimensional activity space are nearly orthogonal with high probability. What, then, is the source of the winner-take-all interference observed in Fig. 2C? In nonlinear networks with finite P , Eq. (3) still describes each task's dynamics, but the gain factor now depends on the dynamics associated with all tasks [29]. Increased activity in any task subspace reduces the gain factor and suppresses the dynamics in other subspaces (see Methods). Consequently, at the steady state, the strongest task's latent variables dominate and evolve as they would in a single-task network, with all others suppressed. This form of interference persists even in the limit $N \rightarrow \infty$ and is unique to nonlinear networks that exhibit mixed selectivity and thus share neurons across tasks. In the example of Fig. 2C, the growth of task 1's oscillatory latent variables suppresses dynamics in the task-2 subspace and causes trajectories to approach the origin, eliminating bistability.

These properties are not specific to networks with our particular choice of nonlinearity, as we observe similar winner-take-all behavior for other saturating nonlinearities, as well as in networks with non-saturating nonlinearities that are stabilized through recurrent inhibition [35] (Fig. S3). However, the fact that at most one task can dominate, rather than multiple tasks dominating simultaneously, is a consequence of Gaussian-distributed loadings, a simplification we use for mathematical tractability (see Discussion).

Chaotic fluctuations in networks with large P

Higher cortical areas participate in many distinct tasks, so our primary focus is on networks with large P . We therefore simulate networks with increasing P . We configure the overlap matrix of task 1 to produce a stable limit cycle as before, while the overlap matrices of other tasks are randomly generated. We also set $D^{(1)}$ to be large enough that task 1 would dominate any other task in isolation.

As P increases, the average magnitude of the projection of neural activity onto the task-1 subspace first decreases gradually, then drops sharply (Fig. 2D). This drop corresponds to a transition from a state in which task 1 is dominant and exhibits a stable limit cycle to a state in which no task




Figure 2. Dynamics of networks with varying P .

See Methods for parameter values. **(A)** Latent-variable flow fields (black) and a latent-variable trajectory (purple) in a single-task network with **A** chosen to produce limit-cycle dynamics. Dashed line: limit cycle. **(B)** Latent-variable flow fields (black) and latent variable trajectories for two separate runs with different initial conditions (red) in a single-task network with **A** chosen to produce two stable fixed points. Black X's: final time points for all trajectories used to generate flow fields. **(C)** A two-task network with $D^{(1)} > D^{(2)}$. Left: flow fields (black) of projections $\bar{z}^{(1)}(t)$ onto task-1 subspace and a latent-variable trajectory (purple). Right: flow fields (black) of projections $\bar{z}^{(2)}(t)$ onto task-2 subspace simulated over many initial conditions, as well as two example trajectories (red). Black X's: final time points for all trajectories. **(D)** Time-averaged latent variable norm $\langle \|\bar{z}^{(1)}(t)\| \rangle_t$ for task $\mu = 1$ as a function of the number of total tasks $\alpha = P/N$ normalized by N . Same overlap parameters as (C) for $\mu = 1, 2$, and beyond $\mu > 2$ we randomly sampled $A_{rr}^{(\mu)}$. **(E)** Same plot of task-1 dynamics as in left panel of (C) but for a network with $P = 1000$ and smaller axis scales. **(F)** Time-averaged power $\langle \|\bar{z}^{(\mu)}(t)/D^{(\mu)}\|^2 \rangle_t$ of the projections onto all task subspaces, normalized by $D^{(\mu)}$ and sorted in descending order, for a network with $P = 500$ tasks (black) and a network with $P = 2500$ tasks (gray) from panel (D). **(G)** Top: average gain as a function of time for two different networks, one with no task dominant (gray) and one with task 1 dominating (black). Middle: normalized latent-variable trajectories $\bar{z}_r^{(\mu)}(t)/D^{(\mu)} = \phi(t) \cdot \mathbf{n}^{(\mu,r)}$ for each task in the task-1-dominant network. Bottom: same as middle for the network with no task dominant; note the difference in y-axis scale.

is dominant. The point at which this transition occurs corresponds to a fixed ratio of number of tasks to number of neurons, denoted by $\alpha = P/N$ (Fig. 2D [↗](#), light vs. dark cyan).

In an example network with large P and no task dominant, the task-1 latent variables fluctuate near the origin and are smaller in magnitude than those of networks with smaller P (Fig. 2E [↗](#), purple trajectory). However, signatures of the task's oscillatory structure remain visible in an average flow field computed over many samples of network activity (Fig. 2E [↗](#), black). As we will show later, the irregular trajectories of latent variables in networks with large P are driven by chaotic fluctuations that arise from many small overlaps among distinct tasks' subspaces, an effect that is vanishing for small P but significant for large P . In this state, activity is distributed across task subspaces (Fig. 2F [↗](#), gray), rather than residing primarily in a dominant-task subspace (Fig. 2F [↗](#), black). Additionally, the gain factor equilibrates to a time-independent value (Fig. 2G [↗](#) top, gray), rather than being dynamically entrained to any one task's latent variables (Fig. 2G [↗](#) top, black).

The simulations shown in Fig. 2 [↗](#) demonstrate that nonlinear multi-task networks exhibit interference. When P is small, interference leads to winner-take-all dynamics in which one fixed task always dominates, regardless of the initial condition. As P is increased, interference additionally leads to chaotic fluctuations that eventually suppress task-related dynamics. Unlike previous models in which chaos arises due to unstructured, random connectivity [\[36\]](#), or weight matrices with both unstructured and low-rank components [\[28\]](#), in our model all components of the synaptic weight matrix are task-related (Eq. 4 [↗](#)). This suggests that chaotic fluctuations, which are often proposed as an explanation for irregular neural activity recorded *in vivo*, can arise solely as a consequence of learning many tasks (see Discussion).

Modulation of effective connectivity enables task selection

The results so far show that multi-task networks with fixed weights of the form of Eq. 4 [↗](#) cannot flexibly transition between distinct task-related activity patterns. Instead, at most one task dominates the winner-take-all dynamics (Fig. 2 [↗](#)). We investigated mechanisms to circumvent this limitation. One possibility is to provide an input aligned with the task subspace of the desired task. Surprisingly, for networks with Gaussian loading vectors, this strategy does not enable task selection (Fig. S4 [↗](#)) because such inputs modulate the gain factor by an equal amount across all tasks [\[29\]](#).

We therefore hypothesize that a mechanism that modulates $D^{(\mu)}$, the strength of the component of the weight matrix associated with task μ , could enable selective activation of a desired task's dynamics. Such a modulation can be interpreted as increasing the effective recurrent strength of one task-specific latent dynamical system. Targeted gain control in which inputs lead to changes in effective connectivity through a nonlinearity, neuromodulatory signals, or short-term synaptic plasticity are potential biological substrates (see Discussion).

We develop an analytical theory that characterizes the dynamical regimes of multi-task networks with large P and predicts when task selection due to modulation of $D^{(\mu)}$ is successful. The theory applies to networks for which P and N are large, but with a fixed ratio $\alpha = P/N$. It also assumes that R , the dimension of the manifold associated with a single task, is fixed. This scaling is motivated by the hypothesis that neural circuits perform many distinct tasks, each of which is associated with a low-dimensional manifold (Fig. 3A [↗](#)).

The theory reveals that modulation of $D^{(\mu)}$ indeed supports selection of task-related dynamics in multi-task networks defined as Eq. 4 [↗](#). These connectivity modulations can place the dynamics of such networks into one of four distinct regimes. In the first, weights are small in magnitude and no activity is produced, so we do not consider this case further. In the second, the *spontaneous* state, chaotic fluctuations occur and no task is dominant (as in Fig. 2E [↗](#)). The other two states are the chaotic and nonchaotic *task-selected* states. The network transitions from the spontaneous state to a state in which task μ is selected if $D^{(\mu)}$ is increased sufficiently (Fig. 3B [↗](#)). Of course, task selection could be trivially achieved by setting $D^{(\mu)}$ to zero for all but one task, or by increasing it




Figure 3. Task-selected states.

(A) Top: Illustration of task selection as promoting activity in the chosen task's subspace, which contains task-specific dynamics. Bottom: Phase diagram as a function of the task strength $D^{(\mu)}$, comprising the spontaneous, chaotic task-selected, and nonchaotic task-selected states. Schematic of task selection at the level of connectivity: modulating the strength of the chosen task's connectivity component through the associated $D^{(\mu)}$ value. (C) Trajectory of a task-1 latent variable (normalized by $D^{(1)}$), $z_1^{(1)}(t)/D^{(1)}$, when task-1 limit-cycle dynamics are selected (left) vs. when task-2 bistable dynamics are selected (right), in all 3 regimes. To illustrate bistability, multiple trajectories with different initial conditions are plotted on the right. (D) Same as (C) but for a task-2 latent variable $z_1^{(2)}(t)/D^{(2)}$. (E) Same as (C) and (D) but for an example neuron's activity.

for a single task so dramatically that it dominates the entire weight matrix. However, our model requires only small increases in $D^{(\mu)}$ to produce large changes in task-related dynamics (see Discussion).

The simulations thus far have depicted the transition from a task-selected state to the spontaneous state as α is varied (Fig. 2D). Henceforth, we consider transitions from the spontaneous state to task-selected states due to modulation of $D^{(\mu)}$ for fixed α . Such a transition can be induced for any task by modulating its associated $D^{(\mu)}$ value (assuming the task does not involve purely decaying dynamics; see Methods). Thus, there are many distinct task-selected states, for which activity is large in the selected task's subspace and small in the others (Fig. 3C,D). This contrasts with networks without modulation of $D^{(\mu)}$, in which winner-take-all dynamics do not permit the selection of different tasks (Fig. 2C).

Theory of the spontaneous state

The spontaneous state occurs when two conditions on the connectivity are met. First, for the network not to be quiescent, the average synaptic weight magnitude, $\langle J_{ij}^2 \rangle$, must be sufficiently large (see Methods). This is similar to the condition for spontaneous activity to occur in networks with independent and identically distributed weights [36] and requires that $\alpha > 0$, i.e. the number of tasks indeed scales with the network size. The second condition is that there must be no task μ whose strength $D^{(\mu)}$ is so large that it dominates the other tasks, leading to a task-selected state.

Here we describe latent-variable dynamics, and in later sections, we describe single-neuron activity (the theory describes both; see Methods). In the spontaneous state, the latent variables $z_r^{(\mu)}(t)$ ($r = 1, \dots, R$) for each task ($\mu = 1, \dots, P$) follow

$$z_r^{(\mu)}(t) = \eta_r^{(\mu)}(t) + D^{(\mu)} \langle \phi' \rangle \sum_{r'=1}^R A_{rr'}^{(\mu)} [F * z_{r'}^{(\mu)}](t), \quad (6)$$

where $\eta_r^{(\mu)}(t)$ is a temporally correlated Gaussian noise process and $F(\tau)$ is a temporal filter. Eq. (6) describes the dynamics in previous simulations when no task is dominant (e.g. Fig. 2E; or Fig. 2G, bottom).

Eq. (6) can be compared to Eq. (3), which describes a single-task network (Fig. 2A). The first difference is the presence of an additional term $\eta_r^{(\mu)}(t)$, which represents interference due to effective noise from the many small overlaps with other tasks' dynamics. The second is that the gain factor $\langle \phi' \rangle$ is time-independent. This reflects the fact that, in the spontaneous state, the gain is not entrained to any single task's latent variables. Instead, it is determined by an average over many vanishingly correlated latent variables of different tasks, producing a time-independent quantity (Fig. 2G, black vs. gray). Finally, the linear filter $F(\tau)$ is in general different from that of Eq. (3); it depends on the statistics of the ensemble of $D^{(\mu)}$ and $\mathbf{A}^{(\mu)}$ (see Methods).

Therefore, in the spontaneous state, the dynamics of $z_r^{(\mu)}(t)$, for any task μ , are noise-driven and linear. As such, nonlinear dynamics such as asymptotically stable limit cycles or multistability cannot occur. Nevertheless, the spontaneous-state dynamics within a single task subspace are related to those that would be produced if the task were dominant (Fig. 2A,B). For example, the latent variables associated with limit-cycle dynamics exhibit an oscillatory component (Fig. S5A). Similarly, the latent variables associated with bistability exhibit a slow activity mode along the same axis that contains the fixed points in a single-task network (Fig. S5A). Thus, in the spontaneous state, latent variables for all tasks simultaneously follow noisy, linearized, and weakened versions of their task-selected dynamics.

Theory of the task-selected state

When a specific task μ 's strength $D^{(\mu)}$ is increased sufficiently, the linear dynamics of the spontaneous state (Eq. 6) become unstable, leading to rapid growth of the task's latent variables. When this occurs, the gain factor $\langle \phi'(t) \rangle$ is no longer time-independent and instead is dynamically entrained to task μ 's latent variables. This change amounts to the selection of the nonlinear

dynamics associated with these variables and the suppression of unselected task dynamics (Fig. 3A). Just above this value of $D^{(\mu)}$, task-selected dynamics coexist with chaotic fluctuations. Increasing $D^{(\mu)}$ further leads to a second transition, beyond which the dominant task completely suppresses chaotic fluctuations, creating a nonchaotic task-selected state. Previous studies have identified related transitions in other network models [28, 37–40].

How does the activity of single neurons and of task-related latent variables change across these transitions? Increasing $D^{(1)}$ or $D^{(2)}$ (Fig. 3B) beyond the first transition threshold—to place the network in the corresponding chaotic task-selected state—leads to the associated latent variables growing in magnitude (from $\mathcal{O}(1)$ in the spontaneous state to $\mathcal{O}(\sqrt{N})$), while latent variables for any non-selected task remain $\mathcal{O}(1)$ (Fig. 3C,D). In the chaotic task-selected state just above the transition, the activity of each neuron exhibits a combination of tuning to the selected task's latent variables and chaotic fluctuations, with the two effects contributing to single-neuron activity on the same order (Fig. 3E, darker gray). In contrast, the selected task's latent variables are noiseless to leading order, as single-neuron fluctuations average out in the projection onto the task subspace. Thus, the chaotic task-selected state is characterized by variability at the single-neuron level but not at the level of task-relevant variables. Increasing $D^{(1)}$ or $D^{(2)}$ further to enter the nonchaotic task-selected state suppresses this variability at the single-neuron level as well (Fig. 3E, black).

The value of $D^{(\mu)}$ required for task selection depends on task μ 's structure, captured by its overlap matrix $\mathbf{A}^{(\mu)}$, as well as on the number and structure of other tasks (see Methods). Fig. 4 illustrates these dependencies for an analytically tractable class of networks in which each task corresponds to the generation of a stable oscillation with a different frequency (see Methods). A larger value of $D^{(\mu)}$ is required for task selection when the number of tasks is increased (horizontal slices in Fig. 4A; compare with Fig. 2D). The dependence of the critical value on the overlap matrix of the selected task reduces, in this case, to a dependence on the selected task's oscillation frequency (Fig. 4B).

Single-neuron statistics in spontaneous and task-selected states

Thus far, we have focused on describing latent variables defined by projections of network activity onto task-specific subspaces. The theory also describes activities of individual neurons (see Methods). In the spontaneous state, the activity of a given neuron i is described by

$$\dot{x}_i(t) = -x_i(t) + \eta_i^x(t) + [G * \phi_i](t), \quad (7)$$

where $\eta_i^x(t)$ is a temporally correlated Gaussian noise process and $G(\tau)$ is a filter that produces a nonlinear self-coupling of each neuron to its firing rate. Each neuron's activity fluctuates with negligible tuning to any single task's latent variables (Fig. S5A). The statistics of these fluctuations are homogeneous across neurons. The theory allows analytic calculation of the correlation functions of $\eta_r^{(\mu)}(t)$ and $\eta^x(t)$ and the kernels $F(\tau)$ and $G(\tau)$. This leads to predictions for the single-task and single-neuron correlation functions, $C_{rr}^{(\mu)}(\tau)$ and $C^\Phi(\tau)$, respectively, which we verify with simulations (Fig. S5).

In task-selected states, the equation describing single-neuron activity is similar to Eq. (7), with an additional term corresponding to the neuron's tuning to the dominant task's latent variables:

$$\dot{x}_i(t) = -x_i(t) + \eta_i^x(t) + [G * \phi_i](t) + \sum_{r=1}^R m_i^{(\mu^*, r)} z_r^{(\mu^*)}(t), \quad (8)$$

where μ^* is the index of the dominant task. The $\eta_i^x(t)$ term is still a Gaussian noise process, but its kernel is not time-translation-invariant and instead depends on the dynamics of the dominant task (see Methods).




Figure 4. Phase diagrams for networks generated as a sum of single-task components associated with oscillatory dynamics (see Methods).

(A) Network state as a function of normalized number of tasks $\alpha = P/N$ and the strength of one task $D^{(\mu)}$. All other tasks have $D^{(\nu \neq \mu)} = 2$ (solid line). Dashed line: Boundary between spontaneous and chaotic task-selected state for task μ . Dot-dashed line: Boundary between chaotic and nonchaotic task-selected state for task μ . **(B)** Same as (A) but for fixed $\alpha = 1$ and instead varying the frequency of the limit cycle task that is selected. Black boxes: Parameter values whose associated single-neuron properties are examined in Fig. 5.

In task-selected states, neurons are statistically heterogeneous, following different trajectories depending on their tuning to the dominant task's latent variables (Fig. 5A). This heterogeneity implies that the extent to which single-neuron activity can be explained by one task's latent variables varies across neurons in the population. The mean of this distribution is low in the spontaneous state, higher in the chaotic task-selected state, and highest in the nonchaotic task-selected state (Fig. 5B).

The average autocovariance function $C^X(\tau) = \langle x_i(t)x_i(t + \tau) \rangle_{t,i}$ illustrates the differences between the temporal statistics of single-neuron firing rates in the spontaneous and task-selected states (Fig. 5C). The first transition occurs when task-related structure, in this case oscillations, becomes visible at the single-neuron level (Fig. 5C, left vs. middle). The second transition occurs when all firing-rate variance is accounted for by dominant-task activity, or equivalently, in the case of a limit-cycle task, when $C^X(0) = C^X(T)$ for an oscillation with period T (Fig. 5C, middle vs. right).

Dimensionality of spontaneous and task-selected states

To analyze the structure of population activity in multi-task networks, we also calculate its dimension. We quantify dimension using the participation ratio, a measure of linear dimensionality that approximates the number of orthogonal directions in activity space that are explored by the network's dynamics. This quantity has been the subject of numerous experimental and theoretical studies [37, 41–45].

In the spontaneous state, the dimension scales with the network size N , as previously observed in networks with random connectivity (43; Fig. 5D, left). Although this scaling produces dimension values that are large, they are always much lower than the ambient dimension N due to weak correlations among neurons [45], a common observation in neural data [22, 46, 47]. At the other extreme, in the nonchaotic task-selected state, the dimension is determined solely by the dominant task's dynamics and is therefore proportional to R (Fig. 5D, right). Between these extremes, in the chaotic task-selected state, the dimension is larger than this limiting value by a factor that does not grow with N and can be predicted by the theory (Fig. 5D, middle; see Methods). This increase arises from chaotic fluctuations in directions orthogonal to the dominant task's subspace. The magnitudes of these fluctuations depend on the values of $D^{(\mu)}$ for both the selected and unselected tasks (Fig. 56).

These results suggest that experimental measurements of neural dimensionality should be interpreted in the context of the underlying network state. In particular, the observed dimension may depend on which task-selected states are present, the statistics of transitions between these states, and the dimension of the manifolds associated with each individual task. Consistent with previous studies of neural activity during specific behavioral tasks [21, 48, 49], our model predicts that the dimension of activity is low when activity is measured within a single task-selected state, even when fluctuations are present at the single-neuron level. An experimentally observed increase beyond this value can arise through two distinct mechanisms. First, the network may transition into the spontaneous state, whose dimension scales with network size, consistent with analyses of large-scale neural recordings [20]. Second, the dimension of activity may increase through the sequential engagement of multiple task-selected states, each associated with a different low-dimensional manifold.

These mechanisms make distinct predictions for how the dimensionality of activity grows with recording time (Fig. 5E). In the spontaneous state, dimensionality increases rapidly at short times and then saturates (Fig. 5E, gray). When the network sequentially switches across task-selected states, the rate of dimensionality growth depends on the amount of time spent in each task state. Even in cases in which this time is short (as in our simulations), dimensionality grows more slowly than in the spontaneous state (Fig. 5E, black).

