Short Term Memory Capacity in Networks via the Restricted Isometry Property

Adam S. Charles, Han Lun Yap, Christopher J. Rozell

School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA

Abstract

Cortical networks are hypothesized to rely on transient network activity to support short term memory (STM). In this paper we study the capacity of randomly connected recurrent linear networks for performing STM when the input signals are approximately sparse in some basis. We leverage results from compressed sensing to provide rigorous non-asymptotic recovery guarantees, quantifying the impact of the input sparsity level, the input sparsity basis, and the network characteristics on the system capacity. Our analysis demonstrates that network memory capacities can scale superlinearly with the number of nodes, and in some situations can achieve STM capacities that are much larger than the network size. We provide perfect recovery guarantees for finite sequences and recovery bounds for infinite sequences. The latter analysis predicts that network STM systems may have an optimal recovery length that balances errors due to omission and recall mistakes. Furthermore, we show that the conditions yielding optimal STM capacity can be embodied in several network topologies, including networks with sparse or dense connectivities.

1 Introduction

————————-

Short term memory (STM) is critical for neural systems to understand non-trivial environments and perform complex tasks. While individual neurons could potentially account for very long or very short stimulus memory (e.g., through changing synaptic weights or membrane dynamics, respectively), useful STM on the order of seconds is conjectured to be due to transient *network* activity. Specifically, stimulus perturbations can cause activity in a recurrent network long after the input has been removed, and recent research hypothesizes that cortical networks may rely on transient activity to support STM (Jaeger & Haas, 2004; Maass et al., 2002; Buonomano & Maass, 2009).

Understanding the role of memory in neural systems requires determining the fundamental limits of STM capacity in a network and characterizing the effects on that capacity of the network size, topology, and input statistics. Various approaches to quantifying the STM capacity of linear (Jaeger, 2001; White et al., 2004; Ganguli et al., 2008; Hermans & Schrauwen, 2010) and nonlinear (Wallace et al., 2013) recurrent networks have been used, often assuming Gaussian input statistics (Jaeger, 2001; White et al., 2004; Hermans & Schrauwen, 2010; Wallace et al., 2013). These analyses show that even under optimal conditions, the STM capacity (i.e., the length of the stimulus able to be recovered) scales only linearly with the number of nodes in the network. While conventional wisdom holds that signal structure could be exploited to achieve more favorable capacities, this idea has generally not been the focus of significant rigorous study.

Recent work in computational neuroscience and signal processing has shown that many signals of interest have statistics that are strongly non-Gaussian, with low-dimensional structure that can be exploited for many tasks. In particular, sparsity-based signal models (i.e., representing a signal using relatively few non-zero coefficients in a basis) have recently been shown to be especially powerful. In the computational neuroscience lit-

erature, sparse encodings increase the capacity of associative memory models (Baum et al., 1988) and are sufficient neural coding models to account for several properties of neurons in primary visual cortex (i.e., response preferences (Olshausen & Field, 1996) and nonlinear modulations (Zhu & Rozell, 2013)). In the signal processing literature, the recent work in compressed sensing (CS) (Candes et al., 2006; Ganguli & Sompolinsky, 2012) has established strong guarantees on sparse signal recovery from highly undersampled measurement systems.

Ganguli & Sompolinsky (2010) have previously conjectured that the ideas of CS can be used to achieve STM capacities that exceed the number of network nodes in an orthogonal recurrent network when the inputs are sparse in the canonical basis (i.e. the input sequences have temporally localized activity). While these results are compelling and provide a great deal of intuition, the theoretical support for this approach remains an open question as the results in (Ganguli & Sompolinsky, 2010) use an asymptotic analysis on an approximation of the network dynamics to support empirical findings. In this paper we establish a theoretical basis for CS approaches in network STM by providing rigorous non-asymptotic recovery error bounds for an exact model of the network dynamics and input sequences that are sparse in any general basis (e.g., sinusoids, wavelets, etc.). Our analysis shows conclusively that the STM capacity can scale superlinearly with the number of network nodes, and quantifies the impact of the input sparsity level, the input sparsity basis, and the network characteristics on the system capacity. We provide both perfect recovery guarantees for finite inputs, as well as bounds on the recovery performance when the network has an arbitrarily long input sequence. The latter analysis predicts that network STM systems based on CS may have an optimal recovery length that balances errors due to omission and recall mistakes. Furthermore, we show that the structural conditions yielding optimal STM capacity in our analysis can be embodied in many different network topologies, including networks with both sparse and dense connectivities.

2 Background

2.1 Short Term Memory in Recurrent Networks

Since understanding the STM capacity of networked systems would lead to a better understanding of how such systems perform complex tasks, STM capacity has been studied in several network architectures, including discrete-time networks (Jaeger, 2001; White et al., 2004; Ganguli et al., 2008), continuous-time networks (Hermans & Schrauwen, 2010; Büsing et al., 2010), and spiking networks (Maass et al., 2002; Mayor & Gerstner, 2005; Legenstein & Maass, 2007; Wallace et al., 2013). While many different analysis methods have been used, each tries to quantify the amount of information present in the network states about the past inputs. For example, in one approach taken to study echo state networks (ESNs) (White et al., 2004; Ganguli et al., 2008; Hermans & Schrauwen, 2010), this information preservation is quantified through the correlation between the past input and the current state. When the correlation is too low, that input is said to no longer be represented in the state. The results of these analyses conclude that for Gaussian input statistics, the number of previous inputs that are significantly correlated with the current network state is bounded by a linear function of the network size.

In another line of analysis, researchers have sought to directly quantify the degree to which different inputs lead to unique network states (Jaeger, 2001; Maass et al., 2002; Legenstein & Maass, 2007; Strauss et al., 2012). In essence, the main idea of this work is that a one-to-one relationship between input sequences and the network states should allow the system to perform an inverse computation to recover the original input. A number of specific properties have been proposed to describe the uniqueness of the network state with respect to the input. In spiking liquid state machines (LSMs), in work by Maass et al. (2002), a separability property is suggested that guarantees distinct network states for distinct inputs and follow up work (Legenstein & Maass, 2007) relates the separability property to practical computational tasks through the Vapnik-

Chervonenkis (VC) dimension (Vapnik & Chervonenkis, 1971). More recent work analyzing similar networks using separation properties (Wallace et al., 2013; Büsing et al., 2010) gives an upper bound for the STM capacity that scales like the logarithm of the number of network nodes.

In discrete ESNs, the echo-state property (ESP) ensures that every network state at a given time is uniquely defined by some left-infinite sequence of inputs (Jaeger, 2001). The necessary condition for the ESP is that the maximum eigenvalue magnitude of the system is less than unity (an eigenvalue with a magnitude of one would correspond to a linear system at the edge of instability). While the ESP ensures uniqueness, it does not ensure robustness and output computations can be sensitive to small perturbations (i.e., noisy inputs). A slightly more robust property looks at the conditioning of the matrix describing how the system acts on an input sequence (Strauss et al., 2012). The condition number describes not only a one-to-one correspondence, but also quantifies how small perturbations in the input affect the output. While work by Strauss et al. (2012) is closest in spirit to the analysis in this paper, it ultimately concludes that the STM capacity still scales linearly with the network size.

Determining whether or not a system abides by one of the separability properties depends heavily on the network's construction. In some cases, different architectures can yield very different results. For example, in the case of randomly connected spiking networks, low connectivity (each neuron is connected to a small number of other neurons) can lead to large STM capacities (Legenstein & Maass, 2007; Büsing et al., 2010), whereas high connectivity leads to chaotic dynamics and smaller STM capacities (Wallace et al., 2013). In contrast, linear ESNs with high connectivities (appropriately normalized) (Büsing et al., 2010) can have relatively large STM capacities (on the order of the number of nodes in the network) (Ganguli et al., 2008; Strauss et al., 2012). Much of this work centers around using systems with orthogonal connectivity matrices, which leads to a topology that robustly preserves information. Interestingly, such systems can be constructed to have arbitrary connectivity while preserving the information
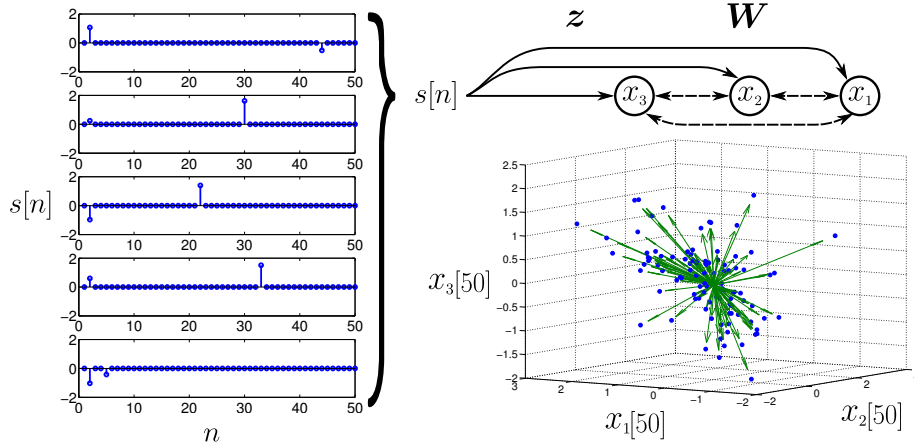
5

Figure 1: The current state of the network encodes information about the stimulus history. Different stimuli (examples shown to the left), when perturbing the same system (in this figure a three neuron orthogonal network) result in distinct states $x = [x_1, x_2, x_3]^T$ at the current time ($n = 50$). The current state is therefore informative for distinguishing between the input sequences.

preserving properties (Strauss et al., 2012).

While a variety of networks have been analyzed using the properties described above, these analyses ignore any structure of the inputs sequences that could be used to improve the analysis (Jaeger, 2001; Mayor & Gerstner, 2005). Conventional wisdom has suggested that STM capacities could be increased by exploiting structure in the inputs, but formal analysis has rarely addressed this case. For example, work by Ganguli & Sompolinsky (2010) builds significant intuition for the role of structured inputs in increasing STM capacity, specifically proposing to use the tools of CS to study the case when the input signals are temporally sparse. However, the analysis by Ganguli & Sompolinsky (2010) is asymptotic and focuses on an annealed (i.e., approximate) version of the system that neglects correlations between the network states over time. The present paper can be viewed as a generalization of this work to provide formal guarantees for STM capacity of the exact system dynamics, extensions to arbitrary orthogonal sparsity bases, and recovery bounds when the input exceeds the capacity of the system (i.e., the input is arbitrarily long).

