

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 TEMPORAL SUPERPOSITION AND FEATURE GEOMETRY OF RNNs UNDER MEMORY DEMANDS

Anonymous authors

Paper under double-blind review

ABSTRACT

Understanding how populations of neurons represent information is a central challenge across machine learning and neuroscience. Recent work in both fields has begun to characterize the representational geometry and functionality underlying complex distributed activity. For example, artificial neural networks trained on data with more features than neurons compress data by representing features non-orthogonally in so-called *superposition*. However, the effect of time (or memory), an additional capacity-constraining pressure, on underlying representational geometry in recurrent models is not well understood. Here, we study how memory demands affect representational geometry in recurrent neural networks (RNNs), introducing the concept of temporal superposition. We develop a theoretical framework in RNNs with linear recurrence trained on a delayed serial recall task to better understand how properties of the data, task demands, and network dimensionality lead to different representational strategies, and show that these insights generalize to nonlinear RNNs. Through this, we identify an effectively linear, dense regime and a sparse regime where RNNs utilize an interference-free space, characterized by a phase transition in the angular distribution of features and decrease in spectral radius. Finally, we analyze the interaction of spatial and temporal superposition to observe how RNNs mediate different representational tradeoffs. Overall, our work offers a mechanistic, geometric explanation of representational strategies RNNs learn, how they depend on capacity and task demands, and why.

1 INTRODUCTION

A major goal in both machine learning and neuroscience is to understand how populations of neurons represent information, and why certain representational geometries are preferred in different settings. Characterizing different strategies can lead to increased interpretability of artificial neural networks (ANNs), as well as help us better understand the functionality of certain brain regions.

A prevalent theme in neuroscience has been the pursuit of highly-specialized functional components in the brain. Much classical work aimed to demonstrate the functions of individual neurons (Yuste, 2015), such as simple and complex cells in the primary visual cortex (Hubel & Wiesel, 1962) or place cells in the hippocampus (O’Keefe & Dostrovsky, 1971). More recently however, large scale population recordings have revealed that neural computation is often distributed across many neurons (Yuste, 2015; Saxena & Cunningham, 2019; Ebitz & Hayden, 2021). It’s thought that such coding strategies are important for flexible, complex behavior (Fusi et al., 2016; Tye et al., 2024).

A parallel line of work has investigated representational geometry in ANNs. While earlier work tried to develop models where neurons have specialized functions (disentangled representations or monosemantic neurons) (Bengio et al., 2014; Olah et al., 2020; Cammarata et al., 2020), recent work has focused on *superposition*, a phenomenon that characterizes a class of distributed representations.

In particular, ANNs are often challenging to interpret due to *polysemanticity*, where neurons respond to many unrelated mixtures of inputs. One explanation for polysemanticity is the superposition hypothesis (Elhage et al., 2022), which posits that ANNs utilize a compression strategy where they represent more features (e.g., directions in activation space) than there are neurons by allowing some interference (non-zero dot products) between them. Superposition is particularly effective when features are sparse, although it leads to features no longer mapping onto individual neurons. In the era of big data, even in models with millions of parameters, superposition appears ubiquitous (Bricken et al., 2023; Templeton et al., 2024) – an effective strategy arising from the inherent sparsity of features in real-world data. A major focus of mechanistic interpretability is now finding new

054 techniques to extract such features from superposition. Fewer studies, however, have focused on
 055 understanding the feature geometry induced by superposition (Sharkey et al., 2025).
 056

057 Previous work in interpretability has not considered the effect of *time* (or memory) as a capacity-
 058 constraining pressure. While less of an issue in feedforward and transformer architectures, this
 059 pressure does arise with recurrence. Recurrent neural networks (RNNs) are important for learning
 060 tasks with temporal dependencies and models of dynamical systems. They are commonly used
 061 as cognitive models in neuroscience (Barak, 2017), as they often replicate the neural activity of
 062 animals when trained on similar tasks (Vyas et al., 2020). Additionally, RNNs are increasingly
 063 popular for long-range sequence modeling due to their computational and memory efficiency (Gu
 064 et al., 2022b; Orvieto et al., 2023; Gu & Dao, 2024). Understanding the effect of capacity through
 065 time on representational geometry is important, as it affects maintenance of long-range dependencies
 066 and could have implications for cognition with limited resources, such as working memory.
 067

068 In this work, we study how time, or task (memory) demands, affects representational geometry and
 069 capacity in RNNs through the lens of feature superposition. We develop a theoretical framework
 070 to better understand representational strategies employed by RNNs and how data and network di-
 071 mensionality interacts with memory demands. Our results characterize and explain behavior across
 072 different recurrent architectures and task settings, providing insight into what RNNs learn and why.
 073 Our contributions are as follows:

- 074 • We introduce the concept of temporal superposition in RNNs.
 075
- 076 • We distinguish two forms of interference – projection and composition interference – and
 077 show how they impact behavior.
 078
- 079 • We derive an expression for the loss on a simple recall task that decomposes into four
 080 interpretable terms, allowing us to explain the geometric strategy employed by the RNN.
 081
- 082 • We study the learning dynamics of RNNs in terms of changing representational geometry.
 083
- 084 • In RNNs with ReLU nonlinearities, we identify the existence of an interference-free space
 085 into which many feature directions can be tightly packed.
 086
- 087 • We identify a phase transition in the geometry between the dense and sparse regimes.
 088
- 089 • We study the interaction of spatial and temporal superposition, and how RNNs mediate this
 090 tradeoff depending on task requirements.

091 2 KEY INTUITIONS: SPATIAL & TEMPORAL SUPERPOSITION

092 Here we introduce the key ideas we will develop formally in the remainder of the paper. First, we
 093 review *spatial* superposition, which was studied extensively by Elhage et al. (2022) in feedforward
 094 networks, and then extend these concepts to characterize *temporal* superposition in RNNs.
 095

096 **Spatial superposition.** We follow the linear representation hypothesis and assume that features
 097 (loosely defined as interpretable properties of the input; see Elhage et al. (2022); Park et al. (2024))
 098 are represented as directions in activation space. When there are more features of the data than
 099 neurons and features are sparse (i.e., do not commonly co-occur), it becomes optimal for neurons to
 100 represent features non-orthogonally (in shared dimensions of the activation space) so as to compress
 101 more in (Figure 1a). This comes at the cost of possible interference between co-occurring features,
 102 but if features are sufficiently sparse, this can be outweighed by the benefit of representing more
 103 features. Spatial superposition only becomes viable in the presence of a nonlinearity that reduces
 104 interference between features; linear networks instead learn a PCA of the most important features.
 105

106 **Temporal superposition.** In addition to spatial superposition, we claim that RNNs exhibit another
 107 form of superposition due to the axis of time and that this phenomenon is fundamentally different
 108 from the feedforward case. In particular, in addition to having some spatial component, we can think
 109 of each input feature as also having a temporal component dependent on its sequential position. This
 110 means that features are represented differently depending on which timestep they occur at, even if
 111 the input itself is the same. In other words, representations of features are now determined by
 112 “when” just as much as “what”. For example, if an RNN receives an impulse of some feature A
 113 as input at timestep t , the representation of this input will move through a set of distinct feature
 114 directions as time goes on and the impulse grows older, until it ceases to be task-relevant (as in

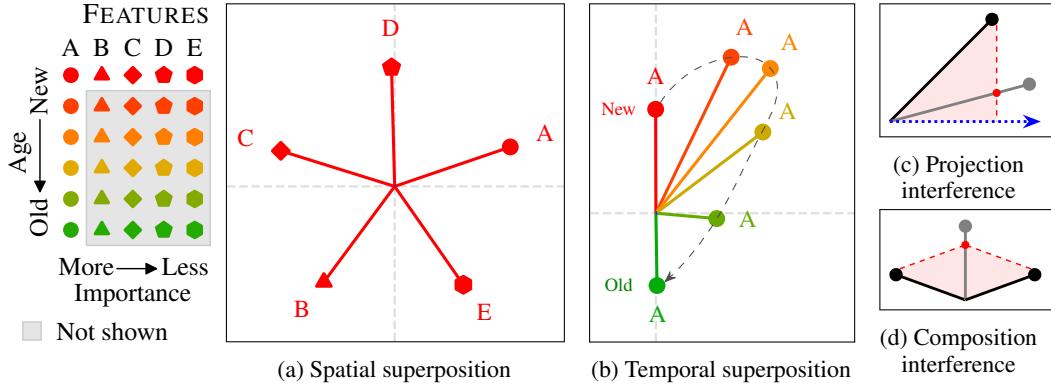


Figure 1: **Spatial and temporal superposition both occur in RNNs and are characterized by different representational strategies.** (a) When there are more input features (5: A-E) than hidden dimensions (2) and features are sparse, it becomes favorable to compress these features into the activation space non-orthogonally in spatial superposition. (b) When an input feature (A) must be held in memory for more time (red: 0 timesteps old; green: 5 timesteps old) than hidden dimensions (2) and features are temporally sparse, it becomes favorable to utilize temporal superposition to exploit the interference-free space opposite the output feature. (c) Projection interference (red dashed line) occurs when the activation of a feature (black) is read-out (blue arrow) as the activation of a different feature (gray). (d) Composition interference (red dashed line) occurs when the activation of multiple features (black) is linearly combined into an activation that imitates another feature (gray).

Figure 1b). If features are held over longer periods of time (because they remain task-relevant), more feature directions are compressed into the hidden state. The hidden state therefore acts as a bottleneck, such that RNNs are forced to either forget features or represent more features than dimensions (in superposition) as the length of the task-relevant input window increases. In addition to compressing more features due to memory demands, features must also be read out at particular timesteps, rather than being immediately available as in feedforward models, giving rise to different behavior from that of spatial superposition. We distinguish *temporal* superposition as resulting from representing features across a longer period of time (higher memory demand) in a lower-dimensional activation space, whereas we refer to *spatial* superposition as representing more input features (higher-dimensional data feature space) in a lower-dimensional activation space.

3 MATHEMATICAL SETUP

3.1 RNN MODEL

We study a RNN parameterized by matrices $W_x \in \mathbb{R}^{N_h \times N_x}$, $W_h \in \mathbb{R}^{N_h \times N_h}$, $W_y \in \mathbb{R}^{N_h \times N_y}$ with a hidden state $\mathbf{h}_t \in \mathbb{R}^{N_h}$ that receives an input $\mathbf{x}_t \in \mathbb{R}^{N_x}$ at each timestep t and produces an output $\hat{\mathbf{y}}_t \in \mathbb{R}^{N_y}$. The RNN is given by

$$\mathbf{h}_t = W_x \mathbf{x}_t + W_h \sigma_h(\mathbf{h}_{t-1}) \quad \hat{\mathbf{y}}_t = \sigma_y(W_y^\top \mathbf{h}_t) \quad (1)$$

where σ represents either a linear activation or ReLU, depending on the setting we consider (see Appendix F.1 for discussion on applying the activation directly to \mathbf{h}_{t-1}). We refer to the model with linear recurrence and readout as a linear RNN, the model with linear recurrence and nonlinear readout as a state space model (SSM), and the model with nonlinear recurrence and readout as a nonlinear RNN. We first focus on the setting where both σ_h, σ_y are linear and later consider cases of nonlinearity. We initialize the hidden state \mathbf{h}_0 to $\mathbf{0}$. In the case where σ is linear, this yields

$$\mathbf{h}_t = \sum_{i=1}^t W_h^{t-i} W_x \mathbf{x}_i \quad (2)$$

3.2 FEATURE DIRECTIONS

We define $s := t - i$ to indicate how many ‘recurrences’ an input from timestep i has undergone at timestep t (or the ‘time window’ it has been in the RNN for). This allows us to account for a feature

162 in terms of how long it has been in the RNN relative to a particular timestep. Letting $W_s := W_h^s W_x$,
 163 we can rewrite the expression for the hidden state above as

$$165 \quad \mathbf{h}_t = \sum_{s=0}^{t-1} W_h^s W_x \mathbf{x}_{t-s} = \sum_{s=0}^{t-1} W_s \mathbf{x}_{t-s} \quad (3)$$

167 To isolate the effect of temporal superposition and make visualization easier, we first consider scalar
 168 inputs and outputs ($N_x = N_y = 1$) and a two-dimensional hidden size ($N_h = 2$). After building a
 169 thorough understanding of temporal superposition, in Section 4.5 we introduce spatial superposition
 170 by increasing the dimensionality of the data ($N_x = N_y = 5$) and study its interaction with temporal
 171 superposition. We note that our framework and main results naturally extend to multidimensional
 172 inputs and outputs, and higher dimensional hidden spaces (Appendix G). For clarity, we refer to
 173 $W_y, W_s \in \mathbb{R}^{2 \times 1}$ as $\mathbf{w}_y, \mathbf{w}_s$ to indicate that they are vectors, and to $\mathbf{x}, \mathbf{y}, \hat{\mathbf{y}} \in \mathbb{R}^1$ as x, y, \hat{y} to
 174 indicate they are scalars (the equations above remain otherwise the same).

175 In this form, we can see how each input feature x_{t-s} is independently and linearly represented in
 176 the hidden state in the direction given by \mathbf{w}_s (the feature direction). This makes it clear that, at time
 177 t , the model has access to the entire history of t input features, but that for $t > N_h$, the hidden
 178 state inevitably becomes bottlenecked. Further, this form illustrates the role of the readout \mathbf{w}_y in
 179 the interference between different features. Although there are up to t feature directions ($\mathbf{w}_{s=0:t-1}$)
 180 contained within the hidden state, only the feature directions that project onto \mathbf{w}_y at timestep t will
 181 affect the RNN output \hat{y}_t . For visual clarity in our main figures, we tie the readout $\mathbf{w}_y := \mathbf{w}_{s=k}$ (see
 182 definition of k below; discussion and figures replicated without weight tying in Appendix I.2; full
 183 experimental details in Appendix I).

184 3.3 K-DELAY TASK

186 In order to directly control for the time span for which features remain task-relevant, we consider
 187 the k -delay task (Jaeger, 2002), in which the model is trained to reproduce the input sequence after
 188 a fixed delay of k timesteps. In particular, the RNN is tasked with producing the target output
 189 $y_t = x_{t-k}$ at each timestep ($y_t = 0$ for $t \leq k$). Therefore k acts as a control parameter that specifies
 190 how long an input feature must be maintained in the hidden state for successful task performance.
 191 We note that this task is essentially an extension of the setup in Elhage et al. (2022) to the temporal
 192 domain and that the two are identical for $k = 0$. We use a squared-error loss given by

$$193 \quad \mathcal{L} = \sum_{t=1}^T \|y_t - \hat{y}_t\|^2 = \sum_{t=1}^k \|0 - \hat{y}_t\|^2 + \sum_{t=k+1}^T \|x_{t-k} - \hat{y}_t\|^2 \quad (4)$$

196 **Task-relevant and irrelevant features.** To provide some intuition about the k -delay task, we can
 197 refer to features in terms of their utility. At each timestep $t \geq k+1$, there is an *output feature*
 198 *direction* ($\mathbf{w}_{s=k}$) which functions to produce $\hat{y}_t \approx x_{t-k}$ by projecting onto the readout \mathbf{w}_y . There
 199 are also *intermediate feature directions* (\mathbf{w}_s for $0 \leq s < k$), which represent input features ($x_{t-k+1} : x_t$) from the relevant k -length memory window. These intermediate features are held in memory to
 200 be read out at future timesteps. Both output features and intermediate features ($k+1$ features in
 201 total) are *task-relevant* features at timestep t . These features contribute to task performance either
 202 immediately or in the future, so the model is incentivized to represent them as faithfully as possible.

204 There also potentially exist historical, *task-irrelevant* features. These correspond to features (x_{t-s}
 205 represented by \mathbf{w}_s) from the more distant ($s > k$) past. Being beyond the k (shift) window, these
 206 features cannot contribute to current or future task performance. The model is therefore incentivized
 207 to forget these features (which may otherwise interfere with other task-relevant features).

209 3.4 PROJECTION AND COMPOSITION INTERFERENCE

211 Now that we've defined feature directions, we can understand how they interact and potentially
 212 interfere with each other. Here we introduce and define two forms of interference that occur in
 213 RNNs: projection interference and composition interference.

214 **Projection interference.** Projection interference occurs when the activation of a feature is read-
 215 out as an activation of a different feature, as shown in Figure 1c. This occurs when a feature is

216 represented (\mathbf{w}_s) non-orthogonally to the readout (\mathbf{w}_y), causing an unintended non-zero projection
 217 onto that read-out. For example, say feature direction $\mathbf{w}_{s=A}$ and feature direction $\mathbf{w}_{s=B}$ ($A \neq B$)
 218 are both non-orthogonal to the read-out \mathbf{w}_y . Then, an input (x_{t-A}) from timestep $t - A$ can be
 219 mistakenly read-out as having been an input (x_{t-B}) from timestep $t - B$, and vice versa.
 220

221 **Composition interference.** Composition interference occurs when the activation of multiple features
 222 is linearly combined into an activation that imitates another feature (Figure 1d). In the example
 223 above, if the RNN receives inputs x_{t-A} and x_{t-B} , their feature directions ($\mathbf{w}_{s=A}$, $\mathbf{w}_{s=B}$) are co-
 224 activated and linearly combine to form $\mathbf{w}_{s=A} + \mathbf{w}_{s=B}$. A consequence is that if there is another
 225 feature direction $\mathbf{w}_{s=C}$ similar to this combination, the activation for inputs x_{t-A} and x_{t-B} will be
 226 indistinguishable from that of a single input x_{t-C} from timestep $t - C$, which never occurred.
 227

228 4 RESULTS

229 4.1 LOSS DECOMPOSITION REVEALS GEOMETRIC STRATEGY

230 In order to better understand how the learning problem incentivizes certain geometric strategies, we
 231 begin by studying the loss of linear RNNs. By assuming temporal independence and sparsity of the
 232 data (Appendix A), we derive (Appendix B) a form of the loss comprised of four terms.
 233

$$234 \mathbb{E}[\mathcal{L}] = \sum_{t=k+1}^T \underbrace{\left(p\nu \|\mathbf{w}_y^\top \mathbf{w}_{s=k} - 1\|^2 - 2p^2\mu^2 \sum_{s \neq k}^{t-1} \mathbf{w}_y^\top \mathbf{w}_s \right)}_{\text{task benefit}} \underbrace{\left(\sum_{s \neq k}^{t-1} \mathbf{w}_y^\top \mathbf{w}_s \right)}_{\text{mean correction}} \\ 235 + \sum_{t=1}^T \underbrace{\left(p\nu \sum_{s \neq k}^{t-1} (\mathbf{w}_y^\top \mathbf{w}_s)^2 \right)}_{\text{projection interference cost}} \underbrace{\left(p^2\mu^2 \sum_{s \neq s'}^{t-1} (\mathbf{w}_y^\top \mathbf{w}_s) \cdot (\mathbf{w}_y^\top \mathbf{w}_{s'}) \right)}_{\text{composition interference}} \quad (5) \\ 236 \\ 237 \\ 238 \\ 239 \\ 240 \\ 241$$

242 p controls temporal sparsity (how frequently features occur in time; smaller p corresponds to higher
 243 sparsity), and μ and ν are the mean and variance of the input distribution, respectively. By studying
 244 the terms above, we can understand the competing incentives in the loss.
 245

246 First, the task benefit term is the value of successfully performing the task by aligning the feature
 247 direction $\mathbf{w}_{s=k}$ of the input x_{t-k} to the readout \mathbf{w}_y . This corresponds to producing x_{t-k} at time
 248 t , as required by the task. The mean correction term acts to offset any non-zero mean of the input
 249 distribution by exploiting projection interference – the RNN uses projection interference ($\mathbf{w}_y^\top \mathbf{w}_{s \neq k}$)
 250 as a bias in the absence of one and, in fact, the term disappears if we include a bias term in the output
 251 or trivially if the mean μ is 0.

252 Next, we can see the effects of interference in the loss. The projection interference cost introduces
 253 a penalty on feature directions \mathbf{w}_s that project onto the readout \mathbf{w}_y at the incorrect time ($s \neq k$).
 254 Additionally, the composition interference term comes into play when there are multiple features
 255 simultaneously active. Geometrically, this term penalizes positive correlations between \mathbf{w}_s vectors
 256 while rewarding negative correlations, with respect to their projection onto \mathbf{w}_y , such that negative
 257 (destructive) interference is preferred over positive (constructive) interference. This essentially en-
 258 courages \mathbf{w}_s vectors to spread out in activation space as much as possible, ideally forming antipodal
 259 pairs (similar to Elhage et al. (2022) and analogous to the form in Saxe et al. (2014)).
 260

261 We train linear RNNs on the k -delay task and see that our expected value of the loss closely pre-
 262 dicted the empirical loss (Figure 2 *top*; Appendix C for more discussion). We also observe several
 263 different stages of learning, corresponding to unique changes in each of the loss terms that map
 264 onto specific geometric configurations (Figure 2 *bottom*). In particular, the RNN initially aligns all
 265 feature directions (\mathbf{w}_s) to the readout (\mathbf{w}_y) (quantified by output projection $\mathbf{w}_y^\top \mathbf{w}_s$), corresponding
 266 to a decrease in the task error term and increase in magnitude in the other three terms. After this
 267 initial alignment, feature directions begin to spread out in activation space based on their temporal
 268 ordering, causing the task error, mean correction, and composition interference terms to decrease
 269 in magnitude, while projection interference increases. This appears to be analogous to the learning
 270 phases of data eigenvector alignment (Atanasov et al., 2022) and scaling of eigenvalues (Saxe et al.,
 271 2014; Proca et al., 2025). Furthermore, the ‘staircase’ loss is indicative of saddle-to-saddle dynamics
 272 (Jacot et al., 2022) and geometric restructuring (Haputhanthri et al., 2024). The final arrangement

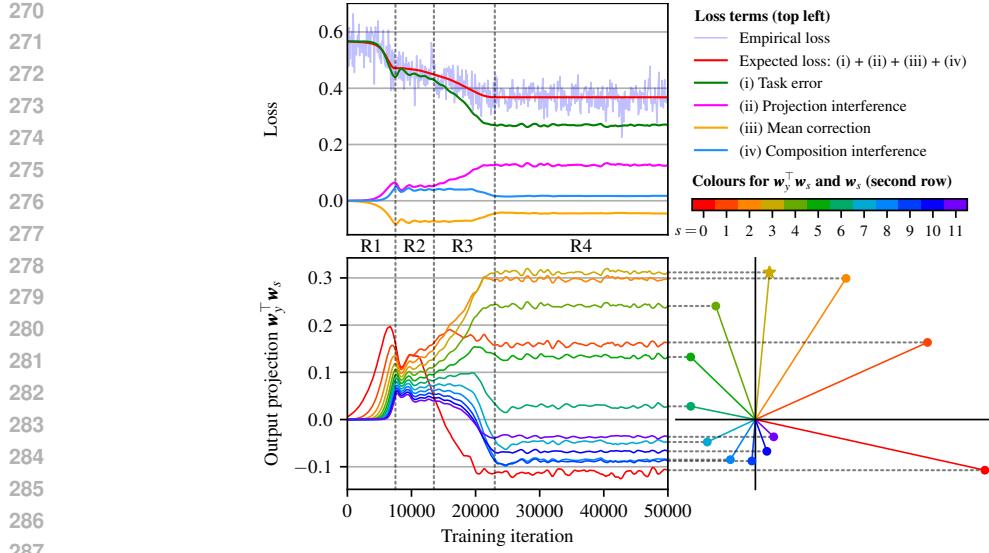


Figure 2: **Learning dynamics characterized by initial readout alignment, then separation of feature directions.** For a 2D linear RNN on the 3-delay task. (Top) The expected loss $\mathbb{E}[\mathcal{L}]$ matches the empirical loss curve. The four loss terms exhibit different dynamics during training corresponding to the geometric configuration of feature directions. (Bottom left) Quantifying the projection onto the readout for each feature direction via the output projection $w_y^\top w_s$ (for $0 \leq s < 12$), there is an initial readout alignment followed by temporally-ordered separation. (Bottom right) The final arrangement of w_s vectors converges to a spiral, matched to the output projections. The output feature $w_{s=3}$ (star) has the highest projection onto w_y .

of feature directions the RNN converges to is shown in Figure 3b. Here, in the linear case, the RNN learns to downscale and rotate features with each recurrence (such that old features spiral into the origin), implementing a ‘smooth’ forgetting. Indeed, in the 2D case, a spiral sink (e.g., spectral radius $\|\mathbf{w}_{s=0}\| < 1$) is the optimal solution for a linear RNN in the k -delay task (Appendices D.2 and D.3), as old features gradually fade from the hidden state.