The two mechanisms also differ in their saturation values. In the spontaneous state, network activity explores a chaotic attractor in which projections onto certain directions in neural activity space contain more variance than others, limiting dimensionality. In contrast, sequential task selection produces a sequence of trajectories that are highly decorrelated from one another. As a




Figure 5. Single-neuron and population properties of a multi-task networks.

Parameters are the same as Fig. 4 (see Methods). **(A)** Example neurons (top) and latent variables for the selected task (bottom) across network states. **(B)** Histogram R^2 values for linear regression of $x_i(t)$ against task- μ latent-variable activity. **(C)** Average single-neuron autocovariance functions $C^X(t) = \langle x_i(t)x_i(t + \tau) \rangle_{t,i}$. **(D)** Top: Principal component trajectories across same network states from (A), (B), and (C). For currents $x_i(t)$ and firing rates $\phi_i(t)$, principal component axes and projections onto these axes are computed separately, with $\phi_i(t)$ projections doubled for visibility. Axis scales: 150 (left), 300 (middle), 500 (right). Bottom: Linear embedding dimension computed as participation ratio of the firing rates $\phi_i(t)$, as a function of $D^{(\mu)}$. Light dots: individual network simulations. Heavy dots: averages over networks. **(E)** Firing-rate dimension as a function of recording time for a network in the spontaneous state and switching across task-selected states. For the latter, one of 150 tasks is randomly selected for 80 time constants, followed by 3 time constants of spontaneous activity, and this process is repeated. Thin lines: individual network simulations. Thick lines: average. Inset: The same plot but for networks with different N , showing dependence on N in the spontaneous state.

result, the saturation value for sequential task selection is approximately equal to the number of visited task states multiplied by the within-task dimension. This value can be either higher or lower than the spontaneous-state dimension. For a sufficiently large number of tasks, the dimensionality associated with sequential task selection can greatly exceed that of the spontaneous state (Fig. 5E, inset). Altogether, the theory demonstrates that either a fixed high-dimensional covariance matrix or state-dependent switching among different low-dimensional covariance matrices are both potential models for high-dimensional activity in neural recordings.

Discussion

We identified limitations to multi-task capabilities in nonlinear low-rank networks due to interference among tasks. Such limitations are not evident in analyses of networks constructed to perform a single task or in linear networks that lack the winner-take-all interference we described. We proposed a solution to the problem of interference that relies on modulation of effective connectivity. Networks with such a mechanism transition between a spontaneous state in which no task is selected and task-selected states in which activity lies primarily on a task-specific low-dimensional manifold. Our theory thus retains the analytic tractability of low-rank networks [28, 29, 34] while describing the geometry of neural dynamics associated with multiple tasks.

Neuronal architectures for multi-task networks

The connectivity of our model, and its modulation during task selection, may be interpreted in multiple ways. If the network is viewed as a single recurrently connected circuit, modulation of $D^{(\mu)}$ corresponds to a change in the strength of a single task-related component of the weight matrix (Fig. 6A). Such a change requires selective modulation of synapses depending on their participation in this weight matrix component. Potential biological substrates include rapid synaptic plasticity or targeted neuromodulation [51–56].

Alternatively, the loadings $\mathbf{m}^{(\mu,r)}$ and $\mathbf{n}^{(\mu,r)}$ may be interpreted as synaptic weights between a network and an auxiliary population of neurons, for instance thalamocortical and corticothalamic weights, respectively (Fig. 6B). In this case, task selection is accomplished through adjustment of the gain of a subset of task-associated thalamic neurons. A linear network model of such a mechanism has been proposed [50]. More generally, because only small modulations of effective connectivity are required in our nonlinear model, its weights may be interpreted as a combination of fixed corticocortical and adjustable thalamocortical synapses (Fig. 6C).

We focused on how a single network can generate many distinct sets of low-dimensional dynamics, without explicitly modeling task-related sensory inputs or how these dynamics are transformed to produce behavior. In low-rank single-task networks, task-related inputs define additional dimensions of the task subspace [28, 34]. Similarly, we expect our qualitative results on interference and winner-take-all dynamics to generalize to cases where distinct tasks operate with distinct input and output channels. Future work could examine the consequences of sharing neural modes across tasks, whether input channels, latent dynamical modes, or output channels [2, 3, 5, 7].

Related models of multi-tasking

Previous studies have examined multi-tasking in recurrent neural networks that are trained end-to-end to perform individual tasks that depend on a contextual input [3, 27], or sets of tasks commonly used in neuroscience experiments [4, 14]. These models produce diverse dynamical behaviors but their weight matrices have complex statistical structure from the training process that limits analytical tractability. It would be interesting to examine whether the tonic inputs commonly used to signal task context in trained multi-task networks effectively drive similar task-selection mechanisms via gain modulation, perhaps in task-specific populations [14, 29, 34], by shifting the operating point of the nonlinearity.

Other works have studied nonlinear models with closed-form expressions for the weight matrix, including summations over a large number of low-rank components, as in our model. In certain cases, such connectivity leads to chaotic activity [38, 45, 55, 57, 58]. These studies considered




Figure 6. Illustrations of three interpretations of the network's effective connectivity.

(A) Modulation of $D^{(\mu)}$ modifies the recurrent connections J_{ij}^{CC} of a cortical network. **(B)** Interactions between neurons are driven by effective connectivity J_{ij}^{CTC} defined by thalamocortical loops. Modulation of $D^{(\mu)}$ adjusts the gain of thalamic neurons [50]. **(C)** Model in which the effective connectivity is a combination of (A) and (B).

networks that are designed to produce specific dynamical features, including fixed points [38], continuous attractors [58], and sequential activity [55, 57]. Our model, on the other hand, permits each task to be associated with any low-dimensional dynamical system described by Eq. (3) [4], that is, any dynamical system that can be described within the framework of single-population low-rank networks [29]. Thus, although we limited our analysis to limit cycles and fixed points, our model also accommodates sequential activity and non-normal latent dynamics [59].

Modulating connectivity has been proposed in previous works. A previous study developed a linear network model of thalamocortical interactions (Fig. 6B [4]) that sequences transient motor motifs [50]. Other models have used statistical inference to learn time-varying scaling factors that adjust low-rank connectivity components in nonlinear networks [6, 56]. Our theory provides an analytical account of such modulation and its consequences for task selection and spontaneous activity.

In biological and artificial systems, complex behaviors are often built by combining simpler behavioral components [14, 60]. In our model, each “task” should be interpreted as a minimal computational component associated with a behavior, which may only be part of what is often referred to as a task in behavioral experiments. For example, fixation, stimulus, delay, and response periods may each demand different network dynamics [6]. To retain analytic tractability, here we assumed no overlap or reuse of dynamics across tasks, but including such structure is an interesting direction for future work.

Related theoretical work suggests that selecting among low-rank connectivity components, each suited to a particular subtask, is a more effective strategy for continual learning than treating each compound behavior as its own learning problem [6]. Animal experiments indicate that complex motor control arises from sequencing different low-dimensional dynamics, each associated with a particular manifold of state space, to generate flexible behavioral sequences [2]. Our theory provides an analytically tractable mechanism for such sequential composition. In contrast, the winner-take-all dynamics of our model does not support multiple simultaneously active tasks, due to the Gaussian-distributed loadings. Such dynamics require additional circuit structure, such as multiple neuronal populations (e.g. from mixture-of-Gaussians loadings) with population-specific gain factors [29, 34]. Nonetheless, the same gain-dependent mechanisms that enable task selection and suppression in the Gaussian case may also facilitate simultaneous selection of multiple tasks—and suppression of all unselected tasks—in more complex models, which we leave to future work. Such models would additionally allow for more complex single-task dynamics beyond what can be achieved with one population [34].

Magnitude of effective connectivity modulation

Our theory shows that, in nonlinear networks, a modest modulation of effective connectivity is often sufficient for task selection. Adjustment of all values of $D^{(\mu)}$, such that each value is either zero for unselected tasks, or nonzero for the selected task, would achieve task selection trivially by completely overwriting the synaptic weight matrix each time a task is selected. In contrast, our model requires adjusting the $D^{(\mu)}$ value for only the selected task, which corresponds to changes in synaptic weights that are $\mathcal{O}(1/P)$, whereas each synaptic weight is $\mathcal{O}(1/\sqrt{P})$. Thus a small task-related connectivity component can dramatically influence a network’s activity, an observation that has been made in models with both structured low-rank and unstructured weight matrix components [28, 33, 39, 61]. Our model demonstrates that this feature holds in networks for which all connectivity components are task-related.

Moreover, the degree of modulation, even relative to $1/P$, can be made small in our model. The change in $D^{(\mu)}$ required for task selection depends on its distance from the critical value at which task selection occurs (Fig. 4B [4], black solid line vs. dashed). When α is small, a plausible limit in which the number of tasks is large but much smaller than the number of neurons, these distances can be decreased (Fig. 4A [4]). Rapid reconfiguration of dynamics thus occurs with only small connectivity adjustments, as each task is close to selection. In such networks, dynamics are near the edge of stability in the spontaneous state [62], providing a functional explanation for the presence of irregular spontaneous activity.

Dimensionality and the relationship between spontaneous and evoked activity

Spontaneous activity has historically been defined as activity in the absence of sensory stimulation [63]. Classic theoretical work showed how such activity could be self-generated by a nonlinear network without external inputs [36]. However, recent high-yield recordings during periods of unstimulated, uninstructed behavior have revealed high-dimensional activity whose high-variance components are tuned to kinematic variables [19, 20], which is not well described by unstructured chaotic activity. The spontaneous state of our model has similarities to the chaotic regime of unstructured networks [36] but describes both the irregular activity of neurons and also of task-relevant latent variables. The spontaneous state of our model may explain the idiosyncratic behaviors that have been observed to accompany periods of “quiet wakefulness” [15], while the chaotic task-selected state may explain how such behaviors coexist with coherent goal-directed actions [64].

When freely exploring, rodents perform a variety of self-initiated behaviors, such as whisking or grooming [1]. Many of these behavioral “syllables” feature precise kinematic control that is likely not explained by the spontaneous state of our model. Rather, recordings during long periods of free behavior may correspond to switching among different task-selected states of our model. This predicts that individual neurons should have heterogeneous tuning to kinematic variables for an active behavioral syllable, and this tuning will change across syllables (Fig. 5B). A second prediction is that, when pooling neural activity across distinct syllables, the observed dimension should increase with the diversity of behavioral syllables observed. Conversely, dimensionality measured over many repetitions of a single behavioral syllable should be small, even if there is trial-to-trial variability at the single-neuron level (Fig. 5D). A final prediction is that the initial rate of dimension growth relative to recording time should be faster during quiet wakefulness than during behavior, and the magnitude of this effect depends on the average syllable duration (Fig. 5E). These predictions can be examined in neural recordings during free behavior [65].

Methods and materials

Dynamical mean-field theory

We now describe the full theory in greater detail, leaving derivations of the results to the Appendix. The dynamics describing the recurrent network (Eq. 1) can be interpreted as a bidirectionally coupled system of neurons and latent variables (Fig. 7A). Each latent variable is defined as a sum over neuronal firing rates, and in turn the input to each neuron can be written as a sum over latent variables:

$$z_r^{(\mu)}(t) = D^{(\mu)} \sum_i n_i^{(\mu,r)} \phi_i(t) \quad (9)$$

$$\dot{x}_i(t) = -x_i(t) + \sum_{\mu=1}^P \sum_{r=1}^R m_i^{(\mu,r)} z_r^{(\mu)}(t). \quad (10)$$

The magnitudes of these latent variables can be large $\mathcal{O}(\sqrt{N})$, as in Fig. 2A or small ($\mathcal{O}(1)$), as in Fig. 2E. We refer to any task whose latent variables are large as a “dominant” task, and the others as “non-dominant” tasks. This distinction is central to the mean-field description, because each neuron’s activity is significantly affected by any dominant task, meaning that a finite fraction of its variance can be explained by that task’s variables. This is not true for the non-dominant tasks, although they collectively provide non-negligible input as an ensemble. A task can be dominant for two reasons: 1) initial conditions or inputs pushing the state of the network to preferentially align with the task subspace, or 2) self-




Figure 7. Conceptual illustration of the DMFT.

(A) Cartoon illustrating an alternative interpretation of the neuron-to-neuron recurrent dynamics, as continuous-time, bidirectional interactions between a reservoir of neurons (blue: currents; yellow: firing rates) and a reservoir of latent variables (dark and light purple). Here we have a single set of dominant latent variables ($P^* = 1$), illustrated as larger purple disks. **(B)** Types of effective input a single neuron receives in the DMFT: input from dominant task latent variables with tuning determined by $m_i^{(\mu, r)}$ (top); summed input from weak-task latent variables, which is well described as a sample $\eta_i^x(t)$ from a Gaussian process (middle); and nonlinear effective self-interaction that arises through coherent summation of the neuron's perturbative effect on latent-variable trajectories (bottom). **(C)** Types of effective input a single weak-task latent variable receives in the DMFT: summed input from neurons, which is well described as a sample $\eta_r^{(\mu)}(t)$ from a Gaussian process (top); and effective self-interactions that are mediated by single-neuron response properties as well as task-specific parameters $D^{(\mu)}$ and $\mathbf{A}^{(\mu)}$ (bottom).

amplifying dynamics within the task subspace, as shown in Fig. 3. A random network state results in weak latent variables for every task.

We set aside a separate index $\mu^* = 1^*, \dots, P^*$, with $P^* \geq 0$ small, to enumerate tasks $z_r^{(\mu^*)}(t)$ with dominant $\mathcal{O}(\sqrt{N})$ activity, if there are any. (Although in the main text we emphasized that there can be at most one dominant task, this is true for sufficiently long times (to forget initial conditions) and when there are no external inputs. Generally, initial conditions or external input drive can set multiple task-related latent variables to be large. Once inputs are released, the network eventually equilibrates to having at most one dominant task. In all cases, only finitely many tasks can be dominant since the state vector norm is at most $\sqrt{N}$.) Then $\mu = 1, \dots, P$ continues to enumerate the many non-dominant tasks with $\mathcal{O}(1)$ activity. In the $N, P \rightarrow \infty$ limit with fixed ratio $\alpha = P/N$, fixed P^* , and fixed R , we can statistically describe the trajectories of all neurons and weak-task-related latent variables as decoupled processes, rather than as solutions to the high-dimensional dynamics of Eq. (1). In this mean-field description, the activity of the i th neuron $x_i(t)$ and the r th latent variable $z_r^{(\mu)}(t)$ of the μ th non-dominant task evolve as

$$\begin{aligned} \dot{x}_i(t) = & -x_i(t) + \eta_i^x(t) + \sum_{\mu^*=1^*}^{P^*} \sum_{r=1}^R m_i^{(\mu^*,r)} z_r^{(\mu^*)}(t) \\ & + \int_{-\infty}^t dt' G^z(t, t') \phi_i(t') \quad \eta_i^x(t) \sim \mathcal{GP}(0, Q^z(t, t')) \end{aligned} \quad (11)$$

$$\begin{aligned} z_r^{(\mu)}(t) = & \eta_r^{(\mu)}(t) + D^{(\mu)} \sum_{r'=1}^R A_{rr'}^{(\mu)} \int_{-\infty}^t dt' S^\phi(t, t') z_{r'}^{(\mu)}(t') \quad \eta_r^{(\mu)}(t) \\ & \sim \mathcal{GP}\left(0, \left(D^{(\mu)}\right)^2 C^\phi(t, t')\right) \end{aligned} \quad (12)$$

In deriving the theory, several functions of time emerge that are shared by all neurons and tasks: four kernel functions $C^\phi(t, t')$, $S^\phi(t, t')$, $Q^z(t, t')$, $G^z(t, t')$ that describe effective noise statistics and interactions, plus the set of dominant-task latent variable trajectories $\{z_r^{(\mu^*)}(t)\}_{\mu^*, r}$. (In the main text, we wrote $S^\phi(t, t') = \langle \Phi'(t) F(t, t') \rangle$ and $G^z(t, t') = G(t, t')$ to simplify notation—see section on self-consistency conditions.) All these functions' values depend collectively on $\{(D^{(\mu)}, \mathbf{A}^{(\mu)})\}_{\mu=1}^P$ and α , and they can be calculated by enforcing self-consistency, as we describe later.

Each neuron i is described in Eq. (11) as leaky integration of 3 current sources (Fig. 7B):

1. A Gaussian process noise term $\eta_i^x(t)$ arising from interactions with the large reservoir of weak-task-related latent variables and capturing the effects of chaotic fluctuations. Its temporal statistics $Q^z(t, t')$ are inherited from average statistics of non-dominant-task latent variables.
2. Drive from the dominant-task latent variables $z_r^{(\mu^*)}(t)$ via the neuron-specific loadings $m_i^{(\mu^*, r)}$.
3. A nonlinear self-coupling to the neuron's own firing rate, which is generally a consequence of introducing task-related overlaps. The kernel mediating this coupling $G^z(t, t')$ is determined by weighted averages of the induced network connectivity motifs of all orders.

Similarly, each weakly active task μ is described in Eq. (12) as an integral equation with two input sources (Fig. 7C):

1. A Gaussian process noise term $\eta_r^{(\mu)}(t)$ that arises from the sum over many irregular neuronal activity trajectories via Eq. (9). Its temporal statistics $(D^{(\mu)})^2 C^\phi(t, t')$ are determined by the average neuronal autocovariance function and scaled by $D^{(\mu)}$.
2. A coupling among task μ 's latent variables mediated by $D^{(\mu)} \mathbf{A}^{(\mu)}$ and the kernel $S^\phi(t, t')$. The kernel $S^\phi(t, t')$ is determined by average impulse-response properties of the neuronal population. The average gain is a crucial factor in determining $S^\phi(t, t')$, since (un)saturated neurons are less (more) responsive to inputs.

The dominant-task latent variables satisfy

$$z_r^{(\mu^*)}(t) = \left\langle D^{(\mu^*)} n_i^{(\mu^*, r)} \langle \phi(x_i(t)) \rangle_{W_i^*} \right\rangle_i, \quad (13)$$

where $\langle \cdot \rangle_{W_i^*}$ indicates averaging over the random process described by Eq. (11) for a given neuron i . In general, this expression does not have a simple closed-form solution, and calculating the forward dynamics requires iterative numerical approaches to maintain self-consistency. However, in special cases where currents $x_i(t)$ are Gaussian, the dynamics of $z_r^{(\mu^*)}(t)$ essentially reduce to the simple form of Eq. (3) from the single-task network case, by Stein's Lemma. (With a small difference: $\langle \Phi(t) \rangle$ is determined by $z_r^{(\mu^*)}(t)$ as well as a constant interference term from non-dominant tasks. But since this term has Gaussian statistics, Stein's Lemma still applies.)

In summary, our theory is able to completely describe network activity for an arbitrary α and ensemble of task dynamics, at the level of both single neurons and task-specific subspaces. Inputs to each can be decomposed into regular, task-related signals and effectively noisy fluctuations due to chaos, with mutually determined statistics that can be analytically solved for.

Note on N-scaling

Here we summarize how various structural and dynamical quantities scale with the network size N . Components of loading vectors ($\mathbf{m}^{(\mu, r)}$, $\mathbf{n}^{(\mu, r)}$ for $r = 1, \dots, R$, $\mu = 1, \dots, P$) are $\mathcal{O}(1/\sqrt{N})$, resulting in elements of each single-task weight matrix ($D^{(\mu)} \sum_{r=1}^R \mathbf{m}^{(\mu, r)} (\mathbf{n}^{(\mu, r)})^T$ for $\mu = 1, \dots, P$) being $\mathcal{O}(1/N)$. Because the number of tasks P is $\mathcal{O}(N)$, each single-task weight matrix has zero mean, and the single-task weight matrices are sampled independently, the elements of the full weight matrix that sums over all single-task matrices is $\mathcal{O}(1/\sqrt{N})$.

