

MEASURING AND CONTROLLING SOLUTION DEGENERACY ACROSS TASK-TRAINED RECURRENT NEURAL NETWORKS

Anonymous authors

Paper under double-blind review

ABSTRACT

Task-trained recurrent neural networks (RNNs) are versatile models of dynamical processes widely used in machine learning and neuroscience. While RNNs are easily trained to perform a wide range of tasks, the nature and extent of the degeneracy in the resultant solutions (i.e., the variability across trained RNNs) remain poorly understood. Here, we provide a unified framework for analyzing degeneracy across three levels: behavior, neural dynamics, and weight space. We analyzed RNNs trained on diverse tasks across machine learning and neuroscience domains, including N-bit flip-flop, sine wave generation, delayed discrimination, and path integration. Our key finding is that the variability across RNN solutions, quantified on the basis of neural dynamics and trained weights, depends primarily on network capacity and task characteristics such as complexity. We introduce information-theoretic measures to quantify task complexity and demonstrate that increasing task complexity consistently reduces degeneracy in neural dynamics and generalization behavior while increasing degeneracy in weight space. These relationships hold across diverse tasks and can be used to control the degeneracy of the solution space of task-trained RNNs. Furthermore, we provide several strategies to control solution degeneracy, enabling task-trained RNNs to learn more consistent or diverse solutions as needed. We envision that these insights will lead to more reliable machine learning models and could inspire strategies to better understand and control degeneracy observed in neuroscience experiments.

1 INTRODUCTION

Recurrent neural networks (RNNs) are widely used across machine learning and computational neuroscience for modeling dynamic processes. They can be efficiently trained using standard nonconvex optimization techniques and have proven useful for understanding neural dynamics during task performance (Sussillo, 2014; Rajan et al., 2016; Barak; Mastrogiuseppe & Ostojic, 2018; Vyas et al., 2020; Driscoll et al., 2024). While it is well-known that in feedforward networks, differences in weight initialization and randomness during training (stochastic gradients, mini-batch sample variability, etc.) can cause different final solutions, less is known about how such differences manifest when training RNNs (Das & Fiete, 2020; Turner & Barak, 2024; Collins et al., 2022; Glorot & Bengio, 2010; Fort et al., 2019; Goodfellow et al., 2015; Li et al., 2018; Jastrzebski et al., 2018; Chaudhari et al., 2017; Frankle & Carbin, 2019; Kornblith et al., 2019).

Traditionally, the study of task-trained RNNs has often focused on analyzing models trained using a single approach, implicitly assuming that multiple RNNs trained on the same task would converge to similar solutions, even when trained differently or when starting from different initial conditions. However, recent work has shown that this assumption may not hold universally. For instance, Maheswaranathan et al. (2019) found that while trained RNNs may share certain topological features, their internal representation geometry can vary widely. Similarly, Turner et al. (2021) discovered that task-trained networks can develop qualitatively distinct dynamics.

These contrasting findings raise a fundamental question about the degeneracy of task-trained RNN solutions: when do networks trained on the same task converge to similar dynamics and representations, and when do they exhibit substantial degeneracy? Previous studies offer conflicting per-

054 spectives, with some suggesting that task-trained RNNs exhibit dynamics that are universal across
055 instances (Maheswaranathan et al., 2019), while others emphasize the degeneracy across individual
056 solutions (Turner et al., 2021; Galgali et al., 2023; Gholamrezaei & Whishaw; Gao et al.; Mehrer
057 et al.).

058 In this paper, we reconcile these differing views by providing a unified framework for analyzing
059 degeneracy at three levels: behavior, neural dynamics, and weight space. We hypothesize that the
060 variation in solutions of trained RNNs is influenced by *task under-specification*. In other words,
061 when constraints do not uniquely or adequately determine the network’s solution to a given task, we
062 observe greater variability across trained networks (D’Amour et al., 2020). To test this hypothesis,
063 we investigate how different task characteristics, particularly at the level of complexity of the inputs
064 and outputs, affect degeneracy of the solutions found by task-trained RNNs at three levels (across
065 behavioral, neural, and weight space). We simulated four tasks – N-Bits Flip Flip, Delayed Discrimi-
066 nation, Sine Wave Generation, and Path Integration 1. Our key finding is that as task complexity
067 increases, the solution space becomes more constrained, reducing degeneracy in both behavior and
068 neural dynamics but increasing degeneracy in the underlying network weights.

069 Our simulations confirm the above hypothesis at the behavioral and neural-dynamical levels: with
070 increasing task complexity, the degeneracy of solutions – measured at the behavioral level by the
071 coefficient of variation of the out-of-distribution performance and at the dynamical level by the Dy-
072 namical Similarity Analysis distance – decreases consistently across all four tasks we tested 3 5
073 (Ostrow et al., 2023). Interestingly, at the weight level, we observe the opposite trend: the degener-
074 acy of solutions increases with task complexity.

075 After quantifying degeneracy at behavioral, neural, and weight-levels, next we propose practical
076 strategies for controlling the level of degeneracy in task-trained RNNs, including altering task com-
077 plexity, incorporating auxiliary loss functions, and applying structural constraints during training.
078 These methods provide flexibility for researchers aiming to tailor the training process to their spe-
079 cific needs, whether they are seeking more consistent (Kepple et al., 2022) or more diverse RNN
080 solutions (Liebana Garcia et al., 2023; Fascianelli et al., 2024; Pan-Vazquez et al., 2024; Kepple
081 et al., 2022).

082 083 084 1.1 CONTRIBUTIONS

085
086 Our contributions with this paper are as follows:

- 087
088 • **Information-theoretic quantification of task complexity:** We introduce a new frame-
089 work for assessing task complexity by quantifying the information content in the input and
090 target output time series of a task. This measure effectively captures the extent to which
091 input-output signals constrain neural dynamics. Furthermore, this measure correlates well
092 with the neural-dynamical degeneracy we observed both across different complexity levels
093 within the same task and across different tasks.
- 094
095 • **Quantifying the degeneracy of solution spaces:** We measure degeneracy at the behav-
096 ioral, dynamical, and weight levels in populations of trained networks across four different
097 tasks of wide applicability to both ML and neuroscience. Our findings reveal that, across
098 all tasks, increased task complexity leads to lower degeneracy at the behavioral and neural
099 dynamical levels, but concomitantly higher degeneracy at the weight level.
- 100
101 • **Methods for controlling degeneracy:** We propose five techniques for controlling degen-
102 eracy in task-trained RNNs, including manipulating task complexity, applying specific reg-
103 ularization during training through auxiliary losses, and structural constraints during train-
104 ing. Notably, we demonstrate that both neural dynamical degeneracy and weight degeneracy
105 can be manipulated to change either in the same (“covariant”) or opposite directions
106 (“contravariance”), providing flexibility for researchers to tailor network solutions to their
107 specific needs (Cao & Yamins, 2024; Kepple et al., 2022; Fascianelli et al., 2024; Mah-
eswaranathan et al., 2019; Durstewitz et al., 2023; Gilpin, 2024) .

2 METHODS

2.1 MODEL ARCHITECTURE

We used discrete-time nonlinear *vanilla* recurrent neural networks (RNNs), where the update rule is defined as: $\mathbf{h}_t = F(\mathbf{h}_{t-1}, \mathbf{x}_t) = \tanh(\mathbf{W}_h \mathbf{h}_{t-1} + \mathbf{W}_x \mathbf{x}_t + \mathbf{b})$, where $\mathbf{h}_t \in \mathbb{R}^n$ is the hidden state at time t , $\mathbf{x}_t \in \mathbb{R}^m$ is the input at time t , $\mathbf{W}_h \in \mathbb{R}^{n \times n}$ is the recurrent weight matrix, $\mathbf{W}_x \in \mathbb{R}^{n \times m}$ is the input weight matrix, and $\mathbf{b} \in \mathbb{R}^n$ is the bias vector.

A linear readout layer is applied on top of the RNN hidden state to produce the model’s prediction at each time step. The RNNs use Backpropagation Through Time (BPTT) as the learning rule where the RNN is unrolled over time, allowing gradients to be computed for each time step in the sequence (Werbos, 1990). All networks are trained using supervised learning via the Adam optimizer, with a learning rate specific to each task determined via hyperparameter tuning (a list of all training-related hyperparameters can be found in Appendix A). For each task, we train 50 networks using different initializations, until they achieve near-asymptotic loss on the test set. For the N-Bit Flip-Flop and Path Integration tasks 1, we use RNNs with 64 hidden units and input/output dimensions appropriate to each task’s specific requirements. For the Delayed Discrimination and Sine Wave Generation tasks 1, we use RNNs with 128 hidden units.