2.2 Compressed Sensing

In the CS paradigm, a signal $s \in \mathbb{R}^N$ is sparse in a basis Ψ so that it can be approximately written as $s \approx \Psi a$, where most of the entries in $a \in \mathbb{R}^N$ are zero. This signal is observed through measurements $x \in \mathbb{R}^M$ taken via a compressive (e.g., $M \ll N$) linear operator:

$$x = As + \epsilon. \tag{1}$$

Sparse coefficients representing the signal are recovered by solving the convex optimization

$$\widehat{a} = \arg \min_{a} ||a||_1 \quad \text{such that} \quad ||x - A\Psi a||_2 \leq ||\epsilon||_2, \tag{2}$$

where $||\epsilon||_2$ is the magnitude of the measurement noise.

There is substantial evidence from the signal processing and computational neuroscience communities that many natural signals are sparse in an appropriate basis (Olshausen & Field, 1996; Elad et al., 2008). The recovery problem above requires that the system knows the sparsity basis Ψ to perform the recovery, which neural systems may not know a priori. We note that recent work has shown that appropriate sparsity bases can be learned from example data (Olshausen & Field, 1996), even in the case where the system only observes the inputs through compressed measurements (Isely et al., 2011). While the analysis doesn't depend on the exact method for solving the optimization in equation (2), we also note that this type of optimization can be solved in biologically plausible network architectures (e.g., (Rozell et al., 2010; Rhen & Sommer, 2007; Hu et al., 2012; Balavoine et al., 2012, 2013; Shapero et al., 2011)).

The most common sufficient condition in CS for stable recovery is known as the Restricted Isometry Property (RIP) (Candes & Tao, 2006). Formally, we say that RIP-$(2K, \delta)$ holds for A in the basis Ψ if for any vector s that is $2K$-sparse in Ψ we have

that

$$C\left(1 - \delta\right) \leq ||\boldsymbol{A}\boldsymbol{s}||_2^2 / ||\boldsymbol{s}||_2^2 \leq C\left(1 + \delta\right) \tag{3}$$

holds for constants $C > 0$ and $0 < \delta < 1$. Said another way, the RIP guarantees that all pairs of vectors that are K-sparse in $\boldsymbol{\Psi}$ have their distances preserved after projecting through the matrix $\boldsymbol{A}$. This can be seen by observing that for a pair of K-sparse vectors, their difference has at most $2K$ nonzeros. In this way, the RIP can be viewed as a type of separation property for sparse signals that is similar in spirit to the separation properties used in previous studies of network STM (Jaeger, 2001; White et al., 2004; Maass et al., 2002; Hermans & Schrauwen, 2010).

When $\boldsymbol{A}$ satisfies the RIP-$(2K, \delta)$ in the basis $\boldsymbol{\Psi}$ with 'reasonable' δ (e.g. $\delta \leq \sqrt{2} - 1$) and the signal estimate is $\widehat{\boldsymbol{s}} = \boldsymbol{\Psi}\widehat{\boldsymbol{a}}$, canonical results establish the following bound on signal recovery error:

$$||\boldsymbol{s} - \widehat{\boldsymbol{s}}||_2 \leq \alpha\,||\boldsymbol{\epsilon}||_2 + \beta\frac{\left|\left|\boldsymbol{\Psi}^T\left(\boldsymbol{s} - \boldsymbol{s}_K\right)\right|\right|_1}{\sqrt{K}}, \tag{4}$$

where α and β are constants and $\boldsymbol{s}_K$ is the best K-term approximation to $\boldsymbol{s}$ in the basis $\boldsymbol{\Psi}$ (i.e., using the K largest coefficients in $\boldsymbol{a}$) (Candes, 2006). Equation (4) shows that signal recovery error is determined by the magnitude of the measurement noise and sparsity of the signal. In the case that the signal is exactly K-sparse and there is no measurement noise, this bound guarantees perfect signal recovery.

While the guarantees above are deterministic and non-asymptotic, the canonical CS results state that measurement matrices generated randomly from "nice" independent distributions (e.g., Gaussian, Bernoulli) can satisfy RIP with high probability when $M = O(K \log N)$ (Rauhut, 2010). For example, random Gaussian measurement matrices (perhaps the most highly used construction in CS) satisfy the RIP condition for any sparsity basis with probability $1 - O(1/N)$ when $M \geq C\delta^{-2}K \log\left(N\right)$. This extremely

favorable scaling law (i.e., linear in the sparsity level) for random Gaussian matrices is in part due to the fact that Gaussian matrices have many degrees of freedom, resulting in M statistically independent observations of the signal. In many practical examples, there exists a high degree of structure in $\boldsymbol{A}$ that causes the measurements to be correlated. Structured measurement matrices with correlations between the measurements have been recently studied due to their computational advantages. While these matrices can still satisfy the RIP, they typically require more measurement to reconstruct a signal with the same fidelity and the performance may change depending on the sparsity basis (i.e., they are no longer "universal" because they don't perform equally well for all sparsity bases). One example which arises often in the signal processing community is the case of random circulant matrices (Krahmer et al., 2012), where the number of measurements needed to assure that the RIP holds with high probability for temporally sparse signals (i.e., $\boldsymbol{\Psi}$ is the identity) increases to $M \geq C\delta^{-2} K \log^4{(N)}$. Other structured systems analyzed in the literature include Toeplitz matrices (Haupt et al., 2010), partial circulant matrices (Krahmer et al., 2012), block diagonal matrices (Eftekhari et al., 2012; Park et al., 2011), subsampled unitary matrices (Bajwa et al., 2009), and randomly subsampled Fourier matrices (Rudelson & Vershynin, 2008). These types of results are used to demonstrate that signal recovery is possible with highly undersampled measurements, where the number of measurements scales linearly with the "information level" of the signal (i.e., the number of non-zero coefficients) and only logarithmically with the ambient dimension.

3 STM Capacity using the RIP

3.1 Network Dynamics as Compressed Sensing

We consider the same discrete-time ESN model used in previous studies (Jaeger, 2001; Ganguli et al., 2008; Ganguli & Sompolinsky, 2010; White et al., 2004):

$$\boldsymbol{x}[n] = f\left(\boldsymbol{W}\boldsymbol{x}[n-1] + \boldsymbol{z}s[n] + \widetilde{\boldsymbol{\epsilon}}[n]\right), \tag{5}$$

where $\boldsymbol{x}[n] \in \mathbb{R}^M$ is the network state at time n, $\boldsymbol{W}$ is the $(M \times M)$ recurrent (feedback) connectivity matrix, $s[n] \in \mathbb{R}$ is the input sequence at time n, $\boldsymbol{z}$ is the $(M \times 1)$ projection of the input into the network, $\widetilde{\boldsymbol{\epsilon}}[n]$ is a potential network noise source, and $f : \mathbb{R}^M \to \mathbb{R}^M$ is a possible pointwise nonlinearity. As in previous studies (Jaeger, 2001; White et al., 2004; Ganguli et al., 2008; Ganguli & Sompolinsky, 2010), this paper will consider the STM capacity of a linear network (i.e., $f(\boldsymbol{x}) = \boldsymbol{x}$).

The recurrent dynamics of Equation (5) can be used to write the network state at time N:

$$\boldsymbol{x}[N] = \boldsymbol{A}\boldsymbol{s} + \boldsymbol{\epsilon}, \tag{6}$$

where $\boldsymbol{A}$ is a $M \times N$ matrix, the k^{th} column of $\boldsymbol{A}$ is $\boldsymbol{W}^{k-1}\boldsymbol{z}$, $\boldsymbol{s} = [s[N], \ldots, s[1]]^T$, the initial state of the system is $\boldsymbol{x}[0] = 0$, and $\boldsymbol{\epsilon}$ is the node activity not accounted for by the input stimulus (e.g. the sum of network noise terms $\boldsymbol{\epsilon} = \sum_{k=1}^{N} \boldsymbol{W}^{N-k}\widetilde{\boldsymbol{\epsilon}}[k]$). With this network model, we assume that the input sequence $\boldsymbol{s}$ is K-sparse in an orthonormal basis $\boldsymbol{\Psi}$ (i.e., there are only K nonzeros in $\boldsymbol{a} = \boldsymbol{\Psi}^T\boldsymbol{s}$).

3.2 STM Capacity of Finite-Length Inputs

We first consider the STM capacity of a network with finite-length inputs, where a length-N input signal drives a network and the current state of the M network nodes

at time N is used to recover the input history via Equation (2). If $\boldsymbol{A}$ derived from the network dynamics satisfies the RIP for the sparsity basis $\boldsymbol{\Psi}$, the bounds in Equation (4) establish strong guarantees on recovering $\boldsymbol{s}$ from the current network states $\boldsymbol{x}[N]$. Given the significant structure in $\boldsymbol{A}$, it is not immediately clear that any network construction can result in $\boldsymbol{A}$ satisfying the RIP. However, the structure in $\boldsymbol{A}$ is very regular and in fact only depends on powers of $\boldsymbol{W}$ applied to $\boldsymbol{z}$:

$$
\boldsymbol{A} = \begin{bmatrix} \boldsymbol{z} & | & \boldsymbol{W}\boldsymbol{z} & | & \boldsymbol{W}^2\boldsymbol{z} & | & \ldots & | & \boldsymbol{W}^{N-1}\boldsymbol{z} \end{bmatrix}.
$$

Writing the eigendecomposition of the recurrent matrix $\boldsymbol{W} = \boldsymbol{U}\boldsymbol{D}\boldsymbol{U}^{-1}$, we re-write the measurement matrix as

$$
\boldsymbol{A} = \boldsymbol{U} \begin{bmatrix} \tilde{\boldsymbol{z}} & | & \boldsymbol{D}\tilde{\boldsymbol{z}} & | & \boldsymbol{D}^2\tilde{\boldsymbol{z}} & | & \ldots & | & \boldsymbol{D}^{N-1}\tilde{\boldsymbol{z}} \end{bmatrix},
$$

where $\tilde{\boldsymbol{z}} = \boldsymbol{U}^{-1}\boldsymbol{z}$. Rearranging, we get

$$
\boldsymbol{A} = \boldsymbol{U}\widetilde{\boldsymbol{Z}} \begin{bmatrix} \boldsymbol{d}^0 & | & \boldsymbol{d} & | & \boldsymbol{d}^2 & | & \ldots & | & \boldsymbol{d}^{N-1} \end{bmatrix} = \boldsymbol{U}\widetilde{\boldsymbol{Z}}\boldsymbol{F} \tag{7}
$$

where $\boldsymbol{F}_{k,l} = d_k^{l-1}$ is the k^{th} eigenvalue of $\boldsymbol{W}$ raised to the $(l-1)^{\text{th}}$ power and $\widetilde{\boldsymbol{Z}} = \text{diag}\,(\boldsymbol{U}^{-1}\boldsymbol{z})$.