Neural activities ($x_i(t)$, $\Phi_i(t)$ for $i = 1, \dots, N$) are $\mathcal{O}(1)$. Projections of neural activity onto single-task subspaces $z_r^{(\mu)}(t) = \sum_{i=1}^N n_i^{(\mu, r)} \phi_i(t)$ for tasks $\mu = 1, \dots, P$ and within-task components $r = 1, \dots, R$ scale in two possible ways depending on whether a task μ is dominant or non-dominant. In particular, projections onto non-dominant task subspaces are $\mathcal{O}(1)$, while projections onto dominant-task subspaces are $\mathcal{O}(\sqrt{N})$. This is because, in the dominant-task case, the network state preferentially aligns with their associated subspaces, and the overall state norm $\|\Phi(t)\|$ is $\mathcal{O}(\sqrt{N})$ while the norms of loadings vectors are approximately 1. In the language of associative-memory networks that store patterns as fixed points, these forms of (non-)alignment are referred to as uncondensed or condensed.

These distinctions in scaling are visible when comparing networks of different sizes (Fig. 2D light vs. dark cyan). The activity level of dominant task $\mu = 1$ is scaled by approximately $\sqrt{2}$ when the number of neurons is doubled, but past the transition to chaos—when task 1 loses dominance—its activity level drops to a value that is the same for networks of different sizes.

Connectivity statistics

The free parameters of the model are P , R , and $\{(D^{(\mu)}, \mathbf{A}^{(\mu)})\}_{\mu=1}^P$. Statistics of J_{ij} can be written in terms of these parameters. The overall strength of the connectivity, which we call g_{eff} is given by the variance of the off-diagonal elements of $\mathbf{J}$ and depends on the second moment of the task weights (but not the overlaps):

$$g_{\text{eff}}^2 = N \langle J_{ij}^2 \rangle_{\mathbf{J}} = \alpha R \langle D^2 \rangle_{D, \mathbf{A}}. \quad (14)$$

Like the analogous g parameter of [36], holding other parameters fixed while increasing this degree of “disorder” eventually effects a transition to chaotic activity, with stronger chaotic fluctuations with increasing g_{eff} . However, in this model with connection symmetries, g_{eff} does not alone determine the network dynamics.

The mean connection strength is

$$\langle J_{ij} \rangle_{\mathbf{J}} = \delta_{ij} \alpha \langle D \text{Tr}(\mathbf{A}) \rangle_{D, \mathbf{A}}, \quad (15)$$

implying strong autapses $J_{ii} \sim \mathcal{O}(1)$ when the overlaps have a nonzero average trace. This homogeneous mean autapse strength $\langle J_{ii} \rangle_{\mathbf{J}}$ can be interpreted as a zeroth-order connectivity motif ρ_0 . The first-order connectivity motif strength ρ_1 is the correlation between transposed elements for $i \neq j$,

$$\rho_1 = N \langle J_{ij} J_{ji} \rangle_{\mathbf{J}} = \alpha \langle D^2 \text{Tr}(\mathbf{A}^2) \rangle_{D, \mathbf{A}}. \quad (16)$$

Higher-order connectivity motifs, e.g. $\rho_2 = \langle J_{ij} J_{jk} J_{ki} \rangle_{\mathbf{J}}$, arise from nonzero average traces of higher matrix powers of $\mathbf{A}$.

If present, these symmetries will lead to effective self-coupling between each neuron and itself, whether through strong self-connections or multi-synaptic pathways. In a Neumann expansion of the the spontaneous state self-coupling kernel $G^z(\omega)$ (see section on self-consistency conditions),

$$G^z(\omega) = \alpha \left\langle D \text{Tr} \left(\mathbf{A} (\mathbf{I} - D S^\phi(\omega) \mathbf{A})^{-1} \right) \right\rangle_{D, \mathbf{A}} \quad (17)$$

$$= \alpha \sum_{k=0}^{\infty} (S^\phi(\omega))^k \langle D^{k+1} \text{Tr}(\mathbf{A}^{k+1}) \rangle_{D, \mathbf{A}} \quad (18)$$

$$= \alpha \sum_{k=0}^{\infty} (S^\phi(\omega))^k \rho_k, \quad (19)$$

which matches our intuition that neurons self-interact through strong network motifs, modulated by the appropriate power of the impulse-response function.

This self-coupling is responsible for non-Gaussian currents $x(t)$, which makes calculations difficult. However, designing a non-trivial ensemble of dynamics that eliminates ρ_k for every k is possible, and in such cases we get Gaussian $x(t)$.

Stability condition for $D_{\text{crit},1}^{(\mu)}$

Computing the stability condition for an integral equation of the form

$$z_r(t) = \eta^{z_r}(t) + D \sum_{r'=1}^R A_{rr'} [S^\phi * z_{r'}](t) \quad (20)$$

requires the Laplace transform of the transfer function $\mathbf{M}^{(\mu)}(s) = (\mathbf{I} - D^{(\mu)} S^\phi(s) \mathbf{A}^{(\mu)})^{-1}$. All singularities s^* of this inversion must satisfy $\Re(s^*) < 0$ for a value of $D^{(\mu)}$ to produce stable dynamics. This condition is generally difficult to calculate without a closed-form expression for $S^\phi(\tau)$, as it requires numerical estimation of the Laplace transform. If we assume vanishing connectivity motifs and recover Gaussian currents, implying $S^\phi(\tau) = \Theta(\tau) \langle \phi' \rangle e^{-\tau}$, the dynamics become equivalent to a leaky, noise-driven ODE, and the stability condition simplifies to

$$D_{\text{crit},1}^{(\mu)} = \frac{1}{\langle \phi' \rangle \mathcal{S}^{(\mu)}}, \quad (21)$$

where $S^{(\mu)} = \max[\Re(\lambda) \mid \lambda \text{ eigenvalue of } \mathbf{A}^{(\mu)}]$ is the spectral abscissa of $\mathbf{A}^{(\mu)}$.

Self-consistency conditions

For the rest of the methods section and appendices, we often drop the neuron index i and non-dominant task index μ , because in the DMFT they are thought of as ensembles from which N and P samples are taken, respectively, rather than explicit lists of elements. A given non-dominant task μ has an associated set of D , $\mathbf{A}$ values that are sampled from a joint distribution $\Pr(D, \mathbf{A})$, and a given neuron i has an associated set of normalized (to $\mathcal{O}(1)$) couplings $c = \{\bar{m}^{(\mu^*, r)}, \bar{n}^{(\mu^*, r)}\}_{\mu^*, r}$ to the dominant-task latent variables. We denote the random process for a single neuron with given couplings c as W_c^x , and the random process for a single non-dominant task with given dynamics D , $\mathbf{A}$ as $W_{D, \mathbf{A}}^z$. The fully general set of self-consistency conditions are

$$S^\phi(t, t') = \left\langle \left\langle \frac{\delta \phi(t)}{\delta I(t')} \right\rangle_{W_c^x} \right\rangle_c \quad (22)$$

$$C^\phi(t, t') = \left\langle \langle \phi(t) \phi(t') \rangle_{W_c^x} \right\rangle_c \quad (23)$$

$$Q^z(t, t') = \alpha \left\langle \left\langle \sum_r z_r(t) z_r(t') \right\rangle_{W_{D, \mathbf{A}}^z} \right\rangle_{D, \mathbf{A}} \quad (24)$$

$$G^z(t, t') = \alpha \left\langle D \sum_{rr'} A_{rr'}^T \left\langle \frac{\delta z_r(t')}{\delta I_{r'}(t)} \right\rangle_{W_{D, \mathbf{A}}^z} \right\rangle_{D, \mathbf{A}} \quad (25)$$

$$z_r^{(\mu^*)}(t) = D^{(\mu^*)} \langle \bar{n}^{(\mu^*, r)} \langle \phi(t) \rangle_{W_c^x} \rangle_c. \quad (26)$$

Each kernel function (or order parameter) is computed as a within-process average parameterized by some value, e.g. $\langle \cdot \rangle_{W_c^x}$ or $\langle \cdot \rangle_{W_{D, \mathbf{A}}^z}$, that is itself averaged over the ensemble of all such parameters, e.g. $\langle \cdot \rangle_c$ or $\langle \cdot \rangle_{D, \mathbf{A}}$.

Thus, the DMFT amounts to a bipartite ensemble of neurons and task-related latent variables with mutually referential statistics. The kernels $Q^z(t, t')$ and $G^z(t, t')$ defining the single-neuron stochastic dynamics are inherited from averages over the ensemble of latent variables, hence their superscripts referencing their definition. Conversely, the kernels $C^\phi(t, t')$ and $S^\phi(t, t')$ defining single-task-latent-variable stochastic dynamics are inherited from averages over the ensemble of neurons. Again, to avoid explaining the bipartite structure implicit in this notation in the main text, we simplified notation and replaced $S^\phi(t, t') = \langle \phi(t) \rangle F(t, t')$ and $G^z(t, t') = G(t, t')$.

Spontaneous state self-consistency conditions

In the spontaneous state, the self-consistency conditions simplify tremendously. First, there are no $z_r^{(\mu^*)}(t)$ values that need to be solved for, and thus the single-neuron dynamics simplify. Second, we can assume stationarity of the kernel functions, i.e. that each kernel function depends only on $t - t'$, not t and t' separately. (This is a fairly mild additional assumption, but one that may be broken with sufficiently strong connectivity motifs [66]). Neurons are now statistically homogeneous, so we do not need to average over the ensemble of processes W_c^x and can instead consider a single representative process W^x . The non-dominant-task latent variables are linear, and their associated autocovariance functions and impulse-response functions can be written compactly in Fourier space. With all of these simplifications, the self-consistency equations take the form

$$C^\phi(\tau) = \langle \phi(t)\phi(t+\tau) \rangle_{W^x} \quad S^\phi(\tau) = \langle \delta\phi(t)/\delta I(t-\tau) \rangle_{W^x} \quad (27)$$

$$Q^z(\omega) = \alpha \langle D^2 \| \mathbf{M}(\omega) \|^2_{\mathcal{F}} \rangle_{D, \mathbf{A}} C^\phi(\omega) \quad G^z(\omega) = \alpha \langle D \text{Tr}(\mathbf{A} \mathbf{M}(\omega)) \rangle_{D, \mathbf{A}}, \quad (28)$$

where $\mathbf{M}(\omega) = (\mathbf{I} - D S^\phi(\omega) \mathbf{A})^{-1}$ is the transfer function for a task with dynamics according to D , $\mathbf{A}$.

Methods for solving the DMFT

To solve the DMFT means to calculate the kernel functions $C^\phi(t, t')$, $S^\phi(t, t')$, $Q^z(t, t')$, $G^z(t, t')$, plus the set of dominant-task latent variable trajectories $\{z_r^{(\mu^*)}(t)\}_{\mu^*, r}$ for an ensemble of task dynamics in a given discretization of time.

Fully general t, t' case

Solving the theory in this most general case is in principle possible by iteratively enforcing self-consistency in sampling a large number of single-neuron trajectories for provisional guesses of each kernel function and the dominant-task latent variable trajectories. A summary of the process:

1. Start with initial guesses for $C^\phi(t, t')$, $S^\phi(t, t')$, and $\{z_r^{(\mu^*)}(t)\}_{\mu^*, r}$.
2. From $C^\phi(t, t')$ and $S^\phi(t, t')$, compute $Q^z(t, t')$ and $G^z(t, t')$ by explicitly averaging over the P non-dominant tasks in the ensemble. Since the non-dominant latent dynamics are linear, we can take advantage of having analytic expressions for their within-task cross-covariances and impulse-response functions, which are necessary ingredients for computing $Q^z(t, t')$ and $G^z(t, t')$.
3. Simulate N_{batch} samples of Eq. (11) in parallel, each requiring an independent sample of the Gaussian noise $\eta^x(t)$ (by the Cholesky decomposition of $Q^z(t, t')$), as well as a sample of the dominant-task coupling strengths $m_i^{(\mu^*, r)}$.
4. Empirically measure the single-unit autocovariance, impulse-response functions, and dominant-task latent variable trajectories based on these samples. (The autocovariance and dominant latent variables are straightforward to calculate from their definitions, while the impulse-response function requires differentiating the individual $x_i(t)$ trajectories.)
5. Smoothly update the values of $C^\phi(t, t')$, $S^\phi(t, t')$, and $\{z_r^{(\mu^*)}(t)\}_{\mu^*, r}$ based on these empirical measurements.
6. Repeat 2-5 until convergence.

Details for a similar method applied to modern Hopfield networks are outlined in [67].

Although this method, as well as some other methods we present below, is numerical in nature, it still does not require ever constructing or simulating a *network* according to Eq. (1) with a synaptic weight matrix—only sampling neural activity and task-related latent variables as random processes with mutually referential statistics.

We did not ultimately present any results with this method, as it is prohibitively computationally expensive with difficult boundary conditions. However, we summarize it to the reader to emphasize that, as written, the DMFT simply encodes a set of mutually referential statistics between a reservoir of neurons and a reservoir of non-dominant latent variables, with a special set of dominant-task latent variable trajectories. In principle, a self-consistent set of such statistics can be found by iterative sampling and updating. However, calculation becomes much easier in certain special cases, as we articulate below.

Spontaneous state case

At a high-level, the general method for solving the theory for the spontaneous state is similar to the fully general method, but now we have much simpler Fourier representations of the kernel functions, rather than explicit t and t' dependence, and we do not have to worry about dominant-

task latent variable trajectories. The process is simplified to:

1. Begin with an initial guess for $C^\Phi(\tau)$ and $S^\Phi(\tau)$.
2. Based on these values, compute $Q^Z(\omega)$ and $G^Z(\omega)$ in Fourier space using Eq.(28) with an explicit average over the $D^{(\mu)}, \mathbf{A}^{(\mu)}$ ensemble.
3. Sample N_{batch} Gaussian noise samples of $\eta_i^x(t)$ based on the kernel $Q^Z(\tau)$, and compute $x_i(t)$ and $\Phi_i(t)$ from Eq. (11) without dominant-task latent variables.
4. Empirically measure $C^\Phi(\tau)$ and $S^\Phi(\tau)$, leveraging the Furutsu-Novikov theorem for estimating $S^\Phi(\tau)$: $S^\Phi(\omega) = C^{\Phi\eta}(\omega)/C^\eta(\omega)$, where $C^{\Phi\eta}(\tau) = \langle \phi_i(t)\eta_i^x(t+\tau) \rangle_{t,i}$ is the cross-covariance of $\Phi_i(t)$ and the driving noise $\eta_i^x(t)$ [43].
5. Smoothly update

$$C^\phi(\tau) \leftarrow (1 - \delta_{mem}) C^\phi(\tau) + \delta_{mem} \widehat{C^\phi(\tau)} \quad (29)$$

$$S^\phi(\tau) \leftarrow (1 - \delta_{mem}) S^\phi(\tau) + \delta_{mem} \widehat{S^\phi(\tau)} \quad (30)$$

with memory parameter $\delta_{mem} \in (0, 1)$.

6. Repeat 2-5 until convergence.

More details for a similar method (for $R = 1$ special case) can be found in [45]. This method was used for the theory fits in Fig. S5B,C,D, with $\tau_{max} = 120$, $d\tau = 0.05$, $N_{batch} = 10^4$, $N_{iter} = 400$, and $\delta_{mem} = 0.2$.

Spontaneous state with reservoir of limit cycles

In a special case of an ensemble with constant $D^{(\mu)} = D$ and $\mathbf{A}^{(\mu)}$ i.i.d. sampled as block-Haar matrices (see Appendix), the single-unit self-coupling kernel $G^Z(\tau)$ vanishes. Crucially, this means the currents $x(t)$ now have Gaussian statistics, and we can leverage closed-form Gaussian integrals in our calculations. Moreover, the kernel that determines input noise to each neuron $Q^Z(\tau)$ admits a closed-form expression in Fourier space without the need for an explicit average over tasks μ . In turn, $C^X(\omega)$ has a closed-form relationship with $C^\Phi(\omega)$ that is parameterized by the (constant) average gain $\langle \Phi' \rangle$ which depends on $C^X(\tau = 0)$. This can be translated into a second-order ODE (Eq. 150) and solved for as in [36], with the additional step of self-consistently determining the gain. Alternatively, we can directly solve in the Fourier domain by iteratively enforcing

$$C^x(\omega) = \frac{g_{eff}^2}{1 - (\gamma D \langle \Phi' \rangle)^2 + \omega^2} C^\phi(\omega), \quad (31)$$

where $g_{eff}^2 = \alpha R \langle D^2 \rangle_D$ (defined in Eq. (14)), over the whole spectrum. We iterate on the following process:

1. Start with an initial guess for $C^X(\omega)$.
2. Take the inverse Fourier transform to get $C^X(\tau)$.
3. Calculate $C^\Phi(\tau)$ in terms of $C^X(\tau)$ via Eq. (41) and $\langle \Phi' \rangle$ in terms of $C^X(\tau = 0)$ via Eq. (40).
4. Update $C^X(\omega)$ using Eq. (31) with memory parameter $\delta_{mem} \in (0, 1)$.
5. Repeat steps 2-4 until convergence.

We evaluated over a temporal window of $\tau_{max} = 200$, $d\tau = 0.05$ with $\delta_{mem} = 0.5$.

Task-selected state with reservoir of limit cycles

For notational simplicity, we assume there is one dominant oscillation, which we index by $*$. We first write down the dynamics for the two latent variables $z_r^{(*)}(t)$ of the dominant pattern with pattern strength $D^{(*)}$ and overlap $\mathbf{A}^{(*)}$, with intrinsic frequency $\omega^{(*)} = \tan(\theta^{(*)})$. By the assumption of dominance, $z_r^{(*)}(t) \sim \mathcal{O}(\sqrt{N})$, we can divide by a factor of $\sqrt{N}$, making the $\mathcal{O}(1)$ noise vanish for large N , and get

$$\frac{z_r^{(*)}(t)}{\sqrt{N}} = D^{(*)} \sum_{r'} A_{rr'}^{(*)} \int dt' S^\phi(t, t') \frac{z_{r'}^{(*)}(t')}{\sqrt{N}}, \quad (32)$$

which reduces in the case of Gaussian $x(t)$ to

$$(1 + \partial_t) \left(\frac{z_r^{(*)}(t)}{\sqrt{N}} \right) = D^{(*)} \langle \phi' \rangle \sum_{r'} A_{rr'}^{(*)} \left(\frac{z_{r'}^{(*)}(t)}{\sqrt{N}} \right), \quad (33)$$

because $S^\phi(t, t') = \Theta(t - t') \langle \Phi'(t) \rangle e^{-(t-t')}$. These dynamics are also consistent with the direct definition of the condensed patterns from the self-consistency conditions and applying Stein's lemma to Eq. (26), which the Gaussianity of $x(t)$ permits. Since Φ is saturating, this system has a steady-state oscillatory solution of a particular amplitude ρ . Assuming marginal stability of the oscillation enforces constant $\langle \Phi' \rangle = 1/(\gamma D^{(*)} \cos(\theta^{(*)}))$, which depends on both ρ and the noise variance—though we don't yet know how this is partitioned between them. For now, we write the steady-state solution for the dominant pattern in terms of ρ as

$$\frac{z_1^{(*)}(t)}{\sqrt{N}} = \rho \cos(\omega^{(*)}t) \quad \frac{z_2^{(*)}(t)}{\sqrt{N}} = \rho \sin(\omega^{(*)}t). \quad (34)$$

The dynamics for any one neuron $x_m(t)$ features noise drive from $\eta^x(t)$ sampled iid as a Gaussian process with stationary kernel $Q^x(\tau)$, determined like before as an average over the 2-point functions of the weak-task trajectories, plus deterministic drive from this dominant task with coupling $m^{(*,r)}$:

$$(1 + \partial_t) x_m(t) = \eta^x(t) + m^{(*,1)} \sqrt{N} \rho \cos(\omega^{(*)}t) + m^{(*,2)} \sqrt{N} \rho \sin(\omega^{(*)}t) \quad (35)$$

$$= \eta^x(t) + m_1 \rho \cos(\omega^{(*)}t) + m_2 \rho \sin(\omega^{(*)}t) \quad (36)$$

where we have defined $m^{(*,r)} \sqrt{N} = m_r \sim \mathcal{O}(1)$ as the order-1 set of couplings between this arbitrary unit and the condensed oscillation, which encode both amplitude and relative phase information. The P subscript in $x_m(t)$ indicates this dependence on $m_1, m_2 \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$.