2.2 TASKS

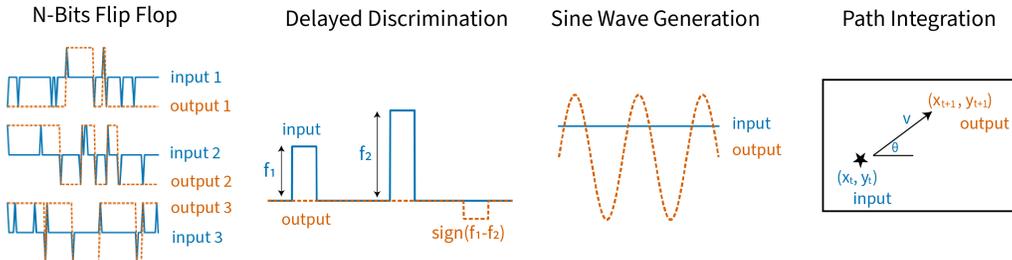


Figure 1: **Task suite:** We train RNNs on a diverse set of four tasks: **N-Bit flip-flop:** Networks are trained to remember the last non-zero input for each of the N channels. **Delayed Discrimination:** Networks compare the magnitude of two temporally separated pulses across N channels. **Sine Wave Generation:** Networks receive static input indicating frequency and produce sine waves of the specified frequency across each of the N channels. **Path Integration:** Networks integrate velocity inputs to track position in a bounded arena in 2D or 3D space. Schematic only shows 2D environment.

N-Bit Flip-Flop Task In this task, RNNs are provided with N independent input channels, each taking discrete values from $\{-1, 0, +1\}$, with a fixed probability of switching p_{switch} . The network has N output channels, each of which is required to remember the value of the last nonzero input received on its respective input channel. We vary the complexity of the task by changing the number of input and/or output channels N .

Delayed Discrimination Task The network is presented with two pulses of varying amplitudes $f_1, f_2 \in [2, 10]$, separated in time by a variable delay $t \in [5, 20]$ time steps. The network is required to output the sign of the difference between the two amplitudes, i.e., $\text{sign}(f_2 - f_1)$. This task can be extended to have N independent input and output channels, where the network must compare f_1 and f_2 within each channel separately. We vary the complexity of the task by varying the number of independent input and/or output channels N .

Sine Wave Generation RNNs receive a static input indicating a target frequency $f \in [1, 30]$ and are required to generate a sine wave of that frequency over time, effectively converting the input f into the output $\sin(2\pi ft)$. We define N_{freq} as the number of possible target frequencies presented in the training set, equally spaced within the range $[1, 30]$. This task can involve N independent input and output channels, where each channel outputs a sine wave with the frequency specified by its respective input channel. We vary the complexity of the task by changing N_{freq} and N , the number of independent input and/or output channels.

Path Integration Task Each network is initialized at a random initial location within a bounded 2D environment. At each time step, the network receives inputs representing the angular direction θ and speed v , and must integrate this information over time to output the updated (x, y) location. To vary the complexity of the task, we introduce a 3D version of the task, where the network receives inputs θ (azimuth angle), ϕ (elevation angle), and v (speed), and must output the updated (x, y, z) location. The network effectively performs path integration by accumulating movement vectors based on the input directions and speeds.

3 RESULTS

3.1 CHARACTERIZING TASKS COMPLEXITY AND NEURAL DYNAMICS (RNN HIDDEN STATES)

We quantified the task complexity by calculating the Shannon entropy of its input and target (output) time-series (Shannon, 1948; Bialek et al., 2001; Crutchfield & Young, 1989). We chose this measure because it directly reflects the amount of information present in the task’s inputs and outputs, which the network must process and represent through its hidden state (Tishby et al., 2000). By quantifying the statistical complexity of the input and output time series, we capture the information-processing demand imposed on the network.

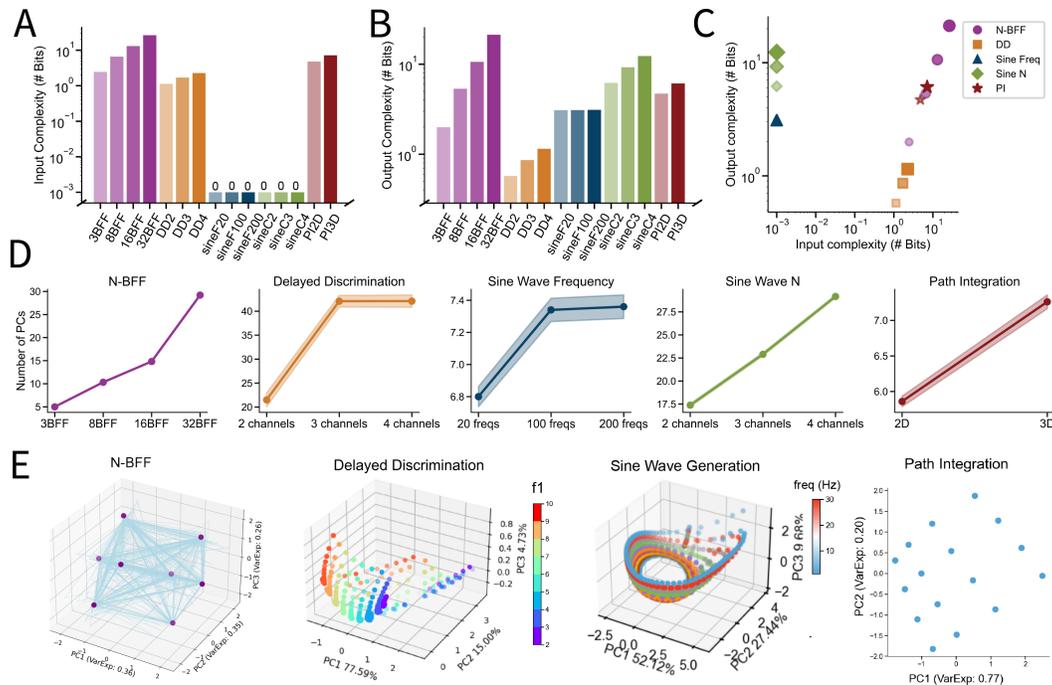


Figure 2: Characterizing tasks and RNN hidden states (neural dynamics): We analyze task complexity and its effects on RNN hidden states using multiple measures: (A) Input complexity, measured as the entropy of the time-averaged input signal. We use the integer in the task name to indicate the number of input-output channels for each task except in “sineF”, where “F” stands for frequency, and the integer indicates the number of frequencies the network learns to produce. (B) Output complexity, measured as the entropy of the time-averaged output signal (C) Our tasks span three quadrants of the Input and Output Complexity space (excluding low input and low output complexity). (D) The number of leading Principal Components (PCs) explaining 95% of the variance, which we use as a measure of the dimensionality of the task, increases consistently across tasks as difficulty increases. (E) RNN hidden state (neural dynamics) trajectories projected onto the leading two PCs for Path Integration and three PCs for the other tasks.

Without loss of generality, given a one-dimensional output time series (y_1, y_2, \dots, y_T) , the entropy of this time series is given by: $H(Y) = -\sum_{k=1}^n p(y_k) \log_2 p(y_k)$, where n is the number of unique values in the time series, and $p(y_k)$ is the probability of observing the value y_k in the time series.

For Sine Wave Generation and Path Integration tasks, the outputs are continuous values rather than discrete ones. Therefore, we binned the output values into discrete intervals before calculating the entropy. For task variants with multiple independent channels, we multiplied the single-channel entropy by the number of channels. Since each channel operates independently and has identical statistical properties, multiplying the single-channel entropy by N provides an accurate measure of the total entropy across all channels. For a training set with M trials, we averaged the entropy across all trials presented in the training set to obtain a representative measure of task complexity:

$$H_{\text{task}} = N \times \left(\frac{1}{M} \sum_{i=1}^M H(Y^{(i)}) \right).$$

Different tasks are characterized by their specific input and output complexity profiles (Figure 2A and 2B). Within each task, increasing the number of input and output channels consistently increases both the input and output complexities. We selected our tasks to span three quadrants of the input-output complexity space (Figure 2C):

- High-input, high-output complexity: N-Bit Flip-Flop and Path Integration tasks
- High-input, low-output complexity: Delayed Discrimination task
- Low-input, high-output complexity: Sine Wave Generation task

Among them, the N-Bit Flip-Flop task exhibits the highest input and output complexities among all tasks, suggesting that it imposes the greatest information-processing demands on the networks.