4.2 FEATURES GROUP INTO INTERFERENCE-FREE SPACE

Building on the previous section, here we consider the impact of adding a nonlinearity to the readout w_y ($\sigma_y(\cdot) = \text{ReLU}(\cdot)$). The nonlinear setting makes deriving analytic solutions to the loss more challenging. We therefore approximate the expectation of the loss in the limit of high temporal sparsity (Appendix E.1), yielding

$$\mathbb{E}[\mathcal{L}] \approx p\nu \left(\underbrace{\sum_{t=k+1}^T (\text{ReLU}(\mathbf{w}_y^\top \mathbf{w}_{s=k}) - 1)^2}_{\text{task benefit}} + \underbrace{\sum_{t=1}^T \sum_{s \neq k}^{t-1} \text{ReLU}(\mathbf{w}_y^\top \mathbf{w}_s)^2}_{\text{projection interference cost}} \right) \quad (6)$$

This expression resembles the one in the linear case, but the inclusion of the ReLU activation has a significant impact on its geometric interpretation. Due to the ReLU, the model only produces output for vectors that have a positive projection onto w_y . Thus, all w_s vectors in the half-space opposite of w_y do not contribute to projection interference. In fact, in the extremely sparse regime (where composition interference becomes negligible), this half-space essentially becomes *interference free*. This reveals a remarkable incentive for the model to take advantage of this phenomenon by packing as many $w_{s \neq k}$ vectors into this half-space as possible (Figure 3a).

To test this prediction, we train linear RNNs with ReLU readouts (SSMs) on the k -delay task at various levels of sparsity. Because the recurrence is still linear, isolation of feature directions $w_{s \neq k}$ into this space is not always perfectly possible as w_s can only be spaced equally (along an elliptical spiral). Despite this, by looking at Figure 3c, we can see that when sparsity is high, models learn to minimize projection interference by grouping the largest feature directions into the interference-free space. The SSM employs an approximation of the strategy in which the largest feature directions w_s occupy the interference-free space while smaller feature directions lie outside of it. Within the

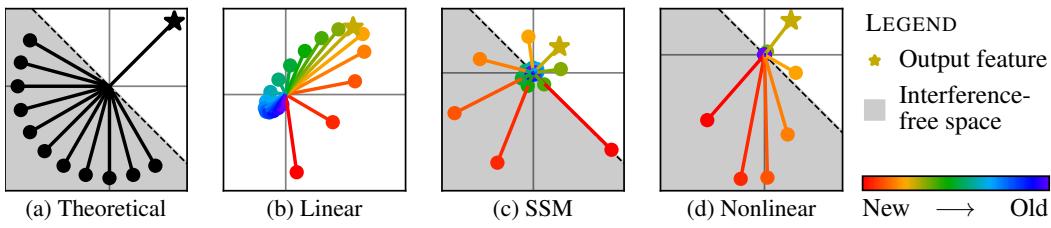


Figure 3: **Theoretical predictions match representational strategies of expressive models.** (a) A ReLU readout creates an interference-free space. In the panel, we display an idealized form, where all intermediate feature directions group into the free space (*gray shading*) and only the output feature direction lies outside of it (*star*). (b) Linear RNNs lack an interference-free space and instead arrange old features to spiral into the origin (c) In the sparse regime, SSMs minimize projection interference by grouping the largest feature directions into the interference-free space. (d) Nonlinear RNNs are expressive enough to fully exploit the interference-free space by grouping all of the intermediate features, separate from the output feature, and implement sharp forgetting.

constraints of linear recurrence, this strategy still minimizes projection interference by exploiting the interference-free space. Additionally, we observe that the spectral radius (i.e., $\|w_{s=0}\|$) increases with k , regardless of sparsity, in order to overcome projection interference (Appendix E.2 for additional simulations).

4.3 PHASE TRANSITION FROM DENSE TO SPARSE CONFIGURATION

By varying sparsity, we can observe the existence of two discrete regimes (Figure 4). When the SSM is trained on data that is dense (low sparsity), it learns a ‘dense-regime’ solution where it arranges features into a spiral sink, similar to the fully-linear RNN. Although it has a nonlinearity on the readout, the SSM does not group the largest feature directions into the interference-free space in this case. We speculate that this is due to the increased likelihood of composition interference occurring when features are dense. In particular, if the SSM groups the largest feature directions into the (negative) interference-free space (e.g., $w_{s=A}$) and they additively sum with smaller (positive) feature directions ($w_{s=B}$), the ReLU will cause the model to output 0 ($\text{ReLU}(w_{s=A} + w_{s=B}) = 0$).

Instead, when inputs are sparse, it is optimal for the SSM to take advantage of the interference-free space, characterizing a ‘sparse-regime’ solution. The most noticeable difference in geometry between dense and sparse regimes is the angle that the task-relevant feature directions span ($k\theta$ is the angle traversed from $w_{s=0}$ to $w_{s=k}$), which exhibits a sharp change as sparsity is varied. In the dense regime, task-relevant feature directions group into a smaller cone ($\approx 90^\circ$). Instead, in the sparse regime, task-relevant feature directions spread out into approximately 270° of the plane, traversing the entirety of the interference-free space to reach the readout direction. By varying sparsity, we can interpolate between these two regimes and observe a phase transition in $k\theta$, accompanied by a decrease in spectral radius ($\rho = \|w_{s=0}\|$). The difference in spectral radius is likely because in the dense regime, the SSM uses larger feature directions to compensate for projection interference (by having a large projection onto the readout that outweighs other projections). In the sparse regime, projection interference is less prevalent and therefore the SSM does not need a spectral radius as large.

4.4 NONLINEAR RNNs EXPLOIT INTERFERENCE-FREE SPACE

We’ve used RNNs with linear recurrence to build our understanding of behavior in analytically tractable settings. However, these models have limited expressivity because their dynamics are constrained to a particular form, such as a spiral. Here, we consider models with nonlinear recurrence and find that, in the sparse regime, they consistently implement our predicted ideal strategy of packing as many w_s feature directions into the interference-free space as possible.

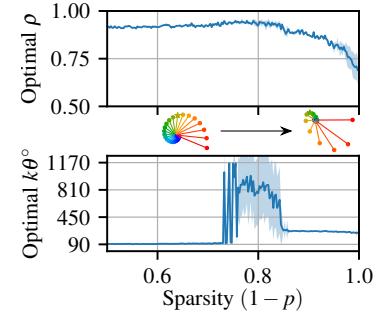


Figure 4: **There exists a phase transition in optimal feature geometry between dense and sparse regimes.** As temporal sparsity increases, (bottom) the angular distribution of feature directions ($k\theta$) exhibits a phase transition, (top) accompanied by a decrease in spectral radius (ρ).

We consider the nonlinear RNN model, now with both $\sigma_h(\cdot), \sigma_y(\cdot) = \text{ReLU}(\cdot)$. Although it's not possible to find a simplified expression for the feature direction \mathbf{w}_s analytically, it still acts as the direction the input x_{t-s} is represented in the limit of high sparsity ($p \rightarrow 0$; Appendix F.2). Furthermore, the interference-free space still exists and can be better exploited as the RNN can form more optimal arrangements of \mathbf{w}_s vectors due to increased expressivity of the recurrence. In Figure 3d, we verify our predictions: nonlinear RNNs learn to pack the k intermediate features into the interference-free space, with only $\mathbf{w}_{s=k}$ lying outside (Appendix F.3 for more simulations). Often, the vectors compress into a single quadrant. This is because the nonlinearity creates a *privileged basis* in the hidden state – positive activations are unaffected by the ReLU, while negative activations will be set to zero in the next timestep. Hence, task-relevant feature directions arrange themselves in the positive quadrant of \mathbf{h}_t .

While a RNN with linear recurrence can only implement smooth forgetting (with spectral radius $\|\mathbf{w}_{s=0}\| < 1$) by shrinking an input's contribution to the hidden state over time, the ReLU activation makes it possible to immediately forget a feature by sending it to the negative quadrant of \mathbf{h}_t . As a result, the nonlinear RNN can implement sharp forgetting, which enables the model to represent only task-relevant features (and remove the possibility of interference from the distant past).

4.5 INTERACTION OF SPATIAL AND TEMPORAL SUPERPOSITION

We've focused on the case of scalar inputs and outputs in 2D space to isolate the effects of temporal superposition and for ease of visualization. We now study the interaction of spatial and temporal superposition by considering vector inputs (to introduce spatial superposition) and changing k (temporal superposition). Recall that for $k = 0$, the RNN is tasked with imitating a feedforward network where it immediately outputs the input; hence, there is no need to represent features from earlier timesteps and we recover the setting of pure spatial superposition from Elhage et al. (2022).

In Figure 5, as k increases, we see how a nonlinear RNN attempts to balance representing the most important features (scaling of loss) across time (A being most important, E being least). Initially, when $k = 1$, it represents features A, B, and C for 2 timesteps and drops D and E altogether. For higher k , the RNN eventually drops all features except A. Hence, we can see a tradeoff between the RNN's incentive to represent multiple input features and the duration of time each must be represented for. The RNN's strategy is 'all-or-none': to gain any advantage from representing a specific feature, the RNN must be able to maintain it in memory for all $k + 1$ timesteps – otherwise it will not meaningfully contribute to decreasing the loss. If the RNN does not have sufficient capacity to represent the feature for all $k + 1$ timesteps, it won't represent the feature at all. This is why we see $k + 1$ feature directions for feature A for all k , while other features only occur for lower k when the RNN has sufficient capacity.

Higher-dimensional hidden states. Up to now, we have restricted the hidden state (N_h) of our models to 2 dimensions for easier visualization and interpretability. To extend our setting to higher-dimensional hidden states, we train nonlinear RNNs on 10-dimensional input ($N_x = 10$) on the 2-delay task, varying hidden size ($N_h = 2, 5, 10$), and measure the projection of each feature direction onto the readout ($W_y^\top W_s$). Based on our previous results, we would expect $W_y^\top W_{s=2}$ to have a diagonal of positive outputs (corresponding to the output feature directions positively projecting onto the readout: for the correct output at the correct time). Moreover, we would expect the rest of the entries in the matrix (as well as all of $W_y^\top W_{s \neq 2}$) to be negative or 0, lying in the interference-free space. Across all hidden sizes, we see this exact strategy (Figure 11 and Appendix G), with RNNs with larger hidden sizes simply capturing more features along the diagonal of $W_y^\top W_{s=2}$ (i.e., the same all-or-none effect described above). Finally, we quantify this behavior by computing the mean of the non-output feature direction projections onto the readout (i.e., $\text{mean}(W_y^\top W_{s \neq k})$) which should be negative in an optimal model, and the mean of the output feature direction projections onto the readout (i.e., $\text{mean}(\text{diag}(W_y^\top W_{s=k}))$), which should be positive. We train RNNs with hidden size 100 on a 2-delay task with 75 features and find that the best performing models group the largest feature directions into an interference-free space and project the output feature onto the readout at the appropriate time, as predicted (Figure 12).

5 RELATED WORK

Representational geometry and interpretability. The study of representational geometry can provide better understanding of how distributed activity encodes information for different behavior. ANNs can exhibit different representational geometries, reflective of the tasks trained on (Johnston

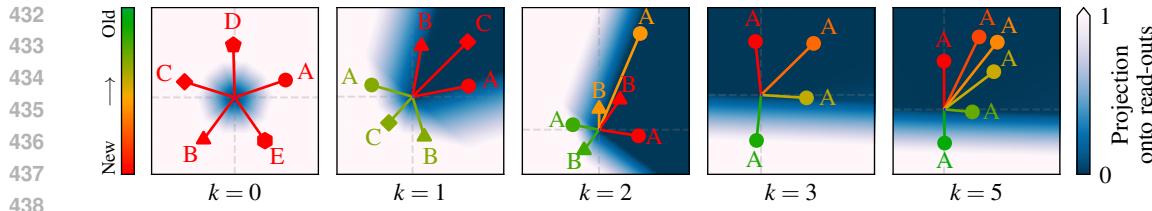


Figure 5: **There exists a tradeoff between spatial and temporal superposition.** By setting $k = 0$, we recover purely spatial superposition (as there is no memory demand) and all 5 features are arranged in a pentagon (feature legend as in Figure 1; interference-free space indicated by blue shading). As k and the corresponding memory demand increases, the RNN prioritizes the most important features in an all-or-none fashion. There is a preference to represent one feature for all $k + 1$ relevant timesteps, as opposed to representing several different features for a shorter duration.

& Fusi, 2023; 2024), dataset size (Henighan et al., 2023), architecture (Jacot et al., 2018; Chizat et al., 2019), and weight parameterization (Flesch et al., 2022; Braun et al., 2022). In addition to advancing our understanding of the solutions ANNs learn in different settings, this can also help neuroscientists identify computational structures in the brain and understand their functional roles (Saxe et al., 2020; Ostojic & Fusi, 2024). In this work, we introduce the concept of temporal superposition and study its effect on representational geometry. Our work is largely inspired by Elhage et al. (2022), one of the first papers to formalize superposition in toy models and study it explicitly. Since then, the field of mechanistic interpretability has centered around the problem of superposition and developing methods to identify meaningful features from model activations (Bricken et al., 2023; Templeton et al., 2024). While these techniques have continued to advance, there has been notably less work devoted to understanding the feature geometry induced by superposition (Sharkey et al., 2025). Here we expand the study of superposition to recurrent architectures to show how memory acts as a capacity constraint, inducing superposition, and how this affects underlying geometry.

Recurrent neural networks. RNNs are important for modeling temporal data and studying dynamic processes. Recently, RNNs (SSMs) with linear recurrence have become popular due to their computational and memory efficiency (Gu et al., 2022b; Orvieto et al., 2023; Gu & Dao, 2024), often initialized with complex-valued parameterizations (Gu et al., 2022a; Orvieto et al., 2023), with increased expressivity (Ran-Milo et al., 2024; Orvieto et al., 2024). We find that complex eigenvalues support temporal superposition by rotating features within an interference-free space. We also study a SSM, showing that it exploits the interference-free space in the sparse regime, but is still constrained in expressivity compared to nonlinear RNNs, resulting in different geometries (and smooth vs sharp forgetting). In neuroscience, RNNs are common for modeling (Barak, 2017), as the brain’s connectivity is highly recurrent and RNNs often replicate neural activity recorded in animals (Vyas et al., 2020; Khona & Fiete, 2021) and behavior (Ji-An et al., 2025) when trained on the same tasks. Similar to the feature geometry we see here, previous work in RNNs has observed rotational dynamics/sequential activity (Rajan et al., 2016; Orhan & Ma, 2019; Cueva et al., 2020; Zhang et al., 2021) for tasks with fixed delay, thought to encode temporal information. Moreover, other work has shown that RNNs trained on tasks with random delays instead exhibit persistent activity, in the form of fixed point attractors (Orhan & Ma, 2019; Liu et al., 2021; Xie et al., 2022b), similar to pure spatial superposition (for example, as we see when $k = 0$). Related to the interference-free space we study in our model, RNNs trained in motor-preparation paradigms similarly develop output-null subspaces where intermediate preparatory activity does not affect behavior (Schimel et al., 2024). There’s been substantial theoretical work on RNNs, both by neuroscientists studying properties of neural computation and by deep learning theorists (Dubreuil et al., 2022; Driscoll et al., 2024; Schuessler et al., 2024; Zucchet & Orvieto, 2024; Proca et al., 2025). One important line of theoretical work has studied low-rank RNNs (Mastrogiovanni & Ostojic, 2018; Schuessler et al., 2020a; Beirán et al., 2020; Dubreuil et al., 2022). These interpretable models have low-dimensional recurrent dynamics, allowing their exact phase portraits to be visualized; furthermore, these dynamics can be directly related to the underlying connectivity statistics. Related to our work, low-rank connectivity also acts as a form of capacity constraint, although the effects of such constraints have not been studied explicitly (but see Beirán et al. (2023) for comparison between low-versus full-rank RNNs). However, there is less work studying feature geometry in the context of capacity constraints induced by memory. Most work implicitly assumes an overparameterized regime (relative to task

486 demands) when studying properties of RNNs (Cohen-Karlik et al., 2023). One exception is François
 487 et al. (2025) which studied the k -delay task in an underparameterized linear RNN in the frequency
 488 domain. While they focused on the dense (linear) regime, here we also study the sparse regime in
 489 nonlinear RNNs, identifying novel behavior.

490 **Memory capacity.** Memory capacity has previously been studied in RNNs (White et al., 2004;
 491 Dambre et al., 2012; Ballarin et al., 2024), most classically in echo-state networks (Jaeger, 2002)
 492 (we refer to short-term memory and not other forms like associative memory). Memory capacity typ-
 493 ically refers to a temporally-dense regime. It's known that linear memory capacity is limited to the
 494 number of neurons in the hidden state N_h , corresponding to a single neuron per feature (orthogonal-
 495 ity/lack of superposition). Here, we consider how RNNs handle capacity constraints, transitioning
 496 from temporally-dense to temporally-sparse regimes. We show how training with different task de-
 497 mands leads to different geometric solutions, aimed at compression and increased (sparse-regime)
 498 capacity. Hence, we provide a mechanistic interpretation of memory capacity under constraints.

499 **Working memory and serial recall.** Our work is related to working memory in cognitive neuro-
 500 science – a cognitive function involving short-term maintenance and manipulation of information
 501 for immediate use. Similar to the k -delay task, working memory is often studied using serial recall,
 502 which has previously been modeled with RNNs (Botvinick & Plaut, 2006; Ganguli et al., 2008).
 503 Two existing theories of sequence working memory involve: (1) activity slots (Luck & Vogel, 1997;
 504 2013; Xie et al., 2022a), in which there exist a set of distinct neural subspaces for different sequence
 505 items, and (2) a resource model (Alvarez & Cavanagh, 2004; Wilken & Ma, 2005; Bays & Husain,
 506 2008), where working memory is a limited-capacity continuous resource that is shared between
 507 items (more items leads to less capacity per item). Interestingly, in the setting we consider we ef-
 508 fectively find both (Soni & Frank, 2025). RNNs arrange features from each timestep into separate
 509 ‘slots’ (directions) along which these features shift through time. In the case of superposition, slots
 510 are not orthogonal (Xie et al., 2022a), but otherwise would be with sufficient capacity. Additionally,
 511 we show that the hidden space is a continuous limited-capacity resource: as memory length (k) or
 512 input features (N_x) increase, there is more demand and features are more likely to interfere. Further,
 513 we show how limited capacity leads RNNs to represent important features, while others are dropped.

514 6 DISCUSSION

515 **Summary of results.** In this work, we study the effect of time on feature representations in RNNs,
 516 introducing the concept of temporal superposition. We identify how features can interfere through
 517 either projection or composition interference, and their corresponding effects. We derive an ana-
 518 lytical form of the loss that decomposes into four interpretable terms, which we use to explain the
 519 resulting learned geometry of RNNs. By deriving an approximation of the loss in the limit of high
 520 temporal sparsity in nonlinear RNNs, we identify the existence of an interference-free space, which
 521 RNNs exploit to minimize projection interference. By varying temporal sparsity, we see how SSMs
 522 exhibit a phase transition from an effectively linear strategy to one that uses the free space, and that
 523 this phase transition is reflected in the angular distribution of features and spectral radius. We further
 524 show that nonlinear RNNs in the sparse regime exploit the interference-free space and implement
 525 sharp forgetting. Finally, we study how spatial and temporal superposition interact as a result of
 526 different task demands and capacity constraints, and how RNNs mediate this tradeoff.

527 **Limitations and future work.** We simplify our theoretical setting by assuming temporal indepen-
 528 dence of features and studying small RNNs. We also study the sparse regime: while the assumption
 529 of sparse input features appears to be reasonable (Elhage et al., 2022), the assumption of temporal
 530 sparsity may be strong. This is dependent on the task and it's an open question how the theoretical
 531 setting considered here extends more generally. Related, we study the k -delay task, which requires
 532 reproduction of a sequence with a delay. An important future direction will be to characterize ge-
 533 ometry and behavior for tasks requiring manipulation of input information and varying memory de-
 534 mands. Finally, we note that one major assumption of the superposition hypothesis is that features
 535 are represented linearly, as directions in activation space (the linear representation hypothesis (Park
 536 et al., 2024)). A possible objection may be to what extent this work captures realistic settings in
 537 seemingly overparameterized modern-day models. Although we study a 2D case for simplicity (but
 538 see Appendix G), memory demands decrease capacity linearly with time (Jelassi et al., 2024) and
 539 consequently finite-width RNNs tasked with learning long-term dependencies will be constrained.
 Moreover, superposition has already been demonstrated in LLMs (Bricken et al., 2023; Templeton
 et al., 2024) and our study indicates that recurrence and memory will exacerbate it.

540
541

REPRODUCIBILITY STATEMENT

542

To ensure reproducibility of our work, we provide code to replicate all of our experiments and figures as supplementary material and include experimental details in Appendix I. For our theoretical results, we include our assumptions about the data in Appendix A and full derivations and proofs in Appendices B, D.1, E.1 and F.2. We also provide additional simulations to those in the main text in Appendices C, D.2, D.3, E.2, F.3 and G.

543

544

REFERENCES

545

546

547

Pau Vilimelis Aceituno, Gang Yan, and Yang-Yu Liu. Tailoring echo state networks for optimal learning. *iScience*, 23(9):101440, 2020. ISSN 2589-0042. doi: <https://doi.org/10.1016/j.isci.2020.101440>. URL <https://www.sciencedirect.com/science/article/pii/S2589004220306325>.

548

549

G. A. Alvarez and P. Cavanagh. The capacity of visual short-term memory is set both by visual information load and by number of objects. *Psychological Science*, 15(2):106–111, 2004.

550

551

552

Alexander Atanasov, Blake Bordelon, and Cengiz Pehlevan. Neural networks as kernel learners: The silent alignment effect. *International Conference on Learning Representations*, 2022.

553

554

555

Giovanni Ballarin, Lyudmila Grigoryeva, and Juan-Pablo Ortega. Memory of recurrent networks: Do we compute it right? *Journal of Machine Learning Research*, 25:1–38, 2024.

556

557

558

Omri Barak. Recurrent neural networks as versatile tools of neuroscience research. *Current Opinion in Neurobiology*, 46:1–6, 2017. ISSN 0959-4388. doi: <https://doi.org/10.1016/j.conb.2017.06.003>. URL <https://www.sciencedirect.com/science/article/pii/S0959438817300429>. Computational Neuroscience.

559

560

561

P. M. Bays and M. Husain. Dynamic shifts of limited working memory resources in human vision. *Science*, 321(5890):851–854, 2008. doi: 10.1126/science.1158023.

562

563

564

Manuel Beirán, Alexis M. Dubreuil, Adrian Valente, Francesca Mastrogioviseppi, and Srdjan Ostojic. Shaping dynamics with multiple populations in low-rank recurrent networks. *Neural Computation*, 33:1572–1615, 2020.

