RNNs of RNNs: Recursive Construction of Stable Assemblies of Recurrent Neural Networks

Anonymous Author(s)
Affiliation
Address
email

Abstract

Recurrent neural networks (RNNs) are widely used throughout neuroscience as models of local neural activity. Many properties of single RNNs are well characterized theoretically, but experimental neuroscience has moved in the direction of studying multiple interacting *areas*, and RNN theory needs to be likewise extended. We take a constructive approach towards this problem, leveraging tools from nonlinear control theory and machine learning to characterize when combinations of stable RNNs will themselves be stable. Importantly, we derive conditions which allow for massive feedback connections between interacting RNNs. We parameterize these conditions for easy optimization using gradient-based techniques, and show that stability-constrained 'network of networks' can perform well on challenging sequential-processing benchmark tasks. Altogether, our results provide a principled approach towards understanding distributed, modular function in the brain.

1 Introduction

The combination and reuse of primitive "modules" has enabled a great deal of progress in computer science, engineering, and biology. Modularity is particularly apparent in the structure of the brain, as different parts are specialized for different functions [Kandel et al., 2000]. Accordingly, most experimental studies throughout the history of neuroscience have focused on a single brain area in association with a single behavior [Abbott and Svoboda, 2020]. Similarly, RNN models of brain function have mostly been limited to a single RNN modeling a single area. However, neuroscience is entering an age where recording from many different brain areas simultaneously during complex behaviors is possible. As experimental neuroscience has shifted towards multi-area recordings, computational techniques for analyzing, modeling, and interpreting these multi-area recordings have blossomed [Abbott and Svoboda, 2020, Perich et al., 2021, Semedo et al., 2019, Yang and Molano-Mazón, 2021]. Despite this, RNN theory has lagged behind.

The theoretical question of RNN stability is crucial for understanding information propagation and manipulation [Vogt et al., 2020, Engelken et al., 2020, Kozachkov et al., 2022]. The conditions under which single, autonomous RNNs are chaotic or stable are well-studied, in particular when the RNN weights are randomly chosen and the number of neurons tends to infinity [Sompolinsky et al., 1988, Engelken et al., 2020]. However, there is very little work addressing the theoretical question of stability in 'networks of networks'. Two facts make this question challenging. Firstly, connecting two stable systems does not, in general, lead to a stable overall system. Secondly, there is a massive amount of feedback between brain areas, so one cannot reasonably assume near-decomposability [Simon, 1962, Abbott and Svoboda, 2020].

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

Given that the brain seems to dynamically reorganize and adapt interareal connectivity to meet task demands and environmental constraints [Miller and Cohen, 2001], this question of how stability is maintained is of the utmost importance. Here we take a bottom-up approach, more specifically asking "what stability properties of the individual modules lend themselves to rapid reorganization?"

We focus on a special type of stability, known as contractive stability [Lohmiller and Slotine, 1998]. Contraction analysis has found wide application in nonlinear control theory [Manchester and Slotine, 2017], synchronization [Pham and Slotine, 2007], and robotics [Chung and Slotine, 2009], but has only recently begun to find application in neuroscience and machine learning [Wensing and Slotine, 2020, Kozachkov et al., 2020, Revay and Manchester, 2020, Kozachkov et al., 2022, Centorrino et al., 2022, Burghi et al., 2022]. Loosely, a contracting system is a dynamical system whose asymptotic behavior is exponentially independent of its initial conditions. Contraction analysis is particularly useful for neuroscience because it is directly applicable to non-autonomous systems, and it allows for modular contraction-preserving combination properties to be derived (Figure 1) [Slotine and Lohmiller, 2001, Slotine, 2003]. See Section A1.2 for a mathematical primer on contraction analysis.

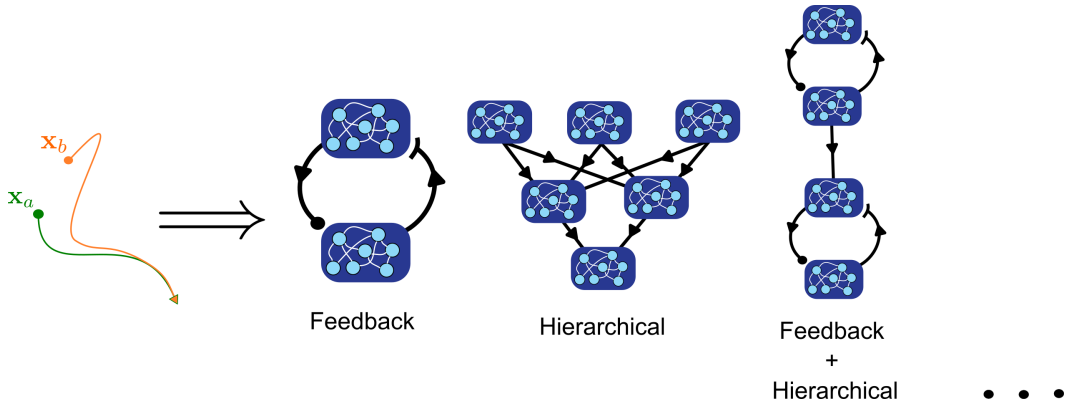

Figure 1: Contractive stability implies a modularity principle. Because contraction analysis tools allow complicated contracting systems to be built up recursively from simpler elements, this form of stability is well suited for understanding biological systems. Contracting combinations can be made between systems with very different dynamics, as long as those dynamics are contracting.

Our contributions are three-fold:

- A novel parameterization for feedback combination of contracting RNNs that enables direct optimization using standard deep learning libraries.
- Novel contraction conditions for continuous-time nonlinear RNNs, to use in conjunction with the combination condition. We also identify flaws in stability proofs from prior literature.
- Experiments demonstrating that our 'network of networks' sets a new state of the art for stability-constrained RNNs on benchmark sequential processing tasks.

2 Network of Networks Model

In this paper we analyze rate-based neural networks. Unlike spiking neural networks, these models are continuous and smooth. We consider the following RNN introduced by Wilson and Cowan [1972], which may be viewed as an approximation to a more biophysically-detailed spiking network:

$$\tau \dot{\mathbf{x}} = -\mathbf{x} + \mathbf{W}\phi(\mathbf{x}) + \mathbf{u}(t) \tag{1}$$

Here $\tau > 0$ is the time-constant of the network, typically between 10-100 milliseconds [Dayan and Abbott, 2005]. The vector $\mathbf{x} \in \mathbb{R}^n$ contains the voltages of all n neurons in the network. The voltages are converted into firing-rates through a static nonlinearity ϕ. We only consider monotonic activation functions with bounded slope: in other words, $0 \leq \phi' \leq g$. We do not restrain the sign of the nonlinearity. Common example nonlinearities $\phi(x)$ that satisfy these constraints are hyperbolic

tangent and softplus. The matrix $\mathbf{W} \in \mathbb{R}^{n \times n}$ contains the synaptic weights of the RNN. It is this matrix that ultimately determines the stability of (1), and will be a main target for our analysis. Finally, $\mathbf{u}(t)$ is the potentially time-varying external input into the network, capturing both explicit input into the RNN from the external world, as well as unmodeled dynamics from other brain areas.

Note that (1) is equivalent to another commonly used class of RNNs where the term $\mathbf{W}\mathbf{x} + \mathbf{u}$ appears inside the nonlinearity. See Section A1.1 or [Miller and Fumarola, 2012] for details. Our mathematical results apply equally well to both types of RNNs.

In order to extend (1) into a model of multiple interacting neural populations, we introduce the index i, which runs from 1 to p, where p is the total number of RNNs in the collection of RNNs. For now we will assume linear interactions between RNNs, because this is the simplest case. Later on we will discuss extensions to nonlinear couplings. This linearity assumptions leads to the following equation for the 'network of networks':

$$\tau \dot{\mathbf{x}}_i = -\mathbf{x}_i + \mathbf{W}_i \phi(\mathbf{x}_i) + \sum_{j=1}^{p} \mathbf{L}_{ij} \mathbf{x}_i + \mathbf{u}_i(t) \qquad \forall i = 1 \cdots p \tag{2}$$

where the matrix $\mathbf{L}_{ij}$ captures the interaction between RNN i and RNN j. Since we allow the RNNs to be of different sizes, $\mathbf{L}_{ij}$ is in general a rectangular matrix. With this notation in hand, we can now formalize the question posed in the abstract. Namely, "what types of connections between stable RNNs automatically preserve stability?" becomes "what restrictions on $\mathbf{L}_{ij}$ must be met in order to ensure overall stability of the network of networks?". We will now derive two combination 'primitives', negative feedback and hierarchical, which allow for recursive connection of contracting modules.

2.1 Generalized Negative Feedback Between RNNs Preserves Stability

We set aside for a moment the problem of determining when (1) is contracting. For now, assume that we have a collection of p contracting RNNs interacting through equation (2). Recall that contraction is defined with respect to a *metric*, a way of measuring distances between trajectories in state space. Thus, the ith RNN is contracting with respect to some metric $\mathbf{M}_i$. We assume for simplicity that this metric is constant. In the case where every RNN receives feedback from every other, we can preserve stability by ensuring these connections are negative feedback. In the simplest case where all RNN modules are contracting in the identity, the negative feedback may be written as:

$$\mathbf{L}_{ij} = -\mathbf{L}_{ji}^T$$

This is a well known result from the contraction analysis literature [Slotine, 2003]. Our first novel contribution is a generalization and parameterization of this negative feedback which allows for direct optimization using gradient-based techniques. In particular, if each $\mathbf{L}_{ij}$ is parameterized as follows:

$$\mathbf{L}_{ij} = \mathbf{B}_{ij} - \mathbf{M}_i^{-1} \mathbf{B}_{ji}^T \mathbf{M}_j \qquad \forall i, j \tag{3}$$

for arbitrary matrix $\mathbf{B}_{ij}$, then the overall network of networks retains the assumed stability properties of the RNN subnetworks. We provide a detailed proof in Section A3.1, but the basic idea relies on ensuring that $\mathbf{L}_{ij} = -\mathbf{L}_{ji}^T$ *in the appropriate metric*. This can be achieved via the constraint:

$$\mathbf{M}_i \mathbf{L}_{ij} = -\mathbf{L}_{ji}^T \mathbf{M}_j$$

Plugging (3) into the above expression verifies that it is indeed satisfied. Because contraction analysis relies on analyzing the symmetric part of Jacobian matrices, the skew-symmetry of $\mathbf{L}_{ij}$ 'cancels out' when computing the symmetric part, and leaves the stability of the subnetwork RNNs untouched.

Recursive Properties of Contracting Combinations The feedback combination (3), taken together with hierarchical combinations, may be used as combination primitives for recursively constructing complicated network of networks while automatically maintaining stability. The recursion comes from the fact that once a modular system is shown to be contracting it may be treated as a single contracting system, which may in turn be combined with other contracting systems, *ad infinitum*. Note that while the feedback (3) requires linear interareal connections, hierarchical interareal connections may be nonlinear [Lohmiller and Slotine, 1998].

3 Many Different Ways to Achieve Local Contraction

In the previous section, we put aside the question of achieving contraction in the subnetwork RNNs. In this section we return to it. Recall that we wish to find restrictions on $\mathbf{W}_i$ such that the ith subnetwork RNN, described by (1), is contracting. Here we derive five such novel conditions (see Section A3 for detailed proofs). As we will discuss, not all contraction conditions are equally useful - for example some conditions are easier to optimize or have higher model capacity than others.

In this section we also point out some flaws in existing stability proofs in the RNN literature, and suggest some pathways towards correcting them.

Theorem 1 (Absolute Value Restricted Weights)**.** *Let $|\mathbf{W}|$ denote the matrix formed by taking the element-wise absolute value of $\mathbf{W}$. If there exists a positive, diagonal $\mathbf{P}$ such that:*

$$\mathbf{P}(g|\mathbf{W}| - \mathbf{I}) + (g|\mathbf{W}| - \mathbf{I})^T\mathbf{P} \prec 0$$

then (1) *is contracting in metric $\mathbf{P}$. Moreover, if $W_{ii} \leq 0$, then $|W|_{ii}$ may be set to zero to reduce conservatism.*

It is easy to find matrices that satisfy Theorem 1, and given a matrix the condition is as easy to check as linear stability is. Moreover, the condition guarantees we can obtain a metric that the system is known to contract in (see Section 4.1 for details). It is less straightforward to enforce this condition during training, however we found that subnetworks constrained by Theorem 1 can achieve high performance in practice by simply fixing $\mathbf{W}$ and only optimizing the connections *between* subnetworks (Section 4.2). As there are fewer parameters to optimize, this training technique is faster.

Theorem 2 (Symmetric Weights)**.** *If $\mathbf{W} = \mathbf{W}^T$ and $g\mathbf{W} \prec \mathbf{I}$, then (1) is contracting.*

It has been known since the early 1990s that if (1) is autonomous (i.e the input $\mathbf{u}$ is constant) and has symmetric weights with eigenvalues less than $1/g$, then there exists a unique fixed point that the network converges to from any initial condition [Matsuoka, 1992]. Theorem 2 generalizes this statement to say that if (1) has symmetric weights with eigenvalues less than $1/g$, it is contracting. This includes previous results as a special case, because an *autonomous* contracting system has a unique fixed point that the network converges to from any initial condition.

Theorem 3 (Product of Diagonal and Orthogonal Weights)**.** *If there exists positive diagonal matrices $\mathbf{P}_1$ and $\mathbf{P}_2$, as well as $\mathbf{Q} = \mathbf{Q}^T \succ 0$ such that*

$$\mathbf{W} = -\mathbf{P}_1\mathbf{Q}\mathbf{P}_2$$

then (1) is contracting in metric $\mathbf{M} = (\mathbf{P}_1\mathbf{Q}\mathbf{P}_1)^{-1}$.

In contrast to the first two contraction conditions, Theorem 3 is very easy to optimize. To meet the constraint that the $\mathbf{P}$ matrices are positive, one can parameterize their diagonal elements as $P_{ii} = d_i^2 + \epsilon$, for some small positive constant ϵ, and optimize d_i directly. To meet the positive definiteness constraint on $\mathbf{Q}$, one may parameterize it as $\mathbf{Q} = \mathbf{E}^T\mathbf{E} + \epsilon\mathbf{I}$ and optimize $\mathbf{E}$ directly.

Theorem 4 (Triangular Weights)**.** *If $g\mathbf{W} - \mathbf{I}$ is triangular and Hurwitz, then (1) is contracting in a diagonal metric.*

Theorem 4 follows from the fact that a hierarchy of contracting systems is also contracting.

Theorem 5 (Singular Value Restricted Weights)**.** *If there exists a positive diagonal matrix $\mathbf{P}$ such that:*

$$g^2\mathbf{W}^T\mathbf{P}\mathbf{W} - \mathbf{P} \prec 0$$

then (1) is contracting in metric $\mathbf{P}$.

In the case of discrete-time RNNs, this contraction condition has been proved by many different authors in many different settings. When $\mathbf{P} = \mathbf{I}$, it is known as the echo-state condition for discrete-time RNNs [Jaeger, 2001]. This was then generalized to diagonal $\mathbf{P}$ in Buehner and Young [2006]. More recently, the original echo-state condition was rediscovered by Miller and Hardt [2018] in the machine learning literature. Following this rediscovery, the condition was generalized to $\mathbf{P} \neq \mathbf{I}$ in Revay and Manchester [2020]. Here we show that it applies to continuous-time RNNs as well.

What do the Jacobian Eigenvalues Tell Us? Several recent papers in machine learning, e.g [Haber and Ruthotto, 2017, Chang et al., 2019], claim that a sufficient condition for stability of the nonlinear system:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, t)$$

is that the associated Jacobian matrix $\mathbf{J}(\mathbf{x}, t) = \frac{\partial \mathbf{f}}{\partial \mathbf{x}}$ has eigenvalues whose real parts are strictly negative, i.e:

$$\max_i \mathrm{Re}(\lambda_i(\mathbf{J}(\mathbf{x}, t)) \leq -\alpha$$

with $\alpha > 0$. Unfortunately, this claim is generally false. For a counter-example, see Section 4.4.2 in [Slotine and Li, 1991].

However, in the *specific* case of the RNN (1), it appears that the eigenvalues of the symmetric part of $\mathbf{W}$ do provide information on global stability in a number of applications. For example, in [Matsuoka, 1992] it was shown that if $\mathbf{W}_s = \frac{1}{2}(\mathbf{W} + \mathbf{W}^T)$ has all its eigenvalues less than unity, and $\mathbf{u}$ is constant, then (1) has a unique, globally asymptotically stable fixed point. This condition also implies that the real parts of the eigenvalues of the Jacobian are uniformly negative. Moreover, in [Chang et al., 2019] it was shown that setting the symmetric part of $\mathbf{W}_s = \frac{1}{2}(\mathbf{W} + \mathbf{W}^T)$ almost equal to zero (yet slightly negative) led to rotational, yet stable dynamics in practice. This leads us to the following theorem, which shows that if the slopes of the activation functions change sufficiently slowly as a function of time, then the condition in [Matsuoka, 1992] in fact implies global contraction of (1).