While the RIP conditioning of $\boldsymbol{A}$ depends on all of the matrices in the decomposition of Equation 7, the conditioning of $\boldsymbol{F}$ is the most challenging because it is the only matrix that is compressive (i.e., not square). Due to this difficulty, we start by specifying a network structure for $\boldsymbol{U}$ and $\widetilde{\boldsymbol{Z}}$ that preserves the conditioning properties of $\boldsymbol{F}$ (other network constructions will be discussed in Section 4). Specifically, as in (White et al., 2004; Ganguli et al., 2008; Ganguli & Sompolinsky, 2010) we choose $\boldsymbol{W}$ to be a random orthonormal matrix, assuring that the eigenvector matrix $\boldsymbol{U}$ has orthonormal columns and preserves the conditioning properties of $\boldsymbol{F}$. Likewise, we choose the feed-forward vector $\boldsymbol{z}$ to be $\boldsymbol{z} = \frac{1}{\sqrt{M}}\boldsymbol{U}\boldsymbol{1}_M$, where $\boldsymbol{1}_M$ is a vector of M ones (the constant

$\sqrt{M}$ simplifies the proofs but has no bearing on the result). This choice for $\boldsymbol{z}$ assures that $\widetilde{\boldsymbol{Z}}$ is the identity matrix scaled by $\sqrt{M}$ (analogous to (Ganguli et al., 2008) where $\boldsymbol{z}$ is optimized to maximize the SNR in the system). Finally, we observe that the richest information preservation apparently arises for a real-valued $\boldsymbol{W}$ when its eigenvalues are complex, distinct in phase, have unit magnitude, and appear in complex conjugate pairs.

For the above network construction, our main result shows that $\boldsymbol{A}$ satisfies the RIP in the basis $\boldsymbol{\Psi}$ (implying the bounds from Equation (4) hold) when the network size scales linearly with the sparsity level of the input. This result is made precise in the following theorem:

Theorem 3.1. *Suppose $N \geq M$, $N \geq K$ and $N \geq O(1)$.[1] Let $\boldsymbol{U}$ be any unitary matrix of eigenvectors (containing complex conjugate pairs) and set $\boldsymbol{z} = \frac{1}{\sqrt{M}}\boldsymbol{U}\mathbf{1}_M$ so that $\widetilde{\boldsymbol{Z}} = diag\left(\boldsymbol{U}^{-1}\boldsymbol{z}\right) = \frac{1}{\sqrt{M}}\boldsymbol{I}$. For M an even integer, denote the eigenvalues of $\boldsymbol{W}$ by $\{e^{jw_m}\}_{m=1}^{M}$. Let the first $M/2$ eigenvalues $\left(\{e^{jw_m}\}_{m=1}^{M/2}\right)$ be chosen uniformly at random on the complex unit circle (i.e., we chose $\{w_m\}_{m=1}^{M/2}$ uniformly at random from $[0, 2\pi)$) and the other $M/2$ eigenvalues as the complex conjugates of these values (i.e., for $M/2 < m \leq M$, $e^{jw_m} = e^{-jw_{m-M/2}}$). Under these conditions, for a given RIP conditioning $\delta < 1$ and failure probability η, if*

$$M \geq C\frac{K}{\delta^2}\mu^2\left(\boldsymbol{\Psi}\right)\log^4\left(N\right)\log(\eta^{-1}), \tag{8}$$

for a universal constant C, then for any $\boldsymbol{s}$ that is K-sparse (i.e., has no more than K non-zero entries)

$$(1 - \delta) \leq \|\boldsymbol{A}\boldsymbol{\Psi}\boldsymbol{s}\|_2^2 / \|\boldsymbol{s}\|_2^2 \leq (1 + \delta)$$

---

[1] The notation $N \geq O(1)$ means that $N \geq C$ for some constant C. For clarity, we do not keep track of the constants in our proofs. The interested reader is referred to (Rauhut, 2010) for specific values of the constants.

with probability exceeding $1 - \eta$.

The proof of this statement is given in Appendix 6.1 and follows closely the approach in (Rauhut, 2010) by generalizing it to both include any basis $\mathbf{\Psi}$ and account for the fact that $\boldsymbol{W}$ is a real-valued matrix. The quantity $\mu(\cdot)$ (known as the coherence) captures the largest inner product between the sparsity basis and the Fourier basis, and is calculated as:

$$\mu(\mathbf{\Psi}) = \max_{n=1,\dots,N} \sup_{t \in [0,2\pi]} \left| \sum_{m=0}^{N-1} \mathbf{\Psi}_{m,n} e^{-jtm} \right|. \tag{9}$$

In the result above, the coherence is lower (therefore the STM capacity is higher) when the sparsity basis is more "different" from the Fourier basis.

The main observation of the result above is that STM capacity scales superlinearly with network size. Indeed, for some values of K and $\mu(\mathbf{\Psi})$ it is possible to have STM capacities much greater than the number of nodes (i.e., $N \gg M$). To illustrate the perfect recovery of signal lengths beyond the network size, Figure 2 shows an example recovery of a single long input sequence. Specifically, we generate a 100 node random orthogonal connectivity matrix $\boldsymbol{W}$ and generate $\boldsymbol{z} = \frac{1}{\sqrt{M}} \boldsymbol{U} \mathbf{1}_M$. We then drive the network with an input sequence that is 480 samples long and constructed using 24 non-zero coefficients (chosen uniformly at random) of a wavelet basis. The values at the non-zero entries were chosen uniformly in the range [0.5,1.5]. In this example we omit noise so that we can illustrate the noiseless recovery. At the end of the input sequence, the resulting 100 network states are used to solve the optimization problem in Equation 2 for recovering the input sequence (using the network architecture in (Rozell et al., 2010)). The recovered sequence, as depicted in Figure 2, is identical to the input sequence, clearly indicating that the 100 nodes were able to store the 480 samples of the input sequence (achieving STM capacity higher than the network size).

Directly checking the RIP condition for specific matrices is NP-hard (one would need to check every possible $2K$-sparse signal). In light of this difficulty in verifying

13

recovery of all possible sparse signals (which the RIP implies), we will explore the qualitative behavior of the RIP bounds above by examining in Figure 3 the average recovery relative MSE (rMSE) in simulation for a network with M nodes when recovering input sequences of length N with varying sparsity bases. Figure 3 uses a plotting style similar to the Donoho-Tanner phase transition diagrams (Donoho & Tanner, 2005) where the average recovery rMSE is shown for each pair of variables under noisy conditions. While the traditional Donoho-Tanner phase transitions plot noiseless recovery performance to observe the threshold between perfect and imperfect recovery, here we also add noise to illustrate the stability of the recovery guarantees. The noise is generated as random additive Gaussian noise at the input ($\widetilde{\epsilon}$ in Equation (5)) to the system with zero mean and variance such that the total noise in the system (ϵ in Equation (6)) has a norm of approximately 0.01. To demonstrate the behavior of the system, the phase diagrams in Figure 3 sweep the ratio of measurements to the total signal length (M/N) and the ratio of the signal sparsity to the number of measurements (K/M). Thus at the upper left hand corner, the system is recovering a dense signal from almost no measurements (which should almost certainly yield poor results) and at the right hand edge of the plots the system is recovering a signal from a full set of measurements (enough to recover the signal well for all sparsity ranges). We generate ten random ESNs for each combination of ratios (M/N, K/M). The simulated networks are driven with input sequences that are sparse in one of four different bases (Canonical, Daubechies-10 wavelet, Symlet-3 wavelet and DCT) which have varying coherence with the Fourier basis. We use the node values at the end of the sequence to recover the inputs.[2]

In each plot of Figure 3, the dashed line denotes the boundary where the system is able to essentially perform perfect recovery (recovery error $\leq 1\%$) up to the noise floor. Note that the area under this line (the white area in the plot) denotes the region where the system is leveraging the sparse structure of the input to get capacities of $N > M$. We

---

[2]For computational efficiency, we use the TFOCS software package (Becker et al., 2011) to solve the optimization problem in Equation (2) for these simulations.
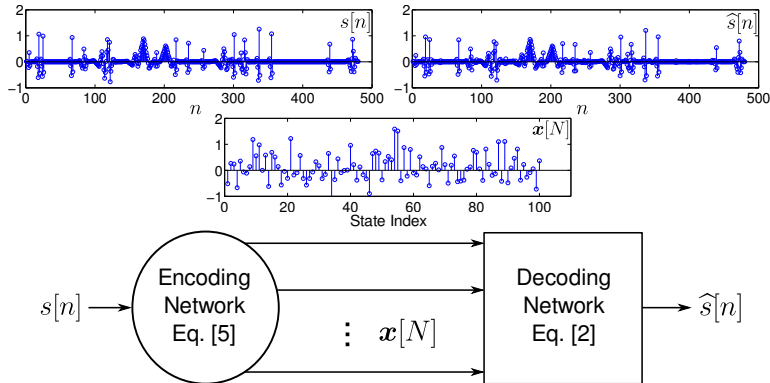
Figure 2: A length 480 stimulus pattern (left plot) that is sparse in a wavelet basis drives the encoding network defined by a random orthogonal matrix W and a feed-forward vector z. The 100 node values (center plot) are then used to recover the full stimulus pattern (right plot) using a decoding network which solves Equation (2).

also observe that the dependence of the RIP bound on the coherence with the Fourier basis is clearly shown qualitatively in these plots, with the DCT sparsity basis showing much worse performance than the other bases.

3.3   STM Capacity of Infinite-Length Inputs

After establishing the perfect recovery bounds for finite-length inputs in the previous section, we turn here to the more interesting case of a network that has received an input beyond its STM capacity (perhaps infinitely long). In contrast to the finite-length input case where favorable constructions for W used random unit-norm eigenvalues, this construction would be unstable for infinitely long inputs. In this case, we take W to have all eigenvalue magnitudes equal to $q < 1$ to ensure stability. The matrix constructions we consider in this section are otherwise identical to that described in the previous section.