As currently setup, this problem is reminiscent of [37], but with a few key differences:

- The noise is not due to an i.i.d. reservoir, but rather a reservoir of low-rank matrices with 2×2 -block Haar overlaps. This has an effect on the temporal shape of the noise statistics (the difference between Eq. 150 and the equation for an unstructured network).
- In [37], neural inputs feature phase heterogeneity across units but uniform amplitude. Our case has heterogeneity in both amplitude and phase. This changes how we write down our average Φ autocovariance function.
- The overall scaling of the (heterogeneous) amplitudes $\rho \|m\|$ is not a free parameter but rather self-consistently determined with $D^{(*)}$, the constant D on the uncondensed patterns, and current variance. This changes the axes on the bifurcation diagram, since $D^{(*)}$ is what we freely vary and not ρ , but once equilibrium ρ is determined, the effect on the dynamics is similar to that of the external current from [37].

We write $x_m(t) = x_m^0(t) + x_m^1(t)$ as a sum of a filtered deterministic part and a filtered noise part, the former having a closed-form expression:

$$x_m^0(t) = \frac{\rho}{\sqrt{1 + (\omega^{(*)})^2}} \left(m_1 \cos(\omega^{(*)}t - \theta_0) + m_2 \sin(\omega^{(*)}t - \theta_0) \right), \quad (37)$$

which has autocovariance $C^{x^0}(\tau) = \tilde{\rho}^2 \cos(\omega^{(*)}\tau)$ for $\tilde{\rho}^2 = \rho^2 / (1 + (\omega^{(*)})^2)$. The filtered noise part $x^1(t)$ has autocovariance with the same relation to $C^\Phi(\omega)$ as in the spontaneous case:

$$C^{x^1}(\omega) = \frac{g_{\text{eff}}^2}{1 + \omega^2 - (\gamma D \langle \phi' \rangle)^2} C^\Phi(\omega). \quad (38)$$

(See Appendix for a derivation of this relation.) If we can write an expression for $C^\Phi(\tau)$ in terms of the statistics of $x_m^0(t)$ and $x_m^1(t)$, then we close the equations. Since they are both Gaussian, their sum is Gaussian, and $C^x(\omega) = C^{x^0}(\omega) + C^{x^1}(\omega)$ determines $C^\Phi(\tau)$ in closed form (after taking the IFFT). That suggests the following iterative process for solving the DMFT:

1. For a given $D^{(*)}$ and $\omega^{(*)}$, first calculate the average gain $\langle \Phi' \rangle$ set by the marginality condition of the dominant oscillation.
2. Set an initial guess $C^{x^1}(\omega)$.
3. Take the IFFT to obtain $C^{x^1}(\tau)$.
4. Invert Eq. (40) to solve for the complementary ρ value that combines with $C^{x^1}(\tau = 0)$ to give the already determined $\langle \Phi' \rangle$ value.
5. Compute $C^x(\tau) = C^{x^0}(\tau) + C^{x^1}(\tau)$ in terms of our provisional $C^{x^1}(\tau)$ and ρ values.
6. Compute $C^\Phi(\tau)$ using Eq. (41).
7. Take the FFT of $C^\Phi(\tau)$ to obtain $C^\Phi(\omega)$.
8. Update estimate for $C^{x^1}(\omega)$ using $C^\Phi(\omega)$ via Eq. (38) with memory parameter $\delta_{\text{mem}} \in (0, 1)$.
9. Repeat steps 2-8 until convergence.

We evaluated over a temporal window of $\tau_{\text{max}} = 200$, $d\tau = 0.05$ with $\delta_{\text{mem}} = 0.5$. Additionally, for a given dominant frequency $\omega^{(*)}$ value, we adjusted the value to precisely match the nearest on-bin frequency value to avoid spectral smearing.

Phase boundaries

Computing $D_{\text{crit},1}^{(\mu)}$ is straightforward from Eq. (21) and only requires solving the spontaneous state DMFT and computing $\langle \Phi' \rangle$. To compute $D_{\text{crit},2}^{(\mu)}$, we used a midpoint solver to identify the value of $D^{(\mu)} \in (D_{\text{crit},1}^{(\mu)}, 1.7 \times D_{\text{crit},1}^{(\mu)})$ that results in a $C^x(\tau)$ curve with

$$\frac{C^x(0) - C^x(2\pi/\omega^{(*)})}{C^x(0)} \approx 0.02, \quad (39)$$

where 0.02 is our criterion for when the relative magnitude of fluctuations is effectively vanished.

Gaussian integrals

Any special case with Gaussian currents (e.g. the reservoir of oscillations) can take advantage of Gaussian integrals to simplify calculation. By our choice of nonlinearity $\phi(x) = \text{erf}(\sqrt{\pi}x/2)$, we can leverage certain closed-form expressions:

$$\langle \phi'(x_i(t)) \rangle_i = \int dy (2\pi)^{-1/2} e^{-y^2/2} \phi'(\sqrt{C^x(t,t)}y) = \sqrt{1 + (\pi/2)C^x(t,t)} \quad (40)$$

$$\phi(x_i(t))\phi(x_i(t+\tau))_i = \frac{2}{\pi} \sin^{-1} \left(\frac{C^x(t, t+\tau)}{2/\pi + C^x(t, t)} \right). \quad (41)$$

Whenever we consider Gaussian currents $x_i(t)$ in this work, they are always zero mean with possibly time-dependent variance $C^X(t, t)$, but more often we are in the stationary case $\tau = t - t'$ with $C^X(0) = C^X(t, t)$.

Dimensionality of task-selected state

See [45] for a detailed walkthrough of dimensionality calculations, which can in principle apply to the spontaneous ($\mathcal{O}(N)$ dimensionality) or task-selected ($\mathcal{O}(1)$ dimensionality) states. However, in practice we found poor numerical fits for dimensionality in the spontaneous state when overlaps were even modestly strong.

The dimensionality calculation simplifies tremendously for the task-selected state. We can leverage the duality between neuron-neuron covariance (used most directly to define the PR) and time-time covariance (which has identical spectrum to the neuron-neuron covariance) as first shown in [45]. For a given variable $a \in \{x, \phi\}$, its PR can be computed in terms of its autocovariance function $C^a(\tau)$ and four-point function $\Psi^a(\tau, \tau') = \frac{1}{N} \sum_{i,j=1}^N C_{ij}^a(\tau) C_{ij}^a(\tau')$

$$\text{PR}^a = \frac{C^a(0)^2}{\Psi^a(0, 0)}. \quad (42)$$

Leveraging the duality discussed above, we can alternatively calculate $\Psi^a(\tau, \tau')$ as

$$\Psi^a(\tau, \tau') = N \langle C^a(t, t') C^a(t + \tau, t' + \tau') \rangle_{t, t'}. \quad (43)$$

For $\tau = \tau' = 0$, we're essentially computing the average square of $C^a(t, t')$ over the whole (t, t') plane. The values near the $t = t'$ diagonal contribute negligibly compared to the whole plane, so we need only examine the asymptotic off-diagonal behavior when $t - t'$ is large. Any fluctuations $\delta C^a(t, t') = C^a(t, t') - C^a(t - t')$ about the stationarized two-point function are negligible in this case, since typical values over the plane are $\mathcal{O}(1)$ because of the dominant dynamics. Thus we can estimate $\Psi^a(0, 0)$ directly from the autocovariance function $C^a(\tau)$ as

$$\Psi^a(0, 0) = \frac{1}{T} \int_{\tau}^{\tau+T} d\tau' (C^a(\tau'))^2 \quad (44)$$

where the domain of integration depends on the nature of the dominant task. For a limit-cycle task, we need only integrate over one period of an oscillation to get the average behavior over the whole plane, as long as we start past when the amplitude decay is finished (Fig. S7 [45]). For a fixed point, we need only integrate over a brief interval when the $C^a(\tau)$ has reached its asymptotic value.

Simulating networks

To validate the theoretical calculations, we simulated networks as well. We varied N depending on how much we needed to reduce finite-size effects in a given context, we always used $dt = 0.05$ and RK4 forward integration. Parameters are in Table 1 [45] for all network simulations in the main and supplementary figures.

Panels	N	P or α	$D^{(\mu)}$	Notes on $A^{(\mu)}$ / sampling
Fig. 2A	2000	$P = 1$	$D = 2.2$	$\mathbf{A} = \begin{pmatrix} 0.8 & 0.4 \\ -0.4 & 0.8 \end{pmatrix}$ (limit cycle)
Fig. 2B	2000	$P = 1$	$D = 2$	$\mathbf{A} = \begin{pmatrix} 0.5 & 0.3 \\ 0.3 & 0.5 \end{pmatrix}$ (bistable)
Fig. 2C	2000	$P = 2$	$D^{(1)} = 2.2, D^{(2)} = 2$	$\mathbf{A}^{(1)}$ limit cycle, and $\mathbf{A}^{(2)}$ bistable
Fig. 2D,E,F	varied	varied	$D^{(1)} = 3, D^{(\mu>1)} = 2$	$\mathbf{A}^{(1)}$ limit cycle; $\mathbf{A}^{(2)}$ bistable; $A_{rr'}^{(\mu>2)}$ sampled i.i.d. $\mathcal{N}(0, 0.5)$ with rejection of non-positive-semidefinite draws
Fig. 2G	4000	$\alpha = 0.5$	$D^{(1)} = 2$ or $9; D^{(\mu>1)} = 2$	$\mathbf{A}^{(1)}$ limit cycle; $\mathbf{A}^{(2)}$ bistable; $A_{rr'}^{(\mu>2)}$ sampled i.i.d. $\mathcal{N}(0, 0.6)$ with rejection of non-positive-semidefinite draws
Fig. 3C,D,E	5000	$\alpha = 0.5$	$D^{(\mu)} = 2$ by default; $D^{(1)}$ modulated from $2 \rightarrow 3.75 \rightarrow 8$; $D^{(2)}$ from $2 \rightarrow 4.2 \rightarrow 8$	$\mathbf{A}^{(1)}$ limit cycle; $\mathbf{A}^{(2)}$ bistable; $A_{rr'}^{(\mu>2)}$ sampled i.i.d. $\mathcal{N}(0, 0.6)$ with rejection of non-positive-semidefinite draws
Fig. 4A	4000	varied	$D^{(\nu \neq \mu)} = 2$, and $D^{(\mu)}$ varied, where μ corresponds to limit cycle of frequency $\omega = 0.25$.	$\mathbf{A}^{(\nu \neq \mu)} = \gamma \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix}$ with $\theta \sim \text{Unif}(0, 2\pi)$ (see Appendix)
Fig. 4B	4000	$\alpha = 1$	$D^{(\mu)}$ varied on y-axis, and μ varied on x-axis by as the intrinsic frequency $\omega = \tan \theta$ for $\mathbf{A}^{(\mu)}$.	same as above
Fig. 5A,B,C	8000	$\alpha = 1$	same $D^{(\mu)}$ as Fig. 4 modulated from $2 \rightarrow 4.15 \rightarrow 5.3$ (black boxes),	same as Fig. 4
Fig. 5D	varied	$\alpha = 1$	same $D^{(\mu)}$ modulated	same as Fig. 4
Fig. 5E	varied	$\alpha = 1$	same as Fig. 4	
Fig. S2	2000	$P = 2$	$D^{(1)} = 2, D^{(2)} = 2.2$	same as Fig. 2C
Fig. S3A	2000	$P = 2$	$D^{(1)} = 3.5, D^{(2)} = 4$ (or swapped)	same as Fig. 2C
Fig. S3B	2000	$P = 2$	$D^{(1)} = 3.3, D^{(2)} = 3$ (or swapped)	same as Fig. 2C
Fig. S4	5000	$P = 2$	$D^{(\mu)} = 2$ or $2 \rightarrow 2.75 \rightarrow 3.5 \rightarrow 4.25$	same as Fig. 3
Fig. S5A-D	8000	$P = 2$	$D^{(\mu)} = 2$	same as Fig. 3
Fig. S6	10000	$P = 2$	$D^{(\nu \neq \mu)} = 2 \rightarrow 4 \rightarrow 8$	same as Fig. 4

Table 1. Summary of network parameters used in simulations across figures.

In all cases, $R = 2$.

Software availability

Core functions for simulating networks and solving the theory are available in <https://github.com/omarschall/mft-theory>, as well as an example notebook showcasing the fits in Fig. 5C.

Additional mathematical details

Connectivity and connection to SVD

Eq. (4) can be equivalently interpreted as the product of three matrices: an $N \times PR$ matrix whose columns comprise the left loading vectors $\mathbf{m}^{(\mu,r)}$, an $PR \times PR$ diagonal matrix whose diagonal elements comprise the task weights $D^{(\mu)}$ each repeated R times, and an $PR \times N$ matrix whose rows comprise the right loading vectors $\mathbf{n}^{(\mu,r)}$ (Fig. S1A). Thus the multi-task network (Eq. 4) is of the same form as a low-rank network with rank PR (Eq. 2).

Just as the Singular Value Decomposition (SVD) decomposes a matrix as a weighted sum of outer products, our model generates a matrix as a weighted sum of outer products (Fig. S1A). The loadings $\mathbf{m}^{(\mu,r)}$ ($\mathbf{n}^{(\mu,r)}$) play the role of the left (right) singular vectors, except they are not perfectly orthonormal but rather orthonormal up to $\mathcal{O}(1/\sqrt{N})$ noise, i.e. $\mathbf{m}^{(\mu,r)} \cdot \mathbf{m}^{(\nu,r')} = \delta_{\mu\nu}\delta_{rr'} + \mathcal{O}(1/\sqrt{N})$. The unrolled index $a = 1, \dots, M \times R$ corresponding to (μ, r) lists the “singular values” $D^{(a)}$ for each of the $P \times R$ modes (Fig. S1B). For a general matrix, these could be any non-negative values, but here we constrain them to be constant within R blocks, i.e. $D^{(\mu,r)} = D^{(\mu,r')} \equiv D^{(\mu)}$. For an arbitrary matrix, the overlaps between left and right singular vectors are generally unconstrained up to basic geometric consistency, and similarly we are free to generate connectivity through a nearly arbitrary pattern of overlaps $O_{ab} = \langle \mathbf{n}^{(a)} \cdot \mathbf{m}^{(b)} \rangle$ (Fig. S1C). We consider the special case of overlap matrices with an $R \times R$ block-diagonal structure (Fig. S1C), because keeping R fixed and allowing $P \rightarrow \infty$ makes the model mathematically tractable. However, the set of overlaps as specified by $A_{rr'}^{(\mu)}$ is arbitrary. (Up to the constraint that the overall $2R \times 2R$ loadings covariance matrix, containing the n - n block $\langle n_i^{(\mu,r)} n_i^{(\mu,r')} \rangle = N^{-1} \delta_{rr'}$ and m - m block $\langle m_i^{(\mu,r)} m_i^{(\mu,r')} \rangle = N^{-1} \delta_{rr'}$ as well as the n - m blocks given by $A_{rr'}^{(\mu)}$, must be positive definite.) This formulation is equivalent to the “random-mode model” with overlaps from [45], except the symmetries have a block-diagonal, rather than merely diagonal (equivalently $R = 1$), structure.

Full details of small- P case

The small- P case is equivalent to the single-task network case, because if P is small then $R \times P$ is also small. In the large- N limit, there is zero direct cross-talk between latent variables from distinct tasks, and Eq. (3) factorizes into self-contained systems of size R

$$z_r^{(\mu)}(t) = D^{(\mu)} \langle \phi'(t) \rangle \sum_{r'=1}^R A_{rr'}^{(\mu)} \left[L * z_{r'}^{(\mu)} \right] \quad (45)$$

up to their sharing of the gain factor via the Gaussian integral of Φ' with variance given by the squared norm of all the low-pass-filtered latent variables $\kappa_r^{(\mu)}(t) = \left[L * z_r^{(\mu)} \right]^2$:

$$\langle \phi'(t) \rangle = \int dy (2\pi)^{-1/2} e^{-y^2/2} \phi' \left(y \sqrt{\frac{1}{N} \sum_{\mu r} \left(\kappa_r^{(\mu)}(t) \right)^2} \right) \quad (46)$$

which can be evaluated by Eq. (40). Here $L(\tau) = \Theta(\tau)e^{-\tau}$ is a low-pass filter, a closed-form special case of the filter $F(\tau)$ that we can leverage with $G(\tau) = 0$, as in particular occurs when $\alpha = 0$ (i.e. small P) but in other cases as well.

Winner-take-all intuition

We present intuition for why there must always be a dominant task in this case. For a task μ in this setting to be dominant, its linearized dynamics must be unstable about the origin, so that it grows until $z_r^{(\mu)}(t)$ is large enough to influence the nonlinearity $\langle \Phi'(t) \rangle$ and self-stabilize. (This necessarily requires $z_r^{(\mu)}(t) \sim \mathcal{O}(\sqrt{N})$ because of the N scaling in the argument of Φ' .) The dynamics are marginally stable when

$$\langle \phi'(t) \rangle D^{(\mu)} \mathcal{S}^{(\mu)} = 1, \quad (47)$$

where $\mathcal{S}^{(\mu)} = \max[\Re(\lambda) \mid \lambda \text{ eigenvalue of } \mathbf{A}^{(\mu)}]$. If no task even achieves $D^{(\mu)} \mathcal{S}^{(\mu)} \geq 1$ with maximal gain $\langle \Phi'(t) \rangle = 1$, then the network is in the quiescent state, with all activity decaying to the origin rapidly. To avoid this trivial case, we assume at least one task is sufficiently strong to avoid the quiescent state. Consider the “strongest” such task, i.e. the one with highest value of $D^{(\mu)} \mathcal{S}^{(\mu)}$. To keep itself balanced, this task must set $\langle \Phi'(t) \rangle \sim D^{(\mu)} \mathcal{S}^{(\mu)}$ so that its dynamics are either marginally stable, or alternate phases of growth and decay cyclically (e.g. an elliptical limit cycle). Since all other tasks share this gain factor and have by definition weaker intrinsic strength, each will have a strictly weaker recurrent drive and therefore feature contracting dynamics. This is true even in cases where multiple tasks are fine-tuned to have the same strength—finite-sampling effects will always result in one task being marginally stronger than another. However, the more fine-tuned two tasks are to have the same strength, the more likely particular network realizations can feature long transient periods of activity where tasks appear “co-dominant,” until one eventually takes over.

What makes Gaussian loadings in particular lead to only one task dominating, is the fact that, in such networks, there is only one “population” in the sense of [29]. Thus there is only one gain factor, which equilibrates to a value that marginally stabilizes the strongest dynamics and suppresses all weaker ones, as described above. In a network with multiple populations, multiple gain factors could in principle separately equilibrate to values that marginally stabilize two different tasks’ dynamics—a trivial example would be adjoining two completely unconnected single-task networks that we call one network. However, nontrivial examples with overlap and interaction among populations also exist, which we leave to future work.

Fixed-point task dominance

In the main text, we showed an example where the limit cycle dynamics dominate the fixed-point dynamics. Here we swap the $D^{(\mu)}$ values such that the bistable dynamics dominate, and the activity in the task-1 subspace decays to the origin (Fig. S2A,B). The two matrices generating the limit-cycle and fixed-point dynamics, respectively,

$$\mathbf{A}^{(1)} = \begin{pmatrix} 0.8 & -0.4 \\ 0.4 & 0.8 \end{pmatrix}, \quad \mathbf{A}^{(2)} = \begin{pmatrix} 0.5 & 0.3 \\ 0.3 & 0.5 \end{pmatrix}, \quad (48)$$

were chosen such that they have identical spectral abscissa of $\mathcal{S}^{(1)} = \mathcal{S}^{(2)} = 0.8$, so that the dominant task is whichever simply has larger $D^{(\mu)}$.

Supplemental figures and tables




Figure S1. More detailed illustrations of connectivity structure. (A) Alternative interpretation of the sum in Fig. 1C as one large matrix product, with left loadings organized in one matrix as columns (blue) and right loadings in another as rows (yellow), with scalings as diagonal elements. The blue or yellow rectangular boundaries group the loading vectors into tasks, so $r = 1, 2$ indexes vectors within a block, while $\mu = 1, \dots, P$ indexes rectangular blocks. (B) Example connectivity strength spectra $D^{(\mu)}$, which can generally be dispersed or concentrated among tasks. (C) Illustration of the overall overlap matrix for all loading vectors, grouped by task by blue or yellow rectangular boundaries, as the matrix product of the left loading vectors as rows of one matrix (blue) with the right loading vectors as columns of one matrix (yellow). (D) Geometric interpretation of an entry of the overlap matrix as the degree of alignment of the associated left and right loading vectors. The $\mathbf{m}$'s (and $\mathbf{n}$'s) are shown to be orthogonal among themselves.