Theorem 6. *Let $\mathbf{D}$ be a positive, diagonal matrix with $D_{ii} = \frac{d\phi_i}{dx_i}$, and let $\mathbf{P}$ be an arbitrary, positive diagonal matrix. If:*

$$(g\mathbf{W} - \mathbf{I})\mathbf{P} + \mathbf{P}(g\mathbf{W}^T - \mathbf{I}) \preceq -c\mathbf{P}$$

and

$$\dot{\mathbf{D}} - cg^{-1}\mathbf{D} \preceq -\beta\mathbf{D}$$

for $c, \beta > 0$, then (1) is contracting in metric $\mathbf{D}$ with rate β.

We stress however, that it is an open question whether or not diagonal stability of $\mathbf{W}$ implies that (1) is contracting. It has been conjectured that diagonal stability of $g\mathbf{W} - \mathbf{I}$ is a sufficient condition for global contraction of (1) [Revay et al., 2020], however this has been difficult to prove. To better characterize this conjecture, we present Theorem 7, which shows by way of counterexample that diagonal stability of $g\mathbf{W} - \mathbf{I}$ does not imply global contraction in a *constant* metric for (1).

Theorem 7. *Satisfaction of the condition $g\mathbf{W}_{sym} - \mathbf{I} \prec 0$ is **not** sufficient to show global contraction of the general nonlinear RNN (1) in any **constant** metric. High levels of antisymmetry in $\mathbf{W}$ can make it impossible to find such a metric, which we demonstrate via a 2×2 counterexample of the following form, with $c \geq 2$: $\mathbf{W} = \begin{bmatrix} 0 & -c \\ c & 0 \end{bmatrix}$*

4 Stability-Constrained Network of Networks Perform Well on Benchmarks

A natural concern is that stability of an RNN may come at the cost of its expressivity, which is particularly relevant for integrating information over long timescales. To investigate whether this might be an issue for our model, we trained a stability-constrained network-of-networks on three benchmark sequential image classification tasks: sequential MNIST, permuted seqMNIST, and sequential CIFAR10. These tasks are often used to measure information processing ability over long sequences [Le et al., 2015]. Images are presented pixel-by-pixel, and the network makes a prediction at the end of the sequence. In permuted seqMNIST, pixels are input in a fixed but random order.

All of our experiments were done on networks governed by (2). The nonlinear subnetwork RNNs were connected to each other via linear all-to-all negative feedback, given by (3). For all subnetworks we use the ReLU activation function. To enforce contraction of each individual subnetwork, we focused on two stability constraints from our theoretical results: Theorems 1 and 5. In the case of Thm 1, we did not train the individual subnetworks' weight matrices, but only trained the connections *between* subnetworks (Figure 2A). For Thm 5, we trained all parameters of the model (Figure 2B).

We refer to networks with subnetworks constrained by Theorem 1 as 'Sparse Combo Nets' and to networks with subnetworks constrained by Theorem 5 as 'SVD Combo Nets'. Throughout the

experimental results we use the notation '$p \times n$ network' - such a network consists of p distinct subnetwork RNNs, with each such subnetwork RNN containing n units.

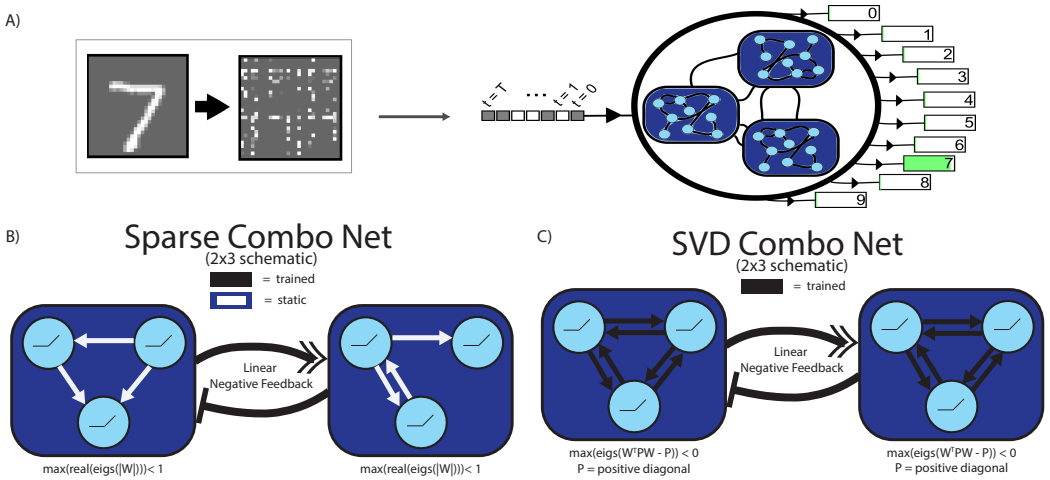

Figure 2: Summary of task structure and network architectures. Images from MNIST (or CIFAR10) were flattened into an array of pixels and fed sequentially into the modular 'network of networks', with classification based on the output at the last time-step. For MNIST, each image was also permuted in a fixed manner (A). The subnetwork 'modules' of our architecture were constrained to meet either Theorem 1 via sparse initialization (B) or Theorem 5 via direct parameterization (C). Linear negative feedback connections were trained between the subnetworks according to (3).

4.1 Network Initialization and Training

For the Sparse Combo Net we were not able to find a parameterization to use during training, but it is easy to randomly generate matrices with a particular likelihood of meeting the Theorem 1 condition, by selecting an appropriate sparsity level and limit on entry magnitude. Sparsity in particular is of interest due to its relevance in neurobiology and machine learning, so it is convenient that the condition makes it easy to verify stability of many different sparse RNNs. As $g = 1$ for ReLU activation, we check potential subnetwork matrices $\mathbf{W}$ by simply verifying linear stability of $|\mathbf{W}| - \mathbf{I}$.

Because every RNN meeting the condition has a corresponding well-defined stable LTI system contracting in the same metric, it is easy to find a metric to use in our training algorithm: solving for $\mathbf{M}$ in $-\mathbf{I} = \mathbf{M}\mathbf{A} + \mathbf{A}^T\mathbf{M}$ will produce a valid metric for any stable LTI system $\mathbf{A}$ [Slotine and Li, 1991]. We also utilize the fact that Hurwitz Metzler matrices are diagonally stable to improve efficiency of computing $\mathbf{M}$ (as well as in our proof of Theorem 1). See Section A4.1.1 for details.

We therefore randomly generated fixed subnetworks satisfying Theorem 1 and trained only the linear connections between them (Figure 3), as well as the linear input and output layers. More information on network initialization, hyperparameter tuning, and training algorithm is provided in Section A4.

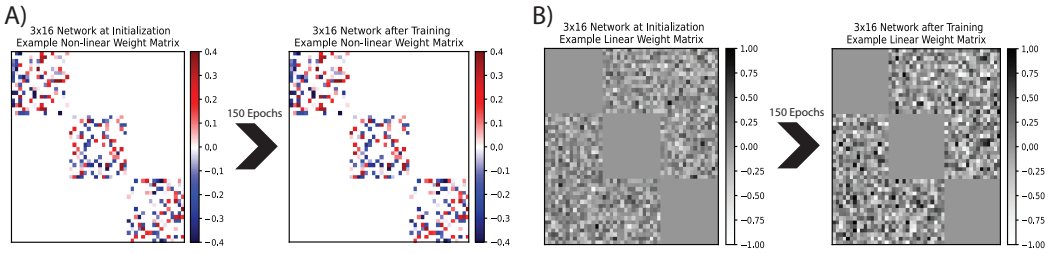

Figure 3: Example 3x16 Sparse Combo Net. Nonlinear intra-subnetwork weights are initialized using a set sparsity, and do not change in training (A). Linear inter-subnetwork connections are constrained to be antisymmetric with respect to the overall network metric, and are updated in training (B).

For the SVD Combo Net on the other hand, we ensured contraction of each subnetwork RNN by direct parameterization (described in Section A5), thus allowing all weights to be trained.

4.2 Results

The Sparse Combo Net architecture achieved the highest overall performance on both permuted seqMNIST and seqCIFAR10, with 96.94% and 65.72% best test accuracies respectively - thereby setting a new SOTA for stable RNNs (Table 1). Furthermore, we were able to reproduce SOTA scores over several repetitions, including 10 trials of seqCIFAR10. Along with repeatability of results, we also show that the contraction constraint on the connections between subnetworks ($\mathbf{L}$ in (3)) is important for performance, particularly in the Sparse Combo Net (Section 4.2.3).

Additionally, we profile how various architecture settings impact performance of our networks. In both networks, we found that increasing the total number of neurons improved task performance, but with diminishing returns (Section 4.2.1). We also found that the sparsity of the hidden-to-hidden weights in Sparse Combo Net had a large impact on the final network performance (Section 4.2.2).

Name	Stable RNN?	Params	sMNIST Repeats Mean (n) [Min]	psMNIST Repeats Mean (n) [Min]	sCIFAR10 Repeats Mean (n) [Min]	Seq MNIST Best	PerSeq MNIST Best	Seq CIFAR Best
LSTM [Chang et al., 2019]		68K	—	—	—	97.3%	92.7%	59.7%
Transformer [Trinh et al., 2018]		0.5M	—	—	—	98.9%	97.9%	62.2%
Antisymmetric [Chang et al., 2019]	?	36K	—	—	—	98%	95.8%	58.7%
Sparse Combo Net	✓	130K	—	**96.85%** (4) [**96.65%**]	64.72% (10) [63.73%]	99.04%	**96.94%**	**65.72%**
Lipschitz [Erichson et al., 2021]	✓	134K	99.2% (10) [99.0%]	95.9% (10) [95.6%]	—	**99.4%**	96.3%	64.2%
CKConv [Romero et al., 2021]		1M	—	—	—	99.32%	98.54%	63.74%
S4 [Gu et al., 2022]		7.9M	—	—	—	**99.63%**	**98.7%**	**91.13%**
Trellis [Bai et al., 2019]		8M	—	—	—	99.2%	98.13%	73.42%

Table 1: Published benchmarks for sequential MNIST, permuted MNIST, and sequential CIFAR10 best test accuracy. Architectures are grouped into three categories: baselines, best performing RNNs with claimed stability guarantee, and networks achieving overall SOTA. Within each grouping, networks are ordered by number of trainable parameters (for CIFAR10 if it differed across tasks). Our network is highlighted. Where possible, we include information on repeatability.

4.2.1 Experiments with Network Size

Understanding the effect of size on network performance is important to practical application of these architectures. For both Sparse Combo Net and SVD Combo Net, increasing the number of subnetworks while holding other settings constant (including fixing the size of each subnetwork at 32 units) was able to increase network test accuracy on permuted seqMNIST to a point (Figure 4).

The greatest performance jump happened when increasing from one module (37.1% Sparse Combo Net, 61.8% SVD Combo Net) to two modules (89.1% Sparse Combo Net, 92.9% SVD Combo Net).

After that the performance increased steadily with number of modules until saturating at $\sim 97\%$ for Sparse Combo Net and $\sim 95\%$ for SVD Combo Net.

As the internal subnetwork weights are not trained in Sparse Combo Net, it is unsurprising that its performance was substantially worse at the smallest sizes. However Sparse Combo Net surpasses SVD Combo Net by the 12×32 network size, which contains a modest 384 total units. Due to the better performance of the Sparse Combo Net, we focused additional analyses there. For more info on SVD Combo Net results, see Section A5.

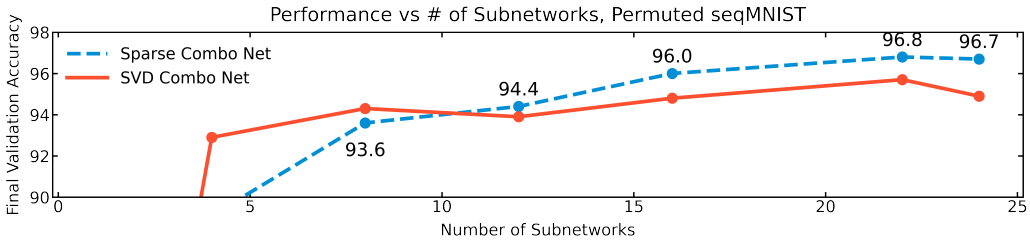

Figure 4: Permuted seqMNIST performance plotted against the number of subnetworks. Each subnetwork has 32 neurons. Results are shown for both Sparse Combo Net and SVD Combo Net.

4.2.2 Experiments with Sparsity

Because of the link between sparsity and stability as well as the biological relevance of sparsity, we explored in detail how subnetwork sparsity affects the performance of Sparse Combo Net. We ran a number of experiments on the permuted seqMNIST task, varying sparsity level while holding network size and other hyperparameters constant. Here we use "$n\%$ sparsity level" to refer to a network with subnetworks that have just $n\%$ of their weights non-zero.

We observed a large (> 5 percentage point) performance boost when switching from a 26.5% sparsity level to a 10% sparsity level in the 11×32 Sparse Combo Net (Figure 5), and subsequently decided to test significantly sparser subnetworks in a 16×32 Sparse Combo Net. We trained networks with sparsity levels of 5%, 3.3%, and 1%, as well as 10% for baseline comparison (Figure S2A). A 3.3% sparsity level produced the best results, leading to our SOTA performance on both permuted seqMNIST and seqCIFAR10. For a component RNN size of just 32 units, this sparsity level is quite small, with only one or two directional connections per neuron on average (Figure S7).

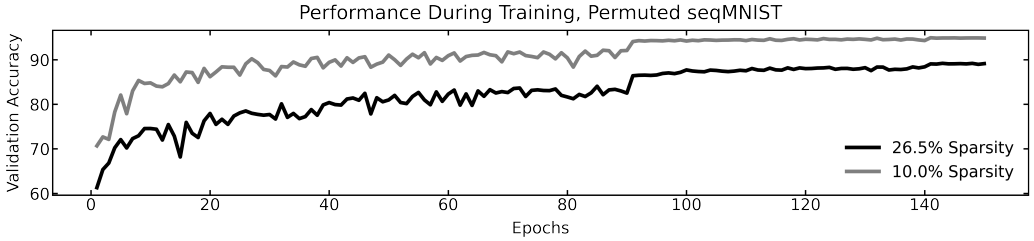

Figure 5: Permuted seqMNIST performance over the course of training for two 11×32 Sparse Combo Nets with different sparsity levels.

The use of sparsity in subnetworks to improve performance suggests another interesting direction that could enable better scalability of total network size - enforcing sparsity in the linear feedback weight matrix ($\mathbf{L}$). We performed pilot testing of this idea, which showed promise in mitigating the saturation effect seen in Figure 4. Those results are detailed in Section A4.2.4 (Table S1). We also discuss related future directions in Section A2.2.

4.2.3 Repeatability and Controls

Sparse Combo Net does not have the connections within its subnetworks trained, so network performance could be particularly susceptible to random initialization. Thus we ran repeatability studies on

permuted sequential MNIST and sequential CIFAR10 using our best network settings (16×32 with subnetwork sparsity level of 3.3%) and an extended training period. Mean performance over 4 trials of permuted seqMNIST was 96.85% with 0.019 variance, while mean performance over 10 trials of seqCIFAR10 was 64.72% with 0.406 variance. Note we also ran a number of additional experiments on size and sparsity settings, described in Section A4.2.

Across the permuted seqMNIST trials, best test accuracy always fell between 96.65% and 96.94%, a range much smaller than the differences seen with changing sparsity settings and network size. Three of the four trials showed best test accuracy $\geq 96.88\%$, despite some variability in early training performance (Figure S3). Similarly, eight of the ten seqCIFAR10 trials had test accuracy $> 64.3\%$, with all results falling between 63.73% and 65.72% (Figure S4). This robustly establishes a new SOTA for stable RNNs, comfortably beating the previously reported (single run) 64.2% test accuracy achieved by Lipschitz RNN [Erichson et al., 2021].

As a control study, we also tested how sensitive the Sparse Combo Net was to the stabilization condition on the interconnection matrix ($\mathbf{L}$ in (3)). To do so, we initialized the individual RNN modules in a 24×32 network as before, but set $\mathbf{L} = \mathbf{B}$ and did not constrain $\mathbf{B}$ at all during training, thus no longer ensuring contraction of the overall system. This resulted in 47.0% test accuracy on the permuted seqMNIST task, a stark decrease from the original 96.7% test accuracy - thereby demonstrating the utility of the contraction condition.

5 Discussion

The question of neural stability is one of the oldest questions in computational neuroscience. Indeed, cyberneticists were concerned with this question before the term 'computational neuroscience' existed [Wiener, 1948, Ashby, 1952]. Stability is a central component in several influential neuroscience theories [Hopfield, 1982, Seung, 1996, Murphy and Miller, 2009], perhaps the most well-known being that memories are stored as stable point attractors [Hopfield, 1982].

Engineers have long noted that stability and modularity are related [Khalil, 2002] - for example, passive dynamical systems can be connected in feedback combinations in ways that automatically preserve stability [Slotine and Li, 1991]. Curiously, this formal connection between stability and modularity was unexplored in the context of computational neuroscience. We begin to fill that gap by providing novel stability conditions for feedback systems consisting of many interacting RNNs.