565

566

567

Manuel Beirán, Nicolas Meirhaeghe, Hansem Sohn, Mehrdad Jazayeri, and Srdjan Ostojic. Parametric control of flexible timing through low-dimensional neural manifolds. *Neuron*, 111(5):739–753, 2023. doi: 10.1101/2021.11.08.467806.

568

569

570

Yoshua Bengio, Patrice Y. Simard, and Paolo Frasconi. Learning long-term dependencies with gradient descent is difficult. *IEEE Transactions on Neural Networks*, 5:2:157–66, 1994.

571

572

573

Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives, 2014. URL <https://arxiv.org/abs/1206.5538>.

574

575

576

Blake Bordelon, Jordan Cotler, Cengiz Pehlevan, and Jacob A. Zavatone-Veth. Dynamically learning to integrate in recurrent neural networks. 2025.

577

578

579

Matthew M. Botvinick and David C. Plaut. Short-term memory for serial order: a recurrent neural network model. *Psychological review*, 113(2):201–233, 2006.

580

581

582

David M. Bradley and Ramesh C. Gupta. On the distribution of the sum of n non-identically distributed uniform random variables. *Annals of the Institute of Statistical Mathematics*, 54(3):689–700, 2002. doi: 10.1023/A:1022483715767. URL <https://doi.org/10.1023/A:1022483715767>. ID: Bradley2002.

583

584

585

586

Lukas Braun, Clémentine Dominé, James Fitzgerald, and Andrew Saxe. Exact learning dynamics of deep linear networks with prior knowledge. *Advances in Neural Information Processing Systems*, 35:6615–6629, 2022.

594 Trenton Bricken, Adly Templeton, Joshua Batson, Brian Chen, Adam Jermyn, Tom Con-
 595 erly, Nick Turner, Cem Anil, Carson Denison, Amanda Askell, Robert Lasenby, Yifan Wu,
 596 Shauna Kravec, Nicholas Schiefer, Tim Maxwell, Nicholas Joseph, Zac Hatfield-Dodds, Alex
 597 Tamkin, Karina Nguyen, Brayden McLean, Josiah E Burke, Tristan Hume, Shan Carter,
 598 Tom Henighan, and Christopher Olah. Towards monosemanticity: Decomposing language
 599 models with dictionary learning. *Transformer Circuits Thread*, 2023. <https://transformer->
 600 [circuits.pub/2023/monosemantic-features/index.html](https://transformer-circuits.pub/2023/monosemantic-features/index.html).

601 Nick Cammarata, Gabriel Goh, Shan Carter, Ludwig Schubert, Michael Petrov, and
 602 Chris Olah. Curve detectors. *Distill*, 2020. doi: 10.23915/distill.00024.003.
 603 <https://distill.pub/2020/circuits/curve-detectors>.

604

605 Lenaic Chizat, Edouard Oyallon, and Francis Bach. On lazy training in differentiable programming.
 606 In *Advances in Neural Information Processing Systems*, 2019.

607

608 Edo Cohen-Karlik, Itamar Menuhin-Gruman, Raja Giryes, Nadav Cohen, and Amir Globerson.
 609 Learning low dimensional state spaces with overparameterized recurrent neural nets. *Inter-*
 610 *national Conference on Learning Representations*, 2023.

611 Christopher J. Cueva, Alex Saez, Encarni Marcos, Aldo Genovesio, Mehrdad Jazayeri, Ranulfo
 612 Romo, C. Daniel Salzman, Michael N. Shadlen, and Stefano Fusi. Low-dimensional dynamics
 613 for working memory and time encoding. *Proceedings of the National Academy of Sciences*, 117
 614 (37):23021–23032, 2020. doi: 10.1073/pnas.1915984117. URL <https://www.pnas.org/>
 615 [doi/abs/10.1073/pnas.1915984117](https://doi.org/10.1073/pnas.1915984117).

616

617 Joni Dambre, David Verstraeten, Benjamin Schrauwen, and Serge Massar. Information processing
 618 capacity of dynamical systems. *Scientific Reports*, 2, 2012.

619

620 Laura Driscoll, Krishna Shenoy, and David Sussillo. Flexible multitask computation in recurrent
 621 networks utilizes shared dynamical motifs. *Nature Neuroscience*, 2024.

622

623 Alexis M. Dubreuil, Adrian Valente, Manuel Beirán, Francesca Mastrogiovanni, and Srdjan Ostojic.
 624 The role of population structure in computations through neural dynamics. *Nature Neuroscience*,
 625 25:783 – 794, 2022.

626

627 R. Becket Ebitz and Benjamin Y. Hayden. The population doctrine in cognitive neuroscience. *Neu-*
 628 *ron*, 109(19):3055–3068, 2021.

629

630 Nelson Elhage, Tristan Hume, Catherine Olsson, Nicholas Schiefer, Tom Henighan, Shauna Kravec,
 631 Zac Hatfield-Dodds, Robert Lasenby, Dawn Drain, Carol Chen, Roger Grosse, Sam McCandlish,
 632 Jared Kaplan, Dario Amodei, Martin Wattenberg, and Christopher Olah. Toy models of superpo-
 633 sition. *Transformer Circuits Thread*, 2022.

634

635 Timo Flesch, Keno Juechems, Tsvetomira Dumbalska, Andrew Saxe, and Christopher Summerfield.
 636 Orthogonal representations for robust context-dependent task performance in brains and neural
 637 networks. *Neuron*, 110:1258–1270.e11, 2022.

638

639 Alexandre François, Antonio Orvieto, and Francis R. Bach. An uncertainty principle for linear
 640 recurrent neural networks. *CoRR*, abs/2502.09287, February 2025. URL <https://doi.org/>
 641 [10.48550/arXiv.2502.09287](https://arxiv.org/abs/2502.09287).

642

643 Stefano Fusi, Earl K. Miller, and Mattia Rigotti. Why neurons mix: high dimensionality for higher
 644 cognition. *Current Opinion in Neurobiology*, 37:66–74, 2016.

645

646 Oded Galor. *Discrete Dynamical Systems*. Springer, Berlin, Germany, December 2007.

647

S. Ganguli, D. Huh, and H. Sompolinsky. Memory traces in dynamical systems. *Proc. Natl. Acad.*
 648 *Sci. U.S.A.*, 105(48):18970–18975, 2008. doi: 10.1073/pnas.0804451105.

649

650 Yoav Ger and Omri Barak. Learning dynamics of rnns in closed-loop environments, 2025. URL
 651 <https://arxiv.org/abs/2505.13567>.

648 Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. In *First*
 649 *Conference on Language Modeling*, 2024. URL <https://openreview.net/forum?id=tEYskw1VY2>.
 650

651 Albert Gu, Karan Goel, Ankit Gupta, and Christopher Ré. On the parameterization and initial-
 652 ization of diagonal state space models. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave,
 653 and Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems*, 2022a. URL
 654 <https://openreview.net/forum?id=yJE7iQSAep>.
 655

656 Albert Gu, Karan Goel, and Christopher Ré. Efficiently modeling long sequences with structured
 657 state spaces, 2022b. URL <https://arxiv.org/abs/2111.00396>.
 658

659 Udit Haputhanthri, Liam Storan, Yiqi Jiang, Adam Shai, Hakki Orhun Akengin, Mark Schnitzer,
 660 Fatih Dinc, and Hidenori Tanaka. Why do recurrent neural networks suddenly learn? bifurcation
 661 mechanisms in neuro-inspired short-term memory tasks. In *ICML 2024 Workshop on Mechanistic*
 662 *Interpretability*, 2024. URL <https://openreview.net/forum?id=njmXdqzHJq>.
 663

663 Tom Henighan, Shan Carter, Tristan Hume, Nelson Elhage, Robert Lasenby, Stanislav Fort,
 664 Nicholas Schiefer, and Christopher Olah. Superposition, memorization, and double descent.
 665 *Transformer Circuits Thread*, 01 2023. URL <https://transformer-circuits.pub/2023/toy-double-descent/>.
 666

667 Sepp Hochreiter, Yoshua Bengio, Paolo Frasconi, and Jurgen Schmidhuber. *Gradient Flow in Re-
 668 current Nets: the Difficulty of Learning Long-Term Dependencies*. IEEE, 2001.
 669

670 David H. Hubel and Torsten N. Wiesel. Receptive fields, binocular interaction and functional archi-
 671 tecture in the cat's visual cortex. *The Journal of Physiology*, 160, 1962.
 672

673 Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and gen-
 674 eralization in neural networks. In *Advances in Neural Information Processing Systems*, pp. 8571–
 675 8580, 2018.
 676

676 Arthur Jacot, François Ged, Berfin Şimşek, Clément Hongler, and Franck Gabriel. Saddle-to-saddle
 677 dynamics in deep linear networks: Small initialization training, symmetry, and sparsity. 2022.
 678 URL <https://arxiv.org/abs/2106.15933>.
 679

680 Herbert Jaeger. Short term memory in echo state networks. *GMD - German National Research
 681 Institute for Computer Science*, 01 2002.
 682

682 Samy Jelassi, David Brandfonbrener, Sham M. Kakade, and Eran Malach. Repeat after me: Trans-
 683 formers are better than state space models at copying. In *Forty-first International Conference on
 684 Machine Learning*, 2024. URL <https://openreview.net/forum?id=duRRoGeoQT>.
 685

686 Li Ji-An, Marcus Benna, and Marcelo Mattar. Discovering cognitive strategies with tiny recurrent
 687 neural networks. *Nature*, 644:993–1001, 07 2025. doi: 10.1038/s41586-025-09142-4.
 688

688 W. Jeffrey Johnston and Stefano Fusi. Abstract representations emerge naturally in neural networks
 689 trained to perform multiple tasks. *Nature Communications*, 14(1):1040, 2023. doi: 10.1038/
 690 s41467-023-36583-0.
 691

692 W. Jeffrey Johnston and Stefano Fusi. Modular representations emerge in neural networks trained
 693 to perform context-dependent tasks, 2024. URL <https://www.biorxiv.org/content/10.1101/2024.09.30.615925v1>.
 694

695 Mikail Khona and Ila Rani Fiete. Attractor and integrator networks in the brain. *Nature Reviews
 696 Neuroscience*, 23:744 – 766, 2021.
 697

698 Yichen Henry Liu, Junda Zhu, Christos Constantinidis, and Xin Zhou. Emergence of pre-
 699 frontal neuron maturation properties by training recurrent neural networks in cognitive
 700 tasks. *iScience*, 24(10):103178, 2021. ISSN 2589-0042. doi: <https://doi.org/10.1016/j.isci.2021.103178>. URL <https://www.sciencedirect.com/science/article/pii/S2589004221011469>.
 701

702 Steven J. Luck and Edward K. Vogel. The capacity of visual working memory for features and
 703 conjunctions. *Nature*, 390(6657):279–281, 1997. doi: 10.1038/36846.

704

705 Steven J. Luck and Edward K. Vogel. Visual working memory capacity: from psychophysics
 706 and neurobiology to individual differences. *Trends in Cognitive Sciences*, 17(8):391–400,
 707 2013. ISSN 1364-6613. doi: <https://doi.org/10.1016/j.tics.2013.06.006>. URL <https://www.sciencedirect.com/science/article/pii/S136461313001265>.

708

709 Francesca Mastrogiuseppe and Srdjan Ostojic. Linking connectivity, dynamics, and computations
 710 in low-rank recurrent neural networks. *Neuron*, 99(3):609–623.e29, 2018. ISSN 0896-6273. doi:
 711 <https://doi.org/10.1016/j.neuron.2018.07.003>. URL <https://www.sciencedirect.com/science/article/pii/S0896627318305439>.

712

713 Kenneth D. Miller and Francesco Fumarola. Mathematical equivalence of two common forms of
 714 firing rate models of neural networks. *Neural Computation*, 24(1):25–31, 2012. doi: 10.1162/NECO_a_00221.

715

716 Chris Olah, Nick Cammarata, Ludwig Schubert, Gabriel Goh, Michael Petrov, and Shan Carter.
 717 Zoom in: An introduction to circuits. *Distill*, 2020. doi: 10.23915/distill.00024.001.
 718 <https://distill.pub/2020/circuits/zoom-in>.

719

720 A. Emin Orhan and Wei Ji Ma. A diverse range of factors affect the nature of neural rep-
 721 resentations underlying short-term memory. *Nature Neuroscience*, 22:275–283, 2019. doi:
 722 10.1038/s41593-018-0314-y.

723

724 Antonio Orvieto, Samuel L. Smith, Albert Gu, Anushan Fernando, Caglar Gulcehre, Razvan Pas-
 725 canu, and Soham De. Resurrecting recurrent neural networks for long sequences. *International
 726 Conference on Machine Learning*, 2023.

727

728 Antonio Orvieto, Soham De, Caglar Gulcehre, Razvan Pascanu, and Samuel L. Smith. Universality
 729 of linear recurrences followed by non-linear projections: Finite-width guarantees and benefits
 730 of complex eigenvalues. In *ICML*, 2024. URL <https://openreview.net/forum?id=47ahB170xb>.

731

732 Srdjan Ostojic and Stefano Fusi. Computational role of structure in neural activity and connec-
 733 tivity. *Trends in Cognitive Sciences*, 28(7):677–690, 2024. ISSN 1364-6613. doi: <https://doi.org/10.1016/j.tics.2024.03.003>. URL <https://www.sciencedirect.com/science/article/pii/S1364661324000561>.

733

734 John O’Keefe and Jonathan O. Dostrovsky. The hippocampus as a spatial map. preliminary evidence
 735 from unit activity in the freely-moving rat. *Brain research*, 34 1:171–5, 1971.

736

737

738 Kiho Park, Yo Joong Choe, and Victor Veitch. The linear representation hypothesis and the geometry
 739 of large language models. In *Forty-first International Conference on Machine Learning*, 2024.
 740 URL <https://openreview.net/forum?id=UGpGkLzwP>.

741

742 Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. On the difficulty of training recurrent neural
 743 networks. In *International Conference on Machine Learning*, 2012.

744

745 Alexandra Maria Proca, Clémentine Carla Juliette Dominé, Murray Shanahan, and Pedro A. M.
 746 Mediano. Learning dynamics in linear recurrent neural networks. In *Forty-second International
 747 Conference on Machine Learning*, 2025. URL <https://openreview.net/forum?id=KG0crIWYnx>.

748

749 Kanaka Rajan, Christopher D. Harvey, and David W. Tank. Recurrent network models of sequence
 750 generation and memory. *Neuron*, 90:128–142, 2016.

751

752 Yuval Ran-Milo, Eden Lumbroso, Edo Cohen-Karlik, Raja Giryes, Amir Globerson, and Nadav
 753 Cohen. Provable benefits of complex parameterizations for structured state space models. In
 754 *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL
 755 <https://openreview.net/forum?id=h15RyEj151>.

756

757 Andrew M. Saxe, James L. McClelland, and Surya Ganguli. Exact solutions to the nonlinear dy-
 758 namics of learning in deep linear neural networks, 2014.

756 Andrew M. Saxe, Stephanie Nelli, and Christopher Summerfield. If deep learning is the answer,
 757 what is the question? *Nature Reviews Neuroscience*, 22:55 – 67, 2020.

758

759 Shreya Saxena and John P Cunningham. Towards the neural population doctrine. *Current Opinion
 760 in Neurobiology*, 55:171–178, 2019. doi: 10.1016/j.conb.2019.02.002.

761

762 Marine Schimel, Ta-Chu Kao, and Guillaume Hennequin. When and why does motor preparation
 763 arise in recurrent neural network models of motor control? 2024. doi: 10.7554/elife.89131.3.
 764 URL <http://dx.doi.org/10.7554/elife.89131.3>.

765

766 Friedrich Schuessler, Alexis Dubreuil, Francesca Mastrogiuseppe, Srdjan Ostojic, and Omri Barak.
 767 Dynamics of random recurrent networks with correlated low-rank structure. *Phys. Rev. Res.*, 2:
 768 013111, Feb 2020a. doi: 10.1103/PhysRevResearch.2.013111. URL <https://link.aps.org/doi/10.1103/PhysRevResearch.2.013111>.

769

770 Friedrich Schuessler, Francesca Mastrogiuseppe, Alexis Dubreuil, Srdjan Ostojic, and Omri Barak.
 771 The interplay between randomness and structure during learning in rnns. In *Advances in Neural
 772 Information Processing Systems*, 2020b.

773

774 Friedrich Schuessler, Francesca Mastrogiuseppe, Srdjan Ostojic, and Omri Barak. Aligned and
 775 oblique dynamics in recurrent neural networks. *eLife*, 2024.

776

777 Lee Sharkey, Bilal Chughtai, Joshua Batson, Jack Lindsey, Jeff Wu, Lucius Bushnaq, Nicholas
 778 Goldowsky-Dill, Stefan Heimersheim, Alejandro Ortega, Joseph Bloom, Stella Biderman, Adria
 779 Garriga-Alonso, Arthur Conmy, Neel Nanda, Jessica Rumbelow, Martin Wattenberg, Nandi
 780 Schoots, Joseph Miller, Eric J. Michaud, Stephen Casper, Max Tegmark, William Saunders,
 781 David Bau, Eric Todd, Atticus Geiger, Mor Geva, Jesse Hoogland, Daniel Murfet, and Tom Mc-
 782 Grath. Open problems in mechanistic interpretability, 2025. URL <https://arxiv.org/abs/2501.16496>.

783

784 Aneri V Soni and Michael J Frank. Adaptive chunking improves effective working memory capacity
 785 in a prefrontal cortex and basal ganglia circuit. February 2025. doi: 10.7554/elife.97894.2. URL
 786 <http://dx.doi.org/10.7554/elife.97894.2>.

787

788 Adly Templeton, Tom Conerly, Jonathan Marcus, Jack Lindsey, Trenton Bricken, Brian Chen,
 789 Adam Pearce, Craig Citro, Emmanuel Ameisen, Andy Jones, Hoagy Cunningham, Nicholas L
 790 Turner, Callum McDougall, Monte MacDiarmid, C. Daniel Freeman, Theodore R. Sumers,
 791 Edward Rees, Joshua Batson, Adam Jermyn, Shan Carter, Chris Olah, and Tom Henighan.
 792 Scaling monosemanticity: Extracting interpretable features from claude 3 sonnet. *Transformer
 793 Circuits Thread*, 2024. URL <https://transformer-circuits.pub/2024/scaling-monosemanticity/index.html>.

794

795 Kay M Tye, Earl K Miller, Felix H Taschbach, and et al. Mixed selectivity: Cellular computations
 796 for complexity. *Neuron*, 112(15):2459–2470, 2024. doi: 10.1016/j.neuron.2024.04.017.

797

798 Loek van Rossem and Andrew M Saxe. Algorithm development in neural networks: Insights from
 799 the streaming parity task. In *Forty-second International Conference on Machine Learning*, 2025.
 800 URL <https://openreview.net/forum?id=3go0lhfxd0>.

801

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

804

805 Olivia L. White, Daniel D. Lee, and Haim Sompolinsky. Short-term memory in orthogonal neural
 806 networks. *Physical Review Letters*, 92(14), 2004. ISSN 1079-7114. doi: 10.1103/physrevlett.92.
 807 148102. URL <http://dx.doi.org/10.1103/PhysRevLett.92.148102>.

808

809 Patrick Wilken and Wei Ma. A detection theory account of change detection. *Journal of vision*, 4:
 810 1120–35, 01 2005. doi: 10.1167/4.12.11.

811

812 Yang Xie, Peiyao Hu, Junru Li, Jingwen Chen, Weibin Song, Xiao-Jing Wang, Tianming Yang,
 813 Stanislas Dehaene, Shiming Tang, Bin Min, and Liping Wang. Geometry of sequence work-
 814 ing memory in macaque prefrontal cortex. *Science*, 375(6581):632–639, 2022a. doi: 10.1126/
 815 science.abm0204. URL <https://www.science.org/doi/abs/10.1126/science.abm0204>.

810 Yuanqi Xie, Yichen Henry Liu, Christos Constantinidis, and Xin Zhou. Neural mech-
811 anisms of working memory accuracy revealed by recurrent neural networks. *Frontiers in*
812 *Systems Neuroscience*, Volume 16 - 2022, 2022b. ISSN 1662-5137. doi:
813 10.3389/fnsys.2022.760864. URL <https://www.frontiersin.org/journals/systems-neuroscience/articles/10.3389/fnsys.2022.760864>.

814
815 Rafael Yuste. From the neuron doctrine to neural networks. *Nature Reviews Neuroscience*, 16(8):
816 487–497, 2015. doi: 10.1038/nrn3962.

817
818 Xiaohan Zhang, Shenquan Liu, and Zhe Sage Chen. A geometric framework for understanding
819 dynamic information integration in context-dependent computation. *iScience*, 24(8), 2021.

820
821 Nicolas Zucchetti and Antonio Orvieto. Recurrent neural networks: vanishing and exploding gradients
822 are not the end of the story. *Advances in Neural Information Processing Systems*, 2024.