Figure S2. Dynamics for the same network as Fig. 2C but with $D^{(1)} = 2$, $D^{(2)} = 2.2$.

(A) Purple: activity in the task-1 subspace for a single run of the dynamics. Black: flow fields over many runs. (B) Red: activity in the task-2 subspace over two runs with different initial conditions. Black: average flow fields over many runs.




Figure S3. Similar simulations as in Fig. 2C, but for networks with different nonlinearities.

(A) A network using a diagonally shifted error function, by 1 in each direction, such that outputs are always positive, and it requires positive current to activate a neuron. (B) A network using a non-saturating nonlinearity, in particular a relu function (shifted leftward by 1). To ensure global stability, an additional term $-g_I/N$ is added to each J_{ij} , with $g_I = 2$.




Figure S4. Comparing modulation of task-specific input drive and task-specific connectivity, for task $\mu = 1$ in the same network shown in Fig. 3.

(A) Left: spontaneous-state activity in the task-1 subspace, featuring weak, noisy oscillations. Top: modulating the additive input drive $I_i^{(1,1)} = I_0 \sqrt{Nm_i^{(1,1)}}$ to each neuron i in the direction of the first task-1 left loading vector. Values of I_0 are (left to right) 0, 0.5, 1, 1.5. Bottom: modulating the strength of the task-1 connectivity component. Values of $D^{(1)}$ are (left to right) 2, 2.75, 3.5, 4.25.




Figure S5. Dynamics in the spontaneous state.

(A) Example activity traces for a network with $\alpha = 0.5$; $D^{(\mu)} = 2$ for all μ ; $\mathbf{A}^{(1)}, \mathbf{A}^{(2)}$ as in Fig. 2C, with $A_{rr}^{(\mu)} \sim \mathcal{N}(0, 0.6)$ i.i.d. with rejection sampling for positive-definiteness violations; and $N = 8000$. Top: latent activity for a limit-cycle task ($\mu = 1$) and a fixed-point task ($\mu = 2$) in the spontaneous state. Bottom: current and firing-rate traces for two example neurons in the network. (B) Lagged cross-covariance functions for $\tilde{z}^{(1)}(t)$, i.e. $\langle \tilde{z}_r^{(1)}(t) \tilde{z}_{r'}^{(1)}(t - \tau) \rangle_t$. Top left: $r = 1, r' = 1$. Top right: $r = 1, r' = 2$, etc. Solid: simulation. Dashed: theory. (D) Single-neuron autocovariance function $C^\Phi(\tau) = \langle \Phi_i(t) \Phi_i(t + \tau) \rangle_t$ (top) and impulse-response function $S^\Phi(\tau) = \langle \Phi_i(t) F(t) \rangle_t$ (bottom) and for the same network. Solid: simulation. Dashed: theory.




Figure S6.

(A) Participation ratio of firing rates for networks with variable $D^{(v)}$ values for all non-selected tasks $v \neq \mu$, as $D^{(\mu)}$ is varied relative to its critical point for task selection. Light dots: individual network simulations. Heavy dots: averages over networks. Dashed line: theory.





Figure S7. Cartoon showing how to read off the dimensionality of the task-selected state from the single-unit autocovariance function.

(A) Black: Theoretically calculated autocovariance function $C^\phi(\tau)$ from Fig. 3F (middle). The value at $\tau = 0$ at the average squared value over one period after full amplitude decay are the relevant quantities for computing the PR. The latter captures asymptotic behavior of the autocovariance function away from $\tau = 0$, which is needed for Eq. (44).



Data availability

The current manuscript is a computational study, so no data have been generated for this manuscript. Code is available on a repository linked in the main text.

Acknowledgements

We thank Manuel Beiran for discussions and for proposing the inhibition-stabilized model of Fig. S3. We thank Lea Duncker for feedback on the manuscript. OM, DGC, and ALK were supported by the Kavli Foundation and the Gatsby Charitable Foundation (grant no. GAT3708). ALK was also supported by the National Science Foundation (grant no. 2443158).

Appendix 1. Derivations of key mathematical results

Path integral for spontaneous state

This section takes advantage of mathematical techniques which are outlined in e.g. [68] or [69]. Similar derivations can be found in [45]. In this section, we derive the DMFT for the spontaneous state, assuming no dominant patterns. We start by writing the Martin-Siggia-Rose-Janssen-de Dominicis path integral for the partition function $Z(\mathbf{J})$ of the network dynamics:

$$Z(\mathbf{J}) \propto \int \mathcal{D}\mathbf{x} \mathcal{D}\hat{\mathbf{x}} \exp \left[i \sum_i \int dt \hat{x}_i(t) \left(T[x_i](t) - \sum_j J_{ij} \phi_j(t) \right) \right] \quad (49)$$

currently keeping the weights as shorthand J_{ij} and using $T = 1 + \partial_t$ as notational convenience. We now compute the average

$$\bar{Z} = \int \prod_{\mu,r,i} dm_i^{(\mu,r)} \prod_{\mu,r,i} dn_i^{(\mu,r)} \Pr \left(\left\{ m_i^{(\mu,r)} \right\}_{r,\mu,i}, \left\{ n_i^{(\mu,r)} \right\}_{r,\mu,i} \right) Z(\mathbf{J})$$

The probability density $\Pr(\cdot \cdot \cdot)$ for a given set of loadings (in general, vector-arrow notation imply a fixed-dimension vector like $\vec{l}_i^{(\mu)} \in \mathbb{R}^{2R}$, while bolded notation implies a vector with extensive (i.e. N -scaling) dimension, e.g. $\mathbf{m}^{(\mu,r)} \in \mathbb{R}^N$) for each neuron i and each task μ ,

$$\vec{l}_i^{(\mu)} = \left(m_i^{(\mu,1)}, \dots, m_i^{(\mu,R)}, n_i^{(\mu,1)}, \dots, n_i^{(\mu,R)} \right) \in \mathbb{R}^{2R}$$

is

$$\Pr \left(\left\{ m_i^{(\mu,r)} \right\}_{r,\mu,i}, \left\{ n_i^{(\mu,r)} \right\}_{r,\mu,i} \right) \propto \exp \left[-\frac{1}{2} \sum_i \sum_{\mu} \vec{l}_i^{(\mu)T} \mathcal{P}^{(\mu)} \vec{l}_i^{(\mu)} \right], \quad (50)$$

since each neuron's loadings are independent of the others', and loadings for different tasks are also independent. Within a task for a given neuron, the loadings are distributed as a multivariate Gaussian with precision matrix

$$\mathcal{P}^{(\mu)} = \text{inv} \left(\Sigma^{(\mu)} \right)$$

We define a set of $R \times P$ latent variables

$$z_r^{(\mu)}(t) = D^{(\mu)} \sum_i n_i^{(\mu,r)} \phi_i(t)$$

and their definition-enforcing conjugates $\hat{z}_r^{(\mu)}(t)$. Then our disorder-averaged partition function is

$$\bar{Z} \propto \int \mathcal{D}\mathbf{x} \mathcal{D}\hat{\mathbf{x}} \mathcal{D}\mathbf{z} \mathcal{D}\hat{\mathbf{z}} d\vec{\Gamma} \exp \left[-\frac{1}{2} \sum_{i\mu} \vec{l}_i^{(\mu)T} \mathcal{P}^{(\mu)} \vec{l}_i^{(\mu)} + i \sum_{\mu r} \hat{z}_r^{(\mu)} z_r^{(\mu)} + i \sum_i \hat{x}_i T[x_i] - i \sum_{i\mu r} m_i^{(\mu,r)} \hat{x}_i z_r^{(\mu)} - i \sum_{i\mu r} D^{(\mu)} n_i^{(\mu,r)} \phi_i \hat{z}_r^{(\mu)} \right]. \quad (51)$$

where suppressed time indices imply time integrals over a shared dummy time index. We can rewrite this as a quadratic in terms of $\vec{l}_i^{(\mu)}$ by defining the vector

$$\vec{b}_{i\mu} = -i \left(\hat{x}_i z_1^{(\mu)}, \dots, \hat{x}_i z_R^{(\mu)}, D^{(\mu)} \phi_i \hat{z}_1^{(\mu)}, \dots, D^{(\mu)} \phi_i \hat{z}_R^{(\mu)} \right) \in \mathbb{R}^{2R}$$

allowing us to perform a matrix version of completing the square for each $i\mu$:

$$\bar{Z} \propto \int \mathcal{D}\mathbf{x} \mathcal{D}\hat{\mathbf{x}} \mathcal{D}\mathbf{z} \mathcal{D}\hat{\mathbf{z}} d\vec{\Gamma} \exp \left[\sum_{i\mu} \left(-\frac{1}{2} \vec{l}_i^{(\mu)T} \mathcal{P}^{(\mu)} \vec{l}_i^{(\mu)} + \vec{l}_i^{(\mu)T} \vec{b}_{i\mu} \right) + i \sum_{\mu r} \hat{z}_r^{(\mu)} z_r^{(\mu)} + i \sum_i \hat{x}_i T[x_i] \right] \quad (52)$$

$$= \int \mathcal{D}\mathbf{x} \mathcal{D}\hat{\mathbf{x}} \mathcal{D}\mathbf{z} \mathcal{D}\hat{\mathbf{z}} d\vec{\Gamma} \exp \left[\sum_{i\mu} \left(-\frac{1}{2} \left(\vec{l}_i^{(\mu)} - \Sigma^{(\mu)} \right)^T \mathcal{P}^{(\mu)} \left(\vec{l}_i^{(\mu)} - \Sigma^{(\mu)} \right) + \frac{1}{2} \vec{b}_{i\mu}^T \Sigma^{(\mu)} \vec{b}_{i\mu} \right) \right] \quad (53)$$

$$+ i \sum_{\mu r} \hat{z}_r^{(\mu)} z_r^{(\mu)} + i \sum_i \hat{x}_i T[x_i] \right] \quad (54)$$

$$\propto \int \mathcal{D}\mathbf{x} \mathcal{D}\hat{\mathbf{x}} \mathcal{D}\mathbf{z} \mathcal{D}\hat{\mathbf{z}} \exp \left[i \sum_{\mu r} \hat{z}_r^{(\mu)} z_r^{(\mu)} + i \sum_i \hat{x}_i T[x_i] + \frac{1}{2} \sum_{i\mu} \vec{b}_{i\mu}^T \Sigma^{(\mu)} \vec{b}_{i\mu} \right]. \quad (55)$$

Recall that $\Sigma^{(\mu)}$ has a block diagonal structure with the identity along each of the two $R \times R$ principal blocks. Then we can rewrite this quadratic term as three terms, reintroducing the time indices briefly:

$$\vec{b}_{i\mu}^T \Sigma^{(\mu)} \vec{b}_{i\mu} = -\frac{1}{N} \sum_r \iint dt dt' \left[\hat{x}_i(t) \hat{x}_i(t') z_r^{(\mu)}(t) z_r^{(\mu)}(t') + \left(D^{(\mu)} \right)^2 \phi_i(t) \phi_i(t') \hat{z}_r^{(\mu)}(t) \hat{z}_r^{(\mu)}(t') \right] \quad (56)$$

$$- \frac{2}{N} \sum_{rr'} D^{(\mu)} A_{rr'}^{(\mu)} \iint dt dt' z_r^{(\mu)}(t) \hat{z}_{r'}^{(\mu)}(t') \hat{x}_i(t) \phi_i(t'). \quad (57)$$

The first two terms are easy to deal with and imply a reciprocal relationship between the two-point function of the latent variables $z_r^{(\mu)}(t)$ and how it appears in the noise driving $x_i(t)$, and the two-point function of the units $\phi_i(t)$ and how it appears in the noise driving $z_r^{(\mu)}(t)$. In the random-mode-model without symmetry [45], these are the only terms, and they lead to a straightforward solution. The third term, which arises due to the latent weight symmetries encoded in the off-diagonal block $A_{rr'}^{(\mu)}$ of $\Sigma^{(\mu)}$, is more difficult to deal with. We have to split the expression into the $t > t'$ and $t < t'$ half-planes. In the former, we can write $z_r^{(\mu)}(t) \hat{z}_r^{(\mu)}(t') = i \delta z_r^{(\mu)}(t) / \delta I_{r'}^{(\mu)}(t')$, and in the latter we can write $\phi_i(t') \hat{x}_i(t) = i \delta \phi_i(t') / \delta I_i(t)$, where $\delta / \delta I_i(t)$ represents the functional derivative with respect to external current. (In the flipped cases, each is zero by causality.)

Then we define a set of fields and their conjugates:

$$C^\phi(t, t') = \frac{1}{N} \sum_i \phi_i(t) \phi_i(t') \quad \text{conjugate: } \hat{Q}^z(t, t') \quad (58)$$

$$Q_r^z(t, t') = \frac{1}{N} \sum_\mu z_r^{(\mu)}(t) z_r^{(\mu)}(t') \quad \text{conjugate: } \hat{C}_r^\phi(t, t') \quad (59)$$

$$S^\phi(t, t') = \frac{1}{N} \sum_i \frac{\delta \phi_i(t)}{\delta I_i(t')} \quad \text{conjugate: } \hat{G}^z(t, t') \quad (60)$$

$$G_{rr'}^Z(t, t') = \frac{1}{N} \sum_\mu D^{(\mu)} A_{rr'}^{(\mu)T} \frac{\delta z_r^{(\mu)}(t)}{\delta I_{r'}^{(\mu)}(t')} \quad \text{conjugate: } \hat{S}_{rr'}^\phi(t, t') \quad (61)$$

Plugging in these definitions, we have

$$\bar{Z} \propto \int \mathcal{D}\mathbf{x} \mathcal{D}\hat{\mathbf{x}} \exp \left[i \sum_i \hat{x}_i T[x_i] - \frac{1}{2} \sum_i \sum_r \hat{x}_i Q_r^z \hat{x}_i - \frac{i}{2} \sum_i \sum_{rr'} \hat{x}_i G_{rr'}^Z \phi_i \right] \quad (62)$$

$$\begin{aligned} & \times \int \mathcal{D}\mathbf{z} \mathcal{D}\hat{\mathbf{z}} \exp \left[i \sum_{\mu r} \hat{z}_r^{(\mu)} z_r^{(\mu)} - \frac{1}{2} \sum_\mu \left(D^{(\mu)} \right)^2 \sum_r \hat{z}_r^{(\mu)} C^\phi \hat{z}_r^{(\mu)} \right. \\ & \left. - \frac{i}{2} \sum_{\mu rr'} D^{(\mu)} A_{rr'}^{(\mu)T} \hat{z}_r^{(\mu)} S^\phi z_r^{(\mu)} \right]. \end{aligned} \quad (63)$$

Note how the interaction term has a half in the neuron path integral and a half in the latent path integral, with kernels defined as population-averaged susceptibilities in the other picture. In other words, a given neuron interacts with itself by a weighted average of latent-latent impulse responses, and a given set of same-task latent variables interact among themselves by an average of single-neuron impulse responses.

Then, integrating over the definitions of these fields and decoupling noninteracting single-neuron and single-task path integrals, we obtain definition-enforcing terms and a global N scaling outside an action 1066 that is purely a function of the fields:

$$\bar{Z} \propto \int \mathcal{D}C^\phi \mathcal{D}\{\hat{C}_r^\phi\} \mathcal{D}S^\phi \mathcal{D}\{\hat{S}_{rr'}^\phi\} \mathcal{D}\{Q_r^z\} \mathcal{D}\hat{Q}^z \mathcal{D}\{G_{rr'}^Z\} \mathcal{D}\hat{G}^z \quad (64)$$

$$\times \exp \left(-NS \left[C^\phi, \{\hat{C}_r^\phi\}, S^\phi, \{\hat{S}_{rr'}^\phi\}, \{Q_r^z\}, \hat{Q}^z, \{G_{rr'}^Z\}, \hat{G}^z \right] \right) \quad (65)$$

with

$$S \left[C^\phi, \left\{ \hat{C}_r^\phi \right\}, S^\phi, \left\{ \hat{S}_{rr'}^\phi \right\}, \left\{ Q_r^z \right\}, \hat{Q}^z, \left\{ G_{rr'}^z \right\}, \hat{G}^z \right] \\ = -\frac{1}{2} C^\phi \hat{Q}^z - \frac{1}{2} \sum_r Q_r^z \hat{C}_r^\phi + \frac{1}{2} S^\phi \hat{G}^z + \frac{1}{2} \sum_{rr'} G_{rr'}^z \hat{S}_{rr'}^\phi \quad (66)$$

$$- \log W^x \left[\left\{ Q_r^z \right\}, Q^z, \left\{ G_{rr'}^z \right\}, G^z \right] \quad (67)$$

$$- \alpha \left\langle \log W_{D,A}^{\vec{z}} \left[C^\phi, \left\{ \hat{C}_r^\phi \right\}, S^\phi, \left\{ \hat{S}_{rr'}^\phi \right\} \right] \right\rangle_{D,A}. \quad (68)$$

Here,

$$W^x = \int \mathcal{D}x \mathcal{D}\hat{x} \exp \left[i \hat{x} T[x] - \frac{1}{2} \sum_r \hat{x} Q_r^z \hat{x} - \frac{1}{2} \phi \hat{Q}^z \phi \right. \\ \left. - \frac{i}{2} \sum_{rr'} \hat{x} G_{rr'}^z \phi - \frac{i}{2} \phi \hat{G}^z \hat{x} \right], \quad (69)$$

$$W_{D,A}^{\vec{z}} = \int \mathcal{D}\vec{z} \mathcal{D}\hat{z} \exp \left[i \sum_r \hat{z}_r z_r - \frac{1}{2} \sum_r z_r \hat{C}_r^\phi z_r - \frac{1}{2} D^2 \sum_r \hat{z}_r C^\phi \hat{z}_r \right. \quad (70)$$

$$\left. - \frac{i}{2} D \sum_{rr'} A_{rr'}^T \hat{z}_{r'} S^\phi z_r - \frac{i}{2} D \sum_{rr'} A_{rr'}^T z_r \hat{S}_{rr'}^\phi \hat{z}_{r'} \right]. \quad (71)$$

We differentiate with respect to each of the $R^2 + R + 2$ fields and their conjugates:

$$\frac{\delta S}{\delta C^\phi(t, t')} = -\hat{Q}^z(t, t') + \alpha \left\langle D^2 \sum_r \langle \hat{z}_r(t) \hat{z}_r(t') \rangle_{W_{D,A}^{\vec{z}}} \right\rangle_{D,A}, \quad (72)$$

$$\frac{\delta S}{\delta \hat{C}_r^\phi(t, t')} = -Q_r^z(t, t') + \alpha \left\langle \langle z_r(t) z_r(t') \rangle_{W_{D,A}^{\vec{z}}} \right\rangle_{D,A}, \quad (73)$$

$$\frac{\delta S}{\delta S^\phi(t, t')} = \hat{G}^z(t, t') + i\alpha \left\langle D \sum_{rr'} A_{rr'}^T \langle \hat{z}_{r'}(t) z_r(t') \rangle_{W_{D,A}^{\vec{z}}} \right\rangle_{D,A}, \quad (74)$$

$$\frac{\delta S}{\delta \hat{S}_{rr'}^\phi(t, t')} = G_{rr'}^z(t, t') + i\alpha \left\langle D A_{rr'}^T \langle z_r(t) \hat{z}_{r'}(t') \rangle_{W_{D,A}^{\vec{z}}} \right\rangle_{D,A}, \quad (75)$$

$$\frac{\delta S}{\delta Q_r^z(t, t')} = -\hat{C}_r^\phi(t, t') + \langle \hat{x}(t) \hat{x}(t') \rangle_{W^x}, \quad (76)$$

$$\frac{\delta S}{\delta \hat{Q}^z(t, t')} = -C^\phi(t, t') + \langle \phi(t) \phi(t') \rangle_{W^x}, \quad (77)$$

$$\frac{\delta S}{\delta G_{rr'}^z(t, t')} = \hat{S}_{rr'}^\phi(t, t') + i \langle \hat{x}(t) \phi(t') \rangle_{W^x}, \quad (78)$$

$$\frac{\delta S}{\delta \hat{G}^z(t, t')} = S^\phi(t, t') + i \langle \phi(t) \hat{x}(t') \rangle_{W^x}. \quad (79)$$