While primarily motivated by neuroscience, our approach is also relevant for machine learning. Deep learning models are often as inscrutable as they are powerful. This opaqueness can be dangerous in safety-critical applications, such as autonomous driving or human-centered robotics, and it limits conceptual progress. In the case of recurrent models, one difficulty is that providing a certificate of *stability* is currently impossible or computationally impractical. Given that stability is a fundamental property of dynamical systems – and is intimately linked to concepts of control, generalization, efficiency, and robustness – the ability to guarantee stability of a recurrent model will be important for ensuring deep models behave as we expect them to.

Moreover, advanced applications of modern machine learning will likely involve combinations of trained networks, as are already used in systems such as DeepMind's AlphaGo. Recursively building such combinations in an effective and stable fashion while also allowing for continual refinement of the individual networks - as nature does for biological networks - will require new analysis tools. Here we have taken a concrete step towards the development of such tools, not only through our theoretical results, but also through their application to create stable combination network architectures that perform well in practice on benchmark tasks.

To summarize, we have theoretically characterized conditions under which 'multi-area RNNs' (i.e RNNs of RNNs) will be stable, given that each individual RNN is stable. We have also provided several novel stability conditions for individual RNNs which are compatible with these stability-preserving interareal connections. We parameterized our stability conditions for easy optimization using gradient-based techniques. We then leveraged this parameterization to successfully train recurrently connected 'network of networks' on three challenging sequential processing benchmarks, demonstrating that stability does not necessarily come at the cost of expressivity. All together, our results contribute progress towards understanding how the brain (i.e a 'RNN of RNNs') maintains stable and accurate function in the presence of massive interareal feedback, as well as external input.

References

Larry Abbott and Karel Svoboda. Brain-wide interactions between neural circuits. *Current Opinion in Neurobiology*, 65:iii–v, Dec 2020. ISSN 09594388. doi: 10.1016/j.conb.2020.12.012.

William Ashby. *Design for a brain: The origin of adaptive behaviour*. Springer Science & Business Media, 1952.

Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. Trellis networks for sequence modeling. *ICLR*, 2019.

Michael Buehner and Peter Young. A tighter bound for the echo state property. *IEEE Transactions on Neural Networks*, 17(3):820–824, 2006.

Thiago B Burghi, Timothy O'Leary, and Rodolphe Sepulchre. Distributed online estimation of biophysical neural networks. *arXiv preprint arXiv:2204.01472*, 2022.

Veronica Centorrino, Francesco Bullo, and Giovanni Russo. Contraction analysis of hopfield neural networks with hebbian learning. *arXiv preprint arXiv:2204.05382*, 2022.

Bo Chang, Minmin Chen, Eldad Haber, and Ed H Chi. Antisymmetricrnn: A dynamical system view on recurrent neural networks. *arXiv preprint arXiv:1902.09689*, 2019.

Soon-Jo Chung and Jean-Jacques E Slotine. Cooperative robot control and concurrent synchronization of lagrangian systems. *IEEE transactions on Robotics*, 25(3):686–700, 2009.

Yimian Dai, Stefan Oehmcke, Fabian Gieseke, Yiquan Wu, and Kobus Barnard. Attention as activation. *arXiv preprint arXiv:2007.07729v2*, 2020.

Peter Dayan and Laurence F Abbott. *Theoretical neuroscience: computational and mathematical modeling of neural systems*. 2005.

Rainer Engelken, Fred Wolf, and Larry F Abbott. Lyapunov spectra of chaotic recurrent neural networks. *arXiv preprint arXiv:2006.02427*, 2020.

N. Benjamin Erichson, Omri Azencot, Alejandro Queiruga, Liam Hodgkinson, and Michael W. Mahoney. Lipschitz recurrent neural networks. *ICLR*, 2021.

Albert Gu, Karan Goel, and Christopher Re. Efficiently modeling long sequences with structured state spaces. *ICLR*, 2022.

Eldad Haber and Lars Ruthotto. Stable architectures for deep neural networks. *Inverse Problems*, 34 (1):014004, 2017.

John J Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, 79(8):2554–2558, 1982.

Herbert Jaeger. The "echo state" approach to analysing and training recurrent neural networks-with an erratum note. 2001.

Eric R Kandel, James H Schwartz, Thomas M Jessell, Steven Siegelbaum, A James Hudspeth, and Sarah Mack. *Principles of neural science*, volume 4. McGraw-hill New York, 2000.

Hassan K Khalil. *Nonlinear systems*, volume 3. 2002.

Leo Kozachkov, Mikael Lundqvist, Jean-Jacques Slotine, and Earl K Miller. Achieving stable dynamics in neural circuits. *PLoS computational biology*, 16(8):e1007659, 2020.

Leo Kozachkov, John Tauber, Mikael Lundqvist, Scott L Brincat, Jean-Jacques Slotine, and Earl K Miller. Robust working memory through short-term synaptic plasticity. *bioRxiv*, 2022.

Dmitry Krotov and John Hopfield. Large associative memory problem in neurobiology and machine learning. *arXiv preprint arXiv:2008.06996*, 2020.

Quoc V Le, Navdeep Jaitly, and Geoffrey E Hinton. A simple way to initialize recurrent networks of rectified linear units. *arXiv preprint arXiv:1504.00941*, 2015.

Mario Lezcano-Casado and David Martınez-Rubio. Cheap orthogonal constraints in neural networks: A simple parametrization of the orthogonal and unitary group. In *International Conference on Machine Learning*, pages 3794–3803. PMLR, 2019.

Winfried Lohmiller and Jean-Jacques E Slotine. On contraction analysis for non-linear systems. *Automatica*, 34(6):683–696, 1998.

Ian R Manchester and Jean-Jacques E Slotine. Control contraction metrics: Convex and intrinsic criteria for nonlinear feedback design. *IEEE Transactions on Automatic Control*, 62(6):3046–3053, 2017.

Kiyotoshi Matsuoka. Stability conditions for nonlinear continuous neural networks with asymmetric connection weights. *Neural networks*, 5(3):495–500, 1992.

Earl K Miller and Jonathan D Cohen. An integrative theory of prefrontal cortex function. *Annual review of neuroscience*, 24(1):167–202, 2001.

John Miller and Moritz Hardt. Stable recurrent models. *arXiv preprint arXiv:1805.10369*, 2018.

Kenneth D Miller and Francesco Fumarola. Mathematical equivalence of two common forms of firing rate models of neural networks. *Neural computation*, 24(1):25–31, 2012.

Brendan K Murphy and Kenneth D Miller. Balanced amplification: a new mechanism of selective amplification of neural activity patterns. *Neuron*, 61(4):635–648, 2009.

Kumpati S. Narendra and Robert Shorten. Hurwitz stability of metzler matrices. *IEEE Transactions On Automatic Control*, 55(6):1484–1487, 2010.

Matthew G. Perich, Charlotte Arlt, Sofia Soares, Megan E. Young, Clayton P. Mosher, Juri Minxha, Eugene Carter, Ueli Rutishauser, Peter H. Rudebeck, Christopher D. Harvey, and Kanaka Rajan. Inferring brain-wide interactions using data-constrained recurrent neural network models. 2021. doi: 10.1101/2020.12.18.423348. URL https://www.biorxiv.org/content/early/2021/03/11/2020.12.18.423348.

Quang-Cuong Pham and Jean-Jacques Slotine. Stable concurrent synchronization in dynamic system networks. *Neural networks*, 20(1):62–77, 2007.

Adityanarayanan Radhakrishnana, Mikhail Belkin, and Caroline Uhler. Overparameterized neural networks implement associative memory. *PNAS*, 117:27162–27170, 2020.

Hubert Ramsauer, Bernhard Schäfl, Johannes Lehner, Philipp Seidl, Michael Widrich, Thomas Adler, Lukas Gruber, Markus Holzleitner, Milena Pavlović, Geir Kjetil Sandve, et al. Hopfield networks is all you need. *arXiv preprint arXiv:2008.02217*, 2020.

Max Revay and Ian Manchester. Contracting implicit recurrent neural networks: Stable models with improved trainability. In *Learning for Dynamics and Control*, pages 393–403. PMLR, 2020.

Max Revay, Ruigang Wang, and Ian R Manchester. Lipschitz bounded equilibrium networks. *arXiv e-prints*, pages arXiv–2010, 2020.

David W. Romero, Anna Kuzina, Erik J. Bekkers, Jakub M. Tomczak, and Mark Hoogendoorn. Ckconv: Continuous kernel convolution for sequential data. *arXiv preprint arXiv:2102.02611v1*, 2021.

João D Semedo, Amin Zandvakili, Christian K Machens, M Yu Byron, and Adam Kohn. Cortical areas interact through a communication subspace. *Neuron*, 102(1):249–259, 2019.

H Sebastian Seung. How the brain keeps the eyes still. *Proceedings of the National Academy of Sciences*, 93(23):13339–13344, 1996.

Herbert A. Simon. The architecture of complexity. *Proceedings of the American Philosophical Society*, 106(6):467–482, 1962. ISSN 0003-049X.

Jean-Jacques Slotine. Modular stability tools for distributed computation and control. *Int. J. Adaptive Control and Signal Processing*, 17(6), 2003.

Jean-Jacques Slotine and Winfried Lohmiller. Modularity, evolution, and the binding problem: a view from stability theory. *Neural Networks*, 14(2):137–145, 2001.

Jean-Jacques E Slotine and Weiping Li. *Applied Nonlinear Control*. Prentice-Hall, 1991.

Haim Sompolinsky, Andrea Crisanti, and Hans-Jurgen Sommers. Chaos in random neural networks. *Physical review letters*, 61(3):259, 1988.

Nicolas Tabareau and Jean-Jacques Slotine. Notes on contraction theory. *arXiv preprint nlin/0601011*, 2006.

Trieu H. Trinh, Andrew M. Dai, Minh-Thang Luong, and Quoc V. Le. Learning longer-term dependencies in rnns with auxiliary losses. *arXiv preprint arXiv:1803.00144v3*, 2018.

Ryan Vogt, Maximilian Puelma Touzel, Eli Shlizerman, and Guillaume Lajoie. On lyapunov exponents for rnns: Understanding information propagation using dynamical systems tools. *arXiv preprint arXiv:2006.14123*, 2020.

Eric W. Weisstein. Positive definite matrix. URL https://mathworld.wolfram.com/PositiveDefiniteMatrix.html.

Patrick M Wensing and Jean-Jacques Slotine. Beyond convexity—contraction and global convergence of gradient descent. *Plos one*, 15(8):e0236661, 2020.

Norbert Wiener. *Cybernetics or Control and Communication in the Animal and the Machine*. MIT press, 1948.

Hugh R. Wilson and Jack D. Cowan. Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical Journal*, 12(1):1–24, 1972. ISSN 0006-3495. doi: https://doi.org/10.1016/S0006-3495(72)86068-5.

Guangyu Robert Yang and Manuel Molano-Mazón. Towards the next generation of recurrent network models for cognitive neuroscience. *Current opinion in neurobiology*, 70:182–192, 2021.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See Supplementary Section A2.1
   (c) Did you discuss any potential negative societal impacts of your work? [No] There are no unique potential negative societal impacts of this work
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] See mathematical details within the paper, Section 2 and Section 3
   (b) Did you include complete proofs of all theoretical results? [Yes] See Supplementary Section A3
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] Our code is linked in Supplementary Section A4.1.2, and can be run directly on Google Colab. We used only extremely common benchmark image datasets in our experiments, which can be utilized directly via the provided Colab link.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] All initialization and training details are reported in Supplementary Section A4.1.1. The results for all hyperparameter tuning trials are included in Supplementary Section A4.4
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] Where applicable in the main text, we reported mean, min, max, and variance of test accuracy across our repeated trials (see Section 4.2.3 in particular). We also provide extended information on repeatability tests in Supplementary Section A4.2.3, and tables that report the results of each individual trial we ran in Supplementary Section A4.4
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See information on the code in Supplementary Section A4.1.2
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] We cite the article that introduced the sequential image classification tasks when introducing them in Section 4
   (b) Did you mention the license of the assets? [Yes] We include the licenses for MNIST and CIFAR10 datasets as part of Supplementary Section A4.1.2
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] Our code is linked in Supplementary Section A4.1.2
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

Appendix

A1 Extended Background **15**

A1.1 Two Different RNNs . . . 15

A1.2 Contraction Math . . . 15

A1.2.1 Feedback and Hierarchical Combinations . . . 15

A2 Extended Discussion **16**

A2.1 Limitations . . . 16

A2.2 Future Directions . . . 16

A3 Proofs for Main Results **17**

A3.1 Proof of Feedback Combination Property . . . 17

A3.2 Proof of Theorem 1 . . . 18

A3.3 Proof of Theorem 2 . . . 19

A3.4 Proof of Theorem 3 . . . 21

A3.5 Proof of Theorem 4 . . . 21

A3.6 Proof of Theorem 5 . . . 22

A3.7 Proof of Theorem 6 . . . 22

A3.8 Proof of Theorem 7 . . . 23

A4 Sparse Combo Net Details **26**

A4.1 Extended Methods . . . 26

A4.1.1 Initialization and Training . . . 26

A4.1.2 Code and Datasets . . . 27

A4.2 Extended Results . . . 28

A4.2.1 Network Size and Modularity Comparison . . . 28

A4.2.2 Sparsity Settings Comparison . . . 29

A4.2.3 Repeatability Tests . . . 30

A4.2.4 Scalability Pilot: Feedback Sparsity Tests . . . 32

A4.3 Architecture Interpretability . . . 33

A4.3.1 Number of Parameters . . . 33

A4.3.2 Inspecting Trained Network Weights . . . 33

A4.4 Tables of Results by Trial (Supplementary Documentation) . . . 34

A4.4.1 Permuted seqMNIST Trials . . . 35

A4.4.2 seqCIFAR10 Trials . . . 37

A5 SVD Combo Net Details **38**

A5.1 Parameterization Information . . . 38

A5.2 Control Experiment . . . 38

A5.3 Model Code . . . 39

A1 Extended Background

A1.1 Two Different RNNs

Note that in neuroscience, the variable $\mathbf{x}$ in equation (1) is typically thought of as a vector of neural membrane potentials. It was shown in [Miller and Fumarola, 2012] that the RNN (1) is equivalent via an affine transformation to another commonly used RNN model,

$$\tau\dot{\mathbf{y}} = -\mathbf{y} + \phi(\mathbf{W}\mathbf{y} + \mathbf{b}(t)) \tag{4}$$

where the variable $\mathbf{y}$ is interpreted as a vector of firing rates, rather than membrane potentials. The two models are related by the transformation $\mathbf{x} = \mathbf{W}\mathbf{y} + \mathbf{b}$, which yields

$$\tau\dot{\mathbf{x}} = \mathbf{W}(-\mathbf{y} + \phi(\mathbf{W}\mathbf{y} + \mathbf{b})) + \tau\dot{\mathbf{b}} = -\mathbf{x} + \mathbf{W}\phi(\mathbf{x}) + \mathbf{v}$$

where $\mathbf{v} \equiv \mathbf{b} + \tau\dot{\mathbf{b}}$. Thus $\mathbf{b}$ is a low-pass filtered version of $\mathbf{v}$ (or conversely, $\mathbf{v}$ may be viewed as a first order prediction of $\mathbf{b}$) and the contraction properties of the system are unaffected by the affine transformation. Note that the above equivalence holds even in the case where $\mathbf{W}$ is not invertible. In this case, the two models are proven to be equivalent, provided that $\mathbf{b}(0)$ and $\mathbf{y}(0)$ satisfy certain conditions–which are always possible to satisfy [Miller and Fumarola, 2012]. Therefore, any contraction condition derived for the x (or y) system automatically implies contraction of the other system. We exploit this freedom freely throughout the paper.

A1.2 Contraction Math

It can be shown that the non-autonomous system

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, t)$$

is contracting if there exists a metric $\mathbf{M}(\mathbf{x}, t) = \mathbf{\Theta}(\mathbf{x}, t)^T\mathbf{\Theta}(\mathbf{x}, t) \succ 0$ such that uniformly

$$\dot{\mathbf{M}} + \mathbf{M}\mathbf{J} + \mathbf{J}^T\mathbf{M} \preceq -\beta\mathbf{M}$$

where $\mathbf{J} = \frac{\partial \mathbf{f}}{\partial \mathbf{x}}$ and $\beta > 0$. For more details see the main reference [Lohmiller and Slotine, 1998]. Similarly, a non-autonomous discrete-time system

$$\mathbf{x}_{t+1} = \mathbf{f}(\mathbf{x}_t, t)$$

is contracting if

$$\mathbf{J}^T\mathbf{M}_{t+1}\mathbf{J} - \mathbf{M}_t \preceq -\beta\mathbf{M}_t$$

A1.2.1 Feedback and Hierarchical Combinations

Consider two systems, independently contracting in constant metrics $\mathbf{M}_1$ and $\mathbf{M}_2$, which are combined in feedback:

$$\begin{aligned} \dot{\mathbf{x}} &= \mathbf{f}(\mathbf{x}, t) + \mathbf{B}\mathbf{y} \\ \dot{\mathbf{y}} &= \mathbf{g}(\mathbf{y}, t) + \mathbf{G}\mathbf{x} \end{aligned} \tag{Feedback Combination}$$

If the following relationship between $\mathbf{B}, \mathbf{G}, \mathbf{M}_1$, and $\mathbf{M}_2$ is satisfied:

$$\mathbf{B} = -\mathbf{M}_1^{-1}\mathbf{G}^T\mathbf{M}_2$$

then the combined system is contracting as well. This may be seen as a special case of the feedback combination derived in [Tabareau and Slotine, 2006]. The situation is even simpler for hierarchical

combinations. Consider again two systems, independently contracting in some metrics, which are combined in hierarchy:

$$\begin{aligned} \dot{\mathbf{x}} &= \mathbf{f}(\mathbf{x}, t) \\ \dot{\mathbf{y}} &= \mathbf{g}(\mathbf{y}, t) + \mathbf{h}(\mathbf{x}, t) \end{aligned} \tag{Hierarchical Combination}$$

where $\mathbf{h}(\mathbf{x}, t)$ is a function with *bounded* Jacobian. Then this combined system is contracting in a diagonal metric, as shown in [Lohmiller and Slotine, 1998]. By recursion, this extends to hierarchies of arbitrary depth.