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

864 **A ASSUMPTIONS ABOUT DATA**
865866 We make a few simplifying assumptions about the form of the data to make our setting analytically
867 tractable:
868869 **Assumption 1 (temporal independence):** Inputs are generated from an IID stochastic process of
870 scalar random variables $\{X_t\}_{t \geq 1}$.
871872 **Assumption 2 (temporal sparsity):** Following Elhage et al. (2022), we let $X_t = B_t U_t$, where
873 $B_t \sim \text{Bernoulli}(p)$, U_t are identically distributed according to any distribution and all
874 $\{B_t\}_{t \geq 1} \cup \{U_t\}_{t \geq 1}$ are mutually independent.
875876 Our second assumption allows us to explicitly control temporal sparsity by varying p : smaller
877 (larger) p corresponds to higher (lower) sparsity. Setting $p = 1$ recovers an arbitrary IID stochastic
878 process.
879880 While temporal independence may be a strong assumption, it is standard in the literature on memory
881 capacity (Jaeger, 2002; Ballarin et al., 2024). Temporal sparsity may also be a strong assumption,
882 depending on the task or setting considered.
883884 **B DECOMPOSING THE k -DELAY LOSS**
885886 Here, we derive a simplified form of the loss of a linear RNN trained on the k -delay task in terms of
887 four interpretable terms.
888889 **B.1 DERIVATION**
890891 Under the assumptions in Appendix A, the expected value of the squared-error loss (Equation (4))
892 incurred by a linear RNN simplifies as follows:
893

894
$$\mathbb{E}[\mathcal{L}] = \mathbb{E} \left[\sum_{t=1}^k \left\| 0 - \hat{Y}_t \right\|^2 + \sum_{t=k+1}^T \left\| X_{t-k} - \hat{Y}_t \right\|^2 \right] \quad (7)$$

895

896
$$= \mathbb{E} \left[\sum_{t=1}^k \hat{Y}_t^2 \right] + \mathbb{E} \left[\sum_{t=k+1}^T \left(X_{t-k} - \hat{Y}_t \right)^2 \right] \quad (8)$$

897

898
$$= \sum_{t=1}^k \mathbb{E} \left[\hat{Y}_t^2 \right] + \sum_{t=k+1}^T \mathbb{E} \left[\left(X_{t-k} - \hat{Y}_t \right)^2 \right] \quad (9)$$

899

900
$$= \sum_{t=1}^T \mathbb{E} \left[\hat{Y}_t^2 \right] + \sum_{t=k+1}^T \mathbb{E} \left[X_{t-k}^2 \right] - 2 \sum_{t=k+1}^T \mathbb{E} \left[X_{t-k} \hat{Y}_t \right] + \sum_{t=k+1}^T \mathbb{E} \left[\hat{Y}_t^2 \right] \quad (10)$$

901

902
$$= \sum_{t=k+1}^T \left(\mathbb{E} \left[X_{t-k}^2 \right] - 2\mathbb{E} \left[X_{t-k} \hat{Y}_t \right] \right) + \sum_{t=1}^T \mathbb{E} \left[\hat{Y}_t^2 \right] \quad (11)$$

903

904
$$= \sum_{t=k+1}^T \left(\mathbb{E} \left[X_{t-k}^2 \right] - 2\mathbb{E} \left[X_{t-k} \sum_{s=0}^{t-1} \mathbf{w}_y^\top \mathbf{w}_s X_{t-s} \right] \right) + \sum_{t=1}^T \mathbb{E} \left[\left(\sum_{s=0}^{t-1} \mathbf{w}_y^\top \mathbf{w}_s X_{t-s} \right)^2 \right] \quad (12)$$

905

906 We proceed by computing each of these expectations separately. Recalling that, by the temporal
907 sparsity assumption, $X_t = B_t U_t$, we let $\mu := \mathbb{E}[U_t]$ and $\nu := \mathbb{E}[U_t^2]$. Then, as $\{X_t\}_{t \geq 1}$ are
908 assumed to be IID, the first expectation becomes
909

910
$$\mathbb{E} \left[X_{t-k}^2 \right] = \mathbb{E} \left[B_{t-k}^2 U_{t-k}^2 \right] = \mathbb{E} \left[B_{t-k}^2 \right] \mathbb{E} \left[U_{t-k}^2 \right] = \mathbb{E} \left[B_{t-k} \right] \mathbb{E} \left[U_{t-k}^2 \right] = p\nu \quad (13)$$

911

918 Next, in computing the second expectation, we must handle the case of $s = k$ separately. This yields
919

920 921
$$\mathbb{E} \left[X_{t-k} \sum_{s=0}^{t-1} \mathbf{w}_y^\top \mathbf{w}_s X_{t-s} \right] = \mathbb{E} [X_{t-k}^2] \mathbf{w}_y^\top \mathbf{w}_{s=k} + \sum_{\substack{s=0 \\ s \neq k}}^{t-1} \mathbb{E} [X_{t-k} X_{t-s}] \mathbf{w}_y^\top \mathbf{w}_s \quad (14)$$

922
923

924 925
$$= p\nu \mathbf{w}_y^\top \mathbf{w}_{s=k} + \sum_{\substack{s=0 \\ s \neq k}}^{t-1} \mathbb{E} [B_{t-k} U_{t-k} B_{t-s} U_{t-s}] \mathbf{w}_y^\top \mathbf{w}_s \quad (15)$$

926
927

928 929
$$= p\nu \mathbf{w}_y^\top \mathbf{w}_{s=k} + p^2 \mu^2 \sum_{\substack{s=0 \\ s \neq k}}^{t-1} \mathbf{w}_y^\top \mathbf{w}_s \quad (16)$$

930
931

932 For the third expectation, we must handle the diagonal terms (where $s = s'$) separately:
933

934 935
$$\mathbb{E} \left[\left(\sum_{s=0}^{t-1} \mathbf{w}_y^\top \mathbf{w}_s X_{t-s} \right)^2 \right] = \mathbb{E} \left[\sum_{s=0}^{t-1} (\mathbf{w}_y^\top \mathbf{w}_s)^2 X_{t-s}^2 + \sum_{\substack{s, s'=0 \\ s \neq s'}}^{t-1} (\mathbf{w}_y^\top \mathbf{w}_s) (\mathbf{w}_y^\top \mathbf{w}_{s'}) X_{t-s} X_{t-s'} \right] \quad (17)$$

936
937
938

939 940
$$= \sum_{s=0}^{t-1} (\mathbf{w}_y^\top \mathbf{w}_s)^2 \mathbb{E} [X_{t-s}^2] + \sum_{\substack{s, s'=0 \\ s \neq s'}}^{t-1} (\mathbf{w}_y^\top \mathbf{w}_s) (\mathbf{w}_y^\top \mathbf{w}_{s'}) \mathbb{E} [X_{t-s} X_{t-s'}] \quad (18)$$

941
942
943

944 945
$$= p\nu \sum_{s=0}^{t-1} (\mathbf{w}_y^\top \mathbf{w}_s)^2 + p^2 \mu^2 \sum_{\substack{s, s'=0 \\ s \neq s'}}^{t-1} (\mathbf{w}_y^\top \mathbf{w}_s) (\mathbf{w}_y^\top \mathbf{w}_{s'}) \quad (19)$$

946
947
948

949 Substituting these expressions into equation 12, we obtain
950

951 952
$$\mathbb{E}[\mathcal{L}] = \sum_{t=k+1}^T \left(p\nu - 2p\nu \mathbf{w}_y^\top \mathbf{w}_{s=k} - 2p^2 \mu^2 \sum_{\substack{s=0 \\ s \neq k}}^{t-1} \mathbf{w}_y^\top \mathbf{w}_s \right) \quad (20)$$

953
954
955

956 957
$$+ \sum_{t=1}^T \left(p\nu \sum_{s=0}^{t-1} (\mathbf{w}_y^\top \mathbf{w}_s)^2 + p^2 \mu^2 \sum_{\substack{s, s'=0 \\ s \neq s'}}^{t-1} (\mathbf{w}_y^\top \mathbf{w}_s) (\mathbf{w}_y^\top \mathbf{w}_{s'}) \right) \quad (21)$$

958
959
960

961 Finally, we move every occurrence of the $p\nu (\mathbf{w}_y^\top \mathbf{w}_{s=k})^2$ term into the first summation:
962

963 964
$$\mathbb{E}[\mathcal{L}] = \sum_{t=k+1}^T \left(p\nu - 2p\nu \mathbf{w}_y^\top \mathbf{w}_{s=k} + p\nu (\mathbf{w}_y^\top \mathbf{w}_{s=k})^2 - 2p^2 \mu^2 \sum_{\substack{s=0 \\ s \neq k}}^{t-1} \mathbf{w}_y^\top \mathbf{w}_s \right) \quad (22)$$

965
966
967

968 969
$$+ \sum_{t=1}^T \left(p\nu \sum_{\substack{s=0 \\ s \neq k}}^{t-1} (\mathbf{w}_y^\top \mathbf{w}_s)^2 + p^2 \mu^2 \sum_{\substack{s, s'=0 \\ s \neq s'}}^{t-1} (\mathbf{w}_y^\top \mathbf{w}_s) (\mathbf{w}_y^\top \mathbf{w}_{s'}) \right) \quad (23)$$

970
971

972 After factorizing, we arrive at the form:
 973

$$\mathbb{E}[\mathcal{L}] = \sum_{t=k+1}^T \left(\underbrace{p\nu \|\mathbf{w}_y^\top \mathbf{w}_{s=k} - 1\|^2}_{\text{task benefit}} - \underbrace{2p^2\mu^2 \sum_{\substack{s=0 \\ s \neq k}}^{t-1} \mathbf{w}_y^\top \mathbf{w}_s}_{\text{mean correction}} \right) \quad (24)$$

$$+ \sum_{t=1}^T \left(\underbrace{p\nu \sum_{\substack{s=0 \\ s \neq k}}^{t-1} (\mathbf{w}_y^\top \mathbf{w}_s)^2}_{\text{projection interference cost}} + \underbrace{p^2\mu^2 \sum_{\substack{s,s'=0 \\ s \neq s'}}^{t-1} (\mathbf{w}_y^\top \mathbf{w}_s) \cdot (\mathbf{w}_y^\top \mathbf{w}_{s'})}_{\text{composition interference}} \right) \quad (25)$$

987 We omit the 0 indexing of summations in the main text to reduce clutter.
 988

989 B.2 DISCUSSION ON MEAN CORRECTION AND COMPOSITION INTERFERENCE TERMS 990

991 By looking at expectation of the loss, we can see that the mean correction and composition inter-
 992 ference terms only exist when the data distribution has non-zero mean. However, we can also see
 993 that the two terms are scaled by p^2 . This implies that, even for data with a non-zero mean, in the
 994 sparse regime where we take the limit of $p \rightarrow 0$, the impact of these two terms on the loss becomes
 995 negligible. The mean correction term disappears because, even though the underlying data distribu-
 996 tion U has a fixed mean $\mathbb{E}[U] = \mu$, the mean of the overall input X_t is proportional to its sparsity:
 997 $\mathbb{E}[X] = \mathbb{E}[B]\mathbb{E}[U] = p\mu$, so $\mathbb{E}[X] \rightarrow 0$ as $p \rightarrow 0$ and sparse data does have an approximately zero
 998 mean. Furthermore, the composition interference term disappears because the probability of two
 999 or more temporal features simultaneously activating is $O(p^2)$, so composition interference becomes
 1000 negligible as $p \rightarrow 0$.

1001 C GEOMETRIC INTERPRETATION OF LEARNING DYNAMICS 1002

1003 Building on earlier work in feedforward linear networks (Saxe et al., 2014), recent work has inves-
 1004 tigated the learning dynamics of RNNs (Schuessler et al., 2020b; Proca et al., 2025; van Rossem
 1005 & Saxe, 2025; Ger & Barak, 2025; Bordelon et al., 2025). It's been shown that neural networks
 1006 initialized with small random weights appear to undergo an initial phase of eigenvector alignment
 1007 (Atanasov et al., 2022), followed by learning of the largest (and latest occurring) data correlation
 1008 singular values/eigenvalues (Saxe et al., 2014; Proca et al., 2025).

1009 Having derived an exact expression of the expected loss in terms of four interpretable terms, here
 1010 we study their learning dynamics to better characterize the geometric evolution of the network (Figure
 1011 2). Additionally, we visualize the dynamics of the output projection ($\mathbf{w}_y^\top \mathbf{w}_s$ for $0 \leq s < 12$)
 1012 to better understand the functional behavior of the network throughout training. Taken together,
 1013 this analysis provides a direct connection between the model's functionality/performance and its
 1014 representational geometry.

1015 First, we observe that our predicted expectation of the loss (in red) closely matches the empirical
 1016 loss (light purple), corroborating our theoretical result. Interestingly, we see that the loss decreases
 1017 in a 'staircase,' consistent with Saxe et al. (2014); Proca et al. (2025). In particular, Saxe et al. (2014)
 1018 identified that in feedforward networks, the largest data correlation singular values are learned first,
 1019 corresponding to staircase-like drops in the loss, while Proca et al. (2025) extended this work to
 1020 RNNs and identified that task dynamics with singular/eigen-values that are large and occur later in
 1021 the trajectory are learned faster (i.e., stronger correlation with the most recent past). **Work in feedfor-**
 1022 **1023** **1024** **1025** **1026** **1027** **1028** **1029** **1030** **1031** **1032** **1033** **1034** **1035** **1036** **1037** **1038** **1039** **1040** **1041** **1042** **1043** **1044** **1045** **1046** **1047** **1048** **1049** **1050** **1051** **1052** **1053** **1054** **1055** **1056** **1057** **1058** **1059** **1060** **1061** **1062** **1063** **1064** **1065** **1066** **1067** **1068** **1069** **1070** **1071** **1072** **1073** **1074** **1075** **1076** **1077** **1078** **1079** **1080** **1081** **1082** **1083** **1084** **1085** **1086** **1087** **1088** **1089** **1090** **1091** **1092** **1093** **1094** **1095** **1096** **1097** **1098** **1099** **1100** **1101** **1102** **1103** **1104** **1105** **1106** **1107** **1108** **1109** **1110** **1111** **1112** **1113** **1114** **1115** **1116** **1117** **1118** **1119** **1120** **1121** **1122** **1123** **1124** **1125** **1126** **1127** **1128** **1129** **1130** **1131** **1132** **1133** **1134** **1135** **1136** **1137** **1138** **1139** **1140** **1141** **1142** **1143** **1144** **1145** **1146** **1147** **1148** **1149** **1150** **1151** **1152** **1153** **1154** **1155** **1156** **1157** **1158** **1159** **1160** **1161** **1162** **1163** **1164** **1165** **1166** **1167** **1168** **1169** **1170** **1171** **1172** **1173** **1174** **1175** **1176** **1177** **1178** **1179** **1180** **1181** **1182** **1183** **1184** **1185** **1186** **1187** **1188** **1189** **1190** **1191** **1192** **1193** **1194** **1195** **1196** **1197** **1198** **1199** **1200** **1201** **1202** **1203** **1204** **1205** **1206** **1207** **1208** **1209** **1210** **1211** **1212** **1213** **1214** **1215** **1216** **1217** **1218** **1219** **1220** **1221** **1222** **1223** **1224** **1225** **1226** **1227** **1228** **1229** **1230** **1231** **1232** **1233** **1234** **1235** **1236** **1237** **1238** **1239** **1240** **1241** **1242** **1243** **1244** **1245** **1246** **1247** **1248** **1249** **1250** **1251** **1252** **1253** **1254** **1255** **1256** **1257** **1258** **1259** **1260** **1261** **1262** **1263** **1264** **1265** **1266** **1267** **1268** **1269** **1270** **1271** **1272** **1273** **1274** **1275** **1276** **1277** **1278** **1279** **1280** **1281** **1282** **1283** **1284** **1285** **1286** **1287** **1288** **1289** **1290** **1291** **1292** **1293** **1294** **1295** **1296** **1297** **1298** **1299** **1300** **1301** **1302** **1303** **1304** **1305** **1306** **1307** **1308** **1309** **1310** **1311** **1312** **1313** **1314** **1315** **1316** **1317** **1318** **1319** **1320** **1321** **1322** **1323** **1324** **1325** **1326** **1327** **1328** **1329** **1330** **1331** **1332** **1333** **1334** **1335** **1336** **1337** **1338** **1339** **1340** **1341** **1342** **1343** **1344** **1345** **1346** **1347** **1348** **1349** **1350** **1351** **1352** **1353** **1354** **1355** **1356** **1357** **1358** **1359** **1360** **1361** **1362** **1363** **1364** **1365** **1366** **1367** **1368** **1369** **1370** **1371** **1372** **1373** **1374** **1375** **1376** **1377** **1378** **1379** **1380** **1381** **1382** **1383** **1384** **1385** **1386** **1387** **1388** **1389** **1390** **1391** **1392** **1393** **1394** **1395** **1396** **1397** **1398** **1399** **1400** **1401** **1402** **1403** **1404** **1405** **1406** **1407** **1408** **1409** **1410** **1411** **1412** **1413** **1414** **1415** **1416** **1417** **1418** **1419** **1420** **1421** **1422** **1423** **1424** **1425** **1426** **1427** **1428** **1429** **1430** **1431** **1432** **1433** **1434** **1435** **1436** **1437** **1438** **1439** **1440** **1441** **1442** **1443** **1444** **1445** **1446** **1447** **1448** **1449** **1450** **1451** **1452** **1453** **1454** **1455** **1456** **1457** **1458** **1459** **1460** **1461** **1462** **1463** **1464** **1465** **1466** **1467** **1468** **1469** **1470** **1471** **1472** **1473** **1474** **1475** **1476** **1477** **1478** **1479** **1480** **1481** **1482** **1483** **1484** **1485** **1486** **1487** **1488** **1489** **1490** **1491** **1492** **1493** **1494** **1495** **1496** **1497** **1498** **1499** **1500** **1501** **1502** **1503** **1504** **1505** **1506** **1507** **1508** **1509** **1510** **1511** **1512** **1513** **1514** **1515** **1516** **1517** **1518** **1519** **1520** **1521** **1522** **1523** **1524** **1525** **1526** **1527** **1528** **1529** **1530** **1531** **1532** **1533** **1534** **1535** **1536** **1537** **1538** **1539** **1540** **1541** **1542** **1543** **1544** **1545** **1546** **1547** **1548** **1549** **1550** **1551** **1552** **1553** **1554** **1555** **1556** **1557** **1558** **1559** **1560** **1561** **1562** **1563** **1564** **1565** **1566** **1567** **1568** **1569** **1570** **1571** **1572** **1573** **1574** **1575** **1576** **1577** **1578** **1579** **1580** **1581** **1582** **1583** **1584** **1585** **1586** **1587** **1588** **1589** **1590** **1591** **1592** **1593** **1594** **1595** **1596** **1597** **1598** **1599** **1600** **1601** **1602** **1603** **1604** **1605** **1606** **1607** **1608** **1609** **1610** **1611** **1612** **1613** **1614** **1615** **1616** **1617** **1618** **1619** **1620** **1621** **1622** **1623** **1624** **1625** **1626** **1627** **1628** **1629** **1630** **1631** **1632** **1633** **1634** **1635** **1636** **1637** **1638** **1639** **1640** **1641** **1642** **1643** **1644** **1645** **1646** **1647** **1648** **1649** **1650** **1651** **1652** **1653** **1654** **1655** **1656** **1657** **1658** **1659** **1660** **1661** **1662** **1663** **1664** **1665** **1666** **1667** **1668** **1669** **1670** **1671** **1672** **1673** **1674** **1675** **1676** **1677** **1678** **1679** **1680** **1681** **1682** **1683** **1684** **1685** **1686** **1687** **1688** **1689** **1690** **1691** **1692** **1693** **1694** **1695** **1696** **1697** **1698** **1699** **1700** **1701** **1702** **1703** **1704** **1705** **1706** **1707** **1708** **1709** **1710** **1711** **1712** **1713** **1714** **1715** **1716** **1717** **1718** **1719** **1720** **1721** **1722** **1723** **1724** **1725** **1726** **1727** **1728** **1729** **1730** **1731** **1732** **1733** **1734** **1735** **1736** **1737** **1738** **1739** **1740** **1741** **1742** **1743** **1744** **1745** **1746** **1747** **1748** **1749** **1750** **1751** **1752** **1753** **1754** **1755** **1756** **1757** **1758** **1759** **1760** **1761** **1762** **1763** **1764** **1765** **1766** **1767** **1768** **1769** **1770** **1771** **1772** **1773** **1774** **1775** **1776** **1777** **1778** **1779** **1780** **1781** **1782** **1783** **1784** **1785** **1786** **1787** **1788** **1789** **1790** **1791** **1792** **1793** **1794** **1795** **1796** **1797** **1798** **1799** **1800** **1801** **1802** **1803** **1804** **1805** **1806** **1807** **1808** **1809** **1810** **1811** **1812** **1813** **1814** **1815** **1816** **1817** **1818** **1819** **1820** **1821** **1822** **1823** **1824** **1825** **1826** **1827** **1828** **1829** **1830** **1831** **1832** **1833** **1834** **1835** **1836** **1837** **1838** **1839** **1840** **1841** **1842** **1843** **1844** **1845** **1846** **1847** **1848** **1849** **1850** **1851** **1852** **1853** **1854** **1855** **1856** **1857** **1858** **1859** **1860** **1861** **1862** **1863** **1864** **1865** **1866** **1867** **1868** **1869** **1870** **1871** **1872** **1873** **1874** **1875** **1876** **1877** **1878** **1879** **1880** **1881** **1882** **1883** **1884** **1885** **1886** **1887** **1888** **1889** **1890** **1891** **1892** **1893** **1894** **1895** **1896** **1897** **1898** **1899** **1900** **1901** **1902** **1903** **1904** **1905** **1906** **1907** **1908** **1909** **1910** **1911** **1912** **1913** **1914** **1915** **1916** **1917** **1918** **1919** **1920** **1921** **1922** **1923** **1924** **1925** **1926** **1927** **1928** **1929** **1930** **1931** **1932** **1933** **1934** **1935** **1936** **1937** **1938** **1939** **1940** **1941** **1942** **1943** **1944** **1945** **1946** **1947** **1948** **1949** **1950** **1951** **1952** **1953** **1954** **1955** **1956** **1957** **1958** **1959** **1960** **1961** **1962** **1963** **1964** **1965** **1966** **1967** **1968** **1969** **1970** **1971** **1972** **1973** **1974** **1975** **1976** **1977** **1978** **1979** **1980** **1981** **1982** **1983** **1984** **1985** **1986** **1987**

1026 By looking at the dynamics of the output projection, we see that initially the RNN tries to decrease
 1027 the loss by aligning all of the \mathbf{w}_s vectors to correlate positively with \mathbf{w}_y (R1 in Figure 2). This de-
 1028 creases task error and takes advantage of the mean correction term, providing an initial improvement
 1029 to the loss. All the vectors become positively correlated with each other (facing similar directions),
 1030 leading to increased interference. This appears closely related to the initial phase of eigenvector
 1031 alignment mentioned earlier (Atanasov et al., 2022).

1032 Eventually, the mean correction term plateaus. Instead, we observe that the output projection for
 1033 $\mathbf{w}_y^\top \mathbf{w}_{s=0}$ (the projection for the most recently-incoming feature) changes dramatically (R2-R3), fol-
 1034 lowed by a separation of the other feature directions approximately converging according in order of
 1035 newest to oldest feature directions (R3-R4). The model converges to its final geometric arrangement,
 1036 spreading the \mathbf{w}_s vectors out across the entire plane (R4). As part of this, $\mathbf{w}_y^\top \mathbf{w}_{s=k}$ becomes much
 1037 larger, causing a significant drop in task error (R3). Additionally, some of the other $\mathbf{w}_y^\top \mathbf{w}_{s \neq k}$ also
 1038 increase in magnitude, causing projection interference to increase (R3-R4). As some of the $\mathbf{w}_y^\top \mathbf{w}_s$
 1039 fall below zero, the model is less able to take advantage of mean correction, but instead, the even
 1040 distribution of vectors causes a drop in composition interference (R4). It’s not clear exactly how
 1041 the dynamics of the changing geometry correspond to the learning of data correlation eigenvalues
 1042 (Proca et al., 2025), but it is likely the two are linked.

1043

1044 D ANALYZING THE SPECTRAL RADIUS OF W_h

1045

1046 D.1 PROVING THE ECHO STATE PROPERTY

1047

1048 In this section, we prove that in our setting, achieving finite loss in the infinite time limit ($T \rightarrow \infty$) is
 1049 only possible in models with a spectral radius of W_h less than 1 (the so-called echo state property),
 1050 to demonstrate the optimality of certain solutions found by the networks we study. While there may
 1051 exist other solutions that perform well on short sequences with a spectral radius greater or equal to
 1052 1, we show that their performance will degrade as sequence length increases.

1053 In particular, we consider models that satisfy

$$1054 \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}[\mathcal{L}(T)] < \infty \quad (26)$$

1055 where T refers to the sequence length and $\mathcal{L}(T)$ is defined as per Equation (4). We prove that this
 1056 condition is only possible if the spectral radius of W_h is $\rho(W_h) < 1$.

1057 *Proof.* Consider the expected loss over an entire sequence, $\mathbb{E}[\mathcal{L}(T)]$, as a sum of terms,

$$1058 \mathbb{E}[\mathcal{L}(T)] = L_{t=1} + L_{t=2} + \dots + L_{t=T} \quad (27)$$

1059 where L_t is the loss incurred by the model on timestep t . Then, for $t > k$, we have

$$1060 L_{t>k} = p\nu \|\mathbf{w}_y^\top \mathbf{w}_{s=k} - 1\|^2 - 2p^2\mu^2 \sum_{\substack{s=0 \\ s \neq k}}^{t-1} \mathbf{w}_y^\top \mathbf{w}_s \quad (28)$$

$$1061 + p\nu \sum_{\substack{s=0 \\ s \neq k}}^{t-1} (\mathbf{w}_y^\top \mathbf{w}_s)^2 + p^2\mu^2 \sum_{\substack{s, s'=0 \\ s \neq s'}}^{t-1} (\mathbf{w}_y^\top \mathbf{w}_s) \cdot (\mathbf{w}_y^\top \mathbf{w}_{s'}) \quad (29)$$

1062 This rearranges to

$$1063 L_{t>k} = (p\nu - p^2\mu^2) \left(\|\mathbf{w}_y^\top \mathbf{w}_{s=k} - 1\|^2 + \sum_{\substack{s=0 \\ s \neq k}}^{t-1} (\mathbf{w}_y^\top \mathbf{w}_s)^2 \right) + p^2\mu^2 \left(\sum_{s=0}^{t-1} \mathbf{w}_y^\top \mathbf{w}_s - 1 \right)^2 \quad (30)$$

1064 whose terms are all non-negative as $p\nu - p^2\mu^2 = \text{Var}[X_t] \geq 0$.