Solving the saddle-point system yields

$$\hat{Q}_\star^z(t, t') = \hat{C}_{\star, r}^\phi(t, t') = 0, \quad (80)$$

$$S_{\star, rr'}^\phi(t, t') = S_\star^\phi(t', t) = \left\langle \frac{\delta \phi(t')}{\delta I(t)} \right\rangle_\star, \quad (81)$$

$$C_\star^\phi(t, t') = \langle \phi(t) \phi(t') \rangle_\star, \quad (82)$$

$$Q_{\star, r}^z(t, t') = \alpha \left\langle \left\langle z_r(t) z_r(t') \right\rangle_{W_{D, A}^z} \right\rangle_{D, A}, \quad (83)$$

$$G_{\star, rr'}^z(t, t') = \alpha \left\langle D A_{rr'}^T \left\langle \frac{\delta z_r(t)}{\delta I_{r'}(t')} \right\rangle_{W_{D, A}^z} \right\rangle_{D, A}, \quad (84)$$

$$\hat{G}_\star^z(t, t') = \alpha \left\langle D \sum_{rr'} A_{rr'}^T \left\langle \frac{\delta z_r(t')}{\delta I_{r'}(t)} \right\rangle_{W_{D, A}^z} \right\rangle_{D, A}, \quad (85)$$

so that

$$\hat{G}_\star^z(t, t') = \sum_{rr'} G_{\star, rr'}^z(t', t). \quad (86)$$

Then, going back into the single-site and latent processes, we obtain a Φ process driven by noise term $\eta^x \sim \mathcal{GP}(0, Q_\star^z)$ where $Q_\star^z = \sum_r Q_{\star, r}^z$, plus self-interactions via self-coupling kernel $\hat{G}_\star^z$. However, the process for the latent variables is more complex and involves R^2 self-couplings. Altogether, we have

$$T[x](t) = \eta^x(t) + \int_{-\infty}^t dt' \hat{G}_\star^z(t, t') \phi(x(t')) \quad \eta^x \sim \mathcal{GP}(0, Q_\star^z(t, t')) \quad (87)$$

$$\begin{aligned} z_r(t) &= \eta^{z_r}(t) + D \sum_{r'=1}^R A_{rr'} \int_{-\infty}^t dt' S_\star^\phi(t, t') z_{r'}(t') \quad \eta^{z_r} \\ &\sim \mathcal{GP}(0, D^2 C_\star^\phi(t, t')) \end{aligned} \quad (88)$$

From now on, we drop the $\star$ notation and assume any kernel as written is the saddle-point solution. Further, we assume that all kernels depend only on $t-t'$, allowing us to leverage the convolution theorem. In Fourier space, we solve for $z_r(\omega)$ and get

$$z_r(\omega) = \eta^{z_r}(\omega) + D \sum_{r'} A_{rr'} S^\phi(\omega) z_{r'}(\omega), \quad (89)$$

$$\Rightarrow \vec{z}(\omega) = (\mathbf{I} - DS^\phi(\omega)\mathbf{A})^{-1} \vec{\eta}^z(\omega) \equiv \mathbf{M}(\omega) \vec{\eta}^z(\omega), \quad (90)$$

$$\begin{aligned} \Rightarrow Q_r^z(\omega) &= \alpha \left\langle \left\langle z_r(\omega) z_r(\omega)^* \right\rangle_{W_{D,A}^z} \right\rangle_{D,A} \\ &= \alpha \left\langle \left\langle \sum_{r'} M_{rr'}(\omega) \eta^{z_{r'}}(\omega) \sum_{r''} M_{rr''}(\omega)^* \eta^{z_{r''}}(\omega)^* \right\rangle_{W_{D,A}^z} \right\rangle_{D,A} \end{aligned} \quad (91)$$

$$= \alpha \left\langle D^2 [\mathbf{M}(\omega) \mathbf{M}^\dagger(\omega)]_{rr} \right\rangle_{D,A} C^\phi(\omega). \quad (92)$$

Then using Furutsu–Novikov we get an expression for $\hat{G}^z(\omega)$. We summarize together with $Q^z(\omega)$ as

$$Q^z(\omega) = \alpha \left\langle D^2 \|\mathbf{M}(\omega)\|_{\mathcal{F}}^2 \right\rangle_{D,A} C^\phi(\omega), \quad (93)$$

$$\hat{G}^z(\omega) = \alpha \langle D \text{Tr}(\mathbf{A} \mathbf{M}(\omega)) \rangle_{D,A}. \quad (94)$$

Path integral for task-selected state

We pick up at Eq. (55) and explicitly separate the μ index into a set of P^* dominant and $P - P^* \approx P$ weak tasks. We perform the disorder integration only over the weak tasks' couplings and keep the average partition function depending on the particular realization of the dominant couplings:

$$\begin{aligned} &\bar{Z} \left[\left\{ m_i^{(\mu^*, r)}, n_i^{(\mu^*, r)} \right\} \right] \\ &\propto \int \mathcal{D}\mathbf{x} \mathcal{D}\hat{\mathbf{x}} \mathcal{D}\mathbf{z} \mathcal{D}\hat{\mathbf{z}} \exp \left[i \sum_i \hat{x}_i T[x_i] - i \sqrt{N} \sum_{i\mu^*r} m_i^{(\mu^*, r)} \hat{x}_i z_r^{(\mu^*)} \right] \end{aligned} \quad (95)$$

$$+ i \sum_{\mu r} \hat{z}_r^{(\mu)} z_r^{(\mu)} - i \sqrt{N} \sum_{i\mu^*r} D^{(\mu^*)} n_i^{(\mu^*, r)} \phi_i \hat{z}_r^{(\mu^*)} \quad (96)$$

$$- \frac{1}{2N} \sum_{\mu ir} \hat{x}_i z_r^{(\mu)} \hat{x}_i z_r^{(\mu)} \quad (97)$$

$$- \frac{1}{2N} \sum_{\mu ir} \left(D^{(\mu)} \right)^2 \phi_i \hat{z}_r^{(\mu)} \phi_i \hat{z}_r^{(\mu)} \quad (98)$$

$$- \frac{1}{N} \sum_{\mu ir r'} D^{(\mu)} A_{rr'}^{(\mu)T} z_r^{(\mu)} \hat{x}_i \hat{z}_{r'}^{(\mu)} \phi_i \quad (99)$$

$$C^\phi(t, t') = \frac{1}{N} \sum_i \phi_i(t) \phi_i(t') \quad \text{conjugate : } \hat{Q}^z(t, t') \quad (100)$$

$$Q_r^z(t, t') = \frac{1}{N} \sum_\mu z_r^{(\mu)}(t) z_r^{(\mu)}(t') \quad \text{conjugate : } \hat{C}_r^\phi(t, t') \quad (101)$$

$$S^\phi(t, t') = \frac{1}{N} \sum_i \frac{\delta \phi_i(t)}{\delta I_i(t')} \quad \text{conjugate : } \hat{G}^z(t, t') \quad (102)$$

$$G_{rr'}^Z(t, t') = \frac{1}{N} \sum_\mu D^{(\mu)} A_{rr'}^{(\mu)T} \frac{\delta z_r^{(\mu)}(t)}{\delta I_{r'}^{(\mu)}(t')} \quad \text{conjugate : } \hat{S}_{rr'}^\phi(t, t') \quad (103)$$

$$z_r^{(\mu^*)}(t) = \frac{1}{\sqrt{N}} D^{(\mu^*)} \sum_i n_i^{(\mu^*, r)} \phi_i(t) \quad \text{conjugate : } \hat{z}_r^{(\mu^*)}(t) \quad (104)$$

For the dominant tasks, their activity is no longer a dynamical or stochastic variable but rather a trajectory that is subject to the saddle-point condition of the overall dynamics. Note that we have rescaled the usual definition of the task latents by a factor of $1/\sqrt{N}$ so that all order parameters are $\mathcal{O}(1)$. We insert definitions of order parameters when possible and separate into integrals over $\mathbf{x}$'s and $\mathbf{z}$'s.

$$\begin{aligned} \bar{Z}^x \left[\left\{ m_i^{(\mu^*, r)}, n_i^{(\mu^*, r)} \right\} \right] &\propto \int \mathcal{D}\mathbf{x} \mathcal{D}\hat{\mathbf{x}} \exp \left[i \sum_i \hat{x}_i T[x_i] \right. \\ &\quad \left. - i\sqrt{N} \sum_{i\mu^*r} m_i^{(\mu^*, r)} \hat{x}_i z_r^{(\mu^*)} - i\sqrt{N} \sum_{i\mu^*r} D^{(\mu)} n_i^{(\mu^*, r)} \phi_i \hat{z}_r^{(\mu^*)} \right] \end{aligned} \quad (105)$$

$$- \frac{1}{2} \sum_i \sum_r \hat{x}_i Q_r^z \hat{x}_i - i \sum_i \sum_{rr'} \hat{x}_i G_{rr'}^z \phi_i \quad (106)$$

(107)

$$\begin{aligned} \bar{Z}^z &\propto \int \mathcal{D}\mathbf{z} \mathcal{D}\hat{\mathbf{z}} \exp \left[i \sum_{\mu r} \hat{z}_r^{(\mu)} z_r^{(\mu)} - \frac{1}{2} \sum_\mu \left(D^{(\mu)} \right)^2 \sum_r \hat{z}_r^{(\mu)} C^\phi \hat{z}_r^{(\mu)} \right. \\ &\quad \left. - i \sum_\mu D^{(\mu)} \sum_{rr'} \hat{z}_r^{(\mu)} S^\phi z_{r'}^{(\mu)} \right] \end{aligned} \quad (108)$$

(109)

Then we can construct one path integral over the fields as a function of the realizations of the normalized dominant couplings $c = \left\{ \bar{m}_i^{(\mu^*, r)}, \bar{n}_i^{(\mu^*, r)} \right\}$:

$$\begin{aligned} \bar{Z} \left[\left\{ m_i^{(\mu^*, r)}, n_i^{(\mu^*, r)} \right\} \right] &\propto \int \mathcal{D}C^\phi \mathcal{D} \left\{ \hat{C}_r^\phi \right\} \mathcal{D}S^\phi \mathcal{D} \left\{ \hat{S}_{rr'}^\phi \right\} \mathcal{D} \left\{ Q_r^z \right\} \mathcal{D} \hat{Q}^z \mathcal{D} \\ &\quad \left\{ G_{rr'}^z \right\} \mathcal{D} \hat{G}^z \mathcal{D} \left\{ z_r^{(\mu^*)} \right\} \mathcal{D} \left\{ \hat{z}_r^{(\mu^*)} \right\} \end{aligned} \quad (110)$$

$$\begin{aligned} &\times \exp \left(\right. \\ &\quad \left. - NS \left[C^\phi, \left\{ \hat{C}_r^\phi \right\}, S^\phi, \left\{ \hat{S}_{rr'}^\phi \right\}, \left\{ Q_r^z \right\}, \hat{Q}^z, \left\{ G_{rr'}^z \right\}, \hat{G}^z, \left\{ z_r^{(\mu^*)} \right\}, \left\{ \hat{z}_r^{(\mu^*)} \right\} \right] \right) \end{aligned} \quad (111)$$

with

$$S[\text{all fields}] = -\frac{1}{2}C^\phi \hat{Q}^z - \frac{1}{2} \sum_r Q_r^z \hat{C}_r^\phi \quad (112)$$

$$+\frac{1}{2}S^\phi \hat{G}^z + \frac{1}{2} \sum_{rr'} G_{rr'}^z \hat{S}_{rr'}^\phi + i \sum_{\mu^* r} \hat{z}_r^{(\mu^*)} z_r^{(\mu^*)} \quad (113)$$

$$-\left\langle \log W_c^x \left[\{Q_r^z\}, \hat{Q}^z, \{G_{rr'}^z\}, \hat{G}^z, \{\hat{z}_r^{(\mu^*)}\}, \{\hat{z}_r^{(\mu^*)}\} \right] \right\rangle_c \quad (114)$$

$$-\alpha \left\langle \log W_{D,A}^{\bar{z}} \left[C^\phi, \{\hat{C}_r^\phi\}, S^\phi, \{\hat{S}_{rr'}^\phi\} \right] \right\rangle_{D,A}. \quad (115)$$

Here,

$$W_c^x = \int \mathcal{D}x \mathcal{D}\hat{x} \exp \left[i\hat{x}T[x] - i\hat{x} \sum_{\mu^* r} \bar{m}^{(\mu^*,r)} z_r^{(\mu^*)} - i\phi \sum_{\mu^* r} D^{(\mu^*)} \bar{n}^{(\mu^*,r)} \hat{z}_r^{(\mu^*)} \right] \quad (116)$$

$$-\frac{1}{2} \sum_r \hat{x} Q_r^z \hat{x} - \frac{1}{2} \phi \hat{Q}^z \phi - \frac{i}{2} \sum_{rr'} \hat{x} G_{rr'}^z \phi - \frac{i}{2} \phi \hat{G}^z \hat{x} \quad (117)$$

$$W_{D,A}^{\bar{z}} = \int \mathcal{D}\bar{z} \mathcal{D}\hat{z} \exp \left[i \sum_r \hat{z}_r z_r - \frac{1}{2} \sum_r z_r \hat{C}_r^\phi z_r - \frac{1}{2} D^2 \sum_r \hat{z}_r C^\phi \hat{z}_r \right] \quad (118)$$

$$-\frac{i}{2} D \sum_{rr'} A_{rr'}^T \hat{z}_{r'} S^\phi z_r - \frac{i}{2} D \sum_{rr'} A_{rr'}^T z_r \hat{S}_{rr'}^\phi \hat{z}_{r'} \quad (119)$$

where we have defined $\bar{m}^{(\mu^*,r)} = \sqrt{N} m^{(\mu^*,r)}$ (and similarly $\bar{n}^{(\mu^*,r)}$) as the $\mathcal{O}(1)$ normalized couplings of an arbitrary unit to the dominant patterns. Then differentiating with respect to each order parameter, we get

$$\frac{\delta S}{\delta C^\phi(t,t')} = -\hat{Q}^z(t,t') + \alpha \left\langle D^2 \sum_r \langle \hat{z}_r(t) \hat{z}_r(t') \rangle_{W_{D,A}^{\bar{z}}} \right\rangle_{D,A}, \quad (120)$$

$$\frac{\delta S}{\delta \hat{C}_r^\phi(t,t')} = -Q_r^z(t,t') + \alpha \left\langle \langle z_r(t) z_r(t') \rangle_{W_{D,A}^{\bar{z}}} \right\rangle_{D,A}, \quad (121)$$

$$\frac{\delta S}{\delta S^\phi(t,t')} = \hat{G}^z(t,t') + i\alpha \left\langle D \sum_{rr'} A_{rr'}^T \langle \hat{z}_{r'}(t) z_r(t') \rangle_{W_{D,A}^{\bar{z}}} \right\rangle_{D,A}, \quad (122)$$

$$\frac{\delta S}{\delta \hat{S}_{rr'}^\phi(t,t')} = G_{rr'}^z(t,t') + i\alpha \left\langle D A_{rr'}^T \langle z_r(t) \hat{z}_{r'}(t') \rangle_{W_{D,A}^{\bar{z}}} \right\rangle_{D,A}, \quad (123)$$

$$\frac{\delta S}{\delta Q_r^z(t, t')} = -\hat{C}_r^\phi(t, t') + \left\langle \left\langle \hat{x}(t) \hat{x}(t') \right\rangle_{W_c^x} \right\rangle_c, \quad (124)$$

$$\frac{\delta S}{\delta \hat{Q}^z(t, t')} = -C^\phi(t, t') + \left\langle \left\langle \phi(t) \phi(t') \right\rangle_{W_c^x} \right\rangle_c \quad (125)$$

$$\frac{\delta S}{\delta G_{rr'}^z(t, t')} = \hat{S}_{rr'}^\phi(t, t') + i \left\langle \left\langle \hat{x}(t) \phi(t') \right\rangle_{W_c^x} \right\rangle_c \quad (126)$$

$$\frac{\delta S}{\delta \hat{G}^z(t, t')} = S^\phi(t, t') + i \left\langle \left\langle \phi(t) \hat{x}(t') \right\rangle_{W_c^x} \right\rangle_c \quad (127)$$

$$\frac{\delta S}{\delta z_r^{(\mu^*)}(t)} = i \hat{z}_r^{(\mu^*)}(t) + i \left\langle \left\langle \hat{x}(t) \right\rangle_{W_c^x} \bar{m}^{(\mu^*, r)} \right\rangle_c \quad (128)$$

$$\frac{\delta S}{\delta \hat{z}_r^{(\mu^*)}(t)} = i z_r^{(\mu^*)}(t) + i \left\langle \left\langle \phi(t) \right\rangle_{W_c^x} D^{(\mu^*)} \bar{n}^{(\mu^*, r)} \right\rangle_c \quad (129)$$

We get the same self-consistency conditions as before, except now the single-site properties are averages over heterogeneous units, where the particular sample of c $= \{\bar{m}^{(\mu^*, r)}, \bar{n}^{(\mu^*, r)}\}_{\mu^*, r}$ for that unit determines its dynamics and properties, plus two additional conditions:

$$\hat{z}_r^{(\mu^*)}(t) = 0 \quad (130)$$

$$z_r^{(\mu^*)}(t) = D^{(\mu^*)} \left\langle \bar{n}^{(\mu^*, r)} \left\langle \phi(t) \right\rangle_{W_c^x} \right\rangle_c. \quad (131)$$

Then we have our final mean-field theory

$$T[x](t) = \eta^x(t) + \sum_{\mu^* r} \bar{m}^{(\mu^*, r)} z_r^{(\mu^*)}(t) + \int_{-\infty}^t dt' G^z(t, t') \phi(t') \quad \eta^x(t) \quad (132)$$

$$\sim \mathcal{GP}(0, Q^z(t, t'))$$

$$z_r(t) = \eta^{z_r}(t) + D \sum_{r'} A_{rr'} \int_{-\infty}^t dt' S^\phi(t, t') z_{r'}(t') \quad \eta^{z_r}(t) \quad (133)$$

$$\sim \mathcal{GP}(0, D^2 C^\phi(t, t'))$$

In the main text a factor of $\sqrt{N}$ gets traded between $\bar{m}^{(\mu^*, r)}$ and the non-normalized dominant pattern $z_r^{(\mu^*)}(t)$.

DMFT simplifications for ensemble of limit-cycle tasks

Consider the following process for generating $R \times R$ matrices, where R must be even. We start by sampling $R/2$ independent 2×2 rotation matrices via $\theta_l \sim U(0, 2\pi)$ for $l = 1, \dots, R/2$:

$$\mathbf{R}_l = \begin{pmatrix} \cos \theta_l & -\sin \theta_l \\ \sin \theta_l & \cos \theta_l \end{pmatrix}, \quad (134)$$

and arrange them in a block-diagonal structure. We then take a random $R \times R$ Haar orthogonal matrix $\mathbf{B}$ and intermix the entries, but preserving the traces of all matrix powers. We finally scale the result by $\gamma < 1$ (e.g. $\gamma = 0.99$) to ensure that $\Sigma = \begin{pmatrix} \mathbf{I} & \mathbf{A} \\ \mathbf{A}^T & \mathbf{I} \end{pmatrix}$ is positive definite. Then

$$\mathbf{A} = \gamma \mathbf{B} \left(\bigoplus_{l=1}^{R/2} \mathbf{R}_l \right) \mathbf{B}^T \quad (135)$$

has a few desirable properties as a process for randomly generating a matrix. (For simplicity, the main text considers just $R = 2$, $\mathbf{B} = \mathbf{I} \Rightarrow \mathbf{A} = \gamma \mathbf{R}_1$, although its useful properties extend to this more general case.) First, for any $k \geq 1$, we have

$$\begin{aligned} \langle \text{Tr}(\mathbf{A}^k) \rangle &= \left\langle \text{Tr} \left(\gamma^k \bigoplus_{l=1}^{R/2} \mathbf{R}_l^k \right) \right\rangle \\ &= \gamma^k \sum_{l=1}^{R/2} \langle \text{Tr}(\mathbf{R}_l^k) \rangle = \gamma^k \sum_{l=1}^{R/2} \langle \text{Tr}(\mathbf{R}_l) \rangle = 0, \end{aligned} \quad (136)$$

using the fact that $\mathbf{R}_l^k$ has the same distribution as $\mathbf{R}_l$. The γ factor conveniently also reduces the SEM when we effectively “measure” the mean trace matrix power for finite samples of such matrices in simulated networks, which can be large for large powers k .