In this scenario, the recurrent application of W in the system dynamics assures that each input perturbation will decay steadily until it has zero effect on the network state. While good for system stability, this decay means that each input will slowly recede

Figure 3: Random orthogonal networks can have a STM capacity that exceeds the number of nodes. These plots depict the recovery relative mean square error (rMSE) for length-1000 input sequences from M network nodes where the input sequences are K-sparse. Each figure depicts recovery for a given set of ratios M/N and K/M. Recovery is near perfect (rMSE $\leq 1\%$; denoted by the dotted line) for large areas of each plot (to the left of the $N = M$ boundary at the right of each plot) for sequences sparse in the canonical basis or various wavelet basis (shown here are 4 level decompositions in Symlet-3 wavelets and Daubechies-10 wavelets). For bases more coherent with the Fourier basis (e.g., discrete cosine transform-DCT), recovery performance above $N = M$ can suffer significantly. All the recovery here was done for noise such that $\|\boldsymbol{\epsilon}\|_2 \approx 0.01$.

into the past until the network activity contains no useable memory of the event. In other words, *any* network with this decay can only hope to recover a proxy signal that accounts for the decay in the signal representation induced by the forgetting factor q. Specifically, we define this proxy signal to be $\boldsymbol{Qs}$, where $\boldsymbol{Q} = \text{diag}\left([1, q, q^2, \ldots]\right)$. Previous work (Ganguli et al., 2008; Jaeger, 2001; White et al., 2004) has characterized recoverability by using statistical arguments to quantify the correlation of the node values to each past input perturbation. In contrast, our approach is to provide recovery bounds on the rMSE for a network attempting to recover the N past samples of $\boldsymbol{Qs}$, which corresponds to the weighted length-N history of $\boldsymbol{s}$. Note that in contrast to the previous section where we established the length of the input that can be perfectly recovered, the amount of time we attempt to recall (N) is now a parameter that can be varied.

Our technical approach to this problem comes from observing that activity due to inputs older than N acts as interference when recovering more recent inputs. In other words, we can group older terms (i.e., from farther back than N time samples ago)

with the noise term, resulting again in $\boldsymbol{A}$ being an M by N linear operation that can satisfy RIP for length-N inputs. In this case, after choosing the length of the memory to recover, the guarantees in Equation (4) hold when considering every input older than N as contributing to the "noise" part of the bound.

Specifically, in the noiseless case where $\boldsymbol{s}$ is sparse in the canonical basis ($\mu\left(\boldsymbol{I}\right) = 1$) with a maximum signal value s_{max}, we can bound the first term of Equation (4) using a geometric sum that depends on N, K and q. For a given scenario (i.e., a choice of q, K and the RIP conditioning of $\boldsymbol{A}$), a network can support signal recovery up to a certain sparsity level K^*, given by:

$$K^* = \frac{M\delta^2}{C\log^\gamma\left(N\right)}, \tag{10}$$

where γ is a scaling constant (e.g., $\gamma = 4$ using the present techniques, but $\gamma = 1$ is conjectured (Rudelson & Vershynin, 2008)). We can also bound the second term of Equation (4) by the sum of the energy in the past N perturbations that are beyond this sparsity level K^*. Together these terms yield the bound on the recovery of the proxy signal:

$$
\begin{aligned}
||\boldsymbol{Q}\boldsymbol{s} - \boldsymbol{Q}\widehat{\boldsymbol{s}}||_2 \quad \leq \quad & \beta s_{\mathrm{max}}||\boldsymbol{U}||_2\left(\frac{q^N}{1-q}\right) \\
+ \quad & \frac{\beta s_{\mathrm{max}}}{\sqrt{\min\left[K^*, K\right]}}\left(\frac{q^{\min[K^*,K]} - q^K}{1-q}\right) \\
+ \quad & \alpha \epsilon_{\mathrm{max}}||\boldsymbol{U}||_2\left|\frac{q}{1-q}\right|.
\end{aligned}
\tag{11}
$$

The derivation of the first two terms in the above bound is detailed in Appendix 6.3, and the final term is simply the accumulated noise, which should have bounded norm due to the exponential decay of the eigenvalues of $\boldsymbol{W}$.

Intuitively, we see that this approach implies the presence of an optimal value for the recovery length N. For example, choosing N too small means that there is useful signal information in the network that the system is not attempting to recover, resulting

in omission errors (i.e., an increase in the first term of Equation (4) by counting too much signal as noise). On the other hand, choosing N too large means that the system is encountering recall errors by trying to recover inputs with little or no residual information remaining in the network activity (i.e., an increase in the second term of Equation (4) from making the signal approximation worse by using the same number of nodes for a longer signal length).

The intuitive argument above can be made precise in the sense that the bound in Equation (11) does have at least one local minimum for some value of $0 < N < \infty$. First, we note that the noise term (i.e., the third term on the right side of Equation (11)) does not depend on N (the choice in origin does not change the infinite summation), implying that the optimal recovery length only depends on the first two terms. We also note the important fact that K^* is non-negative and monotonically decreasing with increasing N. It is straightforward to observe that the bound in equation Equation (11) tends to infinity as N increases (due to the presence of K^* in the denominator of the second term). Furthermore, for small values of N, the second term in Equation (11) is zero (due to $K^* > K$), and the first term is monotonically decreasing with N. Taken together, since the function is continuous in N, has negative slope for small N and tends to infinity for large N, we can conclude that it must have at least one local minima in the range $0 < N < \infty$. This result predicts that there is (at least one) optimal value for the recovery length N.

The prediction of an optimal recovery length above is based on the fact that the error bound in Equation (11)), and it is possible that the error itself will not actually show this behavior (since the bound may not be tight in all cases). To test the qualitative intuition from Equation (11), we simulate recovery of input lengths and show the results in Figure 4. Specifically, we generate 50 ESNs with 500 nodes and a decay rate of q =0.999. The input signals are length-8000 sequences that have 400 nonzeros whose locations are chosen uniformly at random and whose amplitudes are chosen from a Gaussian distribution (zero mean and unit variance). After presenting the full 8000
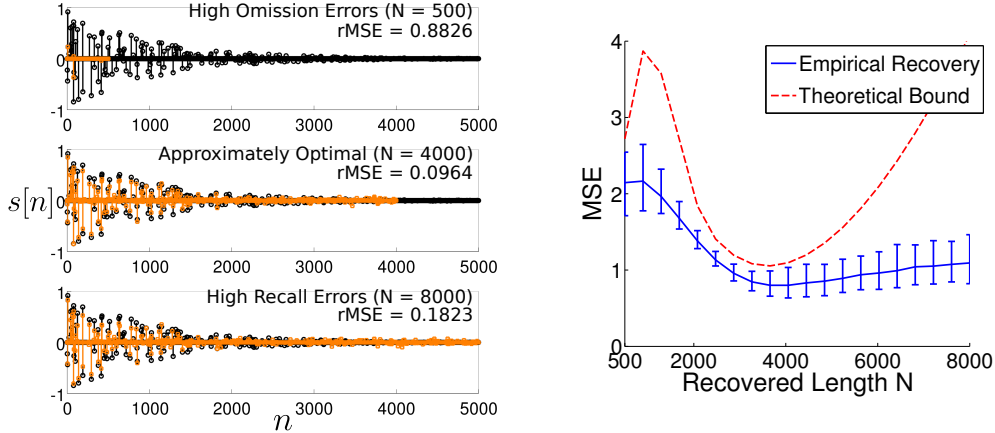
Figure 4: The theoretical bound on the recovery error for the past N perturbations to a network of size M has a minimum value at some optimal recovery length. This optimal value depends on the network size, the sparsity K, the decay rate q, and the RIP conditioning of A. Shown on the right is a simulation depicting the MSE for both the theoretical bound (red dashed line) and an empirical recovery for varying recovery lengths N. In this simulation $K = 400$, $q = 0.999$, $M = 500$. The error bars for the empirical curve show the maximum and minimum MSE. On the left we show recovery (in orange) of a length-8000 decayed signal (in black) when recovering the past 500 (top), 4000 (middle), and 8000 (bottom) most recent perturbations. As expected, at $N = 4000$ (approximately optimal) the recovery has the highest accuracy.

samples of the input signal to the network, we use the network states to recover the input history with varying lengths and compared the resulting MSE to the bound in Equation (11). Note that while the theoretical bound may not be tight for large signal lengths, the recovery MSE matches the qualitative behavior of the bound by achieving a minimum value at $N > M$.

4 Other Network Constructions

4.1 Alternate Orthogonal Constructions

Our results in the previous section focus on the case where W is orthogonal and z projects the signal evenly into all eigenvectors of W. When either W or z deviate from this structure the STM capacity of the network apparently decreases. In this section

we revisit those specifications, considering alternate network structures allowed under these assumptions as well as the consequences of deviating from these assumptions in favor of other structural advantages for a system (e.g., wire length, etc.).

To begin, we consider the assumption of orthogonal network connectivity, where the eigenvalues have constant magnitude and the eigenvectors are orthonormal. Constructed in this way, U exactly preserves the conditioning of $\widetilde{Z}F$. While this construction may seem restrictive, orthogonal matrices are relatively simple to generate and encompass a number of distinct cases. For small networks, selecting the eigenvalues uniformly at random from the unit circle (and including their complex conjugates to ensure real connectivity weights) and choosing an orthonormal set of complex conjugate eigenvectors creates precisely these optimal properties. For larger matrices, the connectivity matrix can instead be constructed directly by choosing W at random and orthogonalizing the columns. Previous results on random matrices (Diaconis & Shahshahani, 1994) guarantee that as the size of W increases, the eigenvalue probability density approaches the uniform distribution as desired. Some recent work in STM capacity demonstrates an alternate method by which orthogonal matrices can be constructed while constraining the total connectivity of the network (Strauss et al., 2012). This method iteratively applies rotation matrices to obtain orthogonal matrices with varying degrees of connectivity. We note here that one special case of connectivity matrices not well-suited to the STM task, even when made orthogonal, are symmetric networks, where the strictly real-valued eigenvalues generates poor RIP conditioning for F.