We then provide a link between the task complexity and the dimensionality of the representation demanded by the tasks. We found that the number of leading Principal Components (PCs) explaining 95% of the variance increases consistently across tasks as difficulty increases, indicating that harder tasks demand higher-dimensional representations and fuller utilization of the networks' representational capacity (Figure 2D) (Gao et al.).

Typical solutions found by converged networks across the different tasks are as follows (Figure 2E): In **N-Bit flip-flop** tasks, networks learn two fixed points corresponding to the output value of $\{-1, +1\}$ for each channel. Networks further factorize the fixed points corresponding to different output channels along different orthogonal dimensions. Specifically, when $N = 3$, networks learn $2^3 = 8$ fixed points on the vertices of a 3D cube. In **Delayed Discrimination** tasks, networks memorize the first input value by storing each input value as a separate fixed point during the delay period, when the inputs to the network are no longer on and working memory is required (Hopfield, 1982; Sussillo & Barak, 2013; Driscoll et al., 2024). In **Sine Wave Generation** tasks, networks develop limit cycles corresponding to different target frequencies indicated by inputs, and traverse different limit cycles to produce different sine wave outputs at the appropriate frequencies. In **2D Path Integration** tasks, without any external input, networks develop 2D maps of their environments using a 2D plane of attractors. Without any inputs, a network's hidden state (neural dynamics) stabilizes on fixed points corresponding to their current location.

3.2 CHARACTERIZING THE DEGENERACY OF TASK-TRAINED RNNs

We now quantify the degeneracy of task-trained RNNs at three levels of analysis: (out of distribution) behavior, neural dynamics, and weight space. We also relate the measured degeneracy back to the task complexity measure introduced in the last section.

3.2.1 DEGENERACY IN DYNAMICS DECREASES WITH TASK COMPLEXITY

We use a recently published measure for comparing dynamical systems, called Dynamical Similarity Analysis (DSA) (Ostrow et al., 2023), to perform pairwise comparisons of the neural dynamics in our task-trained networks. DSA measures the temporal and topological structure of the system's

dynamics and ignores geometric configurations that alter the spatial representation of system trajectories without changing the underlying dynamics. It provides a quantitative framework to assess how different networks may arrive at similar or distinct dynamical implementations of a task (See Appendix C for technical details).

To quantify the relationship between task complexity and dynamical degeneracy, we increased the complexity of the tasks by increasing the number of independent input and output channels. This choice is grounded in information theory, as an entropy measure of task complexity is additive for statistically independent variables. We found that within each task, increasing the number of input and output channels consistently reduced the dynamical degeneracy across the converged networks (Figure 3A). Strikingly, this relationship (“covariant” degeneracy with task complexity) holds across all the tasks we consider (Figure 3C).

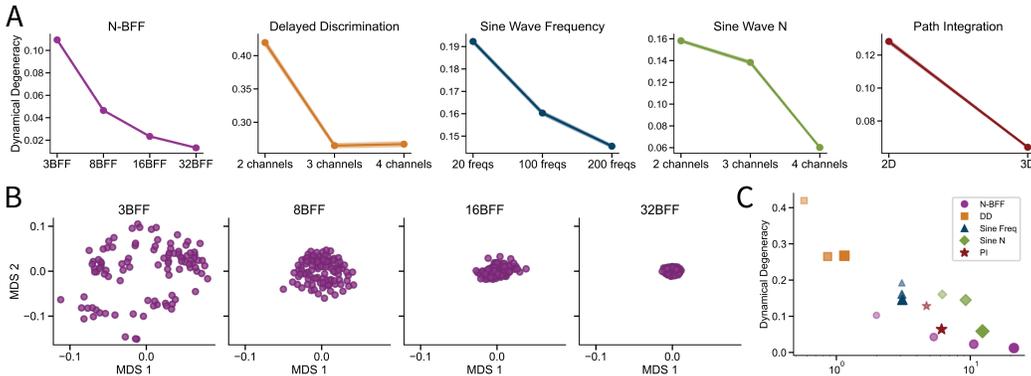


Figure 3: **Dynamical degeneracy decreases with increased task complexity:** (A) Pairwise Dynamical Similarity Analysis (DSA) scores decrease consistently across all tasks as we increase task complexity, suggesting more similar (less degenerate) dynamical solutions for harder tasks. In Sine Wave Frequency, we change the frequency content of the task and in Sine Wave N, we change the number of input-output channels (N) of the task. (B) 2D embedding of pairwise DSA distances using Multi-Dimensional Scaling (MDS) on the N-bit flip-flop task shows how the spread of solutions decreases as task complexity is increased. Each dot represents a network with a different initialization. (C) Output complexity versus the average dynamical degeneracy on a task shows how higher output complexity correlates with lower dynamical degeneracy both within and across tasks. Within each task, larger marker size and less opacity indicate task variant with higher complexity. Smaller DSA score implies dynamics are more similar.

3.2.2 DEGENERACY IN WEIGHTS INCREASES WITH TASK COMPLEXITY

We quantify degeneracy across task-trained RNNs at the level of post-training weights by using a permutation invariant modification of the Frobenius Norm. For a pair of RNNs with recurrent weight matrices given by \mathbf{W}_1 and \mathbf{W}_2 , we define: $d_{PIF}(\mathbf{W}_1, \mathbf{W}_2) = \min_{\mathbf{P}_1, \mathbf{P}_2 \in \mathcal{P}(n)} \|\mathbf{W}_1 - \mathbf{P}_1 \mathbf{W}_2 \mathbf{P}_2\|_F$, where $\mathcal{P}(n)$ is the set of permutation matrices of size $n \times n$, and $\|\cdot\|_F$ denotes the Frobenius norm. (See Appendix D for additional details) For comparing networks of different sizes, we normalize the above norm by the number of parameters in the weight matrix. We found that pairwise distances between weight matrices from converged RNNs increases consistently as task complexity increases across all tasks (Figure 4A).

Inspired by research linking statistical mechanics of random Gaussian landscapes to deep learning theory (Bahri et al., 2020; Fyodorov & Williams, 2007; Bray & Dean, 2007), we hypothesized that increasing task complexity while holding the network size constant could make possible solutions ‘rarer’ at the level of recurrent weights. This implies that the optimization process would have to, on average, search for solutions further away from the initial recurrent weights to converge. We found that the distance between the converged weight matrices and the initialized weight matrices indeed increases consistently with task complexity, and does so across all the tasks (Figure 4B and Table 1). Formally, distance from initial weights is $\|\mathbf{W}_T - \mathbf{W}_0\|_F$, where \mathbf{W}_T is the weight matrix

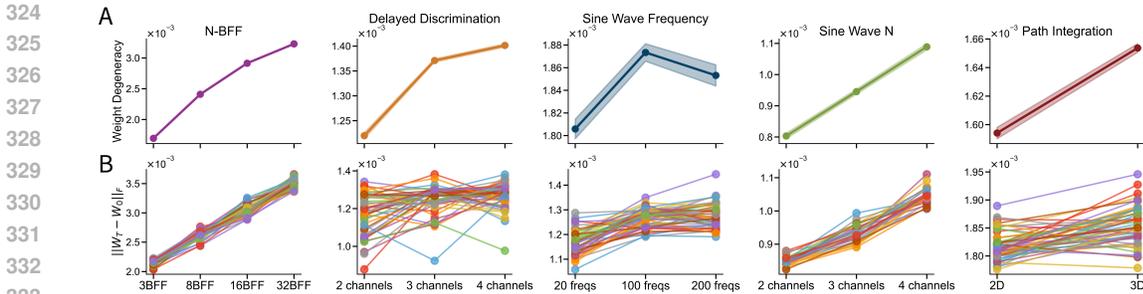


Figure 4: **Degeneracy in weight space (measured using permutation-invariant metrics) increases with increased task complexity:** (A) Pairwise distances between weights, measured by the normalized Frobenius norm of the recurrent connectivity matrix, increase consistently with task complexity, across all tasks. (B) Distances from initialized weights, quantified by the Frobenius Norm, also increases with task complexity, indicating greater divergence in weight space for harder tasks. In other words, we observe greater variation/variability across trained individuals (individual differences) at the connectivity level, even with the opposite trend observed at the dynamical and behavioral levels. Each line represents the distance from the initialized weights for a given network.

after training and W_0 is the initial weight matrix. We normalize this measure by the number of parameters in the weight matrix across tasks.