For independently sampled D , this property eliminates the single-unit self-coupling by setting $G^Z(t, t') = 0$, but without resorting to a trivial distribution of $\mathbf{A} = \mathbf{0}$ as is necessary and sufficient in the $R = 1$ case. We also have a closed-form equation for $Q^Z(\omega)$, which reads

$$Q^Z(\omega) = \alpha \langle D^2 \|\mathbf{M}(\omega)\|_{\mathcal{F}}^2 \rangle_{D, \mathbf{A}} C^\phi(\omega). \quad (137)$$

where $\mathbf{M}(\omega) = (\mathbf{I} - DS^\phi(\omega)\mathbf{A})^{-1}$. Under these assumptions about the distribution of $\mathbf{A}$, we can calculate $Q^Z(\omega)$. First, we note that the similarity transformation pops out of the entire $\mathbf{M}(\omega)$ matrix:

$$\begin{aligned} \mathbf{M}(\omega) &= (\mathbf{I} - DS^\phi \mathbf{A})^{-1} = \left(\mathbf{B} \mathbf{B}^T - \gamma DS^\phi \mathbf{B} \bigoplus_{l=1}^{R/2} \mathbf{R}_l \mathbf{B}^T \right)^{-1} \\ &= \left(\mathbf{B} \left(\mathbf{I} - \gamma DS^\phi(\omega) \bigoplus_{l=1}^{R/2} \mathbf{R}_l \right) \mathbf{B}^T \right)^{-1} \end{aligned} \quad (138)$$

$$\begin{aligned} &= \mathbf{B} \left(\mathbf{I} - \gamma DS^\phi(\omega) \bigoplus_{l=1}^{R/2} \mathbf{R}_l \right)^{-1} \mathbf{B}^T \Rightarrow \|\mathbf{M}(\omega)\|_{\mathcal{F}} \\ &= \left\| \left(\mathbf{I} - \gamma DS^\phi(\omega) \bigoplus_{l=1}^{R/2} \mathbf{R}_l \right)^{-1} \right\|_{\mathcal{F}}, \end{aligned} \quad (139)$$

since the Frobenius norm is invariant under orthogonal similarity transformations. Next we apply the matrix inversion on each block separately:

$$\left(\mathbf{I} - \gamma DS^\phi(\omega) \bigoplus_{l=1}^{R/2} \mathbf{R}_l \right)^{-1} = \bigoplus_{l=1}^{R/2} (\mathbf{I}_{2 \times 2} - \gamma DS^\phi(\omega) \mathbf{R}_l)^{-1}. \quad (140)$$

Each component matrix that gets inverted $\mathbf{I}_{2 \times 2} - \gamma DS^\phi(\omega) \mathbf{R}_l$ has eigenvalues $1 - \gamma DS^\phi(\omega) e^{\pm i\theta_l}$, so their inverses have eigenvalues $1 / (1 - \gamma DS^\phi(\omega) e^{\pm i\theta_l})$. Since this matrix is clearly normal, its Frobenius norm squared is simply the sum of the eigenvalues' squared norms. Putting it all together and assuming D is sampled independently from $\mathbf{A}$, we have

$$\begin{aligned} & \langle D^2 \|\mathbf{M}(\omega)\|_{\mathcal{F}}^2 \rangle_{D, \mathbf{A}} \\ &= \left\langle \sum_{l=1}^{R/2} \left(\frac{D^2}{|1 - \gamma DS^\phi(\omega) e^{i\theta_l}|^2} + \frac{D^2}{|1 - \gamma DS^\phi(\omega) e^{-i\theta_l}|^2} \right) \right\rangle_{D, \{\theta_l\}} \end{aligned} \quad (141)$$

$$= R \left\langle \frac{D^2}{|1 - \gamma DS^\phi(\omega) e^{i\theta}|^2} \right\rangle_{D, \theta} = R \left\langle D^2 \frac{1}{2\pi} \int_0^{2\pi} d\theta \frac{1}{|1 - \gamma DS^\phi(\omega) e^{i\theta}|^2} \right\rangle_D \quad (142)$$

$$= R \left\langle \frac{D^2}{1 - |\gamma DS^\phi(\omega)|^2} \right\rangle_D. \quad (143)$$

The last line requires assuming $|\gamma DS^\phi(\omega)| < 1$, which is a fairly safe assumption for a distribution of D with no outliers. Then, in Fourier space, the statistics of the effective neuron noise follow

$$Q^z(\omega) = \alpha R \left\langle \frac{D^2}{1 - |\gamma DS^\phi(\omega)|^2} \right\rangle_D C^\phi(\omega). \quad (144)$$

With these simplifications, we have now a bipartite mean-field theory

$$(1 + \partial_t) x(t) = \eta^x(t) \quad (145)$$

$$z_r(t) = \eta^{z_r}(t) + D \langle \phi' \rangle \sum_{r'} A_{rr'} \int_{-\infty}^t dt' e^{-(t-t')} z_{r'}(t') \quad (146)$$

because we know $S^\Phi(\tau) = \Theta(\tau) \langle \Phi' \rangle e^{-\tau}$ and $G^Z(\tau) = 0$. We can differentiate the second equation and get

$$(1 + \partial_t) z_r(t) = \dot{\eta}^{z_r}(t) + \eta^{z_r}(t) + D \langle \phi' \rangle \sum_{r'} A_{rr'} z_{r'}(t), \quad (147)$$

exploiting the fact that the derivative of $S^\Phi(\tau)$ is negative of itself, producing a leak term. With this representation of the weak-task latent variable dynamics, it is much more straightforward to determine the stability condition on D : it reduces to whether the largest real part among eigenvalues of $D \langle \Phi' \rangle \mathbf{A}$ is below or above 1.

Let's examine the single-site picture and solve for the temporal dynamics of the $x(t)$ autocovariance:

$$(1 + \omega^2) C^x(\omega) = \alpha R \left\langle \frac{D^2}{1 - |\gamma D S^\phi(\omega)|^2} \right\rangle_D C^\phi(\omega) \quad (148)$$

$$\Rightarrow C^x(\omega) = \alpha R \left\langle \frac{D^2}{1 - (\gamma D \langle \phi' \rangle)^2 + \omega^2} \right\rangle_D C^\phi(\omega). \quad (149)$$

Now let's make yet a further assumption, that D is in fact a constant such that $g_{\text{eff}}^2 = \alpha R D^2$ rather than following some distribution—this is a common choice for the non-selected tasks. Then we have

$$\partial_\tau^2 C^x(\tau) = \left(1 - \frac{\gamma^2 \nu}{\alpha R}\right) C^x(\tau) - g_{\text{eff}}^2 C^\phi(\omega), \quad (150)$$

which is exactly the second-order ODE describing $C^x(\tau)$ for an unstructured network [36], except the $C^x(\tau)$ term is scaled by $1 - \frac{\gamma^2 \nu}{\alpha R}$ of 1, where $\nu \equiv (g_{\text{eff}} \langle \phi' \rangle)^2$ can be thought of as the effective slope for an equivalent linear network [43, 70].

The average gain depends exclusively on $C^x(0)$ (via a Gaussian integral). However, to solve for $C^x(0)$ by the Sompolsky potential-energy trick in this case requires knowing the average gain already, because of its appearance in ν . Therefore we solve for $C^x(0)$ for a provisional value of $\langle \phi' \rangle$, which we iteratively update as we solve for $C^x(0)$ again with our new value of $\langle \phi' \rangle$, until it stops changing. This value of $\langle \phi' \rangle$ then determines the stability for an arbitrary task, based on the spectral abscissa of $D \langle \phi' \rangle A$. However, the purely spectral methods we describe earlier are more effective for small α .

Additional information

Funding

Funder	Grant reference number	Author
Kavli Foundation (Kavli)	GAT3708	Ashok Litwin-Kumar Owen Marschall David G Clark
National Science Foundation (NSF)	2443158	Ashok Litwin-Kumar

Author ORCID iDs

Owen Marschall:  <https://orcid.org/0000-0001-5348-3278>

References

- [1] Markowitz Jeffrey E, Gillis Winthrop F, Jay Maya, Wood Jeffrey, Harris Ryley W, Cieszkowski Robert, Scott Rebecca, Brann David, Koveal Dorothy, Kula Tomasz, *et al.* (2023) Spontaneous behaviour is structured by reinforcement without explicit reward. *Nature* **614**:108-117 <https://doi.org/10.1038/s41586-022-05611-2> | PubMed
- [2] Amematsro Elom A, Trautmann Eric M, Marshall Najja J, Abbott LF, Shadlen Michael N, Wolpert Daniel M, Churchland Mark M (2025) Motor cortex flexibly deploys a high-dimensional repertoire of subskills. *bioRxiv* <https://doi.org/10.1101/2025.09.07.674717> | PubMed

- [3] Mante Valerio, Sussillo David, Shenoy Krishna V, Newsome William T (2013) Context-dependent computation by recurrent dynamics in prefrontal cortex. *nature* **503**:78-84 <https://doi.org/10.1038/nature12742> | PubMed
- [4] Yang Guangyu Robert, Joglekar Madhura R, Song H Francis, Newsome William T, Wang Xiao-Jing (2019) Task representations in neural networks trained to perform many cognitive tasks. *Nature neuroscience* **22**:297-306 <https://doi.org/10.1038/s41593-018-0310-2> | PubMed
- [5] Pagan Marino, Tang Vincent D, Aoi Mikio C, Pillow Jonathan W, Mante Valerio, Sussillo David, Brody Carlos D (2025) Individual variability of neural computations underlying flexible decisions. *Nature* **639**:421-429 <https://doi.org/10.1038/s41586-024-08433-6> | PubMed
- [6] Shan Haozhe, Minni Sun, Duncker Lea (2025) Separating the what and how of compositional computation to enable reuse and continual learning. *arXiv* <https://doi.org/10.52202/085713-3038>
- [7] Tafazoli Sina, Bouchacourt Flora M, Ardalan Adel, Markov Nikola T, Uchimura Motoaki, Mattar Marcelo G, Daw Nathaniel D, Buschman Timothy J (2025) Building compositional tasks with shared neural subspaces. *Nature* **650**:164-172 <https://doi.org/10.1038/s41586-025-09805-2> | PubMed
- [8] Vyas Saurabh, Golub Matthew D, Sussillo David, Shenoy Krishna V (2020) Computation through neural population dynamics. *Annual review of neuroscience* **43**:249-275 <https://doi.org/10.1146/annurev-neuro-092619-094115> | PubMed
- [9] Sadtler Patrick T, Quick Kristin M, Golub Matthew D, Chase Steven M, Ryu Stephen I, Tyler-Kabara Elizabeth C, Yu Byron M, Batista Aaron P (2014) Neural constraints on learning. *Nature* **512**:423-426 <https://doi.org/10.1038/nature13665> | PubMed
- [10] Golub Matthew D, Sadtler Patrick T, Oby Emily R, Quick Kristin M, Ryu Stephen I, Tyler-Kabara Elizabeth C, Batista Aaron P, Chase Steven M, Yu Byron M (2018) Learning by neural reassociation. *Nature neuroscience* **21**:607-616 <https://doi.org/10.1038/s41593-018-0095-3> | PubMed
- [11] Rigotti Mattia, Barak Omri, Warden Melissa R, Wang Xiao-Jing, Daw Nathaniel D, Miller Earl K, Fusi Stefano (2013) The importance of mixed selectivity in complex cognitive tasks. *Nature* **497**:585-590 <https://doi.org/10.1038/nature12160> | PubMed
- [12] Stokes Mark G, Kusunoki Makoto, Sigala Natasha, Nili Hamed, Gaffan David, Duncan John (2013) Dynamic coding for cognitive control in prefrontal cortex. *Neuron* **78**:364-375 <https://doi.org/10.1016/j.neuron.2013.01.039> | PubMed
- [13] Lee Julie J, Krumin Michael, Harris Kenneth D, Carandini Matteo (2022) Task specificity in mouse parietal cortex. *Neuron* **110**:2961-2969 <https://doi.org/10.1016/j.neuron.2022.07.017> | PubMed
- [14] Driscoll Laura N, Shenoy Krishna, Sussillo David (2024) Flexible multitask computation in recurrent networks utilizes shared dynamical motifs. *Nature Neuroscience* **27**:1349-1363 <https://doi.org/10.1038/s41593-024-01668-6> | PubMed
- [15] Lanore Frederic, Cayco-Gajic N Alex, Gurnani Harsha, Coyle Diccon, Silver R Angus (2021) Cerebellar granule cell axons support high-dimensional representations. *Nature neuroscience* **24**:1142-1150 <https://doi.org/10.1038/s41593-021-00873-x> | PubMed
- [16] Natraj Nikhilesh, Silversmith Daniel B, Chang Edward F, Ganguly Karunesh (2022) Compartmentalized dynamics within a common multi-area mesoscale manifold represent a repertoire of human hand movements. *Neuron* **110**:154-174 <https://doi.org/10.1016/j.neuron.2021.10.002> | PubMed
- [17] Sabatini David A, Kaufman Matthew T (2024) Reach-dependent reorientation of rotational dynamics in motor cortex. *Nature Communications* **15**:7007 <https://doi.org/10.1038/s41467-024-51308-7> | PubMed
- [18] Borgognon Simon, Macellari Nicolò, Hickey Alexandra M, Perich Matthew G, Javaheri Houman, Ornelas-Kobayashi Rafael, Delacombaz Maude, Hitz Christopher, Fallegger Florian, Lacour Stéphanie P, et al. (2025) Regional specialization of movement encoding across the primate sensorimotor cortex. *Nature Communications* **16**:5729 <https://doi.org/10.1038/s41467-025-61172-8> | PubMed

- [19] **Stringer Carsen**, Pachitariu Marius, Steinmetz Nicholas, Reddy Charu Bai, Carandini Matteo, Harris Kenneth D (2019) Spontaneous behaviors drive multidimensional, brainwide activity. *Science* **364**:eaav7893 <https://doi.org/10.1126/science.aav7893> | [PubMed](#)
- [20] **Manley Jason**, Lu Sihao, Barber Kevin, Demas Jeffrey, Kim Hyewon, Meyer David, Traub Francisca Martínez, Vaziri Alipasha (2024) Simultaneous, cortex-wide dynamics of up to 1 million neurons reveal unbounded scaling of dimensionality with neuron number. *Neuron* **112**:1694-1709 <https://doi.org/10.1016/j.neuron.2024.02.011> | [PubMed](#)
- [21] **Churchland Mark M**, Cunningham John P, Kaufman Matthew T, Foster Justin D, Nuyujukian Paul, Ryu Stephen I, Shenoy Krishna V (2012) Neural population dynamics during reaching. *Nature* **487**:51-56 <https://doi.org/10.1038/nature11129> | [PubMed](#)
- [22] **Gao Peiran**, Trautmann Eric, Yu Byron, Santhanam Gopal, Ryu Stephen, Shenoy Krishna, Ganguli Surya (2017) A theory of multineuronal dimensionality, dynamics and measurement. *BioRxiv* 214262 <https://doi.org/10.1101/214262>
- [23] **Syeda Atika**, Zhong Lin, Tung Renee, Long Will, Pachitariu Marius, Stringer Carsen (2024) Facemap: a framework for modeling neural activity based on orofacial tracking. *Nature neuroscience* **27**:187-195 <https://doi.org/10.1038/s41593-023-01490-6> | [PubMed](#)
- [24] **Van Vreeswijk Carl**, Sompolinsky Haim (1996) Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* **274**:1724-1726 <https://doi.org/10.1126/science.274.5293.1724> | [PubMed](#)
- [25] **Renart Alfonso**, Rocha Jaime De La, Bartho Peter, Hollender Liad, Parga Néstor, Reyes Alex, Harris Kenneth D (2010) The asynchronous state in cortical circuits. *science* **327**:587-590 <https://doi.org/10.1126/science.1179850> | [PubMed](#)
- [26] **Duncker Lea**, Driscoll Laura, Shenoy Krishna V, Sahani Maneesh, Sussillo David (2020) Organizing recurrent network dynamics by task-computation to enable continual learning. In: *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)*. pp. 14387-14397
- [27] **Langdon Christopher**, Engel Tatiana A (2025) Latent circuit inference from heterogeneous neural responses during cognitive tasks. *Nature neuroscience* **28**:665-675 <https://doi.org/10.1038/s41593-025-01869-7> | [PubMed](#)
- [28] **Mastrogiuseppe Francesca**, Ostojic Srdjan (2018) Linking connectivity, dynamics, and computations in low-rank recurrent neural networks. *Neuron* **99**:609-623 <https://doi.org/10.1016/j.neuron.2018.07.003> | [PubMed](#)
- [29] **Beiran Manuel**, Dubreuil Alexis, Valente Adrian, Mastrogiuseppe Francesca, Ostojic Srdjan (2021) Shaping dynamics with multiple populations in low-rank recurrent networks. *Neural Computation* **33**:1572-1615 https://doi.org/10.1162/neco_a_01381 | [PubMed](#)
- [30] **Cheung Brian**, Terekhov Alexander, Chen Yubei, Agrawal Pulkit, Olshausen Bruno (2019) Superposition of many models into one. *Advances in neural information processing systems* **32** <https://doi.org/10.48550/arxiv.1902.05522>
- [31] **Guo Zengcai V**, Inagaki Hidehiko K, Daie Kayvon, Druckmann Shaul, Gerfen Charles R, Svoboda Karel (2017) Maintenance of persistent activity in a frontal thalamocortical loop. *Nature* **545**:181-186 <https://doi.org/10.1038/nature22324> | [PubMed](#)
- [32] **Sauerbrei Britton A**, Guo Jian-Zhong, Cohen Jeremy D, Mischiati Matteo, Guo Wendy, Kabra Mayank, Verma Nakul, Mensh Brett, Branson Kristin, Hantman Adam W (2020) Cortical pattern generation during dexterous movement is input-driven. *Nature* **577**:386-391 <https://doi.org/10.1038/s41586-019-1869-9> | [PubMed](#)
- [33] **Schuessler Friedrich**, Dubreuil Alexis, Mastrogiuseppe Francesca, Ostojic Srdjan, Barak Omri (2020) Dynamics of random recurrent networks with correlated low-rank structure. *Physical Review Research* **2**:013111 <https://doi.org/10.1103/physrevresearch.2.013111>
- [34] **Dubreuil Alexis**, Valente Adrian, Beiran Manuel, Mastrogiuseppe Francesca, Ostojic Srdjan (2022) The role of population structure in computations through neural dynamics. *Nature neuroscience* **25**:783-794 <https://doi.org/10.1038/s41593-022-01088-4> | [PubMed](#)