While simple to generate in principle, the matrix constructions discussed above are generally densely connected and may be impractical for many systems. However, many other special network topologies that may be more biophysically realistic (i.e., block diagonal connectivity matrices and small-world[3] networks (Mongillo et al., 2008)) can be constructed so that W still has orthonormal columns. For example, consider the case of

---

[3]Small-world structures are typically taken to be networks where small groups of neurons are densely connected amongst themselves, yet sparse connections to other groups reduces the maximum distance between any two nodes.
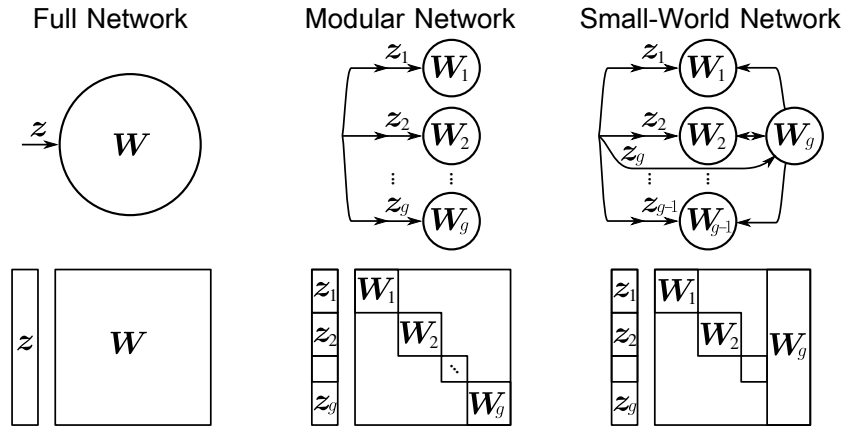
Figure 5: Possible network topologies which have orthogonal connectivity matrices. In the general case, all nodes are connected via non-symmetric connections. Modular topologies can still be orthogonal if each block is itself orthogonal. Small world topologies may also have orthogonal connectivity, especially when a few nodes are completely connected to a series of otherwise disjoint nodes.

a block diagonal connection matrix (illustrated in Figure 5), where many unconnected networks of at least two nodes each are driven by the same input stimulus and evolve separately. Such a structure lends itself to a modular framework, where more of these subnetworks can be recruited to recover input stimuli further in the past. In this case, each block can be created independently as above and pieced together. The columns of the block diagonal matrix will still have unit norm and will be both orthogonal to vectors within its own block (since each of the diagonal sub-matrices are orthonormal) and orthogonal to all columns in other blocks (since there is no overlap in the non-zero indices).

Similarly, a small-world topology can be achieved by taking a few of the nodes in every group of the block diagonal case and allowing connections to all other neurons (either unidirectional or bidirectional connections). To construct such a matrix, a block diagonal orthogonal matrix can be taken, a number of columns can be removed and replaced with full columns, and the resulting columns can be made orthonormal with respect to the remaining block-diagonal columns. In these cases, the same eigenvalue distribution and eigenvector properties hold as the fully connected case, resulting in

the same RIP guarantees (and therefore the same recovery guarantees) demonstrated earlier. We note that this is only one approach to constructing a network with favorable STM capacity and not all networks with small-world properties will perform well.

Additionally, we note that as opposed to networks analyzed in prior work (in particular the work in (Wallace et al., 2013) demonstrating that random networks with high connectivity have short STM), the average connectivity does not play a dominant role in our analysis. Specifically, it has been observed in spiking networks that higher network connectivity can reduce the STM capacity so that is scales only with $\log(M)$ (Wallace et al., 2013)). However, in our ESN analysis, networks can have low connectivity (e.g. 2x2 block-diagonal matrices - the extreme case of the block diagonal structure described above) or high connectivity (e.g. fully connected networks) and have the same performance.

4.2 Suboptimal Network Constructions

Finally, we can also analyze some variations to the network structure assumed in this paper to see how much performance decreases. First, instead of the deterministic construction for $\boldsymbol{z}$ discussed in the earlier sections, there has also been interest in choosing $\boldsymbol{z}$ as i.i.d. random Gaussian values (Ganguli et al., 2008; Ganguli & Sompolinsky, 2010). In this case, it is also possible to show that $\boldsymbol{A}$ satisfies the RIP (with respect to the basis $\boldsymbol{\Psi}$ and with the same RIP conditioning δ as before) by paying an extra $\log(N)$ penalty in the number of measurements. Specifically, we have also established the following theorem:

Theorem 4.1. *Suppose $N \geq M$, $N \geq K$ and $N \geq O(1)$. Let $\boldsymbol{U}$ be any unitary matrix of eigenvectors (containing complex conjugate pairs) and the entries of $\boldsymbol{z}$ be i.i.d. zero-mean Gaussian random variables with variance $\frac{1}{M}$. For M an even integer, denote the eigenvalues of $\boldsymbol{W}$ by $\{e^{jw_m}\}_{m=1}^{M}$. Let the first $M/2$ eigenvalues ($\{e^{jw_m}\}_{m=1}^{M/2}$) be chosen uniformly at random on the complex unit circle (i.e., we chose $\{w_m\}_{m=1}^{M/2}$ uniformly at*

random from $[0, 2\pi)$*) and the other* $M/2$ *eigenvalues as the complex conjugates of these values. Then, for a given RIP conditioning* δ *and failure probability* $N^{-\log^4 N} \le \eta \le \frac{1}{e}$*, if*

$$M \ge C \frac{K}{\delta^2} \mu^2 (\boldsymbol{\Psi}) \log^5 (N) \log(\eta^{-1}), \tag{12}$$

$\boldsymbol{A}$ *satisfies RIP-*(K, δ) *with probability exceeding* $1 - \eta$ *for a universal constant* C.

The proof of this theorem can be found in Appendix 6.2. The additional log factor in the bound in (12) reflects that a random feed-forward vector may not optimally spread the input energy over the different eigen-directions of the system. Thus, some nodes may see less energy than others, making them slightly less informative. Note that while this construction does perform worse that the optimal constructions from Theorem 3.1, the STM capacity is still very favorable (i.e., a linear scaling in the sparsity level and logarithmic scaling in the signal length).

Second, instead of orthogonal connectivity matrices, there has also been interest in network constructions involving non-orthogonal connectivity matrices (perhaps for noise reduction purposes (Ganguli et al., 2008)). When the eigenvalues of $\boldsymbol{W}$ still lie on the complex unit circle, we can analyze how non-orthogonal matrices affect the RIP results. In this case, the decomposition in Equation (7) still holds and Theorem 3.1 still applies to guarantee that $\boldsymbol{F}$ satisfies the RIP. However, the non-orthogonality changes the conditioning of $\boldsymbol{U}$ and subsequently the total conditioning of $\boldsymbol{A}$. Specifically the conditioning of $\boldsymbol{U}$ (the ratio of the maximum and minimum singular values $\sigma_{\max}^2 / \sigma_{\min}^2 = \gamma$) will effect the total conditioning of $\boldsymbol{A}$. We can use the RIP of $\boldsymbol{F}$ and the extreme singular values of $\boldsymbol{U}$ to bound how close $\boldsymbol{UF}$ is to an isometry for sparse vectors, both above by

$$\left\| \boldsymbol{UFs} \right\|_2^2 \le \sigma_{\max}^2 \left\| \boldsymbol{Fs} \right\|_2^2 \le \sigma_{\max}^2 C(1 + \delta) \left\| \boldsymbol{s} \right\|_2^2,$$

and below by

$$\left|\left|\boldsymbol{UFs}\right|\right|_2^2 \geq \sigma_{\min}^2 \left|\left|\boldsymbol{Fs}\right|\right|_2^2 \geq \sigma_{\min}^2 C(1-\delta) \left|\left|\boldsymbol{s}\right|\right|_2^2.$$

By consolidating these bounds, we find a new RIP statement for the composite matrix

$$C'\left(1-\delta'\right)\left|\left|\boldsymbol{s}\right|\right|_2^2 \leq \left|\left|\boldsymbol{UFs}\right|\right|_2^2 \leq C'\left(1+\delta'\right)\left|\left|\boldsymbol{s}\right|\right|_2^2$$

where $\sigma_{\min}^2 C(1-\delta) = C'(1-\delta')$ and $\sigma_{\max}^2 C(1+\delta) = C'(1+\delta')$. These relationships can be used to solve for the new RIP constants:

$$\delta' = \frac{\frac{\gamma-1}{\gamma+1} + \delta}{1 + \delta\frac{\gamma-1}{\gamma+1}}$$

$$C' = \frac{1}{2}C\left(\sigma_{\max}^2 + \sigma_{\min}^2 + \delta(\sigma_{\max}^2 - \sigma_{\min}^2)\right)$$

These expressions demonstrate that as the conditioning of $\boldsymbol{U}$ improves (i.e. $\gamma \to 1$), the RIP conditioning does not change from the optimal case of an orthogonal network ($\delta' = \delta$). However, as the conditioning of $\boldsymbol{U}$ gets worse and γ grows, the constants associated with the RIP statement also get worse (implying more measurements are likely required to guarantee the same recovery performance).

The above analysis primarily concerns itself with constructions where the eigenvalues of $\boldsymbol{W}$ are still unit norm, however $\boldsymbol{U}$ is not orthogonal. Generally, when the eigenvalues of $\boldsymbol{W}$ differ from unity and are not all of equal magnitude, the current approach becomes intractable. In one case, however, there are theoretical guarantees: when $\boldsymbol{W}$ is rank deficient. If $\boldsymbol{W}$ only has $\widetilde{M}$ unit-norm eigenvalues, and the remaining $M - \widetilde{M}$ eigenvalues are zero, then the resulting matrix $\boldsymbol{A}$ is composed the same way, except that the bottom $M - \widetilde{M}$ rows are all zero. This means that the effective

measurements only depend on an $\widetilde{M} \times N$ subsampled DTFT

$$
\begin{aligned}
\boldsymbol{x}[N] \;&=\; \boldsymbol{U}\widetilde{\boldsymbol{Z}}\boldsymbol{F}\boldsymbol{s} + \boldsymbol{\epsilon} \\
&=\; \boldsymbol{U}\widetilde{\boldsymbol{Z}}
\begin{bmatrix}
\widetilde{\boldsymbol{F}} \\
\boldsymbol{0}_{M-\widetilde{M},N}
\end{bmatrix}
\boldsymbol{s} + \boldsymbol{\epsilon} \\
&=\; \boldsymbol{U}\widetilde{\boldsymbol{Z}}_{1:\widetilde{M}}\widetilde{\boldsymbol{F}}\boldsymbol{s} + \boldsymbol{\epsilon}
\end{aligned}
$$

where $\widetilde{\boldsymbol{F}}$ is matrix consisting of the non-zero rows of $\boldsymbol{F}$. In this case we can choose any $\widetilde{M}$ of the nodes and the previous theorems will all hold, replacing the true number of nodes M with the effective number of nodes $\widetilde{M}$.