1065 If the loss incurred at large enough timesteps, $L_{t>k}$, diverges to ∞ as $t \rightarrow \infty$, then the average loss
 1066 per timestep, $\frac{1}{T} \mathbb{E}[\mathcal{L}(T)]$, also diverges to infinity. Hence, for the average loss per timestep to remain
 1067 finite, we require $L_{t>k}$ to remain finite.

Now, assume that $\rho(W_h) \geq 1$ and that $L_{t>k}$ does not diverge to infinity. As all terms in the expression for $L_{t>k}$ are non-negative, the series $\sum_{s \neq k}^{\infty} (\mathbf{w}_y^\top \mathbf{w}_s)^2$ must be convergent. This implies $\lim_{s \rightarrow \infty} \mathbf{w}_y^\top \mathbf{w}_s = 0$. Expanding, we obtain

$$\lim_{s \rightarrow \infty} \mathbf{w}_y^\top \mathbf{w}_s = \lim_{s \rightarrow \infty} \mathbf{w}_y^\top W_h^s \mathbf{w}_x = \mathbf{w}_y^\top \left(\lim_{s \rightarrow \infty} W_h^s \right) \mathbf{w}_x = \mathbf{w}_y^\top W_h^\infty \mathbf{w}_x = 0 \quad (31)$$

where we denote $W_h^\infty = \lim_{s \rightarrow \infty} W_h^s$, which is non-zero as $\rho(W_h) \geq 1$. The equality holds precisely when

$$\mathbf{w}_x \in \ker(W_h^\infty) \quad \text{or} \quad \mathbf{w}_y \in \ker(W_h^{\infty \top}) \quad \text{or} \quad \mathbf{w}_y \perp W_h^\infty \mathbf{w}_x$$

Since $W_h^\infty \neq 0$, we have $\text{rank}(W_h^\infty) = \text{rank}(W_h^{\infty \top}) > 0$. Hence, by the rank-nullity theorem, $\dim(\ker(W_h^\infty)) = N_h - \text{rank}(W_h^\infty) < N_h$ and $\dim(\ker(W_h^{\infty \top})) = N_h - \text{rank}(W_h^{\infty \top}) < N_h$. Thus both $\ker(W_h^\infty)$ and $\ker(W_h^{\infty \top})$ have Lebesgue measure zero.

If $\mathbf{w}_x \notin \ker(W_h^\infty)$, we require the third case, where W_y^\top must lie on the $(N_h - 1)$ -dimensional hyperplane orthogonal to $W_h^\infty \mathbf{w}_x$. This is again a proper subspace of \mathbb{R}^{N_h} with measure zero. Hence, if $\rho(W_h) \geq 1$, then the set of solutions for which $L_{t>k}$ remains finite has measure zero, so $L_{t>k}$ almost surely diverges to infinity. Taking the contrapositive, if $L_{t>k}$ remains finite, then $\rho(W_h) < 1$ almost surely. \square

Intuitively, we have shown that if $\rho(W_h) \geq 1$, then $L_{t>k}$ diverges to infinity except if \mathbf{w}_x , W_h and \mathbf{w}_y *precisely* (not approximately) satisfy certain conditions. We can be confident that these conditions are not satisfied by models in practice: it would require the optimizer to balance the model parameters on an infinitely thin ‘knife edge’, which is practically impossible in floating-point arithmetic. Hence we can safely restrict our attention to models satisfying $\rho(W_h) < 1$.

D.2 OPTIMALITY OF SPIRAL SINKS IN 2D

In the proof above, we have shown that $\rho(W_h) < 1$ in order to have finite loss as sequence length $T \rightarrow \infty$. This condition on the spectral radius is known as the echo state property and can be interpreted as the model forgetting old inputs over time. This constraint guarantees the shrinking aspect of a spiral sink. We now verify that the spiral behavior, specifically, is the optimal solution in 2 dimensions. To do this, we train linear RNNs parameterized such that the trace and determinant of $W_h \in \mathbb{R}^{2 \times 2}$ is fixed. Specifically, for each desired trace-determinant pair (τ, δ) , we optimize over the 2-dimensional manifold

$$\mathcal{M} = \{W_h \in \mathbb{R}^{2 \times 2} : \text{tr}(W_h) = \tau, \det(W_h) = \delta\} \quad (32)$$

which we parameterize by $(\theta_1, \theta_2) \in \mathbb{R}^2$ using the map

$$\varphi(\theta_1, \theta_2) = \begin{bmatrix} \theta_1 & \frac{\theta_1(\tau - \theta_1) - \delta}{\exp(\theta_2)} \\ \exp(\theta_2) & \tau - \theta_1 \end{bmatrix}$$

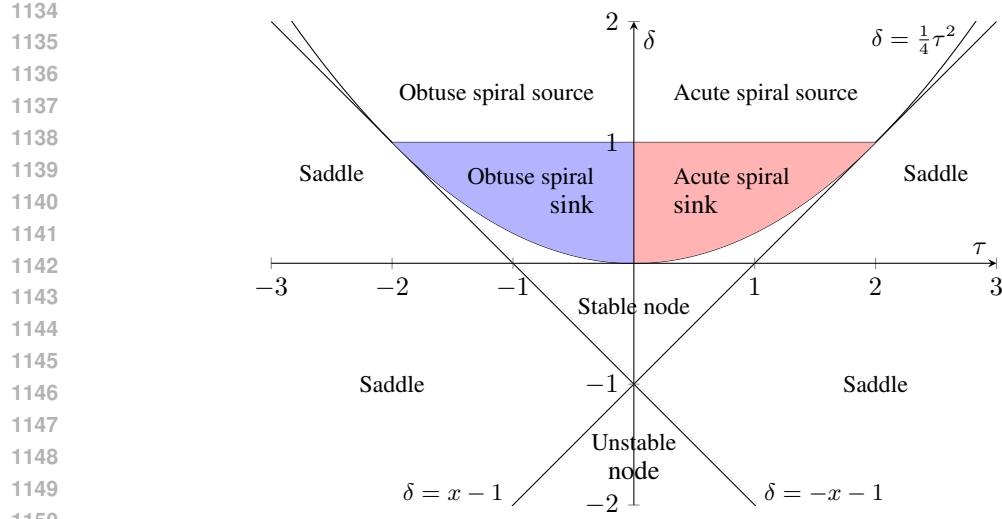
where θ_2 is exponentiated to ensure φ is bijective. We verify that

$$\text{tr}(\varphi(\theta_1, \theta_2)) = \theta_1 + \tau - \theta_1 = \tau \quad (33)$$

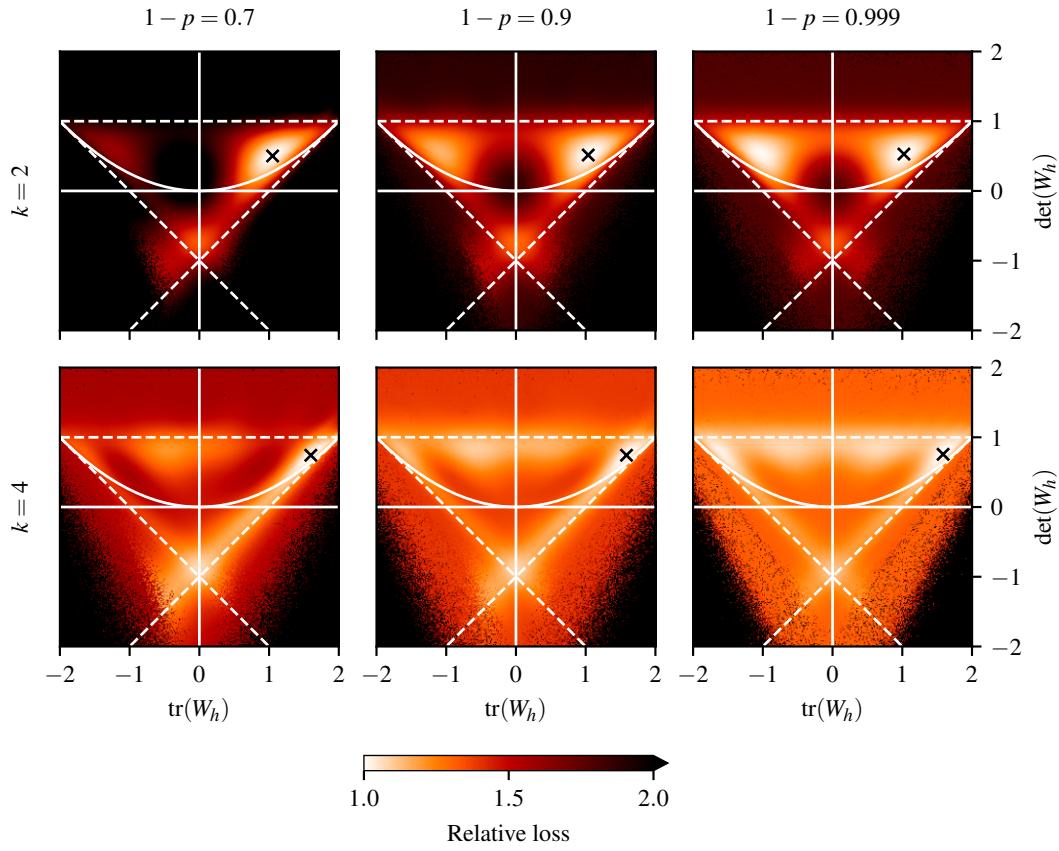
$$\det(\varphi(\theta_1, \theta_2)) = \theta_1(\tau - \theta_1) - \frac{\theta_1(\tau - \theta_1)}{\exp(\theta_2)} \exp(\theta_2) = \delta \quad (34)$$

We sweep through a grid of points in the square $(\tau, \delta) \in [-2, 2]^2$ and for each point, we train a linear RNN parameterized as above. The loss achieved by each model under various task conditions is shown in Figure 7. Each plot can be thought of as a trace-determinant slice of the k -delay loss landscape.

We note that in every case, the optimal solution (brightest point on the plot) is found in the region that corresponds to spiral sinks. This provides strong empirical evidence that the globally optimal solution must be a spiral sink. This makes intuitive sense: rotation is used to implement an approximate delay-line solution by shifting inputs through a sequence of positions in the hidden state, while the gradual shrinking of vectors facilitates the forgetting of old inputs.



1151 **Figure 6: Trace-determinant classification of 2-dimensional discrete linear dynamical systems.**
 1152 Adapted from Galor (2007). Stable systems occupy the triangular region enclosed by the lines $\delta = 1$,
 1153 $\delta = x - 1$ and $\delta = -x - 1$. Within this triangle, spiral sinks are found above the parabola $\delta = \frac{1}{4}\tau^2$.



1182 **Figure 7: Loss landscape of the 2-delay and 4-delay tasks.** Shown in terms of $\text{tr}(W_h)$ and
 1183 $\det(W_h)$ at 3 sparsity levels for linear RNNs with 2-dimensional W_h . At each point on the trace-
 1184 determinant plane a linear model with 2-dimensional hidden state was parameterized, as described
 1185 above, with a fixed trace and determinant. The final training loss is displayed as a multiple of (i.e.
 1186 relative to) the lowest training loss achieved by any of the models. The best-performing model for
 1187 each task is marked by a cross. Standard lines and curves used to classify discrete dynamical sys-
 1188 tems are overlaid in white; refer to Figure 6 for interpretation.

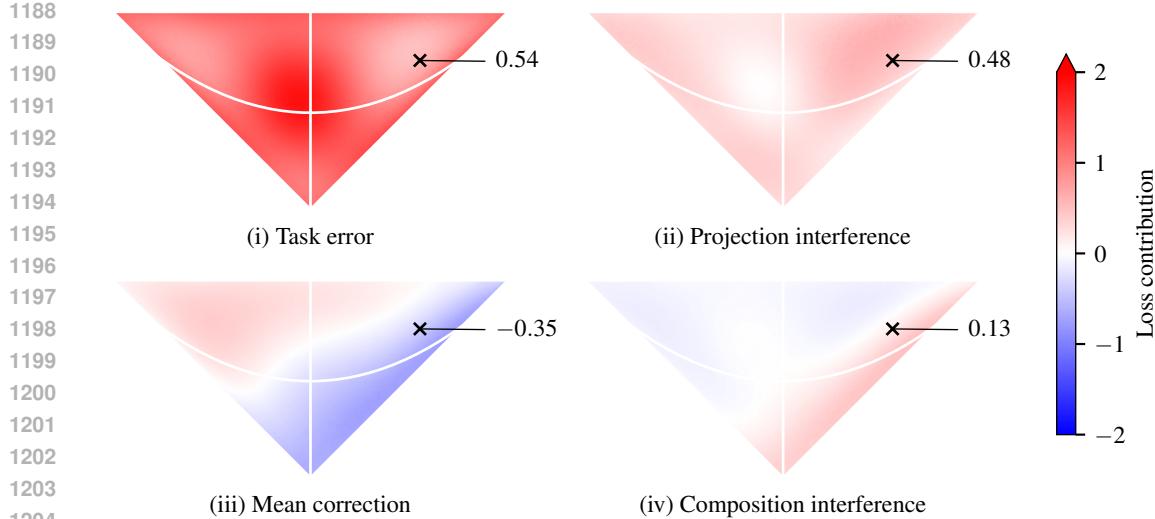


Figure 8: **Decomposed loss landscape of the 2-delay task in the trace-determinant plane.** Trace-determinant loss landscape of the 2-delay task at 0.7 sparsity for stable linear RNNs with 2-dimensional W_h , decomposed into interpretable terms. The best-performing model (in terms of overall empirical loss) is marked with a cross and annotated with its value for each of the loss terms.

D.3 OPTIMALITY OF ACUTE-ANGLED SPIRAL SINKS

Interestingly, the plots become increasingly symmetric with increasing sparsity. In the low-sparsity case for both the 2-delay and 4-delay tasks, the global optimum clearly has $\text{tr}(W_h) > 0$, which implies an acute rotation angle. This can be seen by writing the eigenvalues of W_h as

$$\lambda_{1,2} = \frac{t}{2} \pm \frac{i}{2} \sqrt{4\delta - \tau^2} \quad (35)$$

and observing that $t < 0$ corresponds to $\arg(\lambda_{1,2}) \in (\frac{\pi}{2}, \frac{3\pi}{2})$ (obtuse rotation angle) while $t > 0$ corresponds to $\arg(\lambda_{1,2}) \in (-\frac{\pi}{2}, \frac{\pi}{2})$ (acute rotation angle).

The asymmetry is most clearly visible in the figure for the 2-delay task at 0.7 sparsity. In Figure 8, we decompose this particular loss landscape to investigate the reason for acute-angled rotation being preferred. While none of the terms are perfectly symmetric in $\text{tr}(W_h)$, it is clear that the mean correction term is largely driving this behavior: for acute spirals, it can reduce loss, while for most instances of obtuse spirals, it increases the loss. Unsurprisingly, there exists a trade-off between this term and the others (in particular, composition interference largely seems positive where mean correction is negative, and vice versa), but evidently the optimal balance is firmly in the acute spiral region. Overall, there seems to be a region in which task error, mean correction and composition interference are all relatively low, while projection interference is relatively high – this corresponds precisely with the lowest-loss region in Figure 7 and is exactly the sacrifice we observe in Figure 2.

E STATE SPACE MODEL (LINEAR RECURRENCE, NONLINEAR READOUT)

E.1 APPROXIMATION OF LOSS UNDER HIGH SPARSITY

We now consider the SSM that produces its output through a ReLU activation function, corresponding to $\sigma_y = \text{ReLU}$ and $\sigma_h = \text{id}$. Due to the ReLU activation, such a model can only produce non-negative outputs, so for the k -delay task, it is sensible to restrict the input distribution to be non-negative as well (as we desire $\hat{y}_t = \text{ReLU}(\dots) = x_{t-k}$). Recall that the input distribution is $X_t = B_t U_t$, where $B_t \sim \text{Bernoulli}(p)$ and U_t are identically distributed according to any distribution. Then, since $B_t \geq 0$, we require $U_t \geq 0$ to ensure $X_t \geq 0$. For example, the distribution used by Elhage et al. (2022), where $U_t \sim \text{Uniform}[0, 1]$, would satisfy this requirement.

We now prove some key results about the expected value of the ReLU function applied to functions of a random variable $Z \geq 0$. Firstly, we note that

$$\mathbb{E}[\text{ReLU}(Z)] = \mathbb{E}[Z] \quad (36)$$

We now consider $\mathbb{E}[\text{ReLU}(VZ)]$ for an arbitrary random variable $V \in \mathbb{R}$, which is permitted to be negative. Since $Z \geq 0$, we have

$$\mathbb{E}[\text{ReLU}(VZ)] = \mathbb{E}[\text{ReLU}(V)Z] = \mathbb{E}[\text{ReLU}(V)]\mathbb{E}[Z] \quad (37)$$

In the special case of a constant $V = v$, this simplifies to $\mathbb{E}[\text{ReLU}(vZ)] = \text{ReLU}(v)\mathbb{E}[Z]$.

Now, recalling that $B_t, U_t, X_t \geq 0$ with $\mathbb{E}[B_t] = p$, $\mathbb{E}[U_t] = \mu$ and $\mathbb{E}[X_t] = p\mu$, we apply these results to the expected model output given by

$$\mathbb{E}[Y_t] = \mathbb{E} \left[\text{ReLU} \left(\sum_{s=0}^{t-1} \mathbf{w}_y^\top \mathbf{w}_s X_{t-s} \right) \right] \quad (38)$$

Previously, we were able to apply linearity of expectation to split the expectation of a sum into a sum of expectations, but the ReLU non-linearity precludes using the same approach here. The distribution of the interior weighted sum of uniformly distributed variables is known (Bradley & Gupta, 2002), but its complexity explodes with increasing t . We therefore opt for an approximation.

Let ρ be the spectral radius of W_h . Then the feature x_{t-s} is represented by the vector $\mathbf{w}_s = W_h^s \mathbf{w}_x$ in the hidden state \mathbf{h}_t . For old features, corresponding to large s , we have $\|W_h^s\| \approx \rho^s$ by Gelfand's formula. As argued in Appendix D.1, we are only concerned with the case of $\rho < 1$, so for any $\varepsilon > 0$, there exists a “memory window” of length $T_\varepsilon = \lceil \log(\varepsilon) / \log(\rho) \rceil$ such that for $s \geq T_\varepsilon$, $\|W_h^s\| \approx \rho^s \leq \rho^{T_\varepsilon} \leq \varepsilon$. Hence the contribution of any input older than T_ε has magnitude of order $O(\varepsilon)$. Intuitively, this means that if we set ε small enough (and thus T_ε large enough), we can ignore inputs older than T_ε time steps.

Therefore, the only inputs that can have a significant effect on the model’s behavior are those which arrived in the last T_ε time steps. Since each input is masked by a Bernoulli random variable, the number of non-zero inputs that arrive in T_ε time steps is distributed according to $N_\varepsilon \sim \text{Binomial}(T_\varepsilon, p)$. This quantity essentially counts the number of inputs actually “in play,” meaning that their effect on the hidden state has magnitude larger than ε . Hence, the probability that there are two or more such inputs is given by

$$\Pr[N_\varepsilon \geq 2] = 1 - (1-p)^{T_\varepsilon} - pT_\varepsilon(1-p)^{T_\varepsilon-1} \approx \frac{p^2}{2} T_\varepsilon(T_\varepsilon - 1), \quad (39)$$

where the binomial approximation holds for small p . Therefore, if we are willing to ignore cases that arise with probability less than some $\delta > 0$, we can approximate the behavior of an RNN by its behavior on input sequences with only one non-zero input for $\Pr[N_\varepsilon \geq 2] < \delta$. This occurs when sparsity is high enough to make it vanishingly rare for two or more inputs to be “in play” simultaneously. Specifically, the approximation is valid when

$$p < \sqrt{\frac{2\delta}{T_\varepsilon(T_\varepsilon - 1)}} < \frac{\sqrt{2\delta}}{\log(1/\varepsilon)} \log(1/\rho), \quad (40)$$

or, equivalently,

$$p < \exp \left(-\frac{p \log(1/\varepsilon)}{\sqrt{2\delta}} \right) = \exp \left(\frac{p \log(\varepsilon)}{\sqrt{2\delta}} \right) = \varepsilon^{p/\sqrt{2\delta}}. \quad (41)$$

In particular, for arbitrarily tight $\delta, \varepsilon > 0$, there always exists a p small enough to make the approximation valid for any given model with $\rho < 1$. Under this approximation, our analysis of the ReLU-gated model becomes tractable, as we can ignore all cases that involve two or more non-zero inputs. For instance, the expected model output becomes

$$\mathbb{E} \left[\text{ReLU} \left(\sum_{s=0}^{t-1} \mathbf{w}_y^\top \mathbf{w}_s X_{t-s} \right) \right] \approx \sum_{s=0}^{t-1} \mathbb{E} [\text{ReLU} (\mathbf{w}_y^\top \mathbf{w}_s X_{t-s})] \quad (42)$$

$$1296 \quad = \sum_{s=0}^{t-1} \mathbb{E}[X_{t-s}] \text{ReLU}(\mathbf{w}_y^\top \mathbf{w}_s) \quad (43)$$

$$1297 \quad = p\mu \sum_{s=0}^{t-1} \text{ReLU}(\mathbf{w}_y^\top \mathbf{w}_s). \quad (44)$$

1302 We now follow the steps of Appendix B.1, applying this assumption to derive an interpretable ex-
 1303 pression for the expected squared-error loss $\mathbb{E}[\mathcal{L}]$. We begin with

$$1305 \quad \mathbb{E}[\mathcal{L}] = \sum_{t=k+1}^T \left(\mathbb{E}[X_{t-k}^2] - 2\mathbb{E}\left[X_{t-k} \text{ReLU}\left(\sum_{s=0}^{t-1} \mathbf{w}_y^\top \mathbf{w}_s X_{t-s}\right)\right] \right) \quad (45)$$

$$1308 \quad + \sum_{t=1}^T \mathbb{E}\left[\text{ReLU}\left(\sum_{s=0}^{t-1} \mathbf{w}_y^\top \mathbf{w}_s X_{t-s}\right)^2\right] \quad (46)$$

1312 As before, the first term is $\mathbb{E}[X_{t-k}^2] = p\nu$. To evaluate the second term, we again use the assumption
 1313 that no more than one of the inputs is non-zero. There are two cases: either X_{t-k} is zero, in which
 1314 case the entire term collapses to zero, or X_{t-k} is non-zero, in which case all other X_{t-s} are zero for
 1315 $s \neq k$. Hence the second expectation simplifies to

$$1317 \quad \mathbb{E}\left[X_{t-k} \text{ReLU}\left(\sum_{s=0}^{t-1} \mathbf{w}_y^\top \mathbf{w}_s X_{t-s}\right)\right] \approx \mathbb{E}[X_{t-k} \text{ReLU}(\mathbf{w}_y^\top \mathbf{w}_{s=k} X_{t-k})] \quad (47)$$

$$1319 \quad = \mathbb{E}[X_{t-k}^2] \text{ReLU}(\mathbf{w}_y^\top \mathbf{w}_{s=k}) \quad (48)$$

$$1320 \quad = p\nu \text{ReLU}(\mathbf{w}_y^\top \mathbf{w}_{s=k}) \quad (49)$$

1322 Similarly, in the third expectation, the only non-zero summands are those on the “diagonal” (all
 1323 off-diagonal terms require two inputs to be non-zero, so we ignore them):

$$1325 \quad \mathbb{E}\left[\text{ReLU}\left(\sum_{s=0}^{t-1} \mathbf{w}_y^\top \mathbf{w}_s X_{t-s}\right)^2\right] \approx \sum_{s=0}^{t-1} \mathbb{E}\left[\text{ReLU}(\mathbf{w}_y^\top \mathbf{w}_s X_{t-s})^2\right] \quad (50)$$

$$1329 \quad = \sum_{s=0}^{t-1} \mathbb{E}[X_{t-s}^2] \text{ReLU}(\mathbf{w}_y^\top \mathbf{w}_s)^2 \quad (51)$$

$$1332 \quad = p\nu \sum_{s=0}^{t-1} \text{ReLU}(\mathbf{w}_y^\top \mathbf{w}_s)^2 \quad (52)$$

1334 Putting these together:

$$1336 \quad \mathbb{E}[\mathcal{L}] \approx \sum_{t=k+1}^T (p\nu - 2p\nu \text{ReLU}(\mathbf{w}_y^\top \mathbf{w}_{s=k})) + \sum_{t=1}^T p\nu \sum_{s=0}^{t-1} \text{ReLU}(\mathbf{w}_y^\top \mathbf{w}_s)^2 \quad (53)$$

$$1340 \quad = p\nu \left(\sum_{t=k+1}^T (\text{ReLU}(\mathbf{w}_y^\top \mathbf{w}_{s=k}) - 1)^2 + \sum_{t=1}^T \sum_{\substack{s=0 \\ s \neq k}}^{t-1} \text{ReLU}(\mathbf{w}_y^\top \mathbf{w}_s)^2 \right) \quad (54)$$

1344 E.2 SIMULATIONS ACROSS k AND SPARSITY FOR NONLINEAR READOUT

1346 In the SSM, we observe that the spectral radius (i.e., $\|\mathbf{w}_{s=0}\|$) increases with k , regardless of sparsity
 1347 (Figure 9). This is simply because for larger k , the model must hold inputs in its memory for more
 1348 timesteps. If the spectral radius is too small, then the magnitude of $\mathbf{w}_{s=k}$ in the hidden state will be
 1349 negligible relative to other $\mathbf{w}_{s \neq k}$ and so any task-relevant signal will be overpowered by projection
 interference from other features.