- [35] Ahmadian Yashar, Rubin Daniel B, Miller Kenneth D (2013) Analysis of the stabilized supralinear network. *Neural computation* **25**:1994-2037 https://doi.org/10.1162/neco_a_00472 | PubMed
- [36] Sompolinsky Haim, Crisanti Andrea, Sommers Hans-Jurgen (1988) Chaos in random neural networks. *Physical review letters* **61**:259 <https://doi.org/10.1103/physrevlett.61.259> | PubMed
- [37] Rajan Kanaka, Abbott LF, Sompolinsky Haim (2010) Stimulus-dependent suppression of chaos in recurrent neural networks. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics* **82**:011903 <https://doi.org/10.1103/physreve.82.011903> | PubMed
- [38] Pereira-Obilinovic Ulises, Aljadeff Johnatan, Brunel Nicolas (2023) Forgetting leads to chaos in attractor networks. *Physical Review X* **13**:011009 <https://doi.org/10.1103/physrevx.13.011009>
- [39] Clark David G, Abbott LF (2024) Theory of coupled neuronal-synaptic dynamics. *Physical Review X* **14**:021001 <https://doi.org/10.1103/physrevx.14.021001>
- [40] Clark David G, Beiran Manuel (2025) Structure of activity in multiregion recurrent neural networks. *Proceedings of the National Academy of Sciences* **122**:e2404039122 <https://doi.org/10.1073/pnas.2404039122> | PubMed
- [41] Litwin-Kumar Ashok, Harris Kameron Decker, Axel Richard, Sompolinsky Haim, Abbott LF (2017) Optimal degrees of synaptic connectivity. *Neuron* **93**:1153-1164 <https://doi.org/10.1016/j.neuron.2017.01.030> | PubMed
- [42] Schaffer Evan S, Mishra Neeli, Whiteway Matthew R, Li Wenze, Vancura Michelle B, Freedman Jason, Patel Kripa B, Voleti Venkatakaushik, Paninski Liam, Hillman Elizabeth MC, *et al.* (2023) The spatial and temporal structure of neural activity across the fly brain. *Nature communications* **14**:5572 <https://doi.org/10.1038/s41467-023-41261-2> | PubMed
- [43] Clark David G, Abbott LF, Litwin-Kumar Ashok (2023) Dimension of activity in random neural networks. *Physical Review Letters* **131**:118401 <https://doi.org/10.1103/physrevlett.131.118401> | PubMed
- [44] Morales-Gregorio Aitor, Kurth Anno C, Ito Junji, Kleinhohn Alexander, Barthélemy Frédéric V, Brochier Thomas, Grün Sonja, van Albada Sacha J (2024) Neural manifolds in v1 change with top-down signals from v4 targeting the foveal region. *Cell Reports* **43** <https://doi.org/10.1016/j.celrep.2024.114371> | PubMed
- [45] Clark David G, Marschall Owen, van Meegen Alexander, Litwin-Kumar Ashok (2025) Connectivity structure and dynamics of nonlinear recurrent neural networks. *Physical Review X* **15**:041019 <https://doi.org/10.1103/2jt7-c8cq> | PubMed
- [46] Schneidman Elad, Berry Michael J, Segev Ronen, Bialek William (2006) Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature* **440**:1007-1012 <https://doi.org/10.1038/nature04701> | PubMed
- [47] Panas Dagmara, Amin Hayder, Maccione Alessandro, Muthmann Oliver, van Rossum Mark, Berdondini Luca, Hennig Matthias H (2015) Sloppiness in spontaneously active neuronal networks. *Journal of Neuroscience* **35**:8480-8492 <https://doi.org/10.1523/jneurosci.4421-14.2015> | PubMed
- [48] Cueva Christopher J, Saez Alex, Marcos Encarni, Genovesio Aldo, Jazayeri Mehrdad, Romo Ranulfo, Salzman C Daniel, Shadlen Michael N, Fusi Stefano (2020) Low-dimensional dynamics for working memory and time encoding. *Proceedings of the National Academy of Sciences* **117**:23021-23032 <https://doi.org/10.1073/pnas.1915984117> | PubMed
- [49] Tian Gengshuo John, Zhu Ou, Shirhatti Vinay, Greenspon Charles M, Downey John E, Freedman David J, Doiron Brent (2024) Neuronal firing rate diversity lowers the dimension of population covariability. *bioRxiv* <https://doi.org/10.1101/2024.08.30.610535> | PubMed
- [50] Logiaco Laureline, Abbott LF, Escola Sean (2021) Thalamic control of cortical dynamics in a model of flexible motor sequencing. *Cell reports* **35** <https://doi.org/10.1016/j.celrep.2021.109090> | PubMed
- [51] Mongillo Gianluigi, Barak Omri, Tsodyks Misha (2008) Synaptic theory of working memory. *Science* **319**:1543-1546 <https://doi.org/10.1126/science.1150769> | PubMed

- [52] **Marder Eve** (2012) Neuromodulation of neuronal circuits: back to the future. *Neuron* **76**:1-11 <https://doi.org/10.1016/j.neuron.2012.09.010> | PubMed
- [53] **Tannenbaum Neta Ravid**, Burak Yoram (2016) Shaping neural circuits by high order synaptic interactions. *PLOS Computational Biology* **12**:e1005056 <https://doi.org/10.1371/journal.pcbi.1005056> | PubMed
- [54] **Stroud Jake P**, Porter Mason A, Hennequin Guillaume, Vogels Tim P (2018) Motor primitives in space and time via targeted gain modulation in cortical networks. *Nature neuroscience* **21**:1774-1783 <https://doi.org/10.1038/s41593-018-0276-0> | PubMed
- [55] **Gillett Maxwell**, Pereira Ulises, Brunel Nicolas (2020) Characteristics of sequential activity in networks with temporally asymmetric hebbian learning. *Proceedings of the National Academy of Sciences* **117**:29948-29958 <https://doi.org/10.1073/pnas.1918674117> | PubMed
- [56] **Costacurta Julia**, Bhandarkar Shaunak, Zoltowski David, Linderman Scott (2024) Structured flexibility in recurrent neural networks via neuromodulation. *Advances in Neural Information Processing Systems* **37**:1954-1972 <https://doi.org/10.52202/079017-0062> | PubMed
- [57] **Gillett Maxwell**, Brunel Nicolas (2024) Dynamic control of sequential retrieval speed in networks with heterogeneous learning rules. *eLife* **12**:RP88805 <https://doi.org/10.7554/eLife.88805> | PubMed
- [58] **Battaglia Francesco P**, Treves Alessandro (1998) Attractor neural networks storing multiple space representations: a model for hippocampal place fields. *Physical Review E* **58**:7738 <https://doi.org/10.1103/physreve.58.7738>
- [59] **Pereira-Obilinovic Ulises**, Daie Kayvon, Chen Susu, Svoboda Karel, Darshan Ran (2025) Neural dynamics outside task-coding dimensions drive decision trajectories through transient amplification. *bioRxiv* <https://doi.org/10.1101/2025.11.20.689599> | PubMed
- [60] **Goudar Vishwa**, Peysakhovich Barbara, Freedman David J, Buffalo Elizabeth A, Wang Xiao-Jing (2023) Schema formation in a neural population subspace underlies learning-to-learn in flexible sensorimotor problem-solving. *Nature Neuroscience* **26**:879-890 <https://doi.org/10.1038/s41593-023-01293-9> | PubMed
- [61] **Wakhloo Albert J**, Clark David G, Abbott LF (2025) Associative synaptic plasticity creates dynamic persistent activity. *bioRxiv* <https://doi.org/10.1101/2025.08.20.671131> | PubMed
- [62] **Legenstein Robert**, Maass Wolfgang (2007) Edge of chaos and prediction of computational performance for neural circuit models. *Neural networks* **20**:323-334 <https://doi.org/10.1016/j.neunet.2007.04.017> | PubMed
- [63] **Avitan Lilach**, Stringer Carsen (2022) Not so spontaneous: Multi-dimensional representations of behaviors and context in sensory areas. *Neuron* **110**:3064-3075 <https://doi.org/10.1016/j.neuron.2022.06.019> | PubMed
- [64] **Musall Simon**, Kaufman Matthew T, Juavinett Ashley L, Gluf Steven, Churchland Anne K (2019) Single-trial neural dynamics are dominated by richly varied movements. *Nature neuroscience* **22**:1677-1686 <https://doi.org/10.1038/s41593-019-0502-4> | PubMed
- [65] **Lin Sherry**, Gillis Winthrop F, Weinreb Caleb, Zeine Ayman, Jones Samuel C, Robinson Emma M, Markowitz Jeffrey, Datta Sandeep Robert (2024) Characterizing the structure of mouse behavior using motion sequencing. *Nature Protocols* **19**:3242-3291 <https://doi.org/10.1038/s41596-024-01015-w> | PubMed
- [66] **Montemurro Marcelo A**, Tamarit Francisco A, Stariolo Daniel A, Cannas Sergio A (2000) Out-of-equilibrium dynamics of the hopfield model in its spin-glass phase. *Physical Review E* **62**:5721 <https://doi.org/10.1103/physreve.62.5721> | PubMed
- [67] **Clark David G** (2025) Transient dynamics of associative memory models. *arXiv* <https://doi.org/10.48550/arxiv.2506.05303> | PubMed
- [68] **Crisanti A**, Sompolinsky H (2018) Path integral approach to random neural networks. *Physical Review E* **98**:062120 <https://doi.org/10.1103/physreve.98.062120>

- [69] Helias Moritz, Dahmen David (2020) Functional formulation of stochastic differential equations. In: *Statistical Field Theory for Neural Networks* Springer. pp. 57-67 https://doi.org/10.1007/978-3-030-46444-8_7
- [70] Shen Xuanyu, Hu Yu (2025) Covariance spectrum in nonlinear recurrent neural networks. *arXiv* <https://doi.org/10.48550/arxiv.2508.05288>

Peer reviews

Reviewer #1 (Public review):

Summary:

Marschall et al. develop a theoretical framework for analyzing multi-task dynamics in nonlinear recurrent neural networks (RNNs). In the RNN model, recurrent connectivity is a linear superposition of multiple non-overlapping low-rank components, each corresponding to a separate "task". Each task is an autonomous dynamical system that does not incorporate external inputs. The "multi-task computation" setting examines whether multiple tasks (dynamical systems) can operate concurrently or how a network can switch between tasks.

Within this framework, the authors show that when connectivity consists of two low-rank components implementing two tasks (a limit cycle and a bistable attractor), the network exhibits winner-takes-all dynamics, resulting in only one task being active and the other one suppressed. A task consistently dominates this competition when the magnitude of its low-rank connectivity ("task strength") exceeds that of the other task. When one task is dominant and many other tasks with weaker low-rank connectivity are present, increasing the number of tasks destabilizes the dominant-task dynamics, leading to chaotic fluctuations in network activity. Supported by the dynamical mean-field theory analysis, the authors show that as the dominant-task strength increases, the network transitions through three dynamical regimes: chaotic spontaneous activity, chaotic task-selected dynamics, and non-chaotic task-selected dynamics. The theoretical analysis additionally predicts how latent task-related dynamics manifest in single-neuron activity and how the dimensionality of population activity changes across the three dynamical regimes.

The results are interesting, the analyses and simulations are rigorous, and the text is clear and easy to follow. Overall, this study is a significant and timely contribution to the literature on low-rank RNNs, an influential model class for low-dimensional neural dynamics in computational neuroscience.

Main comments:

(1) In this modeling framework, only one dominant task can be selected while all other tasks are suppressed. In contrast, several previous studies constructed RNNs (either through gradient-descent optimization or reservoir computing) that simultaneously generate outputs for multiple tasks across the corresponding task-specific readouts. Of course, what counts as a task is arbitrary, and one could view the dynamics of a reservoir network as implementing a single high-dimensional "task" with multiple readouts. Nevertheless, it would be helpful to explicitly clarify the distinction and similarities between the current modeling framework and networks that simultaneously solve multiple tasks.

(2) The role of external inputs in task selection appears to be underdeveloped. It is only briefly examined in Fig. S4, with the conclusion that external inputs aligned with the task subspace cannot enable selection of the desired task. However, previous multi-task RNN models (e.g., optimized through gradient descent) are clearly able to switch across many tasks using external inputs. In these models, external inputs modulate RNN activity along specific directions, shaped through gradient descent, to select the relevant task for each input. In

contrast, this study only considers inputs aligned with the m -direction (left loading vector in the low-rank connectivity for a task). Such input cannot enhance the corresponding task's activity through recurrent amplification (Fig. S4). Yet, low-rank RNN theory predicts that inputs aligned with the n -direction (task's right loading vector) are selectively amplified by the recurrent dynamics. Why inputs aligned with the n -direction were not examined? More broadly, focusing only on inputs aligned with either m - or n -direction appears too narrow, as trained multi-task RNNs indicate that task selection through external inputs is possible, but may require input directions different from m and n .

(3) It is unclear what the reason is for the task interference: overlaps between task loading vectors for different tasks or other nonlinear effects? This issue is especially prominent in the analysis of task capacity (Fig. 2D and Fig. 4A). The text states that the loading vectors are independent across tasks, i.e. there is zero expected overlap between loading vectors of different tasks (line 126). For a network of N neurons, a rank- R task requires $2R$ loading vectors. Under the assumption of independence, the maximum number of tasks is $P = N/(2R)$. Then $\alpha = P/N$ can be at most $1/(2R)$. In the simulations, it is chosen $R=2$, such that α can be at most $1/4$ for the loading vectors to remain independent. However, the range of α reaches up to 1 in Fig. 2D and Fig. 4A, suggesting that loading vectors are no longer independent across tasks for larger α . Is the linear dependence between loading vectors (i.e. overlap across tasks) the reason for the task-1 component norm to drop sharply around $\alpha=1/4$ in Fig. 2D? More broadly, is it possible to isolate the contribution of overlap in task loading vectors versus other nonlinear effects?

(4) In the current version of the paper, it is difficult to understand how the analysis based on the dynamical mean-field theory in Methods explains the key observations from the numerical simulations. For example, the mechanism underlying the transition from the spontaneous state to chaotic task-selected state, and to the non-chaotic task-selected state remain opaque. Further, it would be helpful to specify which section in Methods is being referred to at each mention throughout the main text. Finally, Fig. 7 and Eqs. (11-13) clearly state the input sources that drive single unit activity, non-dominant task latent states and dominant task latent state. The decomposition is potentially very informative, but its implications are not discussed in sufficient detail. It would be helpful to provide an intuitive explanation of how contributions of these input sources evolve as connectivity strength of one task increases, and which of them eventually leads to the loss of stability of the previous network state.

(5) The text states that the results can be easily extended to include task-specific inputs and outputs (lines 117-119). However, such an extension does not seem to be straightforward and requires additional explanations. The dynamical mean-field theory analyses here are stationary and describe the steady-state of network dynamics. In contrast, common input-output tasks typically involve transient dynamics, in which time-dependent external inputs keep changing the RNN flow field, and steady-state is never reached. Under these transient conditions, it is unclear whether the same conclusions apply. For example, if a network is at a low-activity baseline when a task-input begins to drive activity in the corresponding task subspace, it is unclear whether the activity in other task-subspaces would grow sufficiently fast to cause interference, or whether such interference would not be observed. Thus, a more detail analyses are necessary to support the extension of the results to input-driven transient tasks beyond autonomous dynamical systems.

(6) When task strength is the same for all tasks, what determines which task will win the competition? Is it frozen noise in connectivity such that one task always wins, or do initial conditions determine which task wins, based on which task's activity grows faster?

(7) Does the theory require the activity of all neurons to operate in the saturating part of nonlinearity? For example, the text states "increased activity reduces the gain factor $\langle \Phi'(t) \rangle$ " (line 169). This statement is only true when $\Phi' < 0$. For sigmoid nonlinearity used in the paper,

$\Phi' > 0$ when firing rate is small. If a substantial fraction of neurons in the network is near the rest state, would the theory still apply? Similarly, this statement does not hold for ReLU nonlinearity, and it is unclear how the theory applies to ReLU networks in Fig. S3. The paper states that the results are not specific to the choice of nonlinearity (lines 176-178). However, the dynamics being studied (bistability and limit cycles) both operate on the saturating part of the nonlinearity. Could the authors clarify the assumptions on the nonlinearity for the theory to apply?

(8) Is it possible to interpret the results in Fig. 4? What does this dependence on the overlap matrix mean? Is there an intuition for this particular dependence, or is it just an observation without general interpretation?

(9) The results in Fig. 5E appear underdeveloped and somewhat arbitrary. It is unclear how the dependence of dimensionality on the recording time would change as a function of time spent in a task. If this time is long, then the curve grows slowly and total dimensionality is high. If this time is very short, then the network may not have sufficient time for all task variables to grow sufficiently large to contribute significantly to the total variance. Thus, the curve may grow faster and the total variance may saturate at a lower value. Hence, it is unclear whether there will be always a qualitative difference from the spontaneous activity curve. Furthermore, since only one of two curves is measured, what quantitative criteria should be used to determine whether it is consistent with task switching or spontaneous state?

(10) On line 368: "For sufficiently large number of tasks, the dimensionality associated with sequential task selection can greatly exceed that of the spontaneous state (Fig. 5E inset)" - it seems that Fig. 5E inset shows the opposite that the dimensionality of spontaneous state can saturate at a very high value for large N , exceeding the dimension of task-switching network in the main plot. Although it is hard to say, since the inset has many lines with only two labels and no ticks on axis, so it is unclear what exactly does it show.

(11) Related to discussion on line 368-370: In a task-switching state, would the switching between tasks also be reflected in behavior? In addition, the time-correlation functions would not be stationary in task-switching state, i.e. they would change over time, whereas they will be stationary in the spontaneous state. Thus, could the two mechanisms be dissociated in experiments using behavior or metrics beyond dimensionality?

(12) On line 385, could the authors provide more details on how they envision the two potential mechanisms-synaptic plasticity and targeted neuromodulation-to reinforce a task-specific low-rank connectivity pattern? If neuromodulation changes the gain of individual neurons, this modulation corresponds to scaling of connectivity by a diagonal matrix, not strengthening of a specific low-rank component. Short-term facilitation or depression also modulates synaptic strength depending on the activity of the presynaptic neuron, thus also scaling connectivity by a diagonal matrix. It is unclear how these two mechanisms could produce a specific low-rank modulation.

<https://doi.org/10.7554/eLife.111734.1.sa1>

Reviewer #2 (Public review):

Summary

The authors ask what recurrent connectivity supports many distinct task-related manifolds when the associated dynamics interfere, how a circuit engages one task while suppressing others, and what produces high-dimensional activity. Extending previous theoretical studies on low-dimensional dynamics in large networks, they use a solvable model whose weight matrix is a weighted sum of many low-rank, task-specific components and develop a

dynamical mean-field theory that relates connectivity, dynamics, and measurable population signatures of multi-tasking.

Strengths

- (1) The question is timely. Low-rank networks are a leading model for low-dimensional latent dynamics, and the composition of dynamical systems has been proposed as a mechanism allowing for rapid, flexible learning; the paper connects these two ideas under a single theoretical framework.
- (2) The proposal that sequential transitions between low-dimensional, low-rank dynamics can account for the _apparent growth of dimensionality with recording time_ is novel and is the paper's most valuable conceptual contribution.
- (3) The mathematical analysis is rigorous, and the spontaneous-state theory is convincingly validated against simulation.
- (4) The model produces concrete, falsifiable predictions - heterogeneous, syllable-dependent single-neuron tuning, low within-state dimensionality despite single-neuron variability, and distinct dimensionality-versus-recording-time signatures for the spontaneous versus task-switching accounts.

Weaknesses - whether the claims are supported by the data

- (1) Chaos is named but not demonstrated._ The large-P and intermediate task-selected regimes are labeled "chaotic," but the manuscript does not establish chaos. In a homogeneous network, it is known that once the fixed point loses stability, the surviving solution is chaotic (Sompolinsky, Crisanti & Sommers 1988); that guarantee does not transfer here. The DMFT noise term is not computed analytically, and the single-neuron correlation functions (Fig. 5) show disorder, not a demonstrated decay of the fluctuation autocorrelation to zero, nor a positive largest Lyapunov exponent. The concern is sharpened by the possibility of _transient_ chaos: orthogonal to a dominant limit cycle, fluctuations may be locally unstable only at certain amplitudes or phases, so the global attractor could remain a stable cycle visited with chaotic excursions. As it stands, the claim of chaos in the intermediate regime is unsupported; it may well hold for some range of the selected-task strength, but this is neither shown numerically nor proven.
- (2) "Analytical theory" overstates what is solved in closed form._ For the task-selected state, the kernels are non-stationary: The DMFT is entrained to the dominant task's dynamics, with an $O(1)$ time-dependent quantity inside the nonlinearity. To my knowledge, there is no closed-form DMFT solution under these conditions. The Methods section supports this, explaining that the general scheme is solved by iterative numerical self-consistency (and described there as prohibitively expensive), and tractability is recovered only in a special block-Haar ensemble with Gaussian currents. This is entirely reasonable, but the main text presents it as an analytical theory; the reliance on numerical solutions of the self-consistency equations should be stated plainly.
- (3) The spontaneous-state transition is the classical critical-gain transition, only reparametrized._ The onset of the no-task-dominant state is governed by $g_{\text{eff}}^2 = \alpha R \angle D^2 \rangle$. It appears to depend on the number of tasks only because per-task strength D is held fixed as tasks accumulate; under a normalization that holds g_{eff} fixed, the transition reduces to a critical-gain point independent of P , as in extensive random networks. Relatedly, the result that chaos "arises solely from learning many tasks" is, mechanistically, random-network chaos: the random task components raise the weight variance and play the role of effective disorder. This is a legitimate and appealing reframing, but it is not a new transition, and the manuscript should make the relationship to the standard criterion explicit.

(4) Significance of the selection mechanism. That boosting a task's gain selects it is intuitive, and the authors note the extreme (one D^μ dominating) is trivial. The non-trivial and genuinely useful contribution is quantitative - that only a small, $O(1/P)$ modulation near criticality is required. This deserves to be foregrounded rather than left to the Discussion.

<https://doi.org/10.7554/eLife.111734.1.sa0>