A2 Extended Discussion

Given the paucity of existing theory on modular networks, our novel stability conditions and proof of concept combination architectures are a significant step in an important new direction. The "network of networks" approach is evident in the biological brain, and has seen early practical success in applications such as AlphaGo. There is much evidence this line of questioning will be critical in the future, and our work is the first on stable *modular* networks.

Furthermore, we develop an architecture based on such combinations of "vanilla" RNNs that is both stable and achieves high performance on benchmark sequence classification tasks using few trainable parameters (small particularly for sequential CIFAR10). When considering just the facts about the network, it really has no business performing anywhere near as well as it does. Note also that without the stability condition in place, the network performance indeed drops substantially.

In order to facilitate the extension of this important line of thinking, we provide additional context on the limitations of our current approach as well as promising ideas for future directions in this section.

A2.1 Limitations

One drawback of our approach is that we parameterize each weight matrix in a special way to guarantee stability. In all cases, this requires us to parameterize matrices as the product of several other matrices. Thus, we gain stability at the cost of increasing the number of parameters, which can slow down training.

Another current drawback is that we only consider constant metrics. In theory, contraction metrics can be state-dependent as well as time-varying. Thus it is possible that we have overly restricted the space of models we consider. Similarly, negative feedback is not the only way to preserve contraction when combining two contracting systems. There are known small-gain theorems in the contraction analysis literature which accomplish the same task [Slotine, 2003]. However, parameterizing these conditions is less straightforward than parameterizing the negative-feedback condition.

A third limitation of the present work is that it does not give a recipe on how to incorporate anatomical knowledge into the building of 'RNNs of RNNs'. Our current approach is 'bottom up', in the sense that we describe complicated networks which can be built from simpler ones while ensuring stability at every level of construction. However, for building biological models of the brain it is important to incorporate known anatomical detail (i.e V4 projects to PFC, PFC projects back to V4, etc). How to do this in a way that preserves stability is an open and interesting question.

Lastly, we only tested our networks on sequential image classification benchmarks. Future work will include other benchmarks such as character or word-level language modeling. Additionally, while we conducted preliminary experiments exploring the role of scale (i.e number of subnetwork RNNs), we did not pursue this at the sizes reached by many modern deep learning applications. Thus it is currently unclear if the performance of these stability-constrained models will scale well enough with the number of subnetworks (or the number of neurons per subnetwork). Testing this correctly will require extensive experimentation with the various initialization and training settings.

A2.2 Future Directions

There are numerous future directions enabled by this work. For example, Theorem 6 suggests that a less restrictive contraction condition on W in terms of the eigenvalues of the symmetric part is

possible and desirable. Meanwhile, Theorem 7 provides important groundwork in finding such a condition, as it shows the need for a time-varying metric. Investigation of input-dependent metrics could be a fruitful next line of research, and would have far-reaching implications in disciplines such as curriculum learning.

Furthermore, the beneficial impact of sparsity on training these stable models suggests a potential avenue for additional experimental work – in particular adding a regularizing sparsity term during training. This could allow Sparse Combo Net to have its internal subnetwork weights trained without losing the stability guarantee, and conversely it could allow SVD Combo Net to reap some of the performance benefits of Sparse Combo Net without giving up the flexibility allowed by training said internal weights.

As described in the limitations above, a major experimental next step will be to test our architectures at greater scale and on more difficult tasks. Since 'network of network' approaches are becoming increasingly popular, our methodology is relevant to a variety of task types, including reinforcement learning applications. Given the biological inspiration, multi-modal learning tasks may also be of particular relevance.

A3 Proofs for Main Results

A3.1 Proof of Feedback Combination Property

Here we apply existing contraction analysis results to derive equation (3). Because (3) is a parameterization of a known contraction conditions [Slotine, 2003], we provide the following statement in the form of a corollary.

Corollary 1 (Network of Networks)**.** *Consider a collection of p subnetwork RNNs governed by* (1)*. Assume that these RNNs each have hidden-to-hidden weight matrices $\{\mathbf{W}_1, \dots, \mathbf{W}_p\}$ and are independently contracting in metrics $\{\mathbf{M}_1, \dots, \mathbf{M}_p\}$. Define the block matrices $\tilde{\mathbf{W}} \equiv$ BlockDiag$(\mathbf{W}_1, \dots, \mathbf{W}_p)$ and $\tilde{\mathbf{M}} \equiv$ BlockDiag$(\mathbf{M}_1, \dots, \mathbf{M}_p)$, as well as the overall state vector $\tilde{\mathbf{x}}^T \equiv (\mathbf{x}_1^T \cdots \mathbf{x}_2^T)$. Then the following 'network of networks' is globally contracting in metric $\tilde{\mathbf{M}}$:*

$$\begin{aligned} \tau \dot{\tilde{\mathbf{x}}} &= -\tilde{\mathbf{x}} + \tilde{\mathbf{W}}\phi(\tilde{\mathbf{x}}) + \mathbf{u}(t) + \mathbf{L}\tilde{\mathbf{x}} \\ \mathbf{L} &\equiv \mathbf{B} - \tilde{\mathbf{M}}^{-1}\mathbf{B}^T\tilde{\mathbf{M}} \end{aligned} \tag{5}$$

Where $\mathbf{B}$ is an arbitrary square matrix.

Proof. Consider the differential Lyapunov function:

$$V = \frac{1}{2}\delta\mathbf{x}^T\tilde{\mathbf{M}}\delta\mathbf{x}$$

The time-derivative of this function is:

$$\dot{V} = \delta\mathbf{x}^T\tilde{\mathbf{M}}\delta\dot{\mathbf{x}} = \delta\mathbf{x}^T\Big(\underbrace{\tilde{\mathbf{M}}\tilde{\mathbf{J}} + \tilde{\mathbf{J}}^T\tilde{\mathbf{M}}}_{\text{Jacobian of RNNs before interconnection}} + \underbrace{\tilde{\mathbf{M}}\mathbf{L} + \mathbf{L}^T\tilde{\mathbf{M}}}_{\text{Interconnection Jacobian}} \Big)\delta\mathbf{x}$$

Since we assume that the RNNs are contracting in isolation, the first term in this sum is less that the slowest contracting rate of the individual RNNs, which we call $\lambda > 0$. The second term in the sum is the zero matrix, by construction. Thus the time-derivative of V is upper-bounded by:

$$\dot{V} \leq -2\lambda V$$

This implies that the network of networks is contracting with rate λ. □

A3.2 Proof of Theorem 1

Our first theorem is motivated by the observation that if the y-system (described in Section A1.1) is to be interpreted as a vector of firing rates, it must stay positive for all time. For a linear, time-invariant system with positive states, diagonal stability is equivalent to stability. Therefore a natural question is if diagonal stability of a linearized y-system implies anything about stability of the nonlinear system. More formally, given an excitatory neural network (i.e $\forall ij, W_{ij} \geq 0$), if the linear system

$$\dot{\mathbf{x}} = -\mathbf{x} + g\mathbf{W}\mathbf{x}$$

is stable, then there exists a positive diagonal matrix P such that:

$$\mathbf{P}(g\mathbf{W} - \mathbf{I}) + (g\mathbf{W} - \mathbf{I})^T\mathbf{P} \prec 0$$

The following theorem shows that the nonlinear system (1) is indeed contracting in metric $\mathbf{P}$, and extends this result to a more general $\mathbf{W}$ by considering only the magnitudes of the weights.

Theorem 1. *Let $|\mathbf{W}|$ denote the matrix formed by taking the element-wise absolute value of $\mathbf{W}$. If there exists a positive, diagonal $\mathbf{P}$ such that:*

$$\mathbf{P}(g|\mathbf{W}| - \mathbf{I}) + (g|\mathbf{W}| - \mathbf{I})^T\mathbf{P} \prec 0$$

then (1) *is contracting in metric $\mathbf{P}$. Moreover, if $W_{ii} \leq 0$, then $|W|_{ii}$ may be set to zero to reduce conservatism.*

This condition is particularly straightforward in the common special case where the network does not have any self weights, with the leak term driving stability. While it can be applied to a more general $\mathbf{W}$, the condition will of course not be met if the network was relying on highly negative values on the diagonal of $\mathbf{W}$ for linear stability. As demonstrated by counterexample in the proof of Theorem 1, it can be impossible to use the same metric $\mathbf{P}$ for the nonlinear RNN in such cases.

Theorem 1 allows many weight matrices with low magnitudes or a generally sparse structure to be verified as contracting in the nonlinear system (1), by simply checking a linear stability condition (as linear stability is equivalent to diagonal stability for Metzler matrices too [Narendra and Shorten, 2010]).

Beyond verifying contraction, Theorem 1 actually provides a metric, with little need for additional computation. Not only is it of inherent interest that the same metric can be shared across systems in this case, it is also of use in machine learning applications, where stability certificates are becoming increasingly necessary. Critically, it is feasible to enforce the condition during training via L2 regularization on $\mathbf{W}$. More generally, there are a variety of systems of interest that meet this condition but do not meet the well-known maximum singular value condition, including those with a hierarchical structure.

Proof. Consider the differential, quadratic Lyapunov function:

$$V = \delta\mathbf{x}^T\mathbf{P}\delta\mathbf{x}$$

where $\mathbf{P} \succ 0$ is diagonal. The time derivative of V is:

$$\dot{V} = 2\delta\mathbf{x}^T\mathbf{P}\dot{\delta\mathbf{x}} = 2\delta\mathbf{x}^T\mathbf{P}\mathbf{J}\delta\mathbf{x} = -2\delta\mathbf{x}^T\mathbf{P}\delta\mathbf{x} + 2\delta\mathbf{x}^T\mathbf{P}\mathbf{W}\mathbf{D}\delta\mathbf{x}$$

where $\mathbf{D}$ is a diagonal matrix such that $\mathbf{D}_{ii} = \frac{d\phi_i}{dx} \geq 0$. We can upper bound the quadratic form on the right as follows:

$$\delta\mathbf{x}^T\mathbf{P}\mathbf{W}\mathbf{D}\delta\mathbf{x} = \sum_{ij} P_i W_{ij} D_j \delta x_i \delta x_j \leq$$

$$\sum_i P_i W_{ii} D_i |\delta x_i|^2 + \sum_{ij, i\neq j} P_i |W_{ij}| D_j |\delta x_i||\delta x_j| \leq g|\delta\mathbf{x}|^T\mathbf{P}|\mathbf{W}||\delta\mathbf{x}|$$

If $W_{ii} \leq 0$, the term $P_i W_{ii} D_i |\delta x_i|^2$ contributes non-positively to the overall sum, and can therefore be set to zero without disrupting the inequality. Now using the fact that $\mathbf{P}$ is positive and diagonal, and therefore $\delta\mathbf{x}^T \mathbf{P} \delta\mathbf{x} = |\delta\mathbf{x}|^T \mathbf{P} |\delta\mathbf{x}|$, we can upper bound $\dot{V}$ as:

$$\dot{V} \leq |\delta\mathbf{x}|^T(-2\mathbf{P} + \mathbf{P}|\mathbf{W}| + |\mathbf{W}|\mathbf{P})|\delta\mathbf{x}| = |\delta\mathbf{x}|^T[(\mathbf{P}(|\mathbf{W}| - \mathbf{I}) + (|\mathbf{W}|^T - \mathbf{I})\mathbf{P})]|\delta\mathbf{x}|$$

where $|W|_{ij} = |W_{ij}|$, and $|W|_{ii} = 0$ if $W_{ii} \leq 0$ and $|W|_{ii} = |W_{ii}|$ if $W_{ii} > 0$. This completes the proof.

Note that $\mathbf{W} - \mathbf{I}$ is Metzler, and therefore will be Hurwitz stable if and only if $\mathbf{P}$ exists [Narendra and Shorten, 2010].

It is also worth noting that highly negative diagonal values in $\mathbf{W}$ will prevent the same metric $\mathbf{P}$ from being used for the nonlinear system. Therefore the method used in this proof cannot feasibly be adapted to further relax the treatment of the diagonal part of $\mathbf{W}$.

The intuitive reason behind this is that in the symmetric part of the Jacobian, $\frac{\mathbf{PWD}+\mathbf{DW}^T\mathbf{P}}{2} - \mathbf{P}$, the diagonal self weights will also be scaled down by small $\mathbf{D}$, while the leak portion $-\mathbf{P}$ remains untouched by $\mathbf{D}$.

Now we actually demonstrate a counterexample, presenting a 2×2 symmetric Metzler matrix $\mathbf{W}$ that is contracting in the identity in the linear system, but cannot be contracting *in the identity* in the nonlinear system (1):

$$\mathbf{W} = \begin{bmatrix} -9 & 2.5 \\ 2.5 & 0 \end{bmatrix}$$

To see that it is not possible for the more general nonlinear system with these weights to be contracting in the identity, take $\mathbf{D} = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}$. Now

$$(\mathbf{WD})_{sym} - \mathbf{I} = \begin{bmatrix} -1 & 1.25 \\ 1.25 & -1 \end{bmatrix}$$

which has a positive eigenvalue of $\frac{1}{4}$.

□

A3.3 Proof of Theorem 2

While regularization may push networks towards satisfying Theorem 1, strictly enforcing the condition during optimization is not straightforward. This motivates the rest of our theorems, which derive contraction results for specially structured weight matrices. Unlike Theorem 1, these results have direct parameterizations which can easily be plugged into modern optimization libraries.

Theorem 2. *If* $\mathbf{W} = \mathbf{W}^T$ *and* $g\mathbf{W} \prec \mathbf{I}$*, then (1) is contracting.*

When $\mathbf{W}$ is symmetric, (1) may be seen as a continuous-time Hopfield network. Continuous-time Hopfield networks with symmetric weights were recently shown to be closely related to Transformer architectures [Krotov and Hopfield, 2020, Ramsauer et al., 2020]. Specifically, the dot-product attention rule may be seen as a discretization of the continuous-time Hopfield network with softmax activation function [Krotov and Hopfield, 2020]. Our results here provide a simple sufficient (and nearly necessary, see above remark) condition for global exponential stability of a given *trajectory* for the Hopfield network. In the case where the input into the network is constant, this trajectory is a fixed point. Moreover, each trajectory associated with a unique input is guaranteed to be unique. Finally, we note that our results are flexible with respect to activation functions so long as they satisfy the slope-restriction condition. This flexibility may be useful when, for example, considering recent work showing that standard activation functions may be advantageously replaced by attention mechanisms [Dai et al., 2020].

Proof. We begin by writing $\mathbf{W} = \mathbf{R} - \mathbf{P}$ for some unknown $\mathbf{R} = \mathbf{R}^T$ and $\mathbf{P} = \mathbf{P}^T \succ 0$. The approach of this proof is to show by construction that the condition $g\mathbf{W} \prec \mathbf{I}$ implies the existence of an $\mathbf{R}$ and $\mathbf{P}$ such that the system is contracting in metric $\mathbf{P}$. We consider the y version of the RNN, which as discussed above is equivalent to the x version via an affine transformation.