5 Discussion

We have seen that the tools of the CS literature can provide a way to quantify the STM capacity in linear networks using rigorous non-asymptotic recovery error bounds. Of particular note is that this approach leverages the non-Gaussianity of the input statistics to show STM capacities that are superlinear in the size of the network and depend linearly on the sparsity level of the input. This work provides a concrete theoretical understanding for the approach conjectured in (Ganguli & Sompolinsky, 2010) along with a generalization to arbitrary sparsity bases and infinitely long input sequences. This analysis also predicts that there exists an optimal recovery length that balances omission errors and recall mistakes.

In contrast to previous work on ESNs that leverage nonlinear network computations for computational power (Jaeger & Haas, 2004), the present work uses a linear network and nonlinear computations for signal recovery. Despite the nonlinearity of the recovery process, the fundamental results of the CS literature also guarantee that the recovery process is stable and robust. For example, with access to only a subset of nodes (due to failures or communication constraints), signal recovery generally degrades gracefully

by still achieving the best possible approximation of the signal using fewer coefficients. Beyond signal recovery, we also note that the RIP can guarantee performance on many tasks (e.g. detection, classification, etc.) performed directly on the network states (Davenport et al., 2010). Finally, we note that while this work only addresses the case where a single input is fed to the network, there may be networks of interest that have a number of input streams all feeding into the same network (with different feed-forward vectors). We believe that the same tools utilized here can be used in the multi-input case, since the overall network state is still a linear function of the inputs.

Acknowledgments

6 Appendix

6.1 Proof of RIP

In this appendix, we show that the matrix $\boldsymbol{A} = \boldsymbol{U}\widetilde{\boldsymbol{Z}}\boldsymbol{F}$ satisfies the RIP under the conditions stated in Equation (8) of the main text in order to prove Theorem 3.1. We note that (Rauhut, 2010) shows that for the canonical basis ($\boldsymbol{\Psi} = \boldsymbol{I}$), the bounds for M can be tightened to $M \geq \max\left\{ C\frac{K}{\delta^2}\log^4 N,\ C'\frac{K}{\delta^2}\log\eta^{-1} \right\}$ using a more complex proof technique than we will employ here. For $\eta = \frac{1}{N}$, the result in (Rauhut, 2010) represents an improvement of several $\log(N)$ factors when restricted to only the canonical basis for $\boldsymbol{\Psi}$. We also note that the scaling constant C found in the general RIP definition of Equation (3) of the main text is unity due to the $\sqrt{M}$ scaling of $\boldsymbol{z}$.

While the proof of Theorem 3.1 is fairly technical, the procedure follows very closely the proof of Theorem 8.1 from (Rauhut, 2010) on subsampled discrete time

Fourier transform (DTFT) matrices. While the basic approach is the same, the novelty in our presentation is the incorporation of the sparsity basis $\boldsymbol{\Psi}$ and considerations for a real-valued connectivity matrix $\boldsymbol{W}$.

Before beginning the proof of this theorem, we note that because $\boldsymbol{U}$ is assumed unitary, $\|\boldsymbol{A\Psi s}\|_2 = \|\widetilde{\boldsymbol{Z}}\boldsymbol{F\Psi s}\|_2$ for any signal $\boldsymbol{s}$. Thus, it suffices to establish the conditioning properties of the matrix $\widehat{\boldsymbol{A}} := \widetilde{\boldsymbol{Z}}\boldsymbol{F\Psi}$. For the upcoming proof, it will be useful to write this matrix as a sum of rank-1 operators. The specific rank-1 operator that will be useful for our purposes is $X_l X_l^H$ with $X_l^H := \boldsymbol{F}_l^H \boldsymbol{\Psi}$, the conjugate of the l-th row of $\boldsymbol{F\Psi}$, where $\boldsymbol{F}_l^H := \begin{bmatrix} 1, & e^{jw_l}, & \cdots, & e^{jw_l(N-1)} \end{bmatrix} \in \mathbb{C}^N$ is the conjugated l-th row of $\boldsymbol{F}$. Because of the way the "frequencies" $\{w_m\}$ are chosen, for any $l > \frac{M}{2}$, $X_l = X_{l-\frac{M}{2}}^*$. The l-th row of $\widehat{\boldsymbol{A}}$ is $\widetilde{z}_l X_l^H$ where $\widetilde{z}_l$ is the l-th diagonal entry of the diagonal matrix $\widetilde{\boldsymbol{Z}}$, meaning that we can use the sum of rank-1 operators to write the decomposition $\widehat{\boldsymbol{A}}^H \widehat{\boldsymbol{A}} = \sum_{l=1}^{M} |\widetilde{z}_l|^2 X_l X_l^H$. If we define the random variable $\boldsymbol{B} := \widehat{\boldsymbol{A}}^H \widehat{\boldsymbol{A}} - \boldsymbol{I}$ and the norm $\|\boldsymbol{B}\|_K := \sup_{y \text{ is } K\text{-sparse}} \dfrac{\boldsymbol{y}^H \boldsymbol{B} \boldsymbol{y}}{\boldsymbol{y}^H \boldsymbol{y}}$, we can equivalently say that $\widehat{\boldsymbol{A}}$ has RIP conditioning δ if

$$\|\boldsymbol{B}\|_K := \left\| \widehat{\boldsymbol{A}}^H \widehat{\boldsymbol{A}} - \boldsymbol{I} \right\|_K = \left\| \sum_{l=1}^{M} |\widetilde{z}_l|^2 X_l X_l^H - \boldsymbol{I} \right\|_K \leq \delta.$$

To aid in the upcoming proof, we make a few preliminary observations and rewrite the quantities of interest in some useful ways. First, because of the correspondences between the summands in $\widehat{\boldsymbol{A}}^H \widehat{\boldsymbol{A}}$ (i.e. $X_l = X_{l-M/2}^*$), we can rewrite $\widehat{\boldsymbol{A}}^H \widehat{\boldsymbol{A}}$ as

$$\widehat{\boldsymbol{A}}^H \widehat{\boldsymbol{A}} = \sum_{l=1}^{M/2} |\widetilde{z}_l|^2 X_l X_l^H + \sum_{l=1}^{M/2} |\widetilde{z}_l|^2 \left(X_l X_l^H \right)^*,$$

making clear the fact that there are only $\frac{M}{2}$ independent w_m's. Under the assumption of

Theorem 3.1, $\widetilde{z}_l = \frac{1}{\sqrt{M}}$ for $l = 1, \cdots, M$. Therefore,

$$\mathcal{E}\left[\sum_{l=1}^{M/2} |\widetilde{z}_l|^2 X_l X_l^H\right] = \sum_{l=1}^{M/2} |\widetilde{z}_l|^2 \mathcal{E}\left[X_l X_l^H\right] = \sum_{l=1}^{M/2} \frac{1}{M} \Psi^H \mathcal{E}\left[F_l F_l^H\right] \Psi = \frac{1}{2} I,$$

where it is straightforward to check that $\mathcal{E}\left[F_l F_l^H\right] = I$. By the same reasoning, we also have $\mathcal{E}\left[\sum_{l=1}^{M/2} |\widetilde{z}_l|^2 \left(X_l X_l^H\right)^*\right] = \frac{1}{2} I$. This implies that we can rewrite B as

$$\begin{aligned}
B &= \sum_{l=1}^{M} \left(|\widetilde{z}_l|^2 X_l X_l^H\right) - I \\
&= \left(\sum_{l=1}^{M/2} |\widetilde{z}_l|^2 X_l X_l^H - \frac{1}{2} I\right) + \left(\sum_{l=1}^{M/2} |\widetilde{z}_l|^2 \left(X_l X_l^H\right)^* - \frac{1}{2} I\right) \\
&=: \; B_1 + B_2.
\end{aligned}$$

The main proof of the theorem has two main steps. First, we will establish a bound on the moments of the quantity of interest $\|B\|_K$. Next we will use these moments to derive a tail bound on $\|B\|_K$, which will lead directly to the RIP statement we seek. The following two lemmas from the literature will be critical for these two steps.

Lemma 6.1 (Lemma 8.2 of (Rauhut, 2010)). *Suppose $M \geq K$ and suppose we have a sequence of (fixed) vectors $Y_l \in \mathbb{C}^N$ for $l = 1, \cdots, M$ such that $\kappa := \max_{l=1,\cdots,M} \|Y_l\|_\infty < \infty$. Let $\{\xi_l\}$ be a Rademacher sequence, i.e., a sequence of i.i.d. ± 1 random variables. Then for $p = 1$ and for $p \in \mathbb{R}$ and $p \geq 2$,*

$$\begin{aligned}
\left(\mathcal{E}\left[\left\|\sum_{l=1}^{M} \xi_l Y_l Y_l^H\right\|_K^p\right]\right)^{1/p} &\leq C' C^{1/p} \kappa \sqrt{p} \sqrt{K} \log(100K) \sqrt{\log(4N)\log(10M)} \\
&\quad \times \sqrt{\left\|\sum_{l=1}^{M} Y_l Y_l^H\right\|_K},
\end{aligned}$$

where C, C' are universal constants.

Lemma 6.2 (Adapted from Proposition 6.5 of (Rauhut, 2010)). *Suppose Z is a random*

variable satisfying

$$(\mathcal{E}\left[|Z|^p\right])^{1/p} \leq \alpha\beta^{1/p}p^{1/\gamma},$$

for all $p \in [p_0, \ p_1]$, and for constants $\alpha, \beta, \gamma, p_0, p_1$. Then, for all $u \in [p_0^{1/\gamma}, \ p_1^{1/\gamma}]$,

$$\mathcal{P}\left[|Z| \geq e^{1/\gamma}\alpha u\right] \leq \beta e^{-u^\gamma/\gamma}.$$

Armed with this notation and these lemmas, we now prove Theorem 3.1:

Proof. We seek to show that under the conditions on M in Theorem 3.1, $\mathcal{P}\left[\|\boldsymbol{B}\|_K > \delta\right] \leq \eta$. Since $\boldsymbol{B} = \boldsymbol{B}_1 + \boldsymbol{B}_2$ and $\{\|\boldsymbol{B}_1\|_K \leq \delta/2\} \cap \{\|\boldsymbol{B}_2\|_K \leq \delta/2\} \subset \{\|\boldsymbol{B}\|_K \leq \delta\}$, then,

$$\mathcal{P}\left[\|\boldsymbol{B}\|_K > \delta\right] \leq \mathcal{P}\left[\|\boldsymbol{B}_1\|_K > \delta/2\right] + \mathcal{P}\left[\|\boldsymbol{B}_2\|_K > \delta/2\right].$$

Thus, it will suffice to bound $\mathcal{P}\left[\|\boldsymbol{B}_1\|_K > \delta/2\right] \leq \eta/2$ since $\boldsymbol{B}_2 = \boldsymbol{B}_1^*$ implies that $\mathcal{P}\left[\|\boldsymbol{B}_2\|_K > \delta/2\right] \leq \eta/2$. In this presentation we let C, C' be some universal constant that may not be the same from line to line.