3.2.3 DEGENERACY IN OUT-OF-DISTRIBUTION GENERALIZATION BEHAVIORS DECREASES WITH TASK COMPLEXITY

Does degeneracy in Out-of-Distribution generalization behaviors change concomitantly with degeneracy in neural dynamics (hidden state) with task complexity? We hypothesize that lower dynamical degeneracy implies that networks respond to input stimuli with more similar dynamics. Therefore, when these networks are probed with out-of-distribution (OOD) inputs, they should produce more similar outputs, leading to lower dispersion in their OOD performance across networks (i.e., lower degeneracy). In other words, as task difficulty increases and dynamical degeneracy decreases, the degeneracy in OOD performance should also decrease.

To test this hypothesis, we measured the OOD performance (mean squared error) of converged networks that all achieved near-asymptotic training loss under the following *length generalization* conditions. For Delayed Discrimination tasks, we doubled the length of the delay period; for all other tasks, we doubled the length of the entire trial. We then calculated the coefficient of variation (CV) of the OOD performance across networks, defined as $CV = \sigma/\mu$, where σ is the standard deviation of the OOD performance across networks, and μ is the mean OOD performance.

Our results show that across all tasks, networks with lower dynamical degeneracy indeed exhibited lower CV in their OOD performance. This finding supports our hypothesis that increased task complexity, which reduces dynamical degeneracy, also leads to more consistent – less degenerate – OOD generalization behaviors across networks. The prediction this result makes is that when trained on harder tasks, “expert” animals in neuroscience labs or task-trained networks/agents in ML should show less individual variation in generalization performance when tested.

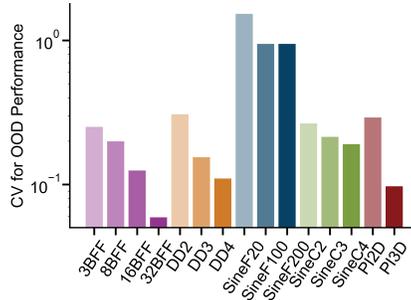


Figure 5: **Behavioral degeneracy, in terms of variability of the OOD generalization performance, decreases with increased task complexity:** Degeneracy of the out of distribution (OOD) generalization behavior, as measured by coefficient of variation of the OOD performance (mean squared error), decreases consistently with increasing task complexity across all tasks.

3.3 CONTROLLING DEGENERACY OF TASK-TRAINED RNNs

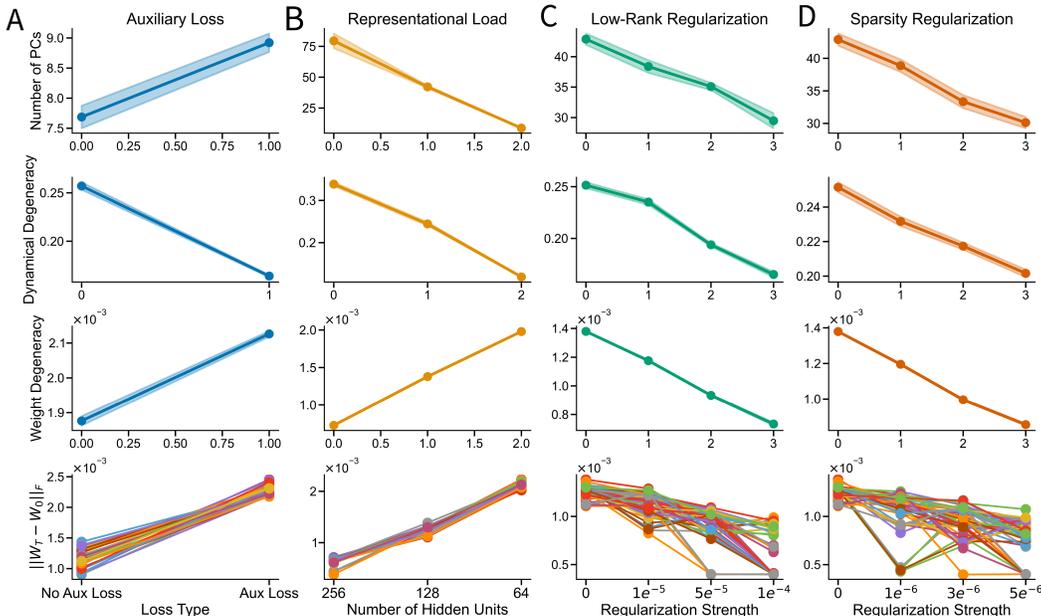


Figure 6: **Differential control of weight and dynamics degeneracy:** Using Delayed Discrimination (DD) tasks as an example, we demonstrate various methods to control the degeneracy of solution spaces found by task-trained RNNs: **(A) Adding an auxiliary loss.** Adding an additional loss term to calculate the signed difference (not just direction) between the two consecutive inputs adds additional structure to the learned dynamics, resulting in reduced dynamical degeneracy and increased weight degeneracy. **(B) Changing network size:** By increasing the networks’ size, we decrease the load on their representational capacity, which in turn, results in more degeneracy in their neural dynamics or hidden states. Simultaneously, we see greater degeneracy in their weight space (i.e., the trend in degeneracy is in the same direction for both dynamics and weights). **(C) Regularizing weights to be low rank:** We see a decrease in *both* dynamical and weight-space degeneracy with increasing low-rank regularization on the recurrent weights and reduced distance from initialized weights, likely due to smaller-magnitude weight updates. **(D) Regularizing weights to be sparse:** Similar to above, we see a decrease in both dynamical and weight degeneracy with increasing sparsity of the recurrent weight matrix.

Table 1: Impact of Various Factors on Degeneracy

	Dynamical Degeneracy	Weight Degeneracy
Increase Task Complexity	↓	↑
Auxiliary Loss	↓	↑
Representational Load	↓	↑
Low-Rank Regularization	↓	↓
Sparsity Regularization	↓	↓

Based on our characterization of task complexity and quantification of degeneracy at multiple levels of analysis, we propose several methods to control degeneracy across task-trained RNNs, making predictions for ML and neuroscience studies:

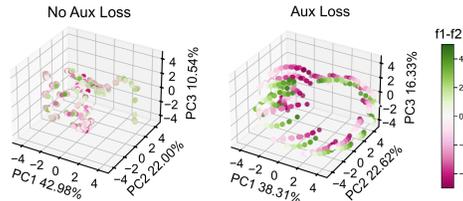
Increase task complexity: As discussed in the previous section, increasing task complexity reduces dynamical degeneracy and increases weight degeneracy in a consistent and measurable way in a suite of tasks of broad interest to ML and neuroscience.

432 **Adding auxiliary loss:** We can directly change the degeneracy of the solution spaces found by task-
 433 trained RNNs by adding additional terms to the objective/loss function that are consistent with but
 434 distinct from the primary task objective. This introduces additional structure to the learned dynam-
 435 ics, which in turn, decreases dynamical degeneracy and increases weight degeneracy, as networks
 436 must satisfy multiple constraints simultaneously. Figure 6A shows an example implementation of
 437 this idea on a Delayed Discrimination task. We demand that the network output $f_1 - f_2$, instead of
 438 only $sign(f_1 - f_2)$. By analyzing the RNN’s neural dynamics or hidden state after both input stimuli
 439 have been presented, we find that the additional objective forces the network to learn “line attrac-
 440 tors” (Strogatz, 2000) with the same sign as network without the additional loss, but with variable
 441 magnitudes reflecting different values of $f_1 - f_2$ (Figure 7).

442 **Modifying representational load (or network capacity):** Changing the network size (or the num-
 443 ber of units in the RNN) affects its representational capacity. Increasing network size typically leads
 444 to increased dynamical degeneracy due to more degrees of freedom. Simultaneously degeneracy
 445 increases across weight space because of more parameters. (Figure 6 B).