Consider the contraction condition:

$$-2\mathbf{M} + \mathbf{M}\mathbf{D}\mathbf{W} + \mathbf{W}^T\mathbf{D}\mathbf{M} \preceq -\beta\mathbf{M}$$

with $\beta > 0$. Substituting in the definitions of $\mathbf{W}$ and $\mathbf{M}$, this condition becomes:

$$-2\mathbf{P} + \mathbf{P}\mathbf{D}(\mathbf{R} - \mathbf{P}) + (\mathbf{R} - \mathbf{P})\mathbf{D}\mathbf{P} \preceq -\beta\mathbf{P}$$

Simplifying the terms and collecting them all on one side, the above may be written as:

$$(\beta - 2)\mathbf{P} + \mathbf{R}\mathbf{D}\mathbf{P} + \mathbf{P}\mathbf{D}\mathbf{R} - 2\mathbf{P}\mathbf{D}\mathbf{P} \preceq 0$$

via the Schur complement, the above term will be satisfied if:

$$(2 - \beta)\mathbf{P} - \mathbf{R}\mathbf{D}\mathbf{P}(2\mathbf{P}\mathbf{D}\mathbf{P})^{-1}\mathbf{P}\mathbf{D}\mathbf{R} =$$
$$(2 - \beta)\mathbf{P} - \frac{1}{2}(\mathbf{R}\mathbf{D}\mathbf{R}) \succeq (2 - \beta)\mathbf{P} - \frac{g}{2}(\mathbf{R}\mathbf{R}) \succeq 0$$

We continue by setting $\mathbf{P} = \gamma^2\mathbf{R}\mathbf{R}$ with $\gamma^2 = \frac{g}{2(2-\beta)}$, so that the above inequality is satisfied. At this point, we have shown that if $\mathbf{W}$ can be written as:

$$\mathbf{W} = \mathbf{R} - \gamma^2\mathbf{R}\mathbf{R}$$

then (1) is contracting in metric $\mathbf{M} = \gamma^2\mathbf{R}\mathbf{R}$. What remains to be shown is that if the condition:

$$g\mathbf{W} - \mathbf{I} \prec 0$$

Is satisfied, then this implies the existence of an $\mathbf{R}$ such that the above is true. To show that this is indeed the case, assume that:

$$\frac{1}{4\gamma^2}\mathbf{I} - \mathbf{W} \succeq 0$$

Substituting in the definition of γ, this is just the statement that:

$$\frac{2(2 - \beta)}{4g}\mathbf{I} - \mathbf{W} \succeq 0$$

Setting $\beta = 2\lambda > 0$, this yields:

$$(1 - \lambda)\mathbf{I} \succeq g\mathbf{W}$$

Since $\mathbf{W}$ is orthogonal, we have the eigendecomposition:

$$\frac{1}{4\gamma^2}\mathbf{I} - \mathbf{W} = \mathbf{V}(\frac{1}{4\gamma^2}\mathbf{I} - \mathbf{\Lambda})\mathbf{V}^T$$

where $\mathbf{V}^T\mathbf{V} = \mathbf{I}$ and $\mathbf{\Lambda}$ is a diagonal matrix containing the eigenvalues of $\mathbf{W}$. Denote the symmetric square-root of this expression as $\mathbf{S}$:

$$\mathbf{S} = \mathbf{V}\sqrt{(\frac{1}{4\gamma^2}\mathbf{I} - \mathbf{\Lambda})}\mathbf{V}^T = \mathbf{S}^T$$

Which implies that:

$$\frac{1}{4\gamma^2}\mathbf{I} - \mathbf{W} = \mathbf{S}^T\mathbf{S}$$

We now define $\mathbf{R}$ in terms of $\mathbf{S}$ as follows:

$$\mathbf{R} = \frac{1}{\gamma}\mathbf{S} + \frac{1}{2\gamma^2}\mathbf{I}$$

Which means that:

$$\frac{1}{4\gamma^2}\mathbf{I} - \mathbf{W} = (\gamma\mathbf{R} - \frac{1}{2\gamma}\mathbf{I})(\gamma\mathbf{R} - \frac{1}{2\gamma}\mathbf{I})$$

Expanding out the right side, we get:

$$\frac{1}{4\gamma^2}\mathbf{I} - \mathbf{W} = \gamma^2\mathbf{R}\mathbf{R} - \mathbf{R} + \frac{1}{4\gamma^2}\mathbf{I}$$

Subtracting $\frac{1}{4\gamma^2}\mathbf{I}$ from both sides yields:

$$\mathbf{W} = \mathbf{R} - \gamma^2\mathbf{R}\mathbf{R}$$

As desired.

□

A3.4 Proof of Theorem 3

Theorem 3. *If there exists positive diagonal matrices $\mathbf{P}_1$ and $\mathbf{P}_2$, as well as $\mathbf{Q} = \mathbf{Q}^T \succ 0$ such that*

$$\mathbf{W} = -\mathbf{P}_1\mathbf{Q}\mathbf{P}_2$$

then (1) is contracting in metric $\mathbf{M} = (\mathbf{P}_1\mathbf{Q}\mathbf{P}_1)^{-1}$.

Proof. Consider again a differential Lyapunov function:

$$V = \delta\mathbf{x}^T\mathbf{M}\delta\mathbf{x}$$

the time derivative is equal to:

$$\dot{V} = -2V + \delta\mathbf{x}^T\mathbf{M}\mathbf{W}\mathbf{D}\delta\mathbf{x}$$

Substituting in the definitions of $\mathbf{W}$ and $\mathbf{M}$, we get:

$$\dot{V} = -2V - \delta\mathbf{x}^T\mathbf{P}_1^{-1}\mathbf{P}_2\mathbf{D}\delta\mathbf{x} \leq -2V$$

Therefore V converges exponentially to zero.

□

A3.5 Proof of Theorem 4

Theorem 4. *If $g\mathbf{W} - \mathbf{I}$ is triangular and Hurwitz, then (1) is contracting in a diagonal metric.*

Note that in the case of a triangular weight matrix, the system (1) may be seen as a feedforward (i.e hierarchical) network. Therefore, this result follows from the combination properties of contracting systems. However, our proof provides a means of explicitly finding a metric for this system.

Proof. Without loss of generality, assume that $\mathbf{W}$ is lower triangular. This implies that $W_{ij} = 0$ if $i \leq j$. Now consider the generalized Jacobian:

$$\mathbf{F} = -\mathbf{I} + \mathbf{\Gamma}\mathbf{W}\mathbf{D}\mathbf{\Gamma}^{-1}$$

with $\mathbf{\Gamma}$ diagonal and $\Gamma_i = \epsilon^i$ where $\epsilon > 0$. Because $\mathbf{\Gamma}$ is diagonal, the generalized Jacobian is equal to:

$$\mathbf{F} = -\mathbf{I} + \mathbf{\Gamma}\mathbf{W}\mathbf{\Gamma}^{-1}\mathbf{D}$$

Now note that:

$$(\mathbf{\Gamma}\mathbf{W}\mathbf{\Gamma}^{-1})_{ij} = \epsilon^i W_{ij}\epsilon^{-j} = W_{ij}\epsilon^{i-j}$$

Where $i \leq j$, we have $W_{ij} = 0$ by assumption. Therefore, the only nonzero entries are where $i \geq j$. This means that by making ϵ arbitrarily small, we can make $\mathbf{\Gamma}\mathbf{W}\mathbf{\Gamma}^{-1}$ approach a diagonal matrix with W_{ii} along the diagonal. Therefore, if:

$$\max_i gW_{ii} - 1 < 0$$

the nonlinear system is contracting. Since $\mathbf{W}$ is triangular, W_{ii} are the eigenvalues of $\mathbf{W}$, meaning that this condition is equivalent to $g\mathbf{W} - \mathbf{I}$ being Hurwitz.

□

A3.6 Proof of Theorem 5

Theorem 5. *If there exists a positive diagonal matrix* $\mathbf{P}$ *such that:*

$$g^2\mathbf{W}^T\mathbf{P}\mathbf{W} - \mathbf{P} \prec 0$$

then (1) is contracting in metric $\mathbf{P}$.

Note that this is equivalent to the discrete-time diagonal stability condition developed in [Revay and Manchester, 2020], for a constant metric. Note also that when $\mathbf{M} = \mathbf{I}$, Theorem 5 is identical to checking the maximum singular value of $\mathbf{W}$, a previously established condition for stability of (1). However a much larger set of weight matrices are found via the condition when $\mathbf{M} = \mathbf{P}$ instead.

Proof. Consider the generalized Jacobian:

$$\mathbf{F} = \mathbf{P}^{1/2}\mathbf{J}\mathbf{P}^{-1/2} = -\mathbf{I} + \mathbf{P}^{1/2}\mathbf{W}\mathbf{P}^{-1/2}\mathbf{D}$$

where $\mathbf{D}$ is a diagonal matrix with $\mathbf{D}_{ii} = \frac{d\phi_i}{dx_i}$. Using the subadditivity of the matrix measure μ_2 of the generalized Jacobian we get:

$$\mu_2(\mathbf{F}) \leq -1 + \mu_2(\mathbf{P}^{1/2}\mathbf{W}\mathbf{P}^{-1/2}\mathbf{D})$$

Now using the fact that $\mu_2(\cdot) \leq ||\cdot||_2$ we have:

$$\mu_2(\mathbf{F}) \leq -1 + ||\mathbf{P}^{1/2}\mathbf{W}\mathbf{P}^{-1/2}\mathbf{D})||_2 \leq -1 + g||\mathbf{P}^{1/2}\mathbf{W}\mathbf{P}^{-1/2}||_2$$

Using the definition of the 2-norm, imposing the condition $\mu_2(\mathbf{F}) \leq 0$ may be written:

$$g^2\mathbf{W}^T\mathbf{P}\mathbf{W} - \mathbf{P} \prec 0$$

which completes the proof.

□

A3.7 Proof of Theorem 6

Theorem 6. *Let* $\mathbf{D}$ *be a positive, diagonal matrix with* $D_{ii} = \frac{d\phi_i}{dx_i}$*, and let* $\mathbf{P}$ *be an arbitrary, positive diagonal matrix. If:*

$$(g\mathbf{W} - \mathbf{I})\mathbf{P} + \mathbf{P}(g\mathbf{W}^T - \mathbf{I}) \preceq -c\mathbf{P}$$

and

$$\dot{\mathbf{D}} - cg^{-1}\mathbf{D} \preceq -\beta\mathbf{D}$$

for $c, \beta > 0$*, then (1) is contracting in metric* $\mathbf{D}$ *with rate* β.

Proof. Consider the differential, quadratic Lyapunov function:

$$V = \delta\mathbf{x}^T\mathbf{P}\mathbf{D}\delta\mathbf{x}$$

where $\mathbf{D} \succ 0$ is as defined above. The time derivative of V is:

$$\dot{V} = \delta\mathbf{x}^T\mathbf{P}\dot{\mathbf{D}}\delta\mathbf{x} + \delta\mathbf{x}^T(-2\mathbf{P}\mathbf{D} + \mathbf{P}\mathbf{D}\mathbf{W}\mathbf{D} + \mathbf{D}\mathbf{W}^T\mathbf{D}\mathbf{P})\delta\mathbf{x}$$

The second term on the right can be factored as:

$$\begin{aligned}
\delta\mathbf{x}^T(-2\mathbf{P}\mathbf{D} + \mathbf{P}\mathbf{D}\mathbf{W}\mathbf{D} + \mathbf{D}\mathbf{W}^T\mathbf{D}\mathbf{P})\delta\mathbf{x} &= \\
\delta\mathbf{x}^T\mathbf{D}(-2\mathbf{P}\mathbf{D}^{-1} + \mathbf{P}\mathbf{W} + \mathbf{W}^T\mathbf{P})\mathbf{D}\delta\mathbf{x} &\leq \\
\delta\mathbf{x}^T\mathbf{D}(-2\mathbf{P}g^{-1} + \mathbf{P}\mathbf{W} + \mathbf{W}^T\mathbf{P})\mathbf{D}\delta\mathbf{x} &= \\
\delta\mathbf{x}^T\mathbf{D}[\mathbf{P}(\mathbf{W} - g^{-1}\mathbf{I}) + (\mathbf{W}^T - g^{-1}\mathbf{I})\mathbf{P}]\mathbf{D}\delta\mathbf{x} &\leq \\
-cg^{-1}\delta\mathbf{x}^T\mathbf{P}\mathbf{D}^2\delta\mathbf{x}
\end{aligned}$$

where the last inequality was obtained by substituting in the first assumption above. Combining this with the expression for $\dot{V}$, we have:

$$\dot{V} \leq \delta\mathbf{x}^T\mathbf{P}\dot{\mathbf{D}}\delta\mathbf{x} - cg^{-1}\delta\mathbf{x}^T\mathbf{P}\mathbf{D}^2\delta\mathbf{x}$$

Substituting in the second assumption, we have:

$$\dot{V} \leq \delta\mathbf{x}^T\mathbf{P}(\dot{\mathbf{D}} - cg^{-1}\mathbf{D}^2)\delta\mathbf{x} \leq -\beta\delta\mathbf{x}^T\mathbf{P}\mathbf{D}\delta\mathbf{x} = -\beta V$$

and thus V converges exponentially to 0 with rate β. □

A3.8 Proof of Theorem 7

Theorem 7. *Satisfaction of the condition*

$$g\mathbf{W}_{sym} - \mathbf{I} \prec 0$$

*is **NOT** sufficient to show global contraction of the general nonlinear RNN (1) in any constant metric. High levels of antisymmetry in* $\mathbf{W}$ *can make it impossible to find such a metric, which we demonstrate via a* 2×2 *counterexample of the form*

$$\mathbf{W} = \begin{bmatrix} 0 & -c \\ c & 0 \end{bmatrix}$$

with $c \geq 2$.

Note that $g\mathbf{W}_{sym} - \mathbf{I} = g\frac{\mathbf{W}+\mathbf{W}^T}{2} - \mathbf{I} \prec 0$ is equivalent to the condition for contraction of the system with *linear* activation in the identity metric.

The main intuition behind this counterexample is that high levels of antisymmetry can prevent a constant metric from being found in the nonlinear system. This is because $\mathbf{D}$ is a diagonal matrix with values between 0 and 1, so the primary functionality it can have in the symmetric part of the Jacobian is to downweight the outputs of certain neurons selectively. In the extreme case of all 0 or 1 values, we can think of this as selecting a subnetwork of the original network, and taking each of the remaining neurons to be single unit systems receiving input from the subnetwork. For a given static configuration of $\mathbf{D}$ (think linear gains), this is a hierarchical system that will be stable if the subnetwork is stable. But as $\mathbf{D}$ can evolve over time when a nonlinearity is introduced, we would need to find a constant metric that can serve completely distinct hierarchical structures simultaneously - which is not always possible.

Put in terms of matrix algebra, D can zero out columns of $\mathbf{W}$, but not their corresponding rows. So for a given weight pair w_{ij}, w_{ji}, which has entry in $\mathbf{W}_{sym} = \frac{w_{ij}+w_{ji}}{2}$, if $D_i = 0$ and $D_j = 1$, the i, j entry in $(\mathbf{WD})_{sym}$ will be guaranteed to have lower magnitude if the signs of w_{ij} and w_{ji} are the same, but guaranteed to have higher magnitude if the signs are different. Thus if the linear system would be stable based on magnitudes alone $\mathbf{D}$ poses no real threat, but if the linear system requires antisymmetry to be stable, $\mathbf{D}$ can make proving contraction quite complicated (if possible at all).

Proof. The nonlinear system is globally contracting in a *constant* metric if there exists a symmetric, positive definite $\mathbf{M}$ such that the symmetric part of the Jacobian for the system, $(\mathbf{MWD})_{sym} - \mathbf{M}$ is negative definite uniformly. Therefore $(\mathbf{MWD})_{sym} - \mathbf{M} \prec 0$ must hold for all possible $\mathbf{D}$ if $\mathbf{M}$ is a constant metric the system *globally* contracts in with any allowed activation function, as some combination of settings to obtain a particular $\mathbf{D}$ can always be found.

Thus to prove the main claim, we present here a simple 2-neuron system that is contracting in the identity metric with linear activation function, but can be shown to have no $\mathbf{M}$ that simultaneously satisfies the $(\mathbf{MWD})_{sym} - \mathbf{M} \prec 0$ condition for two different possible $\mathbf{D}$ matrices.

To begin, take

$$\mathbf{W} = \begin{bmatrix} 0 & -2 \\ 2 & 0 \end{bmatrix}$$

Note that any off-diagonal magnitude ≥ 2 would work, as this is the point at which $\frac{1}{2}$ of one of the weights (found in $\mathbf{W}_{sym}$ when the other is zeroed) will have magnitude too large for $(\mathbf{WD})_{sym} - \mathbf{I}$ to be stable.

Looking at the linear system, we can see it is contracting in the identity because

$$\mathbf{W}_{sym} - \mathbf{I} = \begin{bmatrix} -1 & 0 \\ 0 & -1 \end{bmatrix} \prec 0$$

Now consider $(\mathbf{MWD})_{sym} - \mathbf{M}$ with $\mathbf{D}$ taking two possible values of

$$\mathbf{D}_1 = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \quad and \quad \mathbf{D}_2 = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}$$

We want to find some symmetric, positive definite $\mathbf{M} = \begin{bmatrix} a & m \\ m & b \end{bmatrix}$ such that $(\mathbf{MWD}_1)_{sym} - \mathbf{M}$ and $(\mathbf{MWD}_2)_{sym} - \mathbf{M}$ are both negative definite.

Working out the matrix multiplication, we get

$$(\mathbf{MWD}_1)_{sym} - \mathbf{M} = \begin{bmatrix} 2m - a & b - m \\ b - m & -b \end{bmatrix}$$

and

$$(\mathbf{MWD}_2)_{sym} - \mathbf{M} = \begin{bmatrix} -a & -(a+m) \\ -(a+m) & -2m - b \end{bmatrix}$$

We can now check necessary conditions for negative definiteness on these two matrices, as well as for positive definiteness on $\mathbf{M}$, to try to find an $\mathbf{M}$ that will satisfy all these conditions simultaneously. In this process we will reach a contradiction, showing that no such $\mathbf{M}$ can exist.

A necessary condition for positive definiteness in a real, symmetric $n \times n$ matrix $\mathbf{X}$ is $x_{ii} > 0$, and for negative definiteness $x_{ii} < 0$. Another well known necessary condition for definiteness of a real symmetric matrix is $|x_{ii} + x_{jj}| > |x_{ij} + x_{ji}| = 2|x_{ij}| \;\; \forall i \neq j$. See [Weisstein] for more info on these conditions.