To begin, we use Lemma 6.1 to bound $E_p := (\mathcal{E}\left[\|\boldsymbol{B}_1\|_K^p\right])^{1/p}$ by setting $Y_l = \widetilde{\boldsymbol{z}}_l^* X_l$ for $l = 1, \cdots, \frac{M}{2}$. To meet the conditions of Lemma 6.1 we use a standard "symmetrization" manipulation (see Lemma 6.7 of (Rauhut, 2010)). Specifically, we can write:

$$
\begin{aligned}
E_p &= (\mathcal{E}\left[\|\boldsymbol{B}_1\|_K^p\right])^{1/p} \\
&\leq 2\left(\mathcal{E}\left[\left\|\sum_{l=1}^{M/2}\xi_l Y_l Y_l^H\right\|_K^p\right]\right)^{1/p} \\
&= 2\left(\mathcal{E}\left[\left\|\sum_{l=1}^{M/2}\xi_l|\widetilde{\boldsymbol{z}}_l|^2 X_l X_l^H\right\|_K^p\right]\right)^{1/p},
\end{aligned}
$$

where now the expectation is over the old random sequence $\{w_l\}$, together with a

newly added Rademacher sequence $\{\xi_l\}$. Applying the law of iterated expectation and Lemma 6.1, we have for $p \geq 2$:

$$
\begin{aligned}
E_p^p &:= \mathcal{E}\left[\|\boldsymbol{B}_1\|_K^p\right] && (13) \\
&\leq 2^p \mathcal{E}\left[\mathcal{E}\left[\left\|\sum_{l=1}^{M/2} \xi_l |\widetilde{\boldsymbol{z}}_l|^2 X_l X_l^H\right\|_K^p \mid \{w_l\}\right]\right] && (14) \\
&\leq \left(2C'C^{1/p}\sqrt{p}\kappa\sqrt{K}\log(100K)\sqrt{\log(4N)\log(5M)}\right)^p \mathcal{E}\left[\left\|\sum_{l=1}^{M/2} |\widetilde{\boldsymbol{z}}_l|^2 X_l X_l^H\right\|_K^{p/2}\right] \\
&\leq \left(C^{1/p}\sqrt{p}\kappa\sqrt{C'K\log^4(N)}\right)^p \mathcal{E}\left[\left(\left\|\sum_{l=1}^{M/2}\left(|\widetilde{\boldsymbol{z}}_l|^2 X_l X_l^H - \frac{1}{2}\boldsymbol{I}\right)\right\|_K + \frac{1}{2}\|I\|_K\right)^{p/2}\right] \\
&\leq \left(C^{1/p}\sqrt{p}\kappa\sqrt{C'K\log^4(N)}\right)^p \sqrt{\mathcal{E}\left[\left(\left\|\sum_{l=1}^{M/2}\left(|\widetilde{\boldsymbol{z}}_l|^2 X_l X_l^H - \frac{1}{2}\boldsymbol{I}\right)\right\|_K + \frac{1}{2}\right)^p\right]} \\
&\leq \left(C^{1/p}\sqrt{p}\kappa\sqrt{C'K\log^4(N)}\right)^p \sqrt{\left(E_p + \frac{1}{2}\right)^p}.
\end{aligned}
$$

In the first line above, the inner expectation is over the Rademacher sequence $\{\xi_l\}$ (where we apply Lemma 6.1) while the outer expectation is over the $\{w_l\}$. The third line uses the triangle inequality for the $\|\cdot\|_K$ norm, the fourth line uses Jansen's inequality, and the fifth line uses triangle inequality for moments norm (i.e., $(\mathcal{E}\left[|X+Y|^p\right])^{1/p} \leq (\mathcal{E}\left[|X|^p\right])^{1/p} + (\mathcal{E}\left[|Y|^p\right])^{1/p})$. To get to $\log^4 N$ in the third line, we used our assumption that $N \geq M$, $N \geq K$ and $N \geq O(1)$ in Theorem 3.1. Now using the definition of κ from Lemma 6.1, we can bound this quantity as:

$$
\kappa := \max_l \|Y_l\|_\infty = \max_l |\widetilde{\boldsymbol{z}}_l|\|X_l\|_\infty = \frac{1}{\sqrt{M}}\max_l \|X_l\|_\infty = \frac{1}{\sqrt{M}}\max_{l,n}|\langle \boldsymbol{F}_l, \boldsymbol{\Psi}_n\rangle| \leq \frac{\mu(\boldsymbol{\Psi})}{\sqrt{M}}.
$$

Therefore, we have the following implicit bound on the moments of the random variable

of interest

$$E_p \leq C^{1/p}\sqrt{p}\sqrt{\frac{C'K\mu(\boldsymbol{\Psi})^2\log^4(N)}{M}}\sqrt{E_p + \frac{1}{2}}.$$

The above can be written as $E_p \leq a_p\sqrt{E_p + \frac{1}{2}}$, where $a_p = C^{1/p}\sqrt{p}\sqrt{\frac{4C'K\mu(\boldsymbol{\Psi})^2\log^4(N)}{M}}$. By squaring, rearranging the terms and completing the square, we have $E_p \leq \frac{a_p^2}{2} + a_p\sqrt{\frac{1}{2} + \frac{a_p^2}{4}}$. By assuming $a_p \leq \frac{1}{2}$, this bound can be simplified to $E_p \leq a_p$. Now, this assumption is equivalent to having an upper bound on the range of values of p:

$$a_p \leq \frac{1}{2} \quad \Leftrightarrow \quad \sqrt{p} \leq \frac{1}{2C^{1/p}}\sqrt{\frac{M}{4C'K\mu(\boldsymbol{\Psi})^2\log^4(N)}}$$

$$\Leftrightarrow \quad p \leq \frac{M}{16C^{2/p}C'K\mu(\boldsymbol{\Psi})^2\log^4(N)}.$$

Hence, by using Lemma 6.2 with $\alpha = \sqrt{\frac{C'K\mu(\boldsymbol{\Psi})^2\log^4(N)}{M}}$, $\beta = C$, $\gamma = 2$, $p_0 = 2$, and $p_1 = \frac{M}{16C^{2/p}C'K\mu(\boldsymbol{\Psi})^2\log^4(N)}$ we obtain the following tail bound for $u \in [\sqrt{2}, \sqrt{p_1}]$:

$$\mathcal{P}\left[\|\boldsymbol{B}_1\|_K \geq e^{1/2}\sqrt{\frac{C'K\mu(\boldsymbol{\Psi})^2\log^4(N)}{M}}u\right] \leq Ce^{-u^2/2}.$$

If we pick $\delta < 1$ such that

$$e^{1/2}\sqrt{\frac{C'K\mu(\boldsymbol{\Psi})^2\log^4(N)}{M}}u \leq \frac{\delta}{2} \tag{15}$$

and u such that

$$Ce^{-u^2/2} \leq \frac{\eta}{2} \quad \Leftrightarrow \quad u \geq \sqrt{2\log(2C\eta^{-1})},$$

then we have our required tail bound of $\mathcal{P}\left[\|\boldsymbol{B}_1\|_K > \delta\right] \leq \eta/2$. First, observe that

31

Equation (15) is equivalent to having

$$M \geq \frac{CK\mu(\mathbf{\Psi})^2 \log^4(N) \log(\eta^{-1})}{\delta^2}.$$

Also, because of the limited range of values u can take (i.e., $u \in [\sqrt{2}, \sqrt{p_1}]$), we require that

$$\begin{aligned}
\sqrt{2\log(2C\eta^{-1})} &\leq \sqrt{\frac{M}{16C^{2/p}C'K\mu(\mathbf{\Psi})^2 \log^4(N)}} = \sqrt{p_1} \\
&\Leftrightarrow M \geq CK\mu(\mathbf{\Psi})^2 \log^4(N) \log(\eta^{-1}),
\end{aligned}$$

which, together with the earlier condition on M, completes the proof.

$\square$

6.2 RIP with Gaussian feed-forward vectors

In this appendix we extend the RIP analysis of Appendix 6.1 to the case when $\mathbf{z}$ is chosen to be a Gaussian i.i.d. vector, as presented in Theorem 4.1. It is unfortunate that with the additional randomness in the feed-forward vector, the same proof procedure as in Theorem 3.1 cannot be used. In the proof of Theorem 3.1, we showed that the random variable $\|Z_1\|_K$ has p-th moments that scale like $\alpha\beta^{1/p}p^{1/2}$ (through Lemma 6.1) for a range of p which suggests that it has a sub-gaussian tail (i.e., $\mathcal{P}\left[\|Z_1\|_K > u\right] \leq Ce^{-u^2/2}$) for a range of deviations u. We then used this tail bound to bound the probability that $\|Z_1\|_K$ exceeds a fixed conditioning δ. With Gaussian uncertainties in the feed-forward vector z, Lemma 6.1 will not yield the required sub-gaussian tail but instead gives us moments estimates that result in sub-optimal scaling of M with respect to N. Therefore, we will instead follow the proof procedure of Theorem 16 from (Tropp et al., 2009) that will yield the better measurement rate given in Theorem 4.1.