446 **Imposing structural constraints on weight spaces(or regularization of recurrent weight matri-
 447 ces):** Applying regularization techniques such as low-rank (nuclear norm) (Figure 6 C) or sparsity
 448 (L1) (Figure 6 D) regularization to the recurrent weight matrix generally decreases the degener-
 449 acy across neural dynamical and weight spaces simultaneously (Mastrogiuseppe & Ostojic, 2018;
 450 Narang et al., 2017; Herbert & Ostojic, 2022).

451 These methods provide a toolkit for controlling de-
 452 generacy that naturally follows from our earlier anal-
 453 yses. Importantly, we report that some interventions
 454 (increasing task complexity and modifying network
 455 capacity) induce an inverse or “contravariant” rela-
 456 tionship between weight-space and dynamical degen-
 457 eracy, while others (structural constraints) tend to af-
 458 fect both dynamical and weight-space degeneracy in
 459 the same direction or “covariantly” (Summarized in
 460 Table 1). Overall, the degeneracy of the solution
 461 space is primarily controlled by two key factors: the
 462 task complexity and the representational capacity of
 463 the networks. Among the manipulations to control
 464 degeneracy in RNNs, increasing task complexity and
 465 adding auxiliary losses directly alter the task com-
 466 plexity while representational load and adding struc-
 467 tural constraints directly impact the representational
 468 capacity of the networks.



469 **Figure 7: RNNs learn more structured**
 470 **dynamics with auxiliary losses:** In De-
 471 layed Discrimination tasks trained with aux-
 472 iliary losses (e.g., an auxiliary output $f_1 - f_2$
 473 in addition to main output $sign(f_1 - f_2)$),
 474 RNNs neural dynamics or hidden states
 475 develop more sophisticated representations
 476 to track task-variables with 1-D manifolds
 477 (right) rather than just point-attractors (left).

469 **4 RELATED WORK**

471 Our approach extends a previously proposed measures of task complexity (Meister, 2022), which
 472 defined a task as a mapping from states to actions, and quantified task complexity as the entropy
 473 of all possible state-action mappings. Here, we regard tasks as predefined time-series of inputs and
 474 outputs rather than in terms of state-action mappings. By taking into account the temporal dynamics
 475 explicitly, we quantify the statistical complexity of the input and output time-series using Shannon
 476 entropy. Meister (2022) focuses on evaluating task-complexity from a learning-theory perspective:
 477 specifically, the amount of information an animal needs to acquire a perfect mapping between states
 478 and actions. In contrast, we focus on the representational capacity and the constraints imposed by
 479 the processing demands of a task. Our measure directly captures the amount of information that
 480 must be “carried” through the neural dynamics to produce the desired output. Our measure of task
 481 complexity is also consistent with the notion of *Neural Task Complexity* proposed by Gao et al.
 482 (2017), which provides an upper bound on the dimensionality that neural trajectories can explore
 483 during a task. We have demonstrated that within each type of task, variants with higher output
 484 complexity require higher dimensional representations (Figure 2D).

485 We further show that the output complexity of a task directly shapes the degeneracy of solutions
 found by RNNs. While previous work speculated that the degree to which neural dynamics are

486 attractor-heavy may influence degeneracy (Turner et al., 2021), we have quantified the “attractor-
487 ness” of our tasks and found that it alone does not account for the different levels of dynamical
488 degeneracy across tasks (see Appendix B). Moreover, as we modified the task characteristics to
489 study its effect on degeneracy, we observed that degeneracy remains invariant under certain trans-
490 formations of the task, e.g., changing the delay duration in Delayed Discrimination or altering the
491 environment size in Path Integration tasks. In our framework, these transformations do not affect
492 the output complexity of the task and thus leave the degeneracy unchanged. In the language of
493 complex systems theory, these factors (e.g., delays) are “sloppy dimensions” of the task with respect
494 to degeneracy. In contrast, factors that directly modify output complexity (e.g., adding channels,
495 changing network size) are “stiff dimensions” that significantly control and modulate degeneracy of
496 solutions (Gutenkunst et al., 2007b;a; Daniels et al., 2008; Kepple et al., 2022).

497 Our findings are also consistent with the *Contravariance Principle* from Cao & Yamins (2024): the
498 harder a task, the more constrained the system that solves it will need to be. Another recent study
499 by Huh et al. (2024) hypothesizes that AI models seem to be converging on the same representation
500 spaces independent of modality, because more constraints from the massive-scale data make degen-
501 erate solutions less likely. Notably, the above related findings are based on feedforward networks for
502 visual object categorization and Transformer networks for language modeling, respectively. Here,
503 our study demonstrates that these ideas also generalize to recurrent networks/RNNs.

504 5 DISCUSSION

505 Our paper presents a unified approach to quantifying and controlling the degeneracy of solutions
506 found by task-trained RNNs across behavioral, neural dynamical (or hidden states), and weight-
507 space levels of analysis. We simulated and analyzed a task suite of broad interest to ML and neuro-
508 science, yielding several novel insights.

509 First, information content in inputs and outputs of different tasks, which we use to quantify task
510 complexity, emerges as a crucial determinant of degeneracy. This perspective helps reconcile previ-
511 ously conflicting observations on degeneracy across a range of tasks and measures. Second, across
512 different tasks, as task complexity increases, we observe a robust inverse relationship between neu-
513 ral dynamical and weight degeneracy. As a corollary, while harder tasks lead to more consistent
514 neural dynamics across task-trained networks, their underlying recurrent weight matrices are more
515 degenerate/variable.

516 Third, our study exemplifies several strategies for controlling solution degeneracy, including manip-
517 ulating task complexity, adding auxiliary loss functions, adjusting network capacity, and structural
518 constraints on recurrent weights 1. These prescriptions will allow researchers to tailor the level
519 of degeneracy of networks to suit specific research questions (e.g., probing individual variability
520 across animals in neuroscience experiments) or application needs (Bouthillier et al., 2021; Morik,
521 2005; Yang et al., 2022). Importantly, we observed that certain strategies—increasing task complex-
522 ity or modifying network capacity—induce an inverse (contravariant) relationship between weight
523 and dynamical degeneracy, while others—imposing structural constraints—tend to affect both types
524 of degeneracy in the same or covariant direction.

525 While our study is based on artificial neural network models, some of the above strategies could
526 apply to neuroscience experiments, (Howard, 2002) e.g., introducing an auxiliary sub-task during
527 behavioral shaping of lab animals to learn different tasks may constrain the degeneracy of the so-
528 lutions learned by the biological brain and manifest in a different representational load than in an
529 animal trained using a “curriculum” without an auxiliary sub-task. We note that degeneracy is also
530 ubiquitous in other biological systems (Golub et al., 2018; Prinz et al., 2004; Edelman & Gally,
531 2001); thus, we motivate future theoretical studies based on them.

532 In summary, our work provides insights into the factors that shape the solution landscape of task-
533 trained RNNs. The framework we bridges differing perspectives in prior work, reconciling obser-
534 vations of both universality and individuality in trained networks in ML and neuroscience. Future
535 research could delve deeper into the theoretical underpinnings of the observed relationships in bi-
536 ological and artificial agents, including neural networks, between task complexity, representational
537 capacity, and degeneracy.