Thus we will require a and b to be positive, and can identify the following conditions as necessary for our 3 matrices to all meet the requisite definiteness conditions:

$$2m < a \tag{6}$$

$$-2m < b \tag{7}$$

$$|2m - (a+b)| > 2|b - m| \tag{8}$$

$$|-2m - (a+b)| > 2|a + m| \tag{9}$$

Note that the necessary condition for $\mathbf{M}$ to be PD, $a + b > 2|m|$, is not listed, as it is automatically satisfied if (6) and (7) are.

It is easy to see that if $m = 0$, conditions (8) and (9) will result in the contradictory conditions $a > b$ and $b > a$ respectively, so we will require a metric with off-diagonal elements. To make the absolute values easier to deal with, we will check $m > 0$ and $m < 0$ cases independently.

First we take $m > 0$. By condition (6) we must have $a > 2m$, so between that and knowing the signs of all unknowns are positive, we can reduce many of the absolute values. Condition (8) becomes $a + b - 2m > |2b - 2m|$, and condition (9) becomes $a + b + 2m > 2a + 2m$, which is equivalent to

$b > a$. If $b > a$ we must also have $b > m$, so condition (8) further reduces to $a + b - 2m > 2b - 2m$, which is equivalent to $a > b$. Therefore we have again reached contradictory conditions.

A very similar approach can be applied when $m < 0$. Using condition (7) and the known signs we reduce condition (8) to $2|m| + a + b > 2b + 2|m|$, i.e. $a > b$. Meanwhile condition (9) works out to $a + b - 2|m| > 2a - 2|m|$, i.e. $b > a$.

Therefore it is impossible for a single constant $\mathbf{M}$ to accommodate both $\mathbf{D}_1$ and $\mathbf{D}_2$, so that no constant metric can exist for $\mathbf{W}$ to be contracting in when a nonlinearity is introduced that can possibly have derivative reaching both of these configurations. One real world example of such a nonlinearity is ReLU. Given a sufficiently high negative input to one of the units and a sufficiently high positive input to the other, $\mathbf{D}$ can reach one of these configurations. The targeted inputs could then flip at any time to reach the other configuration.

An additional condition we could impose on the activation function is to require it to be a strictly increasing function, so that the activation function derivative can never actually reach 0. We will now show that a very similar counterexample applies in this case, by taking

$$\mathbf{D}_{1*} = \begin{bmatrix} 1 & 0 \\ 0 & \epsilon \end{bmatrix} \quad and \quad \mathbf{D}_{2*} = \begin{bmatrix} \epsilon & 0 \\ 0 & 1 \end{bmatrix}$$

Note here that the $\mathbf{W}$ used above produced a $(\mathbf{WD})_{sym} - \mathbf{I}$ that just barely avoided being negative definite with the original $\mathbf{D}_1$ and $\mathbf{D}_2$, so we will have to increase the values on the off-diagonals a bit for this next example. In fact anything with magnitude larger than 2 will have some $\epsilon > 0$ that will cause a constant metric to be impossible, but for simplicity we will now take

$$\mathbf{W}_* = \begin{bmatrix} 0 & -4 \\ 4 & 0 \end{bmatrix}$$

Note that with $\mathbf{W}_*$, even just halving one of the off-diagonals while keeping the other intact will produce a $(\mathbf{WD})_{sym} - \mathbf{I}$ that is not negative definite. Anything less than halving however will keep the identity metric valid. Therefore, we expect that taking ϵ in $\mathbf{D}_{1*}$ and $\mathbf{D}_{2*}$ to be in the range $0.5 \geq \epsilon > 0$ will also cause issues when trying to obtain a constant metric.

We will now actually show via a similar proof to the above that $\mathbf{M}$ is impossible to find for $\mathbf{W}_*$ when $\epsilon \leq 0.5$. This result is compelling because it not only shows that ϵ does not need to be a particularly small value, but it also drives home the point about antisymmetry - the larger in magnitude the antisymmetric weights are, the larger the ϵ where we will begin to encounter problems.

Working out the matrix multiplication again, we now get

$$(\mathbf{MW}_*\mathbf{D}_{1*})_{sym} - \mathbf{M} = \begin{bmatrix} 4m - a & 2b - m - 2a\epsilon \\ b - m - 2a\epsilon & -4m\epsilon - b \end{bmatrix}$$

and

$$(\mathbf{MW}_*\mathbf{D}_{2*})_{sym} - \mathbf{M} = \begin{bmatrix} 4m\epsilon - a & -(2a + m - 2b\epsilon) \\ -(2a + m - 2b\epsilon) & -4m - b \end{bmatrix}$$

Resulting in two new main necessary conditions:

$$|4m - a - b - 4m\epsilon| > 2|2b - m - 2a\epsilon| \tag{10}$$

$$|4m\epsilon - a - b - 4m| > 2|2a + m - 2b\epsilon| \tag{11}$$

As well as new conditions on the diagonal elements:

$$4m - a < 0 \tag{12}$$

$$-4m - b < 0 \tag{13}$$

We will now proceed with trying to find a, b, m that can simultaneously meet all conditions, setting $\epsilon = 0.5$ for simplicity.

Looking at $m = 0$, we can see again that $\mathbf{M}$ will require off-diagonal elements, as condition (10) is now equivalent to the condition $a + b > |4b - 2a|$ and condition (11) is similarly now equivalent to $a + b > |4a - 2b|$.

Evaluating these conditions in more detail, if we assume $4b > 2a$ and $4a > 2b$, we can remove the absolute value and the conditions work out to the contradicting $3a > 3b$ and $3b > 3a$ respectively. As an aside, if $\epsilon > 0.5$, this would no longer be the case, whereas with $\epsilon < 0.5$, the conditions would be pushed even further in opposite directions.

If we instead assume $2a > 4b$, this means $4a > 2b$, so the latter condition would still lead to $b > a$, contradicting the original assumption of $2a > 4b$. $2b > 4a$ causes a contradiction analogously. Trying $4b = 2a$ will lead to the other condition becoming $b > 2a$, once again a contradiction. Thus a diagonal $\mathbf{M}$ is impossible

So now we again break down the conditions into $m > 0$ and $m < 0$ cases, first looking at $m > 0$. Using condition (12) and knowing all unknowns have positive sign, condition (10) reduces to $a + b - 2m > |4b - 2(a + m)|$ and condition (11) reduces to $a + b + 2m > |4a - 2(b - m)|$. This looks remarkably similar to the $m = 0$ case, except now condition (10) has $-2m$ added to both sides (inside the absolute value), and condition (11) has $2m$ added to both sides in the same manner. If $4b > 2(a + m)$ the $-2m$ term on each side will simply cancel, and similarly if $4a > 2(b - m)$ the $+2m$ terms will cancel, leaving us with the same contradictory conditions as before.

Therefore we check $2(a + m) > 4b$. This rearranges to $2a > 2(2b - m) > 2(b - m)$, so that from condition (11) we get $b > a$. Subbing condition (12) in to $2(a + m) > 4b$ gives $8b < 4a + 4m < 5a$ i.e. $b < \frac{5}{8}a$, a contradiction. The analogous issue arises if trying $2(b - m) > 4a$. Trying $2(a + m) = 4b$ gives $m = 2b - a$, which in condition (11) results in $5b - a > |6a - 6b|$, while in condition (12) leads to $5a > 8b$, so (11) can further reduce to $5b - a > 6a - 6b$ i.e. $11b > 7a$. But $b > \frac{7}{11}a$ and $b < \frac{5}{8}a$ is a contradiction. Thus there is no way for $m > 0$ to work.

Finally, trying $m < 0$, we now use condition (13) and the signs of the unknowns to reduce condition (10) to $a + b + 2|m| > |4b - 2(a - |m|)|$ and condition (11) to $a + b - 2|m| > |4a - 2(b + |m|)|$. These two conditions are clearly directly analogous to in the $m > 0$ case, where b now acts as a with condition (13) being $b > 4|m|$. Therefore the proof is complete. □

A4 Sparse Combo Net Details

Here we provide comprehensive information on the methodology and results for Sparse Combo Net, including some supplementary experimental results. See the Appendix table of contents for a guide to this section.

A4.1 Extended Methods

A4.1.1 Initialization and Training

As described in the main text, the nonlinear RNN weights for Sparse Combo Net were randomly generated based on given sparsity and entry magnitude settings, and then confirmed to meet the Theorem 1 condition (or discarded if not). For a sparsity level of x and a magnitude limit of y, each subnetwork $\mathbf{W}$ was generated by drawing uniformly from between $-y$ and y with $x\%$ density using scipy.sparse.random, and then zeroing out the diagonal entries. For various potential x and y settings, we quantified both the likelihood that a generated $\mathbf{W}$ would satisfy Theorem 1, and the resulting network performance. Of course this is also dependent on subnetwork size, as larger subnetworks enable greater sparsity. The information we have obtained so far is documented in Section A4.2, in particular A4.2.2.

In training the linear connections between the described nonlinear RNN subnetworks, we constrained the matrix $\mathbf{B}$ in (3) to reflect underlying modularity assumptions. In particular, we only train the off-diagonal blocks of $\mathbf{B}$ and mask the diagonal blocks. We do this to maintain the interpretation of

$\mathbf{L}$ as the matrix containing the connection weights *between* different modules, as diagonal blocks would correspond to self-connections. Furthermore, we only train the lower-triangular blocks of $\mathbf{B}$ while masking the others, to increase training speed.

To obtain the subnetwork RNN metrics necessary for training these linear connections, scipy.integrate.quad was used with default settings to solve for $\mathbf{M}$ in the equation $-\mathbf{I} = \mathbf{M}\mathbf{W} + \mathbf{W}^T\mathbf{M}$, as described in the main text. This was done by integrating $e^{\mathbf{W}^T t}\mathbf{Q}e^{\mathbf{W}t}dt$ from 0 to ∞. For efficiency reasons, and due to the guaranteed existence of a diagonal metric in the case of Theorem 1, integration was only performed to solve for the diagonal elements of $\mathbf{M}$. Therefore a check was added prior to training to confirm that the initialized network indeed satisfied Theorem 1 with metric $\mathbf{M}$. However, it was never triggered by our initialization method.

Initial training hyperparameter tuning was done primarily with 10×16 combination networks on the permuted seqMNIST task, starting with settings based on existing literature on this task, and verifying promising settings using a 15×16 network (Table S2). Initialization settings were held the same throughout, as was later done for the size comparison trials (described in Section A4.2.1).

Once hyperparameters were decided upon, the trials reported on in the main text began. Most of these experiments were also done on permuted seqMNIST, where we characterized performance of networks with different sizes and sparsity levels/entry magnitudes. When we moved to the sequential CIFAR10 task, we began by simply training with the same best settings that were found from these experiments. The results of all attempted trials are reported in Section A4.4.

Unless specified otherwise, all networks reported on in the main text were trained for 150 epochs, using an Adam optimizer with initial learning rate 1e-3 and weight decay 1e-5. The learning rate was cut to 1e-4 after 90 epochs and to 1e-5 after 140. After identifying the most promising settings, we ran repetitions trials on the best networks for 200 epochs with learning rate cuts after epochs 140 and 190.

A4.1.2 Code and Datasets

All Sparse Combo Nets described in the main text were trained using a single GPU on Google Colab. Code to replicate all experiments can be found here: https://colab.research.google.com/drive/1JCT5OMgaMVK_Xh8BDFNRrEsyFOOjvg10?usp=sharing

Runtime for the best performing architecture settings on the sequential CIFAR10 task was ~ 24 hours. A Colab Pro+ account was used to limit timeouts and prioritize GPU access.

The datasets we used for our tasks were MNIST and CIFAR10, downloaded via PyTorch. MNIST is a handwritten digits classification task, consisting of 60,000 training images and 10,000 testing images (each 28x28 and grayscale). It is made available under the terms of the Creative Commons Attribution-Share Alike 3.0 license. CIFAR10 is a dataset of 32x32 color images, split evenly among the following 10 classes: airplane, automobile, bird, cat, deer, dog, frog, horse, ship, and truck. It also contains 60,000/10,000 training/test images, and is distributed under the MIT License.

As mentioned in the main text, we presented these images to our networks pixel by pixel. In the case of MNIST, we used an additional modification to the dataset by permuting the pixels in a randomly determined (but fixed across the dataset) way. The use of these datasets is included in the above link to our code.

Extended results information begins on the next page.

A4.2 Extended Results

In this subsection, we describe additional results we did not get to in the main text, related to scalability, modularity, sparsity, and repeatability. We also add some further discussion of these results, along with more detailed information on the respective experimental set ups.

A4.2.1 Network Size and Modularity Comparison

For the Sparse Combo Net specifically we had additional experiments on architecture size and unit distribution besides what was depicted in the main text. The results of these supplemental experiments both replicated the observed effect in Figure 4 of network size instead using subnetworks with 16 units each as well as higher density (Figure S1A), and evaluated how task performance varies with modularity of a network fixed to have 352 total units (Figure S1B). In the modularity experiment we observed an inverse U shape, with poor performance of a 1×352 net and an 88×4 net, and best performance from a 44×8 net. Note that this experiment compared similar sparsity levels across the different subnetwork sizes. In practice we can achieve better performance with larger subnetworks by leveraging sparsity in a way not possible in smaller subnetworks. These additional experiments are now described in more detail below.

For the initial round of size comparison trials using subnetworks of 16 units each (Figure S1A), the nonlinear RNN weights were set by drawing uniformly from between -0.4 and 0.4 with 40% density using scipy.sparse.random, and then zeroing out the diagonal entries. These settings were chosen because they resulted in $\sim 1\%$ of 16 by 16 weight matrices meeting the Theorem 1 condition. During initialization only the matrices meeting this condition were kept, finishing when the desired number of component RNNs had been set - producing a block diagonal $\mathbf{W}$ like pictured in Figure 3A. This same initialization process was used throughout our experiments. In later experiments we vary the density and magnitude settings.

For the size experiments, we held static the number of units and initialization settings for each component RNN, and tested the effect of changing the number of components in the combination network. 1, 3, 5, 10, 15, 20, 22, 25, and 30 components were tested in this experiment (Figure S1A). Increasing the number of components initially lead to great improvements in test accuracy, but had diminishing returns - test accuracy consistently hit $\sim 93\%$ with a large enough number of subnetworks, but neither loss nor accuracy showed meaningful improvement past the 22×16 network. Interestingly, early training loss and accuracy became substantially worse once the number of components increased past a certain point, falling from 70% to 43% epoch 1 test accuracy between the 22×16 and 30×16 networks. The complete set of results can be found in Table S3.

Note that the size experiment described in the main text (Figure 4A) was a repetition of this original experiment, but now using 32 unit subnetworks with the best performing sparsity settings. The results for the repetition can be found in Table S5.

To better understand how the modularity of the combination networks affects performance, the next experiment held the number of total units constant at 352, selected due to the prior success of the 22×16 network, and tested different allocations of these units amongst component RNNs. Thus 1×352, 11×32, 44×8, and 88×4 networks were trained to compare against the 22×16 (Figure S1B). Increasing the modularity improved performance to a point, with the 44×8 network resulting in final test accuracy of 94.44%, while conversely the 11×32 resulted in decreased test accuracy. However, the 88×4 network was unable to learn, and a 352×1 network would theoretically just be a scaled linear anti-symmetric network.

Because larger networks require different sparsity settings to meet the Theorem 1 condition, these were not held constant between trials in the modularity comparison experiment (Figure S1B), but rather selected in the same way between trials - looking for settings that keep density and scalar balanced and result in $\sim 1\%$ of the matrices meeting the condition. The scalar was applied after sampling non-zero entries from a uniform distribution between -1 and 1. The resulting settings were 7.5% density and 0.077 scalar for 352 unit component RNN, 26.5% density and 0.27 scalar for 32 unit component RNN, 60% density and 0.7 scalar for 8 unit component RNN, and 100% density and 1.0 scalar for 4 unit component RNN. The complete set of results for the modularity experiment can be found in Table S4.