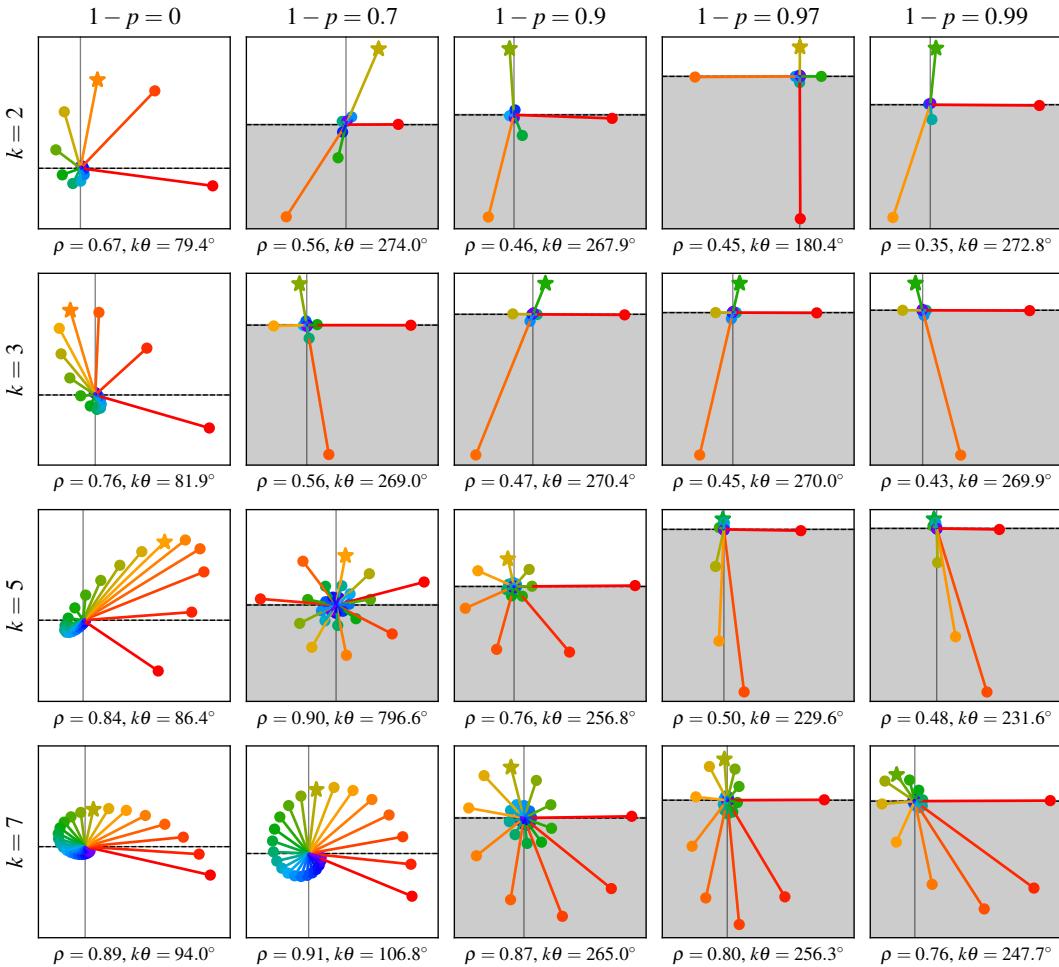


Figure 9: **Solutions to k -delay learned by linear models with ReLU read-out (SSMs).** Rows correspond to different values of k and columns correspond to different sparsity levels. The \mathbf{w}_s are plotted after applying a conformal linear transformation such that the y -component of each \mathbf{w}_s is $\mathbf{w}_y^\top \mathbf{w}_s$, and $\mathbf{w}_x = \mathbf{w}_{s=0}$ points towards positive x . Thus the interference-free half-space is simply given by $y < 0$. As before, the output feature, $\mathbf{w}_{s=k}$, is marked with a star. Note that the plots vary significantly in scale, so it is not meaningful to compare the magnitude of a particular \mathbf{w}_s vector between different plots. The angle θ is calculated as $\arg(\lambda_1)$ where $\lambda_1 \in \mathbb{C}$ is an eigenvalue of W_h .

1404 F NONLINEAR RNN (NONLINEAR RECURRENCE, NONLINEAR READOUT)
14051406 F.1 DISCUSSION ON APPLYING NONLINEARITY TO HIDDEN LAYER
14071408 Recall that we have defined the RNN architecture as
1409

1410
$$\mathbf{h}_t = W_x \mathbf{x}_t + W_h \sigma_h(\mathbf{h}_{t-1}) \quad \hat{\mathbf{y}}_t = \sigma_y(W_y^\top \mathbf{h}_t) \quad (55)$$

1411 While this is a slight departure from traditional RNN architecture in machine learning (with regards
1412 to what we label as \mathbf{h}_t and how the output $\hat{\mathbf{y}}_t$ is computed), this form is standard in computational
1413 neuroscience. Moreover, prior work has shown that these models are mathematically equivalent
1414 (Miller & Fumarola, 2012).

1415
$$\text{ML } \mathbf{h}_t = \sigma_h \left(\underbrace{W_x \mathbf{x}_t + W_h \mathbf{h}_{t-1}}_{\text{our/neuro } \mathbf{h}_t} \right) \quad (56)$$

1416
1417
1418

1419 The motivation for using this form of RNN is that it allows us to study a nonlinear RNN for which
1420 the linear representation hypothesis provably holds (under certain sparsity conditions), as shown in
1421 Appendix F.2, and is the most direct extension of Elhage et al. (2022) to the recurrent setting.
14221423 F.2 NONLINEAR RNNs HAVE LINEAR FEATURE DIRECTIONS IN THE LIMIT OF HIGH
1424 SPARSITY
14251426 Applying a non-linearity to the hidden state immediately breaks the linear representation hypothesis:
1427 each feature would be represented along a (not necessarily smooth) curve rather than a straight line.
1428 This massively complicates the study of non-linear recurrence in general.1429 For the case of linear recurrence, we have shown that for sufficiently sparse input sequences, we can
1430 assume that there is at most one non-zero input “in play” within a model’s hidden state at any given
1431 time. Fundamentally, this was based on the idea that there exists a memory window of length T such
1432 that inputs older than T time steps cannot contribute significantly to the current hidden state. This
1433 essentially arose from the proof in Appendix D.1 that RNNs with linear recurrence must satisfy the
1434 echo state property in order to achieve reasonable loss.1435 We argue that a similar memory window should be expected in models with nonlinear recurrence.
1436 Although it is much harder to prove that the echo state property is a requirement for good perfor-
1437 mance in non-linear models, there exists plenty of evidence for the reverse statement: echo state
1438 networks are, by definition, nonlinear models that satisfy the echo state property (Jaeger, 2002) and
1439 have been shown to achieve strong performance on a variety of sequential processing tasks (Aceituno
1440 et al., 2020). Clearly, the echo state property is not incompatible with strong task performance and
1441 it is not implausible that the linear result – that the echo state property is actually required for tasks
1442 like k -delay – carries over to the case of nonlinear recurrences, based on these empirical observa-
1443 tions. Furthermore, it’s well known that RNNs with spectral radius over 1 cause exploding gradients
1444 and training instability (Bengio et al., 1994; Hochreiter et al., 2001; Pascanu et al., 2012), and that
1445 task dynamics with strong early correlations (producing a spectral radius over 1) result in network
1446 instability (Proca et al., 2025). Therefore, we intuitively expect our reasoning to hold for nonlinear
1447 recurrence: if the input sequence is made sufficiently sparse, we can approximate the model’s behav-
1448 ior by ignoring situations where two or more inputs are non-zero, as historic features are gradually
1449 forgotten by the RNN due to the shrinking effect of the spectral radius ($\rho < 1$).1450 Unrolling the hidden state, we see that an analysis of the general case is intractable:
1451

1452
$$\begin{aligned} \mathbf{h}_t &= \mathbf{w}_x \mathbf{x}_t + W_h \text{ReLU}(\mathbf{h}_{t-1}) \\ &= \mathbf{w}_x \mathbf{x}_t + W_h \text{ReLU}(\mathbf{w}_x \mathbf{x}_{t-1} + W_h \text{ReLU}(\mathbf{h}_{t-2})) \\ &\vdots \\ &= \mathbf{w}_x \mathbf{x}_t + W_h \text{ReLU}(\mathbf{w}_x \mathbf{x}_{t-1} + W_h \text{ReLU}(\cdots W_h \text{ReLU}(\mathbf{w}_x \mathbf{x}_2 + W_h \text{ReLU}(\mathbf{w}_x \mathbf{x}_1)) \cdots)) \end{aligned}$$

1453 Suppose, however, that for some s , only \mathbf{x}_{t-s} is non-zero and all other inputs are assumed to be zero,
1454 as per our approximation. Recalling that in our setup, \mathbf{x}_{t-s} is scalar and non-negative, the hidden
1455

1458 state simplifies as follows:
 1459

$$\mathbf{h}_t = \mathbf{w}_x(0) + W_h \text{ReLU}(\mathbf{w}_x(0) + W_h \text{ReLU}(\cdots W_h \text{ReLU}(\mathbf{w}_x x_{t-s} + W_h(0)) \cdots)) \quad (57)$$

$$= W_h \text{ReLU}(W_h \text{ReLU}(\cdots W_h \text{ReLU}(\mathbf{w}_x x_{t-s}) \cdots)) \quad (58)$$

$$= W_h \text{ReLU}(W_h \text{ReLU}(\cdots x_{t-s} W_h \text{ReLU}(\mathbf{w}_x) \cdots)) \quad (59)$$

$$= x_{t-s} \underbrace{W_h \text{ReLU}(W_h \text{ReLU}(\cdots W_h \text{ReLU}(\mathbf{w}_x) \cdots))}_{\mathbf{w}_s} \quad (60)$$

$$= \mathbf{w}_s x_{t-s} \quad (61)$$

1468 Though it is not analytically possible to find a simplified expression for the vector \mathbf{w}_s , it nevertheless
 1469 is the direction in which the feature x_{t-s} is represented in the limit of sparsity, as $p \rightarrow 0$. Hence,
 1470 in the extremely sparse regime, the linear representation hypothesis holds for this model. This is
 1471 not a trivial result; it relies on both the piecewise linearity of ReLU for non-negative inputs and our
 1472 definition of \mathbf{h}_t as the hidden state prior to application of ReLU.

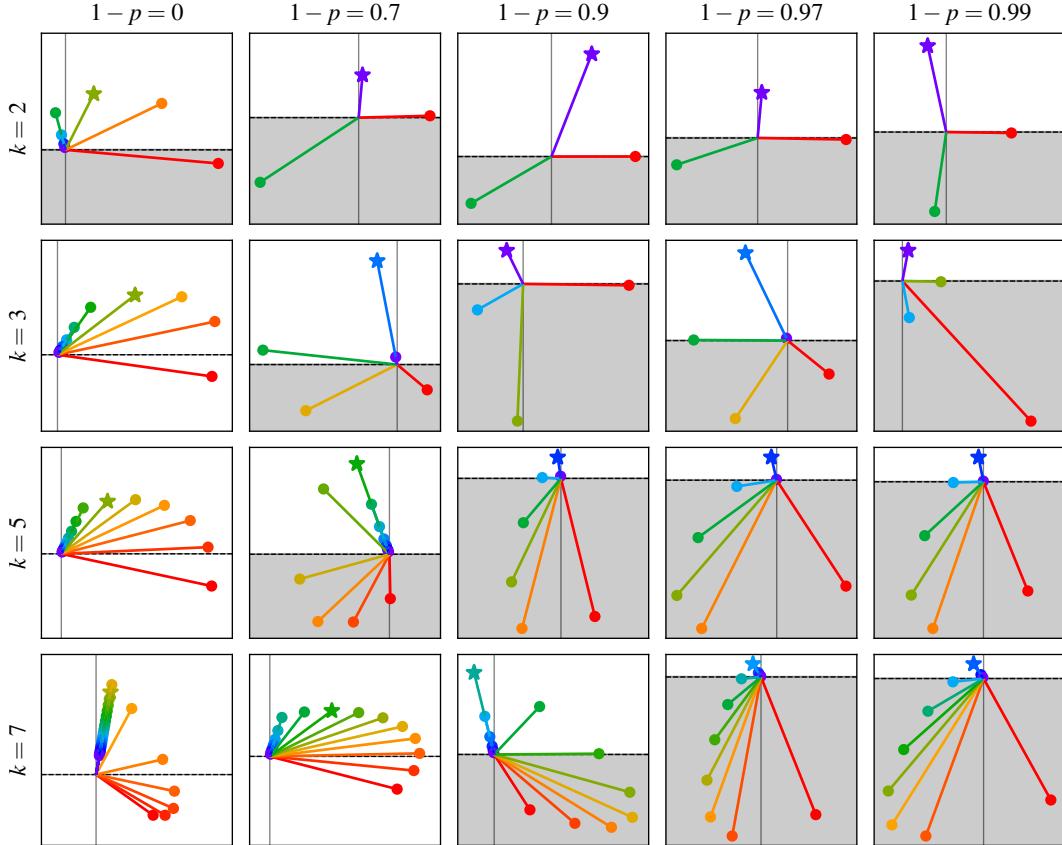
1473

1474 F.3 SIMULATIONS ACROSS k AND SPARSITY FOR NONLINEAR RECURRENCE

1475

1476 In the nonlinear RNN, we observe how the model learns to exploit the interference-free space as
 1477 sparsity increases and to implement sharp forgetting.

1478



1507 **Figure 10: Solutions to k -delay learned by nonlinear RNNs.** Rows correspond to different values
 1508 of k and columns correspond to different sparsity levels. The \mathbf{w}_s are plotted after applying a confor-
 1509 mal linear transformation such that the y -component of each \mathbf{w}_s is $\mathbf{w}_y^\top \mathbf{w}_s$, and $\mathbf{w}_x = \mathbf{w}_{s=0}$ points
 1510 towards positive x . Thus the interference-free half-space is simply given by $y < 0$. As before, the
 1511 output feature, $\mathbf{w}_{s=k}$, is marked with a star. Note that the plots vary significantly in scale, so it is
 not meaningful to compare the magnitude of a particular \mathbf{w}_s vector between different plots.

1512 **G HIGHER DIMENSIONAL RESULTS**
 1513

1514 **G.1 DECOMPOSING THE LOSS FOR VECTOR INPUTS AND OUTPUTS ($N_x > 1$)**
 1515

1516 We can repeat the decomposition in Appendix B to obtain an expression for the expected value of the
 1517 squared-error loss incurred by linear models in the case of vector inputs and outputs ($N_x = N_y > 1$).
 1518 We recall that $X_t, \hat{Y}_t \in \mathbb{R}^{N_x}$, $W_s \in \mathbb{R}^{N_h \times N_x}$ and $W_y \in \mathbb{R}^{N_h \times N_x}$ and proceed as before:
 1519

1520
$$\mathbb{E}[\mathcal{L}] = \mathbb{E} \left[\sum_{t=1}^k \left\| 0 - \hat{Y}_t \right\|^2 + \sum_{t=k+1}^T \|X_{t-k} - \hat{Y}_t\|^2 \right] \quad (62)$$

 1521

1522
$$= \sum_{t=1}^k \mathbb{E} \left[\left\| \hat{Y}_t \right\|^2 \right] + \sum_{t=k+1}^T \mathbb{E} \left[\left\| X_{t-k} - \hat{Y}_t \right\|^2 \right] \quad (63)$$

 1523

1524
$$= \sum_{t=1}^k \mathbb{E} \left[\left\| \hat{Y}_t \right\|^2 \right] + \sum_{t=k+1}^T \mathbb{E} \left[\|X_{t-k}\|^2 - 2X_{t-k}^\top \hat{Y}_t + \left\| \hat{Y}_t \right\|^2 \right] \quad (64)$$

 1525

1526
$$= \sum_{t=k+1}^T \left(\mathbb{E} \left[\|X_{t-k}\|^2 \right] - 2\mathbb{E} \left[X_{t-k}^\top \hat{Y}_t \right] \right) + \sum_{t=1}^T \mathbb{E} \left[\left\| \hat{Y}_t \right\|^2 \right] \quad (65)$$

 1527

1528
$$= \sum_{t=k+1}^T \left(\mathbb{E} \left[\|X_{t-k}\|^2 \right] - 2\mathbb{E} \left[X_{t-k}^\top \sum_{s=0}^{t-1} W_y^\top W_s X_{t-s} \right] \right) + \sum_{t=1}^T \mathbb{E} \left[\left\| \sum_{s=0}^{t-1} W_y^\top W_s X_{t-s} \right\|^2 \right] \quad (66)$$

 1529

1530 We extend the temporal sparsity assumption to assume that each input feature follows the same
 1531 distribution, so that $X_t^{(i)} = B_t^{(i)} U_t^{(i)}$ with $B_t^{(i)} \sim \text{Bernoulli}(p)$, $U_t^{(i)}$ is identically distributed
 1532 according to any distribution and $\{B_t^{(i)}\} \cup \{U_t^{(i)}\}$ are mutually independent.
 1533

1534 Then, with $\mu := \mathbb{E} \left[U_t^{(i)} \right]$ and $\nu := \mathbb{E} \left[\left(U_t^{(i)} \right)^2 \right]$, we can simplify the first expectation to
 1535

1536
$$\mathbb{E} \left[\|X_{t-k}\|^2 \right] = \sum_{i=1}^{N_x} \mathbb{E} \left[\left(X_{t-k}^{(i)} \right)^2 \right] = \sum_{i=1}^{N_x} \mathbb{E} \left[\left(B_{t-k}^{(i)} \right)^2 \left(U_{t-k}^{(i)} \right)^2 \right] = \sum_{i=1}^{N_x} p\nu = N_x p\nu \quad (67)$$

 1537

1538 In computing the second expectation, we must handle the case of $s = k$ separately:
 1539

1540
$$\begin{aligned} & \mathbb{E} \left[X_{t-k}^\top \sum_{s=0}^{t-1} W_y^\top W_s X_{t-s} \right] \\ &= \mathbb{E} \left[\text{tr} \left(\sum_{s=0}^{t-1} X_{t-k}^\top W_y^\top W_s X_{t-s} \right) \right] \end{aligned} \quad (68)$$

 1541

1542
$$= \sum_{s=0}^{t-1} \text{tr} \left(\mathbb{E} \left[X_{t-k}^\top W_y^\top W_s X_{t-s} \right] \right) \quad (69)$$

 1543

1544
$$= \sum_{s=0}^{t-1} \text{tr} \left(W_y^\top W_s \mathbb{E} \left[X_{t-s} X_{t-k}^\top \right] \right) \quad (70)$$

 1545

1546
$$= \text{tr} \left(W_y^\top W_k \mathbb{E} \left[X_{t-k} X_{t-k}^\top \right] \right) + \sum_{s \neq k}^{t-1} \text{tr} \left(W_y^\top W_s \mathbb{E} \left[X_{t-s} X_{t-k}^\top \right] \right) \quad (71)$$

 1547

1548
$$= \text{tr} \left(W_y^\top W_k \mathbb{E} \left[X_{t-k} X_{t-k}^\top \right] \right) + p^2 \mu^2 \sum_{s \neq k}^{t-1} \text{tr} \left(W_y^\top W_s \mathbf{1} \mathbf{1}^\top \right) \quad (72)$$

 1549

$$1566 \quad = [p^2\mu^2\mathbf{1}^\top W_y^\top W_k \mathbf{1} + (p\nu - p^2\mu^2) \operatorname{tr}(W_y^\top W_k)] + p^2\mu^2 \sum_{s \neq k}^{t-1} \mathbf{1}^\top W_y^\top W_s \mathbf{1} \quad (73)$$

$$1569 \quad = (p\nu - p^2\mu^2) \operatorname{tr}(W_y^\top W_k) + p^2\mu^2 \sum_{s=0}^{t-1} \mathbf{1}^\top W_y^\top W_s \mathbf{1} \quad (74)$$

1572 where we make use of the all-ones vector $\mathbf{1} = [1, 1, 1, \dots]^\top \in \mathbb{R}^{N_x}$.