538
539

REFERENCES

- 540
541
542 Samuel K Ainsworth, Jonathan Hayase, and Siddhartha Srinivasa. Git re-basin: Merging models
543 modulo permutation symmetries. *arXiv preprint arXiv:2209.04836*, 2022.
- 544
545 Yasaman Bahri, Jonathan Kadmon, Jeffrey Pennington, Sam S Schoenholz, Jascha Sohl-Dickstein,
546 and Surya Ganguli. Statistical mechanics of deep learning. *Annual Review of Condensed Matter
547 Physics*, 11(1):501–528, 2020.
- 548
549 Omri Barak. Recurrent neural networks as versatile tools of neuroscience research. 46:1–6. ISSN
550 09594388. doi: 10.1016/j.conb.2017.06.003. URL [https://linkinghub.elsevier.
551 com/retrieve/pii/S0959438817300429](https://linkinghub.elsevier.com/retrieve/pii/S0959438817300429).
- 552
553 William Bialek, Ilya Nemenman, and Naftali Tishby. Predictability, complexity and learning, 2001.
554 URL <https://arxiv.org/abs/physics/0007070>.
- 555
556 Xavier Bouthillier, Pierre Delaunay, Mirko Bronzi, Assya Trofimov, Brennan Nichyporuk, Justin
557 Szeto, Nazanin Mohammadi Sepahvand, Edward Raff, Kanika Madan, Vikram Voleti, et al. Ac-
558 counting for variance in machine learning benchmarks. *Proceedings of Machine Learning and
559 Systems*, 3:747–769, 2021.
- 560
561 Alan J Bray and David S Dean. Statistics of critical points of gaussian fields on large-dimensional
562 spaces. *Physical review letters*, 98(15):150201, 2007.
- 563
564 Rosa Cao and Daniel Yamins. Explanatory models in neuroscience, part 2: Functional intelligibility
565 and the contravariance principle. *Cognitive Systems Research*, 85:101200, 2024.
- 566
567 Pratik Chaudhari, Anna Choromanska, Stefano Soatto, Yann LeCun, Carlo Baldassi, Christian
568 Borgs, Jennifer Chayes, Levent Sagun, and Riccardo Zecchina. Entropy-sgd: Biasing gradient
569 descent into wide valleys, 2017. URL <https://arxiv.org/abs/1611.01838>.
- 570
571 Jasmine Collins, Jascha Sohl-Dickstein, and David Sussillo. Capacity and trainability in recurrent
572 neural networks. In *International Conference on Learning Representations*, 2022.
- 573
574 James P. Crutchfield and Karl Young. Inferring statistical complexity. *Phys. Rev. Lett.*, 63:105–
575 108, Jul 1989. doi: 10.1103/PhysRevLett.63.105. URL [https://link.aps.org/doi/
576 10.1103/PhysRevLett.63.105](https://link.aps.org/doi/10.1103/PhysRevLett.63.105).
- 577
578 Alexander D’Amour, Katherine Heller, Dan Moldovan, Ben Adlam, Babak Alipanahi, Alex Beutel,
579 Christina Chen, Jonathan Deaton, Jacob Eisenstein, Matthew D. Hoffman, Farhad Hormozdi-
580 ari, Neil Houlsby, Shaobo Hou, Ghassen Jerfel, Alan Karthikesalingam, Mario Lucic, Yian Ma,
581 Cory McLean, Diana Mincu, Akinori Mitani, Andrea Montanari, Zachary Nado, Vivek Natarajan,
582 Christopher Nielson, Thomas F. Osborne, Rajiv Raman, Kim Ramasamy, Rory Sayres, Jessica
583 Schrouff, Martin Seneviratne, Shannon Sequeira, Harini Suresh, Victor Veitch, Max Vladymy-
584 rov, Xuezhi Wang, Kellie Webster, Steve Yadlowsky, Taedong Yun, Xiaohua Zhai, and D. Sculley.
585 Underspecification presents challenges for credibility in modern machine learning, 2020. URL
586 <https://arxiv.org/abs/2011.03395>.
- 587
588 Bryan C Daniels, Yan-Jiun Chen, James P Sethna, Ryan N Gutenkunst, and Christopher R Myers.
589 Sloppiness, robustness, and evolvability in systems biology. *Current opinion in biotechnology*,
590 19(4):389–395, 2008.
- 591
592 Abhranil Das and Ila R Fiete. Systematic errors in connectivity inferred from activity in strongly
593 recurrent networks. *Nature Neuroscience*, 23(10):1286–1296, 2020.
- 594
595 Laura N. Driscoll, Krishna Shenoy, and David Sussillo. Flexible multitask computation in re-
596 current networks utilizes shared dynamical motifs. *Nature Neuroscience*, 27(7):1349–1363,
597 July 2024. ISSN 1097-6256, 1546-1726. doi: 10.1038/s41593-024-01668-6. URL [https:
598 //www.nature.com/articles/s41593-024-01668-6](https://www.nature.com/articles/s41593-024-01668-6).
- 599
600 Daniel Durstewitz, Georgia Koppe, and Max Ingo Thurm. Reconstructing computational system
601 dynamics from neural data with recurrent neural networks. *Nature Reviews Neuroscience*, 24
602 (11):693–710, 2023.

- 594 Gerald M. Edelman and Joseph A. Gally. Degeneracy and complexity in biological systems. *Pro-*
595 *ceedings of the National Academy of Sciences*, 98(24):13763–13768, November 2001. ISSN
596 0027-8424, 1091-6490. doi: 10.1073/pnas.231499798. URL [https://pnas.org/doi/](https://pnas.org/doi/full/10.1073/pnas.231499798)
597 [full/10.1073/pnas.231499798](https://pnas.org/doi/full/10.1073/pnas.231499798).
598
- 599 Valeria Fascianelli, Aldo Battista, Fabio Stefanini, Satoshi Tsujimoto, Aldo Genovesio, and Ste-
600 fano Fusi. Neural representational geometries reflect behavioral differences in monkeys and
601 recurrent neural networks. *Nature Communications*, 15(1):6479, August 2024. ISSN 2041-
602 1723. doi: 10.1038/s41467-024-50503-w. URL [https://www.nature.com/articles/](https://www.nature.com/articles/s41467-024-50503-w)
603 [s41467-024-50503-w](https://www.nature.com/articles/s41467-024-50503-w).
- 604 Stanislav Fort, Huiyi Hu, and Balaji Lakshminarayanan. Deep ensembles: A loss landscape per-
605 spective. *arXiv preprint arXiv:1912.02757*, 2019.
606
- 607 Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural
608 networks, 2019. URL <https://arxiv.org/abs/1803.03635>.
- 609 Yan V Fyodorov and Ian Williams. Replica symmetry breaking condition exposed by random matrix
610 calculation of landscape complexity. *Journal of Statistical Physics*, 129:1081–1116, 2007.
611
- 612 Aniruddh R Galgali, Maneesh Sahani, and Valerio Mante. Residual dynamics resolves recurrent
613 contributions to neural computation. *Nature Neuroscience*, 26(2):326–338, 2023.
- 614 Peiran Gao, Eric Trautmann, Byron Yu, Gopal Santhanam, Stephen Ryu, Krishna Shenoy, and Surya
615 Ganguli. A theory of multineuronal dimensionality, dynamics and measurement. URL [http:](http://biorxiv.org/lookup/doi/10.1101/214262)
616 [://biorxiv.org/lookup/doi/10.1101/214262](http://biorxiv.org/lookup/doi/10.1101/214262).
617
- 618 Peiran Gao, Eric Trautmann, Byron Yu, Gopal Santhanam, Stephen Ryu, Krishna Shenoy, and Surya
619 Ganguli. A theory of multineuronal dimensionality, dynamics and measurement, November 2017.
620 URL <http://biorxiv.org/lookup/doi/10.1101/214262>.
- 621 Gita Gholamrezaei and Ian Q. Whishaw. Individual differences in skilled reaching for food related
622 to increased number of gestures: Evidence for goal and habit learning of skilled reaching. 123
623 (4):863–874. ISSN 1939-0084, 0735-7044. doi: 10.1037/a0016369. URL [http://doi.apa.](http://doi.apa.org/getdoi.cfm?doi=10.1037/a0016369)
624 [org/getdoi.cfm?doi=10.1037/a0016369](http://doi.apa.org/getdoi.cfm?doi=10.1037/a0016369).
625
- 626 William Gilpin. Generative learning for nonlinear dynamics. *Nature Reviews Physics*, 6(3):194–206,
627 2024.
- 628 Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural
629 networks. In *Proceedings of the thirteenth international conference on artificial intelligence and*
630 *statistics*, pp. 249–256. JMLR Workshop and Conference Proceedings, 2010.
631
- 632 Matthew D. Golub, Patrick T. Sadtler, Emily R. Oby, Kristin M. Quick, Stephen I. Ryu, Eliza-
633 beth C. Tyler-Kabara, Aaron P. Batista, Steven M. Chase, and Byron M. Yu. Learning by neu-
634 ral reassociation. *Nature Neuroscience*, 21(4):607–616, April 2018. ISSN 1097-6256, 1546-
635 1726. doi: 10.1038/s41593-018-0095-3. URL [https://www.nature.com/articles/](https://www.nature.com/articles/s41593-018-0095-3)
636 [s41593-018-0095-3](https://www.nature.com/articles/s41593-018-0095-3).
- 637 Ian J. Goodfellow, Oriol Vinyals, and Andrew M. Saxe. Qualitatively characterizing neural network
638 optimization problems, 2015. URL <https://arxiv.org/abs/1412.6544>.
639
- 640 Ryan N Gutenkunst, Fergal P Casey, Joshua J Waterfall, Christopher R Myers, and James P Sethna.
641 Extracting falsifiable predictions from sloppy models. *Annals of the New York Academy of Sci-*
642 *ences*, 1115(1):203–211, 2007a.
- 643 Ryan N Gutenkunst, Joshua J Waterfall, Fergal P Casey, Kevin S Brown, Christopher R Myers, and
644 James P Sethna. Universally sloppy parameter sensitivities in systems biology models. *PLoS*
645 *computational biology*, 3(10):e189, 2007b.
646
- 647 Elizabeth Herbert and Srdjan Ostojic. The impact of sparsity in low-rank recurrent neural networks.
PLOS Computational Biology, 18(8):e1010426, 2022.