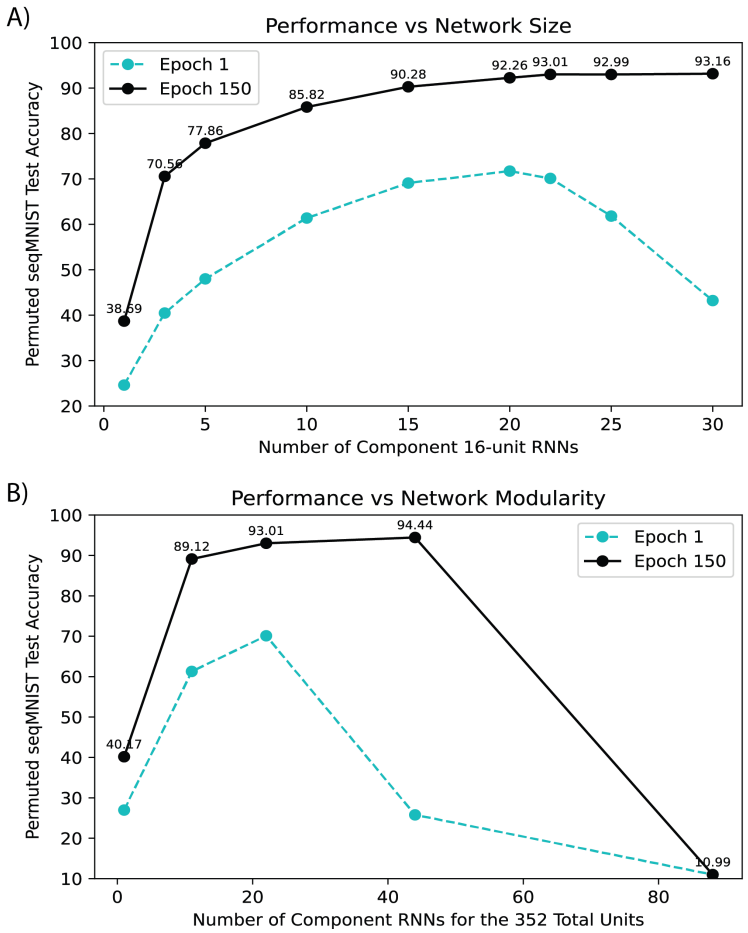

Figure S1: Performance of Sparse Combo Nets on the Permuted seqMNIST task by combination network size. We test the effects on final and first epoch test accuracy of both total network size and network modularity. The former is assessed by varying the number of subnetworks while each subnetwork is fixed at 16 units (A), and the latter by varying the distribution of units across different numbers of subnetworks with the total sum of units in the network fixed at 352 (B). Note that these experiments were run prior to optimizing the sparsity initialization settings. Experiments on total network size were later repeated with the final sparsity settings (Figure 4A). The results of both the size experiments are consistent.

A4.2.2 Sparsity Settings Comparison

Density and scalar settings for the component nonlinear RNNs were initially chosen for each network size using the percentage of random networks that met the Theorem 1 condition. For scalar s, a component network would have non-zero entries sampled uniformly between $-s$ and s.

When we began experimenting with sparsity in the initialization (after seeing the large performance difference in Figure 5), we split the previously described scalar setting into two different scalars - one applied before a random matrix was checked against the Theorem 1 condition, and one applied after a matrix was selected. Of course the latter must be ≤ 1 to guarantee stability is preserved. The scalar was separated out after we noticed that at 5% density, random 32 by 32 weight matrices met the condition roughly 1% of the time whether the scalar was 10 or 100000 - $\sim 85\%$ of sampled matrices using scalar 10 would continue to meet the condition even if multiplied by a factor of 10000. Therefore we wanted a mechanism that could bias selection towards matrices that are stable due to their sparsity and not due to magnitude constraints, while still keeping the elements to a reasonable size for training purposes.

Ultimately, both sparsity and magnitude had a clear effect on performance (Figure S2). Increases in both had a positive correlation with accuracy and loss through most parameters tested. Best test accuracy overall was 96.79%, which was obtained by both a 16×32 network with 5% density and entries between -5 and 5, and a 16×32 network with 3.3% density and entries between -6 and 6. The latter also achieved the best epoch 1 test accuracy observed of 86.79%. Thus we chose to go with these settings for our extended training repetitions on permuted seqMNIST.

It is also worth noting that upon investigation of the subnetwork weight matrices across these trials, the sparser networks had substantially lower maximum eigenvalue of $|\mathbf{W}|$, suggesting that stronger stability can actually correlate with *improved* performance on sequential tasks. This could be due to a mechanism such as that described in [Radhakrishnana et al., 2020].

Results from the initial trials of different sparsity levels across different network sizes can be found in Table S6. Results from the more thorough testing of different sparsity levels and magnitudes in a 16×32 network can be found in Table S7.

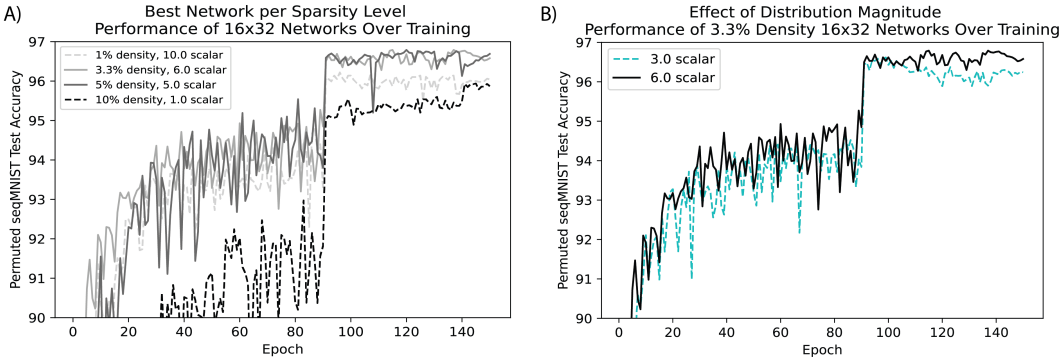

Figure S2: Permuted seqMNIST performance by component RNN initialization settings. Test accuracy is plotted over the course of training for four 16×32 networks with different density levels and entry magnitudes (A), highlighting the role of sparsity in network performance. Test accuracy is then plotted over the course of training for two 3.3% density 16×32 networks with different entry magnitudes (B), to demonstrate the role of the scalar. When the magnitude becomes too high however, performance is out of view of the current axis limits.

A4.2.3 Repeatability Tests

Here we give additional details on the repeatability tests described in the main text, as well as introducing some additional results using other hyperparameters (done after the initial SOTA-setting experiments) on sequential CIFAR10.

To further improve performance once network settings were explored on permuted seqMNIST, an extended training run was tested on the best performing option. Settings were kept the same as above using a 3.3% density 16×32 network, except training now ran for over 200 epochs, with just a single learning rate cut occurring after epoch 200 (exact number of epochs varied based on runtime limit). This experiment was repeated four times and resulted in 96.94% best test accuracy, as described in the main text (Figure S3). Table S8 reports additional details on these trials.

As mentioned, we also characterized Sparse Combo Net performance on the more challenging CIFAR10 sequential image classification task, across a greater number of trials. Ten trials were run for 200 epochs with learning rate scaled after epochs 140 and 190. All other settings were the same as for the permuted seqMNIST repeatability trials. As reported in the main text, the mean test accuracy observed was 64.72%, with variance 0.406, and range 63.73%-65.72%. Training loss and test accuracy over the course of learning are plotted for all ten trials in Figure S4, and exact numbers for each trial are provided in Table S10.

As we encountered difficulties with Colab GPU assignment when we started working on these repetitions, we also trained nine networks over a smaller number of epochs (Figure S5). These networks were trained using the 150 epoch paradigm previously described, although only four of the nine completed training within the 24 hour runtime limit. Complete results for these trials can be

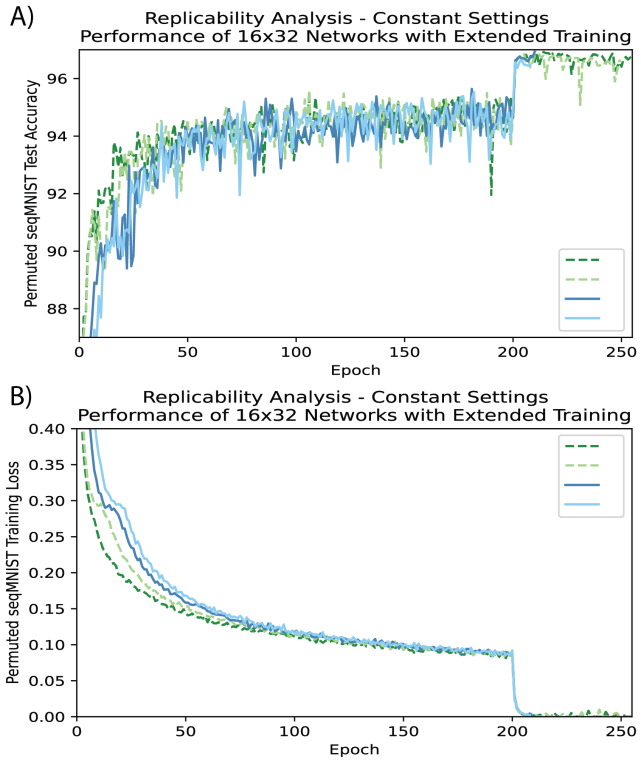

Figure S3: Permuted seqMNIST performance on repeated trials. Four different 16×32 networks with 3.3% density and entries between -6 and 6 were trained for 24 hours, with a single learning rate cut after epoch 200. (A) depicts test accuracy for each of the networks over the course of training. (B) depicts the training loss for the same networks. Exact numbers are reported in Table S8.

found in Table S11. Mean performance among the shorter training trials was 62.82% test accuracy with variance 0.95.

Prior to beginning the repeatability experiments on seqCIFAR10, we explored alternative hyperparameters on this task (Table S9). While we ultimately ended up using the same hyperparameters as for permuted seqMNIST, these results further support the robustness of our architecture.

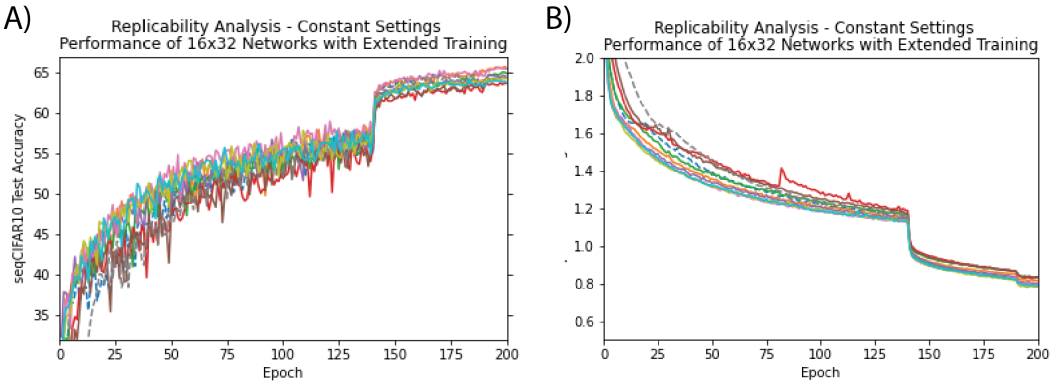

Figure S4: seqCIFAR10 performance on repeated trials. Ten different 16×32 networks with 3.3% density and entries between -6 and 6 were trained for 200 epochs, with learning rate divided by 10 after epochs 140 and 190. (A) depicts test accuracy for each of the networks over the course of training. (B) depicts the training loss for the same networks. Exact numbers are reported in Table S10.

To complete our benchmarking table, we also ran a single 150 epoch trial of our best network settings on the sequential MNIST task. Test accuracy over the course of training for this trial is depicted in Figure S6.

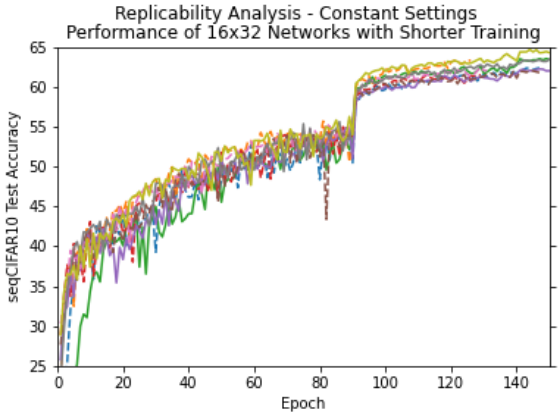

Figure S5: seqCIFAR10 performance on repeated trials with shorter training (done to complete more trials). Nine different 16×32 networks with 3.3% density and entries between -6 and 6 were set up to train for 150 epochs, with learning rate divided by 10 after epochs 90 and 140. Most of these networks hit runtime limit before completing, however they all got through at least 100 epochs and all had test accuracy exceed 61%. This figure depicts test accuracy for each of the networks over the course of training. Networks that completed training are plotted as solid lines, while those that were cut short are dashed.

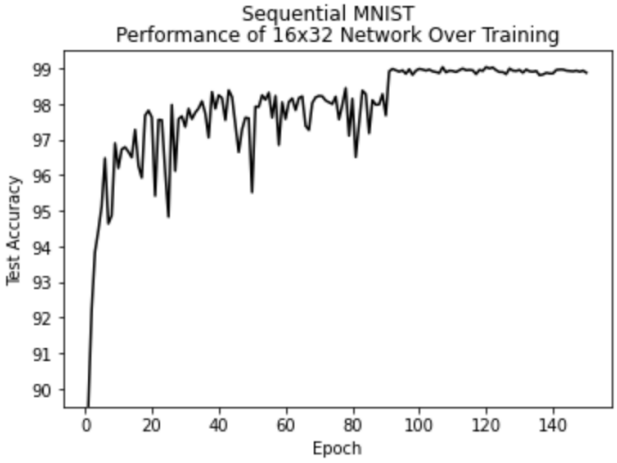

Figure S6: Performance over training on the seqMNIST task for a 16×32 network with best settings (using 150 epoch training protocol). Final test accuracy exceeded 99%.

A4.2.4 Scalability Pilot: Feedback Sparsity Tests

Not only did the 16×32 Sparse Combo Net with 3.3% density achieve test accuracy sequential CIFAR10 that is the highest to date for a provably stable RNN, it was higher than the 1 million parameter CKConv network, which set a recent SOTA for permuted seqMNIST accuracy (Table 1). Our network has $130,000$ trainable parameters by comparison. Thus we achieve very impressive results given the characteristics of our architecture.

However, the results of the network size experiments do raise some concern about the scalability of the approach. Here we provide pilot results suggesting that this issue can likely be addressed via the introduction of sparsity in the linear inter-subnetwork connectivity matrix $\mathbf{L}$. While all results in the main text use all-to-all negative feedback, we have early results suggesting that in larger networks, fewer negative feedback connections may perform better, thus preventing performance saturation with scaling.

Below are the results of testing of this idea in a 24×32 Sparse Combo Net on the sequential CIFAR10 task. We varied the number of feedback connections that were fixed at 0 while all other settings remained static (Table S1). The resulting performance took an inverse U shape, where the network with only 50% of possible feedback connections non-zero had the best test accuracy of 65.14%, achieved in just 124 epochs of training.

Size	Feedback Density	Epochs	Best Overall Test Acc.	Best Test Acc. Through 85 Epochs
24×32	100%	86	52.7%	52.7%
24×32	75%	88	56.49%	56.48%
24×32	66.6%	89	58.84%	58.84%
24×32	50%	124	65.14%	58.01%
24×32	33.3%	129	61.86%	56.05%
24×32	25%	92	54.26%	50.54%
24×32	0%	130	39.8%	38.38%
16×32	100%	150	64.63%	55.82%
16×32	75%	150	64.12%	57.23%
16×32	50%	127	59.87%	54.26%

Table S1: Results from pilot testing on the sparsity of negative feedback connections in a 24×32 Sparse Combo Net and a 16×32 Sparse Combo Net. Feedback Density refers to the percentage of possible subnetwork pairings that were trained in negative feedback, while the remaining inter-network connections were held at 0. All networks were trained with the same 150 epoch training paradigm as mentioned in the main text, but were stopped after hitting a 24 hour runtime limit. Decreasing Feedback Density is a promising path towards further improving performance as the size of Sparse Combo Nets is scaled. The ideal amount of feedback density will likely vary with the size of the combination network.

A4.3 Architecture Interpretability

In this subsection, we give additional information on trainable parameters and properties of the network weights, to assist in interpreting the Sparse Combo Net architecture and its training process.

A4.3.1 Number of Parameters

To report on the number of trainable parameters, we used the following formula:

$\frac{n^2 - M*C^2}{2} + i*n + n*o + n + o$

Where n is the total number of units in the $M \times C$ combination network, o is the total number of output nodes for the task, and i is the total number of input nodes for the task. Thus for the 16×32 networks highlighted here, we have 129034 trainable parameters for the MNIST tasks, and 130058 trainable parameters for sequential CIFAR10.

Note that the naive estimate for the number of trainable parameters would be $n^2 + i*n + n*o + n + o$, corresponding to the number of weights in $\mathbf{L}$, the number of weights in the feedforward linear input layer, the number of weights in the feedforward linear output layer, and the bias terms for the input and output layers, respectively. However, because of the combination property constraints on $\mathbf{L}$, only the lower triangular portion of a block off-diagonal matrix is actually trained, and $\mathbf{L}$ is then defined in terms of this matrix and the metric $\mathbf{M}$. Thus we subtract $M * C^2$ to remove the block diagonal portions corresponding to nonlinear RNN components, and then divide by 2 to obtain only the lower half.

A4.3.2 Inspecting Trained Network Weights

After training was completed, we inspected the state of all networks described in the main text, pulling both the nonlinear ($\mathbf{W}$) and linear ($\mathbf{L}$) weight matrices from both initialization time and the final model. For $\mathbf{W}$, we confirmed it did not change over training, and inspected the max real part of the eigenvalues of $|\mathbf{W}|$ in accordance with Theorem 1. The densest tested matrices tended to have $\lambda_{max}(|\mathbf{W}|) > 0.9$, while the sparsest ones tended to have $\lambda_{max}(|\mathbf{W}|) < 0.1$. For $\mathbf{L}$, we checked the maximum element and the maximum singular value before and after training. In general, both went up over the course of training, but by a modest amount.