1574 Finally, the third expectation simplifies as follows:

$$1575 \quad \mathbb{E} \left[\left\| \sum_{s=0}^{t-1} W_y^\top W_s X_{t-s} \right\|^2 \right] \\ 1576 \quad = \mathbb{E} \left[\left(\sum_{s=0}^{t-1} W_y^\top W_s X_{t-s} \right)^\top \left(\sum_{s'=0}^{t-1} W_y^\top W_{s'} X_{t-s'} \right) \right] \quad (75)$$

$$1582 \quad = \mathbb{E} \left[\left(\sum_{s=0}^{t-1} X_{t-s}^\top W_s^\top W_y \right) \left(\sum_{s'=0}^{t-1} W_y^\top W_{s'} X_{t-s'} \right) \right] \quad (76)$$

$$1585 \quad = \sum_{s=0}^{t-1} \sum_{s'=0}^{t-1} \mathbb{E} [X_{t-s}^\top W_s^\top W_y W_y^\top W_{s'} X_{t-s'}] \quad (77)$$

$$1588 \quad = \sum_{s=0}^{t-1} \sum_{s'=0}^{t-1} \operatorname{tr}(W_s^\top W_y W_y^\top W_{s'} \mathbb{E}[X_{t-s'} X_{t-s}^\top]) \quad (78)$$

$$1591 \quad = \sum_{s=0}^{t-1} \operatorname{tr}(W_s^\top W_y W_y^\top W_s \mathbb{E}[X_{t-s} X_{t-s}^\top]) + \sum_{s \neq s'}^{t-1} \operatorname{tr}(W_s^\top W_y W_y^\top W_{s'} \mathbb{E}[X_{t-s'} X_{t-s}^\top]) \quad (79)$$

$$1594 \quad = \left[p^2\mu^2 \sum_{s=0}^{t-1} \operatorname{tr}(W_s^\top W_y W_y^\top W_s \mathbf{1} \mathbf{1}^\top) + (p\nu - p^2\mu^2) \sum_{s=0}^{t-1} \operatorname{tr}(W_s^\top W_y W_y^\top W_s) \right] \quad (80)$$

$$1597 \quad + p^2\mu^2 \sum_{s \neq s'}^{t-1} \operatorname{tr}(W_s^\top W_y W_y^\top W_{s'} \mathbf{1} \mathbf{1}^\top) \quad (81)$$

$$1600 \quad = (p\nu - p^2\mu^2) \sum_{s=0}^{t-1} \operatorname{tr}(W_s^\top W_y W_y^\top W_s) + p^2\mu^2 \sum_{s=0}^{t-1} \sum_{s'=0}^{t-1} \mathbf{1}^\top W_s^\top W_y W_y^\top W_{s'} \mathbf{1} \quad (82)$$

$$1603 \quad = (p\nu - p^2\mu^2) \sum_{s=0}^{t-1} \|W_y^\top W_s\|_F^2 + p^2\mu^2 \left\| \sum_{s=0}^{t-1} W_y^\top W_s \mathbf{1} \right\|^2 \quad (83)$$

1606 where $\|\cdot\|_F$ is the Frobenius norm.

1608 Putting all the terms together yields

$$1609 \quad \mathbb{E}[\mathcal{L}] = \sum_{t=k+1}^T \left(N_x p\nu - 2 \left[(p\nu - p^2\mu^2) \operatorname{tr}(W_y^\top W_k) + p^2\mu^2 \sum_{s=0}^{t-1} \mathbf{1}^\top W_y^\top W_s \mathbf{1} \right] \right) \quad (84)$$

$$1612 \quad + \sum_{t=1}^T \left[(p\nu - p^2\mu^2) \sum_{s=0}^{t-1} \|W_y^\top W_s\|_F^2 + p^2\mu^2 \left\| \sum_{s=0}^{t-1} W_y^\top W_s \mathbf{1} \right\|^2 \right] \quad (85)$$

G.2 ANALYZING HIGHER-DIMENSIONAL HIDDEN STATES ($N_h > 2$)

1618 So far, for the purposes of feature geometry, we have restricted the hidden state of our models to 2 dimensions. In this section, we demonstrate that the results we have found generalize well to RNNs

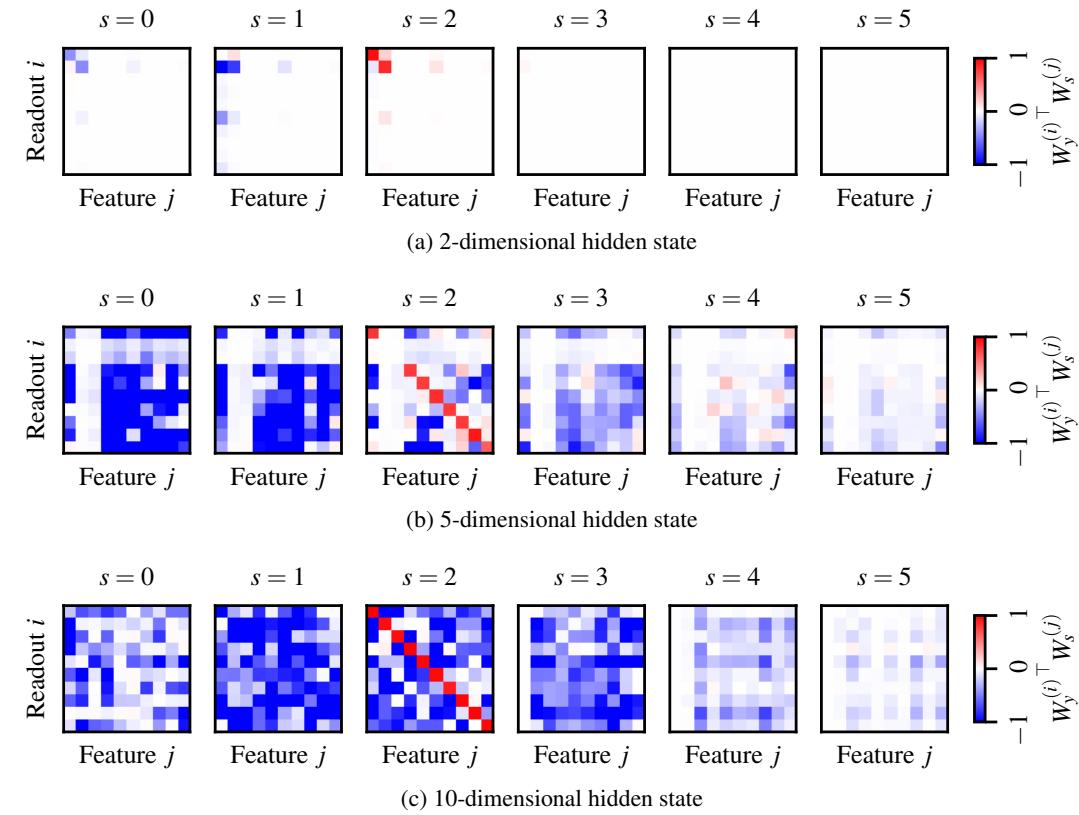


Figure 11: **Projections of feature directions onto readouts in higher-dimensional hidden states.** Each model was trained on the 2-delay task over 10 features (hence there are also 10 readouts). Each plot shows the values of $W_y^\top W_s$ – the projections of feature directions corresponding to s -timestep-old inputs onto each of the readout vectors. Specifically, each cell (i, j) represents the value of $W_y^{(i)\top} W_s^{(j)}$, where $W_y^{(i)}$ is the i -th column of W_y and $W_s^{(j)}$ is the j -th column of W_s . Cells that are red correspond to a positive projection onto the readout – this is desired only on the diagonal of the $s = 2$ panel, where features are being read out into the correct outputs at the correct time; red elsewhere represents projection interference. Blue cells have a negative projection onto the readout and, due to the ReLU, do not contribute to projection interference for their row’s readout. **(a)** A 2-dimensional hidden state is only able to represent 2 features well and ignores the rest – a clear example of the ‘all-or-none’ strategy discussed in Section 4.5. This result is equivalent to the $k = 2$ panel of Figure 5. **(b)** A 5-dimensional hidden state performs much better, representing many more of the features, albeit with some interference. **(c)** A 10-dimensional hidden state performs very well, representing all the features with almost no interference.

with higher-dimensional hidden states. In particular, we find that the interference-free space is not only present, but very well exploited in higher dimensions.

Figure 11 shows the results of training higher-dimensional nonlinear RNNs on the 2-delay task with vector inputs ($N_x = 10$). As such, this is the most generalized form of our results. Even in the 10-dimensional hidden state, there is still significant superposition occurring: a 10-dimensional hidden state can only represent 10 features orthogonally, whereas the task requires 10 features to be held in memory over $k + 1 = 3$ timesteps, equivalent to 30 features being compressed into 10 dimensions.

The 10-dimensional case exhibits very little projection interference – most cells are blue except on the $s = 2$ diagonal, where we expect them to be red. This shows that most activations lie within the interference-free space, where their projection onto every readout vector is negative (and therefore their contribution to projection interference is zero due to the ReLU activation function). We believe this is a strong result that demonstrates that the interference-free space is a significant driver of feature geometry, even in – or perhaps especially in – higher-dimensional hidden states.

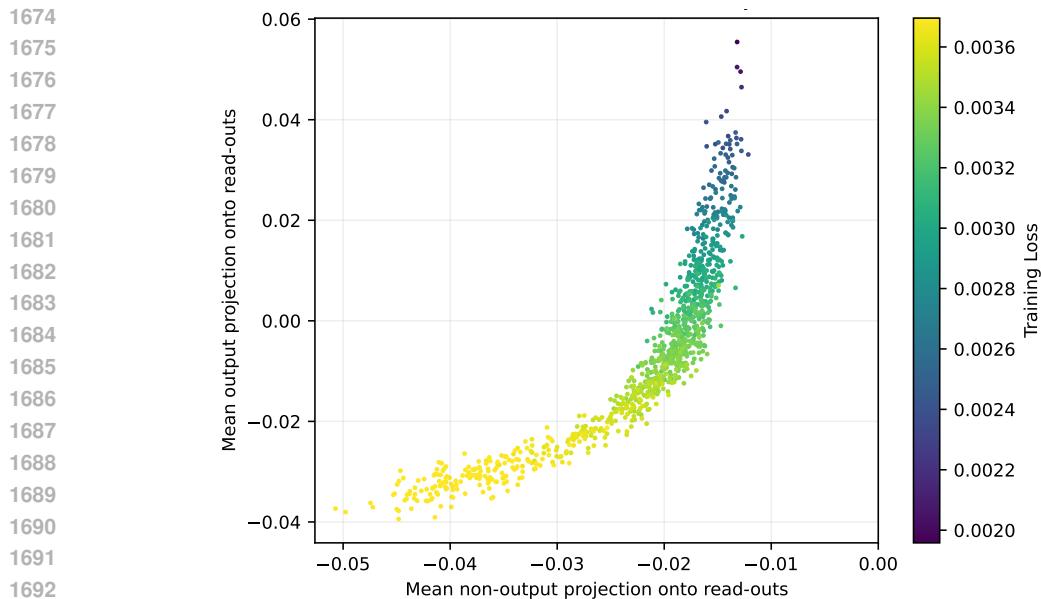


Figure 12: **Mean projections of output and non-output feature directions onto readouts colored by loss in higher-dimensional hidden states.** We train 1000 models with a hidden size of 100 on the 2-delay task with 75 features. The x-axis indicates the mean of the non-output feature direction projections onto the readout (i.e., $W_y^\top W_{s \neq k}$). This is the analogous to the mean of the $s \neq 2$ matrices in Figure 11. The y-axis indicates the mean of the output feature direction projections onto the readout (i.e., $W_y^\top W_{s=k}$), which is analogous to the diagonal of the $s = 2$ matrix in Figure 11. By looking at the plot, we see that all models learn to have non-output feature directions that have negative projections onto the readout (negative x-axis values), utilizing the interference-free space. We also see that performance is correlated with the RNN’s ability to positively project the output feature direction onto the readout, indicating that the optimal models are using the the interference-free space to minimize projection interference with this output feature.

Based on these findings, we extend this idea to larger RNNs, developing metrics to quantify the degree to which these networks exhibit the geometry we would expect based on our study. In particular, we measure the mean non-output feature direction projections onto the readout (i.e., $\text{mean}(W_y^\top W_{s \neq k})$), which allows us to quantify the degree to which non-output feature directions group within the interference-free space. In particular, this value should be negative if the largest non-output feature directions are grouping within the interference-free space. This is analogous to the mean of the $s \neq 2$ matrices in Figure 11.

We also measure the mean of the output feature direction projections onto the readout (i.e., $\text{mean}(\text{diag}(W_y^\top W_{s=k}))$), which quantifies whether the appropriate features are projected onto the readout at the correct time to perform the task. We would expect this value to be positive for models that successfully do this. This is analogous to the mean of the diagonal of the $s = 2$ matrix in Figure 11.

We train 1000 RNNs with a hidden size of 100 on a 2-delay task with 75 features and plot the results in Figure 12, where each point is representative of a single model and is colored by the final loss it achieves. By looking at the figure, we see that all models learn to group the largest non-output feature directions in the interference-free space (indicated by the negative-valued x-axis). We also see that the best-performing models (lowest loss) learn to successfully project the output feature direction onto the readout (indicated by the positive-valued y-axis), indicating that these (optimal) models are outputting the correct features at the correct timestep and using the interference-free space to minimize projection interference with this output feature. Our results pertaining to optimal model geometry are therefore corroborated by this experiment.

In our experiment, there are models that do not learn the optimal solution (i.e., they have a negative output feature direction projection onto the readout; the yellow/light green points in Figure 12).

1728 We suspect that this has to do with the gradient backpropagation through ReLU, as the negative
 1729 projections can get clipped to zero by the ReLU nonlinearity, preventing a learning signal from going
 1730 through. Indeed, we do not use any additional methods to assist with training in this experiment.
 1731 Moreover, if we train instead with a leaky ReLU activation (permitting gradients to backpropagate
 1732 to negative projections), there are no models with a negative mean output projection (all models
 1733 learn the optimal solution metric-wise). We note however that the extension of our work to other
 1734 activation functions is non-trivial and beyond the scope of this work.

H TASKS WITH RANDOM DELAY

1757 In this paper we primarily focus on tasks with a fixed k -delay. Here, we instead consider the effect
 1758 of training on a task with random delay. The task we consider is identical to the k -delay task in that
 1759 the RNN must reproduce the input sequence after k timesteps, but now k is random for each training
 1760 sample ($k \sim \text{Uniform}(0,10)$). One dimension of the input corresponds to the cue, which remains 0
 1761 until the randomly selected k , after which it is set to 1 and the RNN is tasked with outputting the
 1762 sequence, corresponding to the input from the $t - k, \forall t > k$. We train RNNs of each architecture
 1763 (linear: Figure 13, SSM: Figure 14, nonlinear: Figure 15) and visualize how the feature geometry
 1764 changes as the number of input features is varied (the rows) and sparsity is varied (the columns).

1765 Although we are cautious about overinterpreting these plots, we provide a preliminary analysis. The
 1766 results seem to suggest some intermediate geometry between spatial superposition and temporal
 1767 superposition. Indeed, the notion of time-dependency here marks a departure from the rest of the
 1768 paper in that the *sequential ordering* of features is important for the task, but a *time-dependent*
 1769 output is not. We see that most RNNs form solutions where feature directions lie on a shrinking
 1770 line (instead of a spiral sink), with a fixed point at the origin (corresponding to ‘forgetting’). ‘Age’
 1771 (for sequential ordering) is still partially encoded by the magnitude of the feature direction on the
 1772 line. RNNs also appear to be implementing some form of spatial superposition in some cases,
 1773 partitioning the activation space for several different features; this behavior clearly contrasts from
 1774 Figure 5, which, for many input features and delays of up to 10, would only choose to represent one
 1775 feature. However, we also often see a collapse of several feature directions onto the same line. In
 1776 fact, although we study up to 7 input features, the models typically converge to approximately 2-3
 1777 principle directions. We can also see how for 2 features, most RNNs learn to represent these features
 1778 approximately orthogonally.

1779 Remarkably, in the SSM with high sparsity, we exactly recover the pentagon of 5 features charac-
 1780 teristic of spatial superposition. We suspect that the SSMs geometric strategies are a result of each
 1781 feature direction placing itself in the interference-free space of the other feature directions’ readouts
 (hence the spiraling, perpendicular, and pentagon shapes).

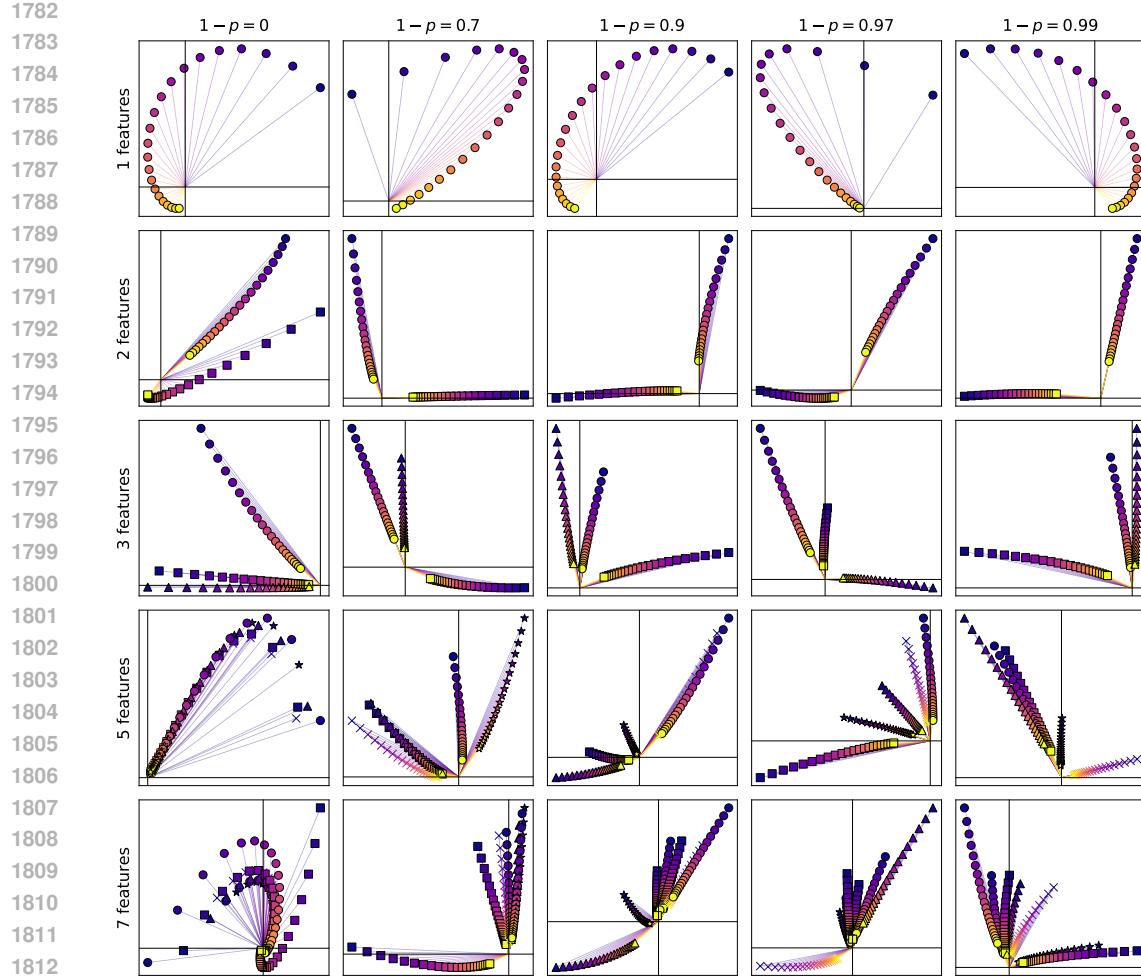


Figure 13: **Solutions to random-delay learned by linear RNNs.** Rows correspond to different numbers of input features (1,2,3,5,7) and columns correspond to different sparsity levels (0, 0.7, 0.9, 0.97, 0.99). Each feature is indicated by a separate marker, and ‘age’ in the network is indicated by color (purple is new; yellow is old).

1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835

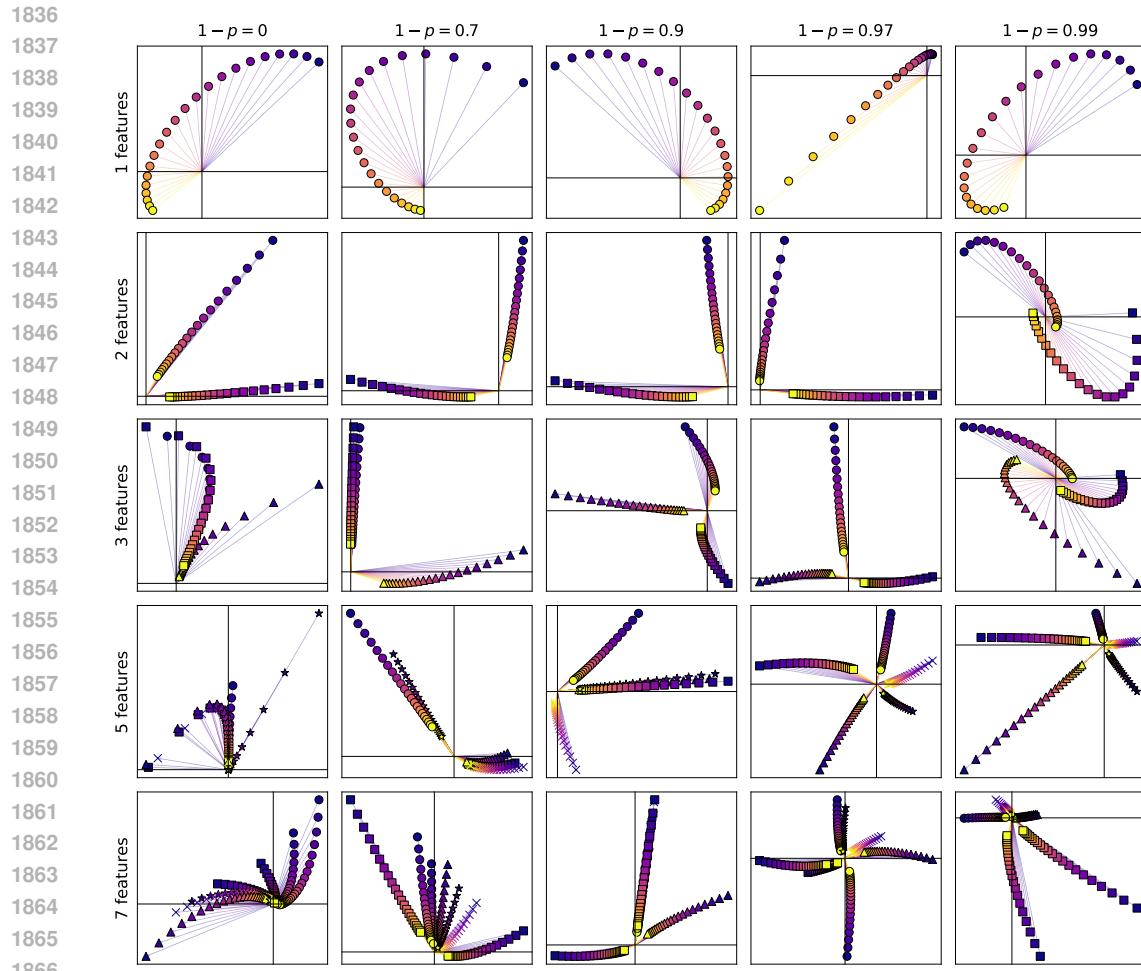


Figure 14: **Solutions to random-delay learned by SSMs.** Rows correspond to different numbers of input features (1,2,3,5,7) and columns correspond to different sparsity levels (0, 0.7, 0.9, 0.97, 0.99). Each feature is indicated by a separate marker, and ‘age’ in the network is indicated by color (*purple is new; yellow is old*).