- 648 JJ Hopfield. Neural networks and physical systems with emergent collective computational abilities.
649 *Proceedings of the National Academy of Sciences*, 79(8):2554–2558, 1982. doi: 10.1073/pnas.
650 79.8.2554. URL <https://www.pnas.org/doi/abs/10.1073/pnas.79.8.2554>.
- 651
- 652 BR Howard. Control of variability. *ILAR journal*, 43(4):194–201, 2002.
- 653
- 654 Minyoung Huh, Brian Cheung, Tongzhou Wang, and Phillip Isola. The platonic representation
655 hypothesis. *arXiv preprint arXiv:2405.07987*, 2024.
- 656
- 657 Stanisław Jastrzebski, Zachary Kenton, Devansh Arpit, Nicolas Ballas, Asja Fischer, Yoshua Ben-
658 gio, and Amos Storkey. Three factors influencing minima in sgd, 2018. URL <https://arxiv.org/abs/1711.04623>.
- 659
- 660 D Kepple, Rainer Engelken, and Kanaka Rajan. Curriculum learning as a tool to uncover learning
661 principles in the brain. In *International Conference on Learning Representations*, 2022.
- 662
- 663 Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of neural
664 network representations revisited, 2019. URL <https://arxiv.org/abs/1905.00414>.
- 665
- 666 Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. Visualizing the loss land-
667 scape of neural nets, 2018. URL <https://arxiv.org/abs/1712.09913>.
- 668
- 669 Samuel Liebana Garcia, Aeron Laffere, Chiara Toschi, Louisa Schilling, Jacek Podlaski, Matthias
670 Fritsche, Peter Zátka-Haas, Yulong Li, Rafal Bogacz, Andrew Saxe, and Armin Lak. Striatal
671 dopamine reflects individual long-term learning trajectories, December 2023. URL [http://](http://biorxiv.org/lookup/doi/10.1101/2023.12.14.571653)
672 biorxiv.org/lookup/doi/10.1101/2023.12.14.571653.
- 673
- 674 Niru Maheswaranathan, Alex Williams, Matthew Golub, Surya Ganguli, and David Sussillo. Uni-
675 versality and individuality in neural dynamics across large populations of recurrent networks. In
676 *Advances in neural information processing systems*, pp. 15629–15641, 2019.
- 677
- 678 Francesca Mastrogiuseppe and Srdjan Ostojic. Linking connectivity, dynamics, and computations
679 in low-rank recurrent neural networks. *Neuron*, 99(3):609–623, 2018.
- 680
- 681 Johannes Mehrer, Courtney J. Spoerer, Nikolaus Kriegeskorte, and Tim C. Kietzmann. In-
682 dividual differences among deep neural network models. 11(1):5725. ISSN 2041-
683 1723. doi: 10.1038/s41467-020-19632-w. URL [http://www.nature.com/articles/](http://www.nature.com/articles/s41467-020-19632-w)
684 [s41467-020-19632-w](http://www.nature.com/articles/s41467-020-19632-w).
- 685
- 686 Markus Meister. Learning, fast and slow, 2022. URL [https://arxiv.org/abs/2205.](https://arxiv.org/abs/2205.02075)
687 [02075](https://arxiv.org/abs/2205.02075).
- 688
- 689 Katharina Morik. Sloppy modeling. *Knowledge representation and organization in machine learn-*
690 *ing*, pp. 107–134, 2005.
- 691
- 692 Sharan Narang, Erich Elsen, Gregory Diamos, and Shubho Sengupta. Exploring sparsity in recurrent
693 neural networks. *arXiv preprint arXiv:1704.05119*, 2017.
- 694
- 695 Mitchell Ostrow, Adam Eisen, Leo Kozachkov, and Ila Fiete. Beyond Geometry: Comparing the
696 Temporal Structure of Computation in Neural Circuits with Dynamical Similarity Analysis, Oc-
697 tober 2023. URL <http://arxiv.org/abs/2306.10168>. arXiv:2306.10168 [cs, q-bio].
- 698
- 699 A Pan-Vazquez, Y Sanchez Araujo, B McMannon, M Louka, A Bandi, L Haetzel, International
700 Brain Laboratory, JW Pillow, ND Daw, and IB Witten. Pre-existing visual responses in a
701 projection-defined dopamine population explain individual learning trajectories, 2024. URL
<https://europepmc.org/article/PPR/PPR811803>.
- 702
- 703 Astrid A Prinz, Dirk Bucher, and Eve Marder. Similar network activity from disparate circuit pa-
704 rameters. *Nature Neuroscience*, 7(12):1345–1352, December 2004. ISSN 1097-6256, 1546-1726.
705 doi: 10.1038/nn1352. URL <https://www.nature.com/articles/nn1352>.
- 706
- 707 Kanaka Rajan, Christopher D Harvey, and David W Tank. Recurrent network models of sequence
708 generation and memory. *Neuron*, 90(1):128–142, 2016.

702 Peter J Schmid. Dynamic mode decomposition and its variants. *Annual Review of Fluid Mechanics*,
703 54(1):225–254, 2022.
704

705 Claude Elwood Shannon. A mathematical theory of communication. *The Bell System Techni-*
706 *cal Journal*, 27:379–423, 1948. URL [http://plan9.bell-labs.com/cm/ms/what/
707 shannonday/shannon1948.pdf](http://plan9.bell-labs.com/cm/ms/what/shannonday/shannon1948.pdf).

708 Steven H. Strogatz. *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chem-*
709 *istry and Engineering*. Westview Press, 2000.
710

711 David Sussillo. Neural circuits as computational dynamical systems. *Current opinion in neurobiol-*
712 *ogy*, 25:156–163, 2014.

713 David Sussillo and Omri Barak. Opening the black box: low-dimensional dynamics in high-
714 dimensional recurrent neural networks. *Neural computation*, 25(3):626–649, 2013.
715

716 Naftali Tishby, Fernando C. Pereira, and William Bialek. The information bottleneck method, 2000.
717 URL <https://arxiv.org/abs/physics/0004057>.

718 Elia Turner and Omri Barak. The simplicity bias in multi-task rnns: shared attractors, reuse of
719 dynamics, and geometric representation. *Advances in Neural Information Processing Systems*,
720 36, 2024.
721

722 Elia Turner, Kabir V Dabholkar, and Omri Barak. Charting and navigating the space of solutions
723 for recurrent neural networks. *Advances in Neural Information Processing Systems*, 34:25320–
724 25333, 2021.

725 Saurabh Vyas, Matthew D Golub, David Sussillo, and Krishna V Shenoy. Computation through
726 neural population dynamics. *Annual Review of Neuroscience*, 43:249–275, 2020.
727

728 Paul J Werbos. Backpropagation through time: what it does and how to do it. *Proceedings of the*
729 *IEEE*, 78(10):1550–1560, 1990.