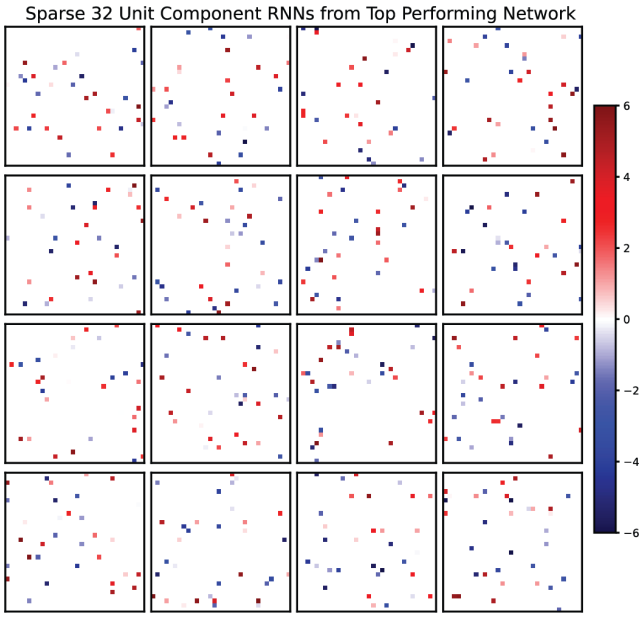

Figure S7: Weight matrices for each of the 32 unit non-linear component RNNs that were used in the best performing 16×32 network on permuted sequential MNIST.

A4.4 Tables of Results by Trial (Supplementary Documentation)

Table S2 shows all trials run on permuted sequential MNIST before beginning the more systematic experiments reported on in the main text. Notably, our networks did not require an extensive hyperparameter tuning process.

Tables S3 and S4 report additional details on the initial size and modularity experiments (Figure S1). Table S5 reports the results of the repeated experiment on Sparse Combo Net size, this time using 32 unit component subnetworks with best sparsity settings (Figure 4A).

Tables S6 and S7 report results from all trials related to our sparsity experiments (Figures 5 and S2).

Table S8 provides further information on the four trials in the permuted seqMNIST repeatability experiments (Figure S3).

Table S9 reports the results of all trials of different hyperparameters on the sequential CIFAR10 task, in chronological order. Ultimately the same settings as those used for permuted seqMNIST were chosen.

Table S10 shows the results of all ten trials in the seqCIFAR10 repeatability experiments (Figure S4). Table S11 shows results from nine additional seqCIFAR10 trials of shorter training duration (Figure S5), run to increase sample size while unable to access the higher quality Colab GPUs.

Finally, Table S1 reports the results of our pilot trial on introducing sparsity into the linear feedback connection matrix - as this table was presented in Section A4.2.4 however, we do not reproduce it here.

A4.4.1 Permuted seqMNIST Trials

Size	Epochs	Adam WD	Initial LR	LR Schedule	Final Test Acc.
10×16	150	5e-5	5e-3	0.1 after 91	84%
10×16	150	1e-5	1e-2	0.1 after 50,100	85%
15×16	150	2e-4	5e-3	0.1 after 50,100	84%
10×16	150	2e-4	1e-2	0.5 every 10	81%
10×16	200	2e-4	1e-2	0.5 after 10 then every 30	81%
10×16	171*	5e-5	1e-2	0.75 after 10,20,60,100 then every 15	84%
15×16	179*	1e-5	1e-3	0.1 after 100,150	90%

Table S2: Training hyperparameter tuning trials, presented in chronological order. * indicates that training was cut short by the 24 hour Colab runtime limit. LR Schedule describes the scalar the learning rate was multiplied by, and at what epochs. The best performing network is highlighted, and represents the training settings we used throughout most of the main text.

Size	Final Test Acc.	Epoch 1 Test Acc.	Final Train Loss
1×16	38.69%	24.61%	1.7005
3×16	70.56%	40.47%	0.9033
5×16	77.86%	47.99%	0.7104
10×16	85.82%	61.38%	0.4736
15×16	90.28%	69.09%	0.3156
20×16	92.26%	71.72%	0.2392
22×16	93.01%	70.11%	0.2073
25×16	92.99%	61.81%	0.2017
30×16	93.16%	43.21%	0.1991

Table S3: Results for combination networks containing different numbers of component 16-unit RNNs. Training hyperparameters and network initialization settings were kept the same across all trials, and all trials completed the full 150 epochs.

Size	Final Test Acc.	Epoch 1 Test Acc.	Final Train Loss
1×352	40.17%	26.97%	1.662
11×32	89.12%	61.29%	0.3781
22×16	93.01%	70.11%	0.2073
44×8	94.44%	25.78%	0.1500
88×4	10.99%	10.99%	2E+35

Table S4: Results for different distributions of 352 total units across a combination network. This number was chosen based on prior 22×16 network performance. For each component RNN size tested, the same procedure was used to select appropriate density and scalar settings. All networks otherwise used the same settings, as in the size experiments.

Name	Size	Final Test Acc.
Sparse Combo Net	1×32	37.1%
Sparse Combo Net	4×32	89.1%
Sparse Combo Net	8×32	93.6%
Sparse Combo Net	12×32	94.4%
Sparse Combo Net	16×32	96%
Sparse Combo Net	22×32	96.8%
Sparse Combo Net	24×32	96.7%
Control	24×32	47%

Table S5: Results for Sparse Combo Nets containing different numbers of component 32-unit RNNs with best found initialization settings, using the standard 150 epoch training paradigm. This experiment was run to demonstrate repeatability of the size results seen in Table S3. All trials were run to completion. A control trial was also run with the largest tested network size - the connections between subnetworks were no longer constrained, and thus this control combination network is not certifiably stable.

Size	Density	Scalar	Final Test Acc.	Epoch 1 Test Acc.	Final Train Loss
11×32	26.5%	0.27	89.12%	61.29%	0.3781
11×32	10%	1.0	94.86%	70.67%	0.1278
22×16	40%	0.4	93.01%	70.11%	0.2073
22×16	20%	1.0	95.27%	76.58%	0.0924
22×16	10%	1.0	94.26%	71.53%	0.1425
44×8	60%	0.7	94.44%	25.78%	0.1500
44×8	50%	1.0	95.05%	30.52%	0.1180

Table S6: Results for different initialization settings - varying sparsity and magnitude of the component RNNs for different network sizes. All other settings remained constant across trials, using our selected 150 epoch training paradigm.

Density	Pre-select Scalar	Post-select Scalar	Final Test Acc.	Epoch 1 Test Acc.	Final Train Loss
10%	1.0	1.0	95.87%	73.67%	0.074
5%	10.0	0.1	95.11%	73.10%	0.1311
5%	10.0	0.2	96.15%	82.50%	0.0051
5%	10.0	0.5	96.69%	75.76%	0.0001
5%	6.0	1.0	96.41%	21.55%	3.3E-5
5%	7.5	1.0	16.75%	11.39%	3068967
3.3%	30.0	0.1	96.24%	83.89%	0.0005
3.3%	30.0	0.2	96.54%	86.79%	4E-5
1%	10.0	1.0	96.04%	81.2%	0.0001

Table S7: Further optimizing the sparsity settings for high performance using a 16×32 network. The final scalar is the product of the pre-selection and post-selection scalars. Note that the 5% density and 7.5 scalar network was killed after 18 epochs due to exploding gradient. All other trials ran for a full 150 epochs.

Epochs	Best Test Acc.	Epoch 1 Test Acc.	Best Train Loss
208	96.65%	61.15%	0.00224
255	96.94%	84.19%	5e-5
250	96.88%	81.08%	5e-5
210	96.93%	67.19%	0.00069

Table S8: Repeatability of the best network settings on permuted seqMNIST. Four trials of 16×32 networks with 3.3% density and entries between -6 and 6, trained for a 24 hour period with a single learning rate cut (0.1 scalar) after epoch 200. All other training settings remained the same as our selected hyperparameters. Trials are presented in chronological order. The mean test accuracy achieved was 96.85% with Variance 0.019.

A4.4.2 seqCIFAR10 Trials

Density	Pre-select Scalar	Post-select Scalar	Epochs	Adam WD	Initial LR	LR Schedule	Best Test Acc.
3.3%	30	0.2	150	1e-5	1e-3	0.1 after 90,140	64.63%
3.3%	30	0.2	34*	1e-5	5e-3	0.1 after 90,140	35.42%
5%	6	1	150	1e-5	1e-3	0.1 after 90,140	60.9%
5%	10	0.5	150	1e-5	1e-4	0.1 after 90,140	54.86%
3.3%	30	0.2	150	1e-5	5e-4	0.1 after 90,140	61.83%
3.3%	30	0.2	200	1e-6	2e-3	0.1 after 140,190	62.31%
3.3%	30	0.2	186*	1e-5	1e-3	0.1 after 140,190	64.75%
3.3%	30	0.2	132*	1e-6	1e-3	0.1 after 140,190	62.31%
5%	10	0.5	195*	1e-5	1e-3	0.1 after 140,190	64.68%

Table S9: Additional hyperparameter tuning for the seqCIFAR10 task, presented in chronological order. * indicates that training was cut short by the 24 hour Colab runtime limit, or in the case of high learning rate was killed intentionally due to exploding gradient. LR Schedule describes the scalar the learning rate was multiplied by, and at what epochs. The best performing network is highlighted. Ultimately we decided on the same network settings and training hyperparameters for further testing, just extending the training period to 200 epochs with the learning rate cuts occurring after epochs 90 and 140.

Epochs	Best Test Acc.	Epoch 1 Test Acc.	Best Train Loss
186	64.75%	10.53%	0.83
200	64.04%	26.35%	0.787
200	64.32%	26.76%	0.778
170	64.88%	20.65%	0.857
200	65.72%	27.28%	0.813
200	63.73%	10.47%	0.826
200	65.03%	18.75%	0.83
200	64.71%	32.36%	0.799
200	64.4%	10.09%	0.83
200	65.63%	30.25%	0.792

Table S10: Repeatability of the best network settings on seqCIFAR10. Ten trials of 16×32 networks with 3.3% density and entries between -6 and 6, trained for 200 epochs with learning rate scaled by 0.1 after epochs 140 and 190. All other training settings remained the same as before. Trials are presented in chronological order. The mean test accuracy achieved was 64.72% with Variance 0.406. Most trials completed all 200 epochs, but two were cut short due to runtime limits.

Epochs	Best Test Acc.	Epoch 1 Test Acc.	Best Train Loss
150	64.63%	28.91%	0.837
150	63.51%	9.7%	0.9
148	62.35%	17.51%	0.89
126	63.33%	23.61%	0.903
129	61.45%	28.93%	0.937
150	62.21%	25.82%	0.9
150	63.42%	23.51%	0.86
147	62.05%	27.67%	0.882
131	62.41%	30.33%	0.912

Table S11: An additional nine trials investigating the repeatability of our results on the seqCIFAR10 task. For these trials we used the same 150 epoch training paradigm as previously, although only four of the networks were able to fully complete training. These trials were done to expand our sample size while access was limited to only slower GPUs. The mean observed test accuracy among the shorter trials was 62.82%, with variance of 0.95.

A5 SVD Combo Net Details

A5.1 Parameterization Information

For the SVD Combo Net, we ensured contraction by directly parameterizing each of the $\mathbf{W}_i$ $(i = 1, 2 \dots p)$ as:

$$\mathbf{W}_i = \mathbf{\Phi}_i^{-1} \mathbf{U}_i \mathbf{\Sigma}_i \mathbf{V}_i^T \mathbf{\Phi}_i \tag{14}$$

where $\mathbf{\Phi}_i$ is diagonal and nonsingular, $\mathbf{U}_i$ and $\mathbf{V}_i$ are orthogonal, and $\mathbf{\Sigma}_i$ is diagonal with $\Sigma_{ii} \in [0, g^{-1})$. We ensure orthogonality of $\mathbf{U}_i$ and $\mathbf{V}_i$ during training by exploiting the fact that the matrix exponential of a skew-symmetric matrix is orthogonal, as was done in [Lezcano-Casado and Martınez-Rubio, 2019]. The network constructed from these subnetworks using (2) is contracting in metric $\tilde{\mathbf{M}} = \text{BlockDiag}(\mathbf{\Phi}_1^2, \dots, \mathbf{\Phi}_p^2)$.

A5.2 Control Experiment

To do the SVD Combo Network stability control experiment, we still constrained the weights of the subnetwork modules using the above parameterization, but we removed any constraint on the connections between modules. Thus each individual module was still guaranteed to be contracting, but the overall system had no such guarantee in this control trial. The experiment was run using a 24×32 combination network and the permuted sequential MNIST task.

In contrast to the primary control experiment run on Sparse Combo Net, for the SVD Combo Net we saw only a very slight performance decrease when the between-module connections were no longer constrained for the overall stability certificate. Test accuracy in the control run was 94.56%, as compared to the 94.9% observed with the standard SVD Combo Net.

This disparity in performance decrease makes sense when considering that the hidden-to-hidden weights of SVD Combo Network are trainable, while those of the Sparse Combo Network are not.

Whatever instabilities are introduced by the lack of constraint on the inter-subnetwork connections likely cannot be adequately compensated for in the Sparse Combo Network.

A5.3 Model Code

```python
#define network
class rnnAssembly_SV(LightningModule):
    '''
    Pytorch module for training the following system:
        tau*dx/dt = -x + W*phi(x) + L*x+ u(t)
    where tau > 0, phi is a nonlinearity, W is block diagonal, L is some 'contracting' combination matrix and u is some input.

    '''

    def __init__(self, input_size, hidden_sizes, output_size,alpha,A):
        super(rnnAssembly_SV, self).__init__()
        self.input_size = input_size
        self.hidden_sizes = hidden_sizes
        self.hidden_size = int(np.sum(hidden_sizes))
        ns = hidden_sizes
        self.output_size = output_size
        self.alpha = alpha

        #input to hidden weights
        self.weight_ih = nn.Parameter(torch.normal(0,1/np.sqrt(self.hidden_size),(self.hidden_size, self.input_size)))
        #hidden to output weights
        self.weight_ho = nn.Parameter(torch.normal(0,1/np.sqrt(self.hidden_size),(self.output_size, self.hidden_size)))

        #mask to modularize overall network
        self.register_buffer("V_mask", create_mask_given_A(torch.eye(len(ns)),ns).bool())

        #parameterization of hidden-to-hidden networks
        self.U = nn.Parameter(torch.eye(self.hidden_size))
        self.V = nn.Parameter(torch.eye(self.hidden_size))
        self.Sigma = nn.Parameter(torch.ones((self.hidden_size,)))
        self.Phi = nn.Parameter(torch.normal(1,1/np.sqrt(self.hidden_size),(self.hidden_size,)))
        self.sing_val_eps = 1e-3
        self.register_buffer("L_mask", create_mask_given_A(A,ns).bool())
        self.register_buffer("S_offset",(self.sing_val_eps)*torch.ones(self.Sigma.shape[0]))

        #trainable inter-subnetwork weights
        self.L_train = nn.Parameter(self.L_mask*torch.normal(0,1/np.sqrt(np.mean(ns)),(np.sum(ns), np.sum(ns))))

        #biases
        self.bias_oh = nn.Parameter(torch.normal(0,1/np.sqrt(self.hidden_size),(1,self.output_size)))
        self.bias_hh = nn.Parameter(torch.normal(0,1/np.sqrt(self.hidden_size),(1,self.hidden_size)))

    def forward(self, input):

        #mask the left and right orth 'stem' matrices to produce diagonal-block structure
        U_masked = self.U*self.V_mask
        V_masked = self.V*self.V_mask

        #take skew-symmetric part of these matrices
        #exponentiate skew-symmetric part to produce left and right orthogonal matrices
        O1 = torch.matrix_exp(U_masked - U_masked.T)
        O2 = torch.matrix_exp(V_masked - V_masked.T)

        #get diagonal of singular value matrix, with elements between [0,1)
        S = torch.exp(-self.Sigma**2 -self.S_offset)

        #construct the weight matrix to be contracting in metric determined by Phi
        W = torch.diag_embed(self.Phi) @ (O1 @ torch.diag_embed(S) @ O2.T) @ torch.diag_embed(1/self.Phi)

        #get metric and metric inverse
        M = torch.diag_embed(1/(self.Phi**2))
        M_inv = torch.diag_embed(self.Phi**2)

        #get mask for network-to-network coupling
        L_masked = self.L_train*self.L_mask

        #initialize hidden state of the network
        inputs = input.unbind(1)
        state = torch.zeros((input.shape[0],self.hidden_size),device = self.device)
        state = state.type_as(state)

        #propagate input through the dynamics and store outputs
        for i in range(len(inputs)):
            fx = -state + F.relu(state @ W.T + inputs[i] @ (self.weight_ih.T) + self.bias_hh) + state @ (L_masked.T - (M @ L_masked) @ M_inv)
            state = state + self.alpha*fx
        hy = state @ (self.weight_ho.T)
        return hy, state
```

Figure S8: Pytorch Lightning code for SVD Combo Net cell.