1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889

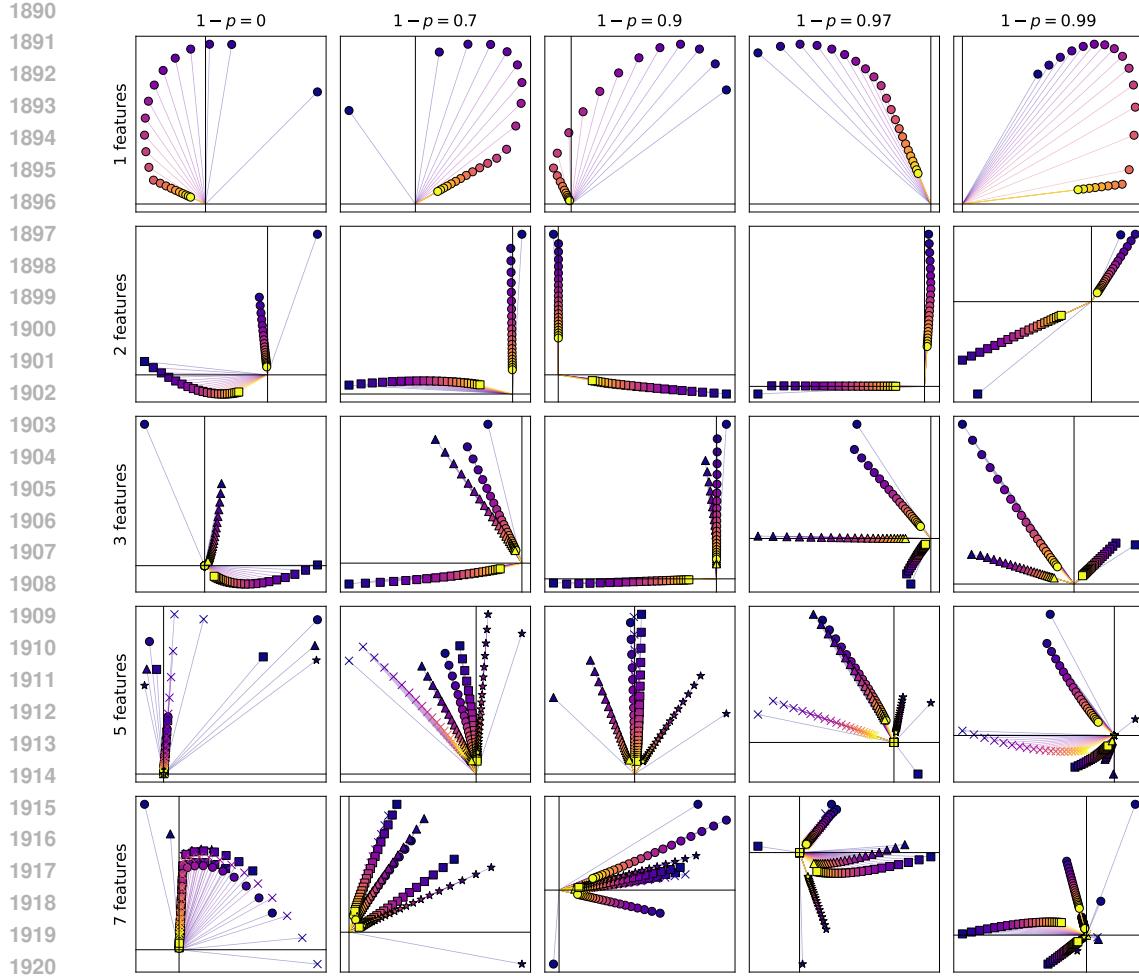


Figure 15: **Solutions to random-delay learned by nonlinear RNNs.** Rows correspond to different numbers of input features (1,2,3,5,7) and columns correspond to different sparsity levels (0, 0.7, 0.9, 0.97, 0.99). Each feature is indicated by a separate marker, and ‘age’ in the network is indicated by color (*purple is new; yellow is old*).

I EXPERIMENTAL DETAILS

I.1 NUMBER OF MODELS THAT DEVELOP THE FEATURE GEOMETRIES DISCUSSED

To verify that a reasonable number of models actually achieve the kinds of feature geometry discussed in our work, we train 1000 models of each architecture for various delay and sparsity values (Table 1). We find that a substantial number of models satisfy our heuristics, designed to identify when a model has achieved the characteristics of the expected feature geometry. We note that optimizing with a two-dimensional hidden space is very challenging for gradient descent. Indeed, in the original work on toy models of superposition (Elhage et al., 2022), the authors also study a 2-dimensional hidden state and report that they fit each model multiple times and take the solution with the lowest loss due to these optimization challenges.

In the case of the linear architecture, the number of models that achieve the expected “spiral sink” feature geometry decreases with increasing k – this is simply because the task becomes too challenging for such a simple architecture to learn well; in many cases, models resort to oscillatory behaviour that achieves suboptimal but lower-than-baseline loss. In contrast, over 40% of SSMs trained on larger k and under high sparsity learn the expected “spiral sink” solution. For the non-

1944 linear model, using a rudimentary heuristic, we find that a reasonable proportion of models can be
 1945 clearly said to have learned a feature geometry in which they exploit the interference-free space as
 1946 discussed in Section 4.4.

1948	1949	1950	1951	1952	1953	1954	1955	1956	1957	1958	1959	1960	Number of models (%)				
													k	$1 - p$	Linear	SSM	Nonlinear
1961	1962	1963	1964	1965	1966	1967	1968	1969	1970	1971	1972	1973	2	0.9	26.5	31.7	31.4
													2	0.97	30.8	20.6	26.8
													2	0.99	27.0	28.5	27.6
													3	0.9	21.0	41.3	18.8
													3	0.97	18.2	32.8	15.6
													3	0.99	19.0	33.8	13.8
													5	0.9	19.4	53.5	8.20
													5	0.97	18.2	48.0	10.3
													5	0.99	16.0	45.3	9.40
													7	0.9	14.3	47.1	12.3
													7	0.97	11.4	42.5	15.6
													7	0.99	11.4	41.5	13.0

Table 1: **Percentage of models that exhibit characteristics of optimal geometry**, as defined by the following heuristics. *Linear*: we expect a reasonable spectral radius and for the feature direction vectors to rotate through less than 180° between input and output, so we count the number of models satisfying $0.5 < \rho(W_h) < 1$ and $5^\circ < |k| \arg(\lambda)| < 180^\circ$, where λ is an eigenvalue of $W_h \in \mathbb{R}^{2 \times 2}$. *SSM*: similar to the linear case, but here we expect the feature direction vectors to rotate through more than 180° between input and output, so we check for $0.5 < \rho(W_h) < 1$ and $|k| \arg(\lambda)| > 180^\circ$. *Nonlinear*: the optimal geometry occurs when all but the output feature direction vector lie in the interference-free space; only one feature direction vector should project positively onto the readout vector, so we check for $w_y^\top w_k > 0$ and $w_y^\top w_s < 0.1$ (tolerating some small positive projections) for $0 < s \leq k$.

1971 I.2 WEIGHT TYING

1974 For visual clarity, in Figure 1, Figure 3 and Figure 5, we set the readout vector $w_y := w_{s=k}$ (or,
 1975 in the case of vector inputs and outputs, $W_y := W_{s=k}$). This means that, in these cases, w_y (or
 1976 W_y) is not a separate trainable parameter of the model and is instead entirely determined by the
 1977 parameters W_h and w_x (or W_x). There are two reasons for doing this: first, it eliminates the need
 1978 for a separate readout vector to be shown, making the plots neater; second, it encourages features
 1979 in spatial superposition to arrange into regular polygons (e.g. Figure 1a). In fact, this is merely an
 1980 extension of the weight tying used by Elhage et al. (2022), where the authors set $W_y := W_x^\top$ for the
 1981 same reasons. Our weight tying is identical to theirs in the $k = 0$ case.

1982 We note that this is just a visualization trick and, to avoid doubt, Figure 16 plots spatial and tem-
 1983 *1984* *1985* *1986* *1987* *1988* *1989* *1990* *1991* *1992* *1993* *1994* *1995* *1996* *1997* *1998* *1999* *2000* *2001* *2002* *2003* *2004* *2005* *2006* *2007* *2008* *2009* *2010* *2011* *2012* *2013* *2014* *2015* *2016* *2017* *2018* *2019* *2020* *2021* *2022* *2023* *2024* *2025* *2026* *2027* *2028* *2029* *2030* *2031* *2032* *2033* *2034* *2035* *2036* *2037* *2038* *2039* *2040* *2041* *2042* *2043* *2044* *2045* *2046* *2047* *2048* *2049* *2050* *2051* *2052* *2053* *2054* *2055* *2056* *2057* *2058* *2059* *2060* *2061* *2062* *2063* *2064* *2065* *2066* *2067* *2068* *2069* *2070* *2071* *2072* *2073* *2074* *2075* *2076* *2077* *2078* *2079* *2080* *2081* *2082* *2083* *2084* *2085* *2086* *2087* *2088* *2089* *2090* *2091* *2092* *2093* *2094* *2095* *2096* *2097* *2098* *2099* *2100* *2101* *2102* *2103* *2104* *2105* *2106* *2107* *2108* *2109* *2110* *2111* *2112* *2113* *2114* *2115* *2116* *2117* *2118* *2119* *2120* *2121* *2122* *2123* *2124* *2125* *2126* *2127* *2128* *2129* *2130* *2131* *2132* *2133* *2134* *2135* *2136* *2137* *2138* *2139* *2140* *2141* *2142* *2143* *2144* *2145* *2146* *2147* *2148* *2149* *2150* *2151* *2152* *2153* *2154* *2155* *2156* *2157* *2158* *2159* *2160* *2161* *2162* *2163* *2164* *2165* *2166* *2167* *2168* *2169* *2170* *2171* *2172* *2173* *2174* *2175* *2176* *2177* *2178* *2179* *2180* *2181* *2182* *2183* *2184* *2185* *2186* *2187* *2188* *2189* *2190* *2191* *2192* *2193* *2194* *2195* *2196* *2197* *2198* *2199* *2200* *2201* *2202* *2203* *2204* *2205* *2206* *2207* *2208* *2209* *2210* *2211* *2212* *2213* *2214* *2215* *2216* *2217* *2218* *2219* *2220* *2221* *2222* *2223* *2224* *2225* *2226* *2227* *2228* *2229* *2230* *2231* *2232* *2233* *2234* *2235* *2236* *2237* *2238* *2239* *2240* *2241* *2242* *2243* *2244* *2245* *2246* *2247* *2248* *2249* *2250* *2251* *2252* *2253* *2254* *2255* *2256* *2257* *2258* *2259* *2260* *2261* *2262* *2263* *2264* *2265* *2266* *2267* *2268* *2269* *2270* *2271* *2272* *2273* *2274* *2275* *2276* *2277* *2278* *2279* *2280* *2281* *2282* *2283* *2284* *2285* *2286* *2287* *2288* *2289* *2290* *2291* *2292* *2293* *2294* *2295* *2296* *2297* *2298* *2299* *2300* *2301* *2302* *2303* *2304* *2305* *2306* *2307* *2308* *2309* *2310* *2311* *2312* *2313* *2314* *2315* *2316* *2317* *2318* *2319* *2320* *2321* *2322* *2323* *2324* *2325* *2326* *2327* *2328* *2329* *2330* *2331* *2332* *2333* *2334* *2335* *2336* *2337* *2338* *2339* *2340* *2341* *2342* *2343* *2344* *2345* *2346* *2347* *2348* *2349* *2350* *2351* *2352* *2353* *2354* *2355* *2356* *2357* *2358* *2359* *2360* *2361* *2362* *2363* *2364* *2365* *2366* *2367* *2368* *2369* *2370* *2371* *2372* *2373* *2374* *2375* *2376* *2377* *2378* *2379* *2380* *2381* *2382* *2383* *2384* *2385* *2386* *2387* *2388* *2389* *2390* *2391* *2392* *2393* *2394* *2395* *2396* *2397* *2398* *2399* *2400* *2401* *2402* *2403* *2404* *2405* *2406* *2407* *2408* *2409* *2410* *2411* *2412* *2413* *2414* *2415* *2416* *2417* *2418* *2419* *2420* *2421* *2422* *2423* *2424* *2425* *2426* *2427* *2428* *2429* *2430* *2431* *2432* *2433* *2434* *2435* *2436* *2437* *2438* *2439* *2440* *2441* *2442* *2443* *2444* *2445* *2446* *2447* *2448* *2449* *2450* *2451* *2452* *2453* *2454* *2455* *2456* *2457* *2458* *2459* *2460* *2461* *2462* *2463* *2464* *2465* *2466* *2467* *2468* *2469* *2470* *2471* *2472* *2473* *2474* *2475* *2476* *2477* *2478* *2479* *2480* *2481* *2482* *2483* *2484* *2485* *2486* *2487* *2488* *2489* *2490* *2491* *2492* *2493* *2494* *2495* *2496* *2497* *2498* *2499* *2500* *2501* *2502* *2503* *2504* *2505* *2506* *2507* *2508* *2509* *2510* *2511* *2512* *2513* *2514* *2515* *2516* *2517* *2518* *2519* *2520* *2521* *2522* *2523* *2524* *2525* *2526* *2527* *2528* *2529* *2530* *2531* *2532* *2533* *2534* *2535* *2536* *2537* *2538* *2539* *2540* *2541* *2542* *2543* *2544* *2545* *2546* *2547* *2548* *2549* *2550* *2551* *2552* *2553* *2554* *2555* *2556* *2557* *2558* *2559* *2560* *2561* *2562* *2563* *2564* *2565* *2566* *2567* *2568* *2569* *2570* *2571* *2572* *2573* *2574* *2575* *2576* *2577* *2578* *2579* *2580* *2581* *2582* *2583* *2584* *2585* *2586* *2587* *2588* *2589* *2590* *2591* *2592* *2593* *2594* *2595* *2596* *2597* *2598* *2599* *2600* *2601* *2602* *2603* *2604* *2605* *2606* *2607* *2608* *2609* *2610* *2611* *2612* *2613* *2614* *2615* *2616* *2617* *2618* *2619* *2620* *2621* *2622* *2623* *2624* *2625* *2626* *2627* *2628* *2629* *2630* *2631* *2632* *2633* *2634* *2635* *2636* *2637* *2638* *2639* *2640* *2641* *2642* *2643* *2644* *2645* *2646* *2647* *2648* *2649* *2650* *2651* *2652* *2653* *2654* *2655* *2656* *2657* *2658* *2659* *2660* *2661* *2662* *2663* *2664* *2665* *2666* *2667* *2668* *2669* *2670* *2671* *2672* *2673* *2674* *2675* *2676* *2677* *2678* *2679* *2680* *2681* *2682* *2683* *2684* *2685* *2686* *2687* *2688* *2689* *2690* *2691* *2692* *2693* *2694* *2695* *2696* *2697* *2698* *2699* *2700* *2701* *2702* *2703* *2704* *2705* *2706* *2707* *2708* *2709* *2710* *2711* *2712* *2713* *2714* *2715* *2716* *2717* *2718* *2719* *2720* *2721* *2722* *2723* *2724* *2725* *2726* *2727* *2728* *2729* *2730* *2731* *2732* *2733* *2734* *2735* *2736* *2737* *2738* *2739* *2740* *2741* *2742* *2743* *2744* *2745* *2746* *2747* *2748* *2749* *2750* *2751* *2752* *2753* *2754* *2755* *2756* *2757* *2758* *2759* *2760* *2761* *2762* *2763* *2764* *2765* *2766* *2767* *2768* *2769* *2770* *2771* *2772* *2773* *2774* *2775* *2776* *2777* *2778* *2779* *2780* *2781* *2782* *2783* *2784* *2785* *2786* *2787* *2788* *2789* *2790* *2791* *2792* *2793* *2794* *2795* *2796* *2797* *2798* *2799* *2800* *2801* *2802* *2803* *2804* *2805* *2806* *2807* *2808* *2809* *2810* *281*

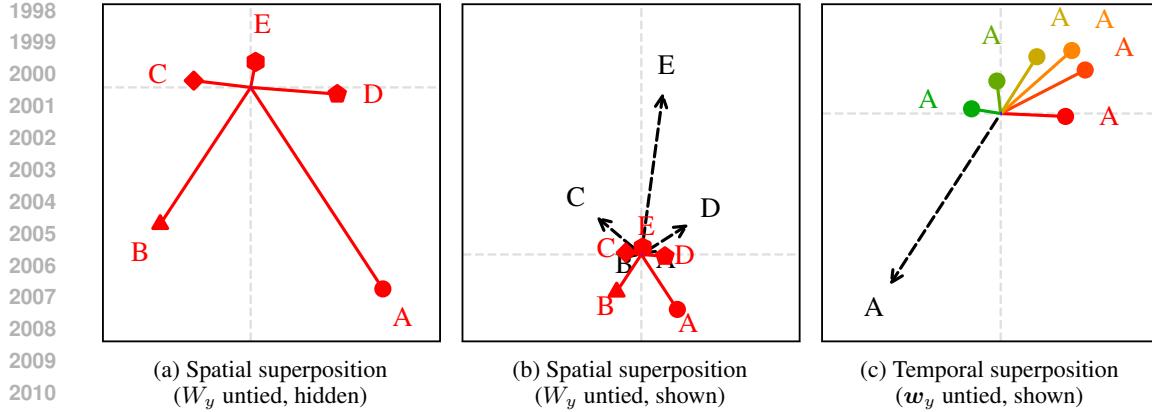


Figure 16: **Spatial and temporal superposition with untied readouts.** Readout vectors are shown as dashed black arrows. (a) If W_y is untied from $W_{s=k}$, spatial superposition of 5 features no longer forms a regular pentagon. This is also true of the results in Elhage et al. (2022). (b) Untied results are more visually cluttered as both the W_s and the W_y vectors need to be plotted, and these sets of vectors can overlap or be at different scales. (c) The difference between tied and untied results is most significant when spatial superposition is involved; for purely temporal superposition, the difference is minimal. Nevertheless, Appendix Figures 9 and 10 each contain many examples of the feature geometry of recurrent models with untied weights.

The specific shape and sparsity of data generated varies by experiment and is discussed below.

I.4 MODEL DEFINITION, INITIALIZATION AND TRAINING

As per Equation (1), there are two activation functions σ_h and σ_y that can be set as follows to achieve a linear RNN, SSM or nonlinear RNN:

- **Linear RNN:** $\sigma_h = \text{id}$, $\sigma_y = \text{id}$
- **SSM:** $\sigma_h = \text{id}$, $\sigma_y = \text{ReLU}$
- **Nonlinear RNN:** $\sigma_h = \text{ReLU}$, $\sigma_y = \text{ReLU}$

In all experiments, the model weights $W_x \in \mathbb{R}^{N_h \times N_x}$ and $W_h \in \mathbb{R}^{N_h \times N_h}$ were initialized using Xavier normal distributions. In cases where the readout weights $W_y \in \mathbb{R}^{N_h \times N_y}$ are *not* tied to $W_h^k W_x$, we initialize W_y using a Xavier normal distribution.

In practice, we often trained many models in parallel to make efficient use of GPU compute. Data was batched into 1000 batches; for each batch, we computed each model’s average training loss over an entire sequence (weighting the contribution of each feature to the loss by the feature’s importance) and backpropagated from this value. We maintained an exponential moving average (EMA) of this value for each model according to the update equation:

$$\ell_{\text{EMA}} \leftarrow \ell_{\text{EMA}} + 0.01(\ell - \ell_{\text{EMA}})$$

After training on all 1000 batches, the final value of ℓ_{EMA} was used to compare models; unless otherwise stated, the model that achieved the lowest EMA training loss was selected for plotting.

In all experiments, we used the AdamW optimizer with a constant learning rate of $\text{lr} = 5 \times 10^{-3}$.

Feature importance is attributed by weighting the loss according to the importance value per feature, as in Elhage et al. (2022). In other words $\mathcal{L} = \sum_{t=1}^T I(\mathbf{x}_{t-k}) \|\mathbf{x}_{t-k} - \hat{\mathbf{y}}_t\|^2$, where $I(\mathbf{x}_{t-k})$ is a scalar sum of the corresponding importance values of \mathbf{x}_{t-k} .

I.5 FIGURE 1

Panel (a): 100 nonlinear RNNs with 2-dimensional hidden states were trained on 10k non-zero sequences of length 10 timesteps each. The task was 0-delay, so $y_t = x_t$, making this equivalent to

2052 the first task used in Elhage et al. (2022). The input sequences had sparsity 0.99 (so $p = 0.01$) and
 2053 contained 5 features, {A, B, C, D, E} with importances $\{1, 0.97, 0.97^2, 0.97^3, 0.97^4\}$ respectively.
 2054

2055 **Panel (b):** Identical to panel (a), except trained on the $k = 5$ task instead of $k = 0$. To decrease
 2056 visual clutter, only the most important feature (A) was plotted. Note that due to the random nature
 2057 of data generation and training, it is not guaranteed that feature A will *always* be prioritized over
 2058 the other, less important features (e.g. B might instead be prioritized), but a single feature is almost
 2059 always prioritized over all others and A is the most common choice.
 2060

2061 I.6 FIGURE 2

2062 A linear RNN with 2-dimensional hidden state was trained on 50k non-zero scalar input sequences,
 2063 each of length 20 time steps and sparsity 0.9. The task was 3-delay. At each step of training, we
 2064 computed the values of $\mathbf{w}_y^\top \mathbf{w}_s$ for $0 \leq s < 12$ and used these to calculate the contribution of each
 2065 of the four terms in equation 5 to the loss, as plotted.
 2066

2067 I.7 FIGURE 3

2068 For each architecture (linear, SSM, nonlinear), 100 models with 2-dimensional hidden states were
 2069 trained on the 5-delay task using 10k non-zero sequences of length 25 timesteps each. We used a
 2070 sparsity level of 0.99, so $p = 0.01$.
 2071

2072 I.8 FIGURE 4

2073 Here we trained models on the $k = 7$ task. We swept through 200 uniformly spaced sparsity values
 2074 in the interval $[0.5, 1]$; at each sparsity level, we trained 1000 SSMs with 2-dimensional hidden
 2075 states on 10k non-zero scalar input sequences of length 25 timesteps each. **For each sparsity level,**
 2076 **we took the best 50 models (top 5%) in terms of lowest EMA training loss.** The mean value and
 2077 standard deviation of ρ and $k\theta$ across these 50 best models for each sparsity level was plotted, thus
 2078 indicating the “optimal” ρ and $k\theta$ at each value of sparsity.
 2079

2080 I.9 FIGURE 5

2081 This experiment is identical to that for Figure 1, except that $k = 1$, $k = 2$ and $k = 3$ were also
 2082 included. The heatmap was computed by taking a 2000×2000 grid of points within the axes and,
 2083 at each point, computing the sum of its non-negative projections onto the readout vectors (columns
 2084 of W_y). The heatmap therefore visualizes the region in activation space within which projection
 2085 interference is zero – the interference-free space.
 2086

2087 I.10 FIGURE 7 AND FIGURE 8

2088 We swept through 300 uniformly spaced values for $\tau \in [-2, 2]$ and 300 uniformly spaced values
 2089 for $\delta \in [-2, 2]$. For each pair (τ, δ) , we parameterized a linear RNN with 2-dimensional hidden
 2090 state as described in Appendix D.2. Each model was trained on 10k non-zero scalar input sequences
 2091 of length 20 timesteps each. Results are plotted for each combination of delay and sparsity level
 2092 $(k, 1 - p) \in \{2, 4\} \times \{0.7, 0.9, 0.999\}$.
 2093

2094 I.11 FIGURE 9

2095 For each combination of delay and sparsity level, $(k, 1 - p) \in \{2, 3, 5, 7\} \times \{0, 0.7, 0.9, 0.97, 0.99\}$,
 2096 100 models with 2-dimensional W_h , linear recurrence and ReLU readout were trained on 10k input
 2097 sequences, each of length 25 timesteps. Each plot shows the vectors of the best-performing model (as
 2098 measured by lowest EMA training loss). The values ρ and $k\theta$ shown under the plots are calculated
 2099 from the final W_h given its eigenvalues $\lambda_1, \lambda_2 \in \mathbb{C}$ as follows:
 2100

$$\rho := \max(|\lambda_1|, |\lambda_2|)$$

$$k\theta := k \arg(\lambda_1)$$

2106 I.12 FIGURE 10
2107

2108 The process was the same as for Figure 9, but with a nonlinear RNN trained instead of an SSM. Due
2109 to the nonlinear recurrence, ρ and $k\theta$ were not meaningful values to compute and so were omitted
2110 from the figure.

2111

2112

2113

2114

2115

2116

2117

2118

2119

2120

2121

2122

2123

2124

2125

2126

2127

2128

2129

2130

2131

2132

2133

2134

2135

2136

2137

2138

2139

2140

2141

2142

2143

2144

2145

2146

2147

2148

2149

2150

2151

2152

2153

2154

2155

2156

2157

2158

2159