730 Rubing Yang, Jialin Mao, and Pratik Chaudhari. Does the data induce capacity control in deep
731 learning? In *International Conference on Machine Learning*, pp. 25166–25197. PMLR, 2022.
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755

756 A TASK AND TRAINING DETAILS

757
758 A.1 N-BIT FLIP FLIP

759
760
761

Hyperparameter	Value
Optimizer	Adam
Learning rate	0.001
Learning rate scheduler	None
Max epochs	300
Steps per epoch	128
Batch size	256
Early stopping threshold	0.0005
Patience	3
Probability of flip	0.3
Number of time steps	100

762
763
764
765
766
767
768
769
770
771
772

773
774
775 A.2 DELAYED DISCRIMINATION

776
777

Hyperparameter	Value
Optimizer	Adam
Learning rate	0.001
Learning rate scheduler	CosineAnnealingWarmRestarts
Max epochs	500
Steps per epoch	128
Batch size	256
Early stopping threshold	0.05
Patience	3
Number of time steps	60
Max delay	20
Lowest stimulus value	2
Highest stimulus value	10

778
779
780
781
782
783
784
785
786
787
788
789
790
791

792
793
794 A.3 SINE WAVE GENERATION

795
796
797

Hyperparameter	Value
Optimizer	Adam
Learning rate	0.0005
Learning rate scheduler	None
Max epochs	500
Steps per epoch	128
Batch size	32
Early stopping threshold	0.05
Patience	3
Number of time steps	100
Lowest frequency	1
Highest frequency	30

798
799
800
801
802
803
804
805
806
807
808
809

810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863

A.4 PATH INTEGRATION

Hyperparameter	Value
Optimizer	Adam
Learning rate	0.001
Learning rate scheduler	ReduceLRonPlateau
Learning rate decay factor	0.5
Learning rate decay patience	40
Max epochs	1000
Steps per epoch	128
Batch size	64
Early stopping threshold (2D)	0.1
Early stopping threshold (3D)	0.3
Patience	3
Number of time steps	100
Lowest frequency	1
Highest frequency	30

B CHARACTERIZING THE ATTRACTORNESS OF THE TASKS

Attractors are associated with stable states in a network’s dynamics. To quantify how attractor-heavy a task is, we measure the **speed of change** in the RNN’s hidden activities. Specifically, we calculate the average normalized difference between the hidden states at two consecutive time steps. This measure reflects the stability of the network’s internal states over time, where slower changes indicate the presence of attractor-like behavior.

For each trial, given a time series of hidden state activities, we calculate the speed of change in the hidden activities as:

$$\Delta h = \frac{\|h_{t+1} - h_t\|}{\frac{1}{T} \sum_{t=1}^T \|h_t\|}$$

where T is the total number of time steps in the trial. We then average Δh across all trials in the training set to obtain a representative measure of the average speed of change in hidden activations for the task. A low Δh indicates that the hidden state is changing slowly over time, suggesting that the network is in or near an attractor state where the hidden activations are relatively stable. Conversely, a high Δh indicates more rapid changes in hidden activity, implying less attractor-heavy dynamics.

As shown in Figure A1, tasks like N-Bits Flip-Flop and Path Integration, which have the highest task complexity and lowest dynamical degeneracy, span opposite ends of the "attractor" spectrum. Interestingly, Delayed Discrimination, which exhibits the highest dynamical degeneracy, shows an intermediate attractor score. This suggests that the attractor score, as measured by the speed of change in hidden activity, does not fully account for the dynamical degeneracy observed in the networks. Therefore, while attractor-like behavior influences network dynamics, it is not the sole factor determining dynamical degeneracy across tasks.

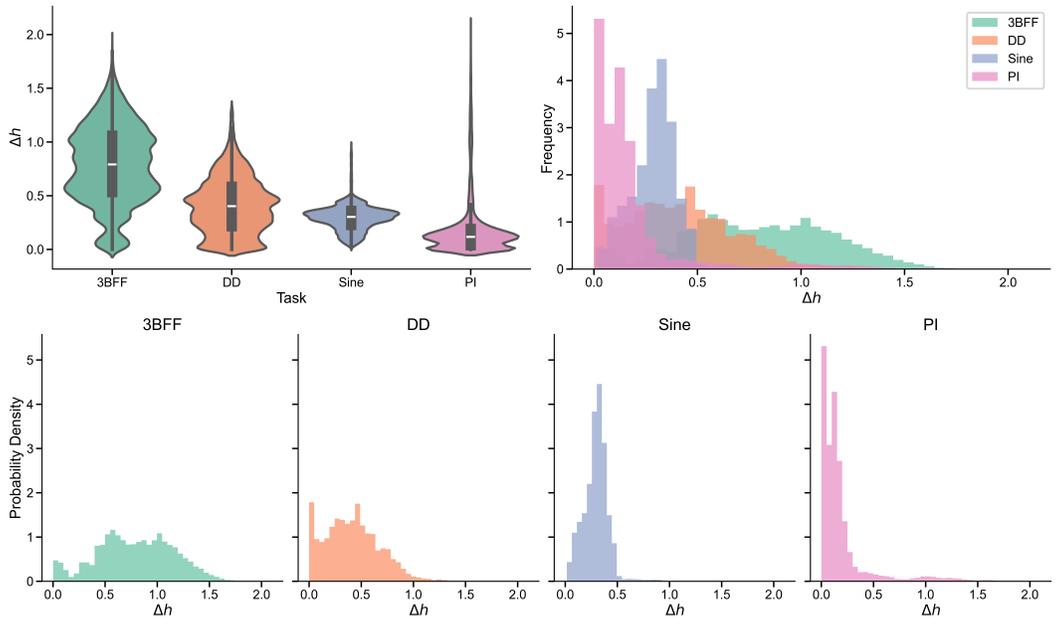


Figure A1: **Normalized speed of change in hidden state activities for each task:** (A) Speed of change in hidden state activity (Δh) for all trials in the training set plotted in violin plot for each task (B) Overlaid histogram of (Δh) for all tasks (C) Individual (Δh) for each task.

C DYNAMICAL SIMILARITY ANALYSIS (DSA)

Briefly, DSA proceeds as follows: Given two RNNs with hidden states $\mathbf{h}_1(t) \in \mathbb{R}^n$ and $\mathbf{h}_2(t) \in \mathbb{R}^n$, we first generate a delay-embedded matrix, \mathbf{H}_1 and \mathbf{H}_2 , by sampling several hidden-state trajectories from each RNN. Next, for each delay-embedded matrix, we use Dynamic Mode Decomposition (DMD) (Schmid, 2022) to extract linear forward operators \mathbf{A}_1 and \mathbf{A}_2 of the two systems' dynamics. Finally, a Procrustes distance between the two matrices \mathbf{A}_1 and \mathbf{A}_2 is used to quantify the dissimilarity between the two dynamical systems and provide an overall DSA score, defined as: $d_{\text{Procrustes}}(\mathbf{A}_1, \mathbf{A}_2) = \min_{\mathbf{Q} \in O(n)} \|\mathbf{A}_1 - \mathbf{Q}\mathbf{A}_2\mathbf{Q}^{-1}\|_F$ where \mathbf{Q} is a rotation matrix from the orthogonal group $O(n)$ and $\|\cdot\|_F$ is the Frobenius norm. This metric quantifies how dissimilar the dynamics of the two RNNs are after accounting for orthogonal transformations. We quantify Dynamical Degeneracy across many RNNs as the average pairwise distance between pairs of RNN neural-dynamics/hidden-state trajectories

Task	Number of delay	Rank
N-Bits Flip Flop	30	1000
Delayed Discrimination	20	100
Sine Wave Generation	30	100
Path Integration	30	100

Table 2: Hyperparameters for DSA

D PERMUTATION-INDEPENDENT DISTANCE BETWEEN WEIGHTS

To quantify the dissimilarity between recurrent weight matrices in a permutation-independent manner, we use a variant of the Frobenius norm that allows for optimal row and column permutations. Given two RNNs with weight matrices $\mathbf{W}_1 \in \mathbb{R}^{n \times n}$ and $\mathbf{W}_2 \in \mathbb{R}^{n \times n}$, we define the Permutation-Independent Frobenius (PIF) distance as:

$$d_{\text{PIF}}(\mathbf{W}_1, \mathbf{W}_2) = \min_{\mathbf{P}_1, \mathbf{P}_2 \in \mathcal{P}(n)} \|\mathbf{W}_1 - \mathbf{P}_1 \mathbf{W}_2 \mathbf{P}_2\|_F$$

where $\mathcal{P}(n)$ is the set of permutation matrices of size $n \times n$, and $\|\cdot\|_F$ denotes the Frobenius norm.

This distance is calculated using an iterative coordinate-descent optimization strategy with multiple random restarts to avoid local minima. At each iteration, both the row permutation \mathbf{P}_1 or the column permutation \mathbf{P}_2 is optimized using a linear-sum-assignment approach as described in (Ainsworth et al., 2022).