Let us begin by recalling a few notations from the proof of Theorem 3.1 and by introducing further notations that will simplify our exposition later. First, recall that we

let X_l^H be the l-th row of $\boldsymbol{F\Psi}$. Thus, the l-th row of our matrix of interest $\widehat{\boldsymbol{A}} = \widetilde{\boldsymbol{Z}}\boldsymbol{F\Psi}$ is $\widetilde{z}_l X_l^H$ where $\widetilde{z}_l$ is the l-th diagonal entry of the diagonal matrix $\widetilde{\boldsymbol{Z}}$. Whereas before, $\widetilde{z}_l = \frac{1}{\sqrt{M}}$ for any $l = 1, \cdots, M$, here it will be a random variable. To understand the resulting distribution of $\widetilde{z}_l$, first note that for the connectivity matrix $\boldsymbol{W}$ to be real, we need to assume that the second $\frac{M}{2}$ columns of $\boldsymbol{U}$ are complex conjugates of the first $\frac{M}{2}$ columns. Thus, we can write $\boldsymbol{U} = [\boldsymbol{U}_R \mid \boldsymbol{U}_R] + j\,[\boldsymbol{U}_I \mid -\boldsymbol{U}_I]$, where $\boldsymbol{U}_R, \boldsymbol{U}_I \in \mathbb{R}^{M \times \frac{M}{2}}$. Because $\boldsymbol{U}^H\boldsymbol{U} = \boldsymbol{I}$, we can deduce that $\boldsymbol{U}_R^T\boldsymbol{U}_I = \boldsymbol{0}$ and that the ℓ_2 norms of the columns of both $\boldsymbol{U}_R$ and $\boldsymbol{U}_I$ are $\frac{1}{\sqrt{2}}$.[4]

With these matrices $\boldsymbol{U}_R, \boldsymbol{U}_I$, let us re-write the random vector $\widetilde{z}$ to illustrate its structure. Consider the matrix $\widehat{\boldsymbol{U}} := [\boldsymbol{U}_R \mid \boldsymbol{U}_I] \in \mathbb{R}^{M \times M}$, which is a scaled unitary matrix (because we can check that $\widehat{\boldsymbol{U}}^T\widehat{\boldsymbol{U}} = \frac{1}{2}\boldsymbol{I}$). Next, consider the random vector $\widehat{z} := \widehat{\boldsymbol{U}}^T z$. Because $\widehat{\boldsymbol{U}}$ is (scaled) unitary and z is composed of i.i.d. zero-mean Gaussian random variables of variance $\frac{1}{M}$, the entries of $\widehat{z}$ are also i.i.d. zero-mean Gaussian random variables, but now with variance $\frac{1}{2M}$. Then, from our definition of $\boldsymbol{U}$ in terms of $\boldsymbol{U}_R$ and $\boldsymbol{U}_I$, for any $l \leq \frac{M}{2}$, we have $\widetilde{z}_l = \widehat{z}_l - j\widehat{z}_{l+\frac{M}{2}}$ and for $l > \frac{M}{2}$, we have $\widetilde{z}_l = \widehat{z}_{l-\frac{M}{2}} + j\widehat{z}_l$. This clearly shows that each of the *first* $\frac{M}{2}$ entries of $\widetilde{z}$ is made up of 2 i.i.d. random variables (one being the real component, the other imaginary), and that the other $\frac{M}{2}$ entries are just complex conjugates of the first $\frac{M}{2}$. Because of this, for $l \leq \frac{M}{2}$, $|\widetilde{z}_l|^2 = |\widetilde{z}_{l+\frac{M}{2}}|^2 = \widehat{z}_l^2 + \widehat{z}_{l+\frac{M}{2}}^2$ is the sum of squares of 2 i.i.d. Gaussian random variables.

From the proof of Theorem 3.1, we also denoted

$$Z := \widehat{\boldsymbol{A}}^H\widehat{\boldsymbol{A}} - \boldsymbol{I} = \left(\sum_{l=1}^{M/2} |\widetilde{z}_l|^2 X_l X_l^H - \frac{1}{2}\boldsymbol{I}\right) + \left(\sum_{l=1}^{M/2} |\widetilde{z}_l|^2 \left(X_l X_l^H\right)^* - \frac{1}{2}\boldsymbol{I}\right) =: Z_1 + Z_2.$$

---

[4] This can be shown by writing

$$\boldsymbol{U}^H\boldsymbol{U} = \left(\begin{bmatrix} \boldsymbol{U}_R^T \\ \boldsymbol{U}_R^T \end{bmatrix} - j\begin{bmatrix} \boldsymbol{U}_I^T \\ -\boldsymbol{U}_I^T \end{bmatrix}\right)([\boldsymbol{U}_R \mid \boldsymbol{U}_R] + j\,[\boldsymbol{U}_I \mid -\boldsymbol{U}_I])$$
$$= \left(\begin{bmatrix} \boldsymbol{U}_R^T\boldsymbol{U}_R & \boldsymbol{U}_R^T\boldsymbol{U}_R \\ \boldsymbol{U}_R^T\boldsymbol{U}_R & \boldsymbol{U}_R^T\boldsymbol{U}_R \end{bmatrix} + \begin{bmatrix} \boldsymbol{U}_I^T\boldsymbol{U}_I & -\boldsymbol{U}_I^T\boldsymbol{U}_I \\ -\boldsymbol{U}_I^T\boldsymbol{U}_I & \boldsymbol{U}_I^T\boldsymbol{U}_I \end{bmatrix}\right) + j\left(\begin{bmatrix} \boldsymbol{U}_R^T\boldsymbol{U}_I & -\boldsymbol{U}_R^T\boldsymbol{U}_I \\ \boldsymbol{U}_R^T\boldsymbol{U}_I & -\boldsymbol{U}_R^T\boldsymbol{U}_I \end{bmatrix} + \begin{bmatrix} \boldsymbol{U}_I^T\boldsymbol{U}_R & \boldsymbol{U}_I^T\boldsymbol{U}_R \\ -\boldsymbol{U}_I^T\boldsymbol{U}_R & -\boldsymbol{U}_I^T\boldsymbol{U}_R \end{bmatrix}\right).$$

Then by equating the above to $I + j0$, we arrive at our conclusion.

It is again easy to check that $\mathcal{E}\left[\sum_{l=1}^{M/2}\left(|\widetilde{z}_l|^2 X_l X_l^H\right)\right] = \mathcal{E}\left[\sum_{l=1}^{M/2}\left(|\widetilde{z}_l|^2 \left(X_l X_l^H\right)^*\right)\right] =$ $\frac{1}{2}I$. Finally, $\widehat{A}$ has RIP conditioning δ whenever $\|Z\|_K \leq \delta$ with $\|Z\|_K := \sup_{y \text{ is } K\text{-sparse}} \frac{y^H Z y}{y^H y}$.

Before moving on to the proof, we first present a lemma regarding the random sequence $|z_l|^2$ that will be useful in the sequel.

Lemma 6.3. *Suppose for* $l = 1, \cdots, \frac{M}{2}$, $|\widetilde{z}_l|^2 = \widehat{z}_l^2 + \widehat{z}_{l+M/2}^2$ *where* $\widehat{z}_l$ *for* $l = 1, \cdots, M$ *is a sequence of i.i.d. zero-mean Gaussian random variables of variance* $\frac{1}{2M}$. *Also suppose that* $\eta \leq 1$ *is a fixed probability. For the random variable* $\max_{l=1,\cdots,M/2} |\widetilde{z}_l|^2$, *we have the following bounds on the expected value and tail probability of this extreme value:*

$$\mathcal{E}\left[\max_{l=1,\cdots,M/2} |\widetilde{z}_l|^2\right] \leq \frac{1}{M}\left(\log\left(\frac{C_1 M}{2}\right) + 1\right), \quad (16)$$

$$\mathcal{P}\left[\max_{l=1,\cdots,M/2} |\widetilde{z}_l|^2 > \frac{C_2 \log\left(C_2' M \eta^{-1}\right)}{M}\right] \leq \eta. \quad (17)$$

Proof. To ease notation, every index l used as a variable for a maximization will be taken over the set $l = 1, \ldots, \frac{M}{2}$ without explicitly writing the index set. To calculate $\mathcal{E}\left[\max_l |\widetilde{z}_l|^2\right]$, we use the following result that allows us to bound the expected value of a positive random variable by its tail probability (see Proposition 6.1 of [Rauhut]):

$$\mathcal{E}\left[\max_l |\widetilde{z}_l|^2\right] = \int_0^\infty \mathcal{P}\left[\max_l |\widetilde{z}_l|^2 > u\right] du. \quad (18)$$

Using the union bound, we have the estimate $\mathcal{P}\left[\max_l |\widetilde{z}_l|^2 > u\right] \leq \frac{M}{2}\mathcal{P}\left[|\widetilde{z}_1|^2 > u\right]$ (since the $|\widetilde{z}_l|^2$ are identically distributed). Now, because $|\widetilde{z}_1|^2$ is a sum of squares of two Gaussian random variables and thus is a (generalized) χ^2 random variable with 2 degrees of freedom (which we shall denote by χ_2),[5] we have

$$\mathcal{P}\left[|\widetilde{z}_1|^2 > u\right] = \mathcal{P}\left[\chi_2 > 2Mu\right] = \frac{1}{\Gamma(1)}e^{\frac{-2Mu}{2}} = C_1 e^{-Mu},$$

---

[5] The pdf of a χ^2 random variable χ_q with q degrees of freedom is given by $p(x) = \frac{1}{2^{q/2}\Gamma(q/2)}x^{q/2-1}e^{-x/2}$. Therefore, it's tail probability can be obtained by integration: $\mathcal{P}\left[\chi_q > u\right] = \int_u^\infty p(x)dx$.

where $\Gamma(\cdot)$ is the Gamma function and the $2Mu$ appears instead of u in the exponential because of the standardization of the Gaussian random variables (initially of variance $\frac{1}{2M}$). To proceed, we break the integral in (18) into 2 parts. To do so, notice that if $u < \frac{1}{M}\log\left(\frac{C_1 M}{2}\right)$, then the trivial upper bound of $\mathcal{P}\left[\max_l |\widetilde{z}_l|^2 > u\right] \leq 1$ is a better estimate than $\frac{C_1 M}{2} e^{-Mu}$. In other words, our estimate for the tail bound of $\max_l |\widetilde{z}_l|^2$ is not very good for small u but gets better with increasing u. Therefore, we have

$$
\begin{aligned}
\mathcal{E}\left[\max_l |\widetilde{z}_l|^2\right] &\leq \int_0^{\frac{1}{M}\log\left(\frac{C_1 M}{2}\right)} 1 \, du + \int_{\frac{1}{M}\log\left(\frac{C_1 M}{2}\right)}^{\infty} \frac{C_1 M}{2} e^{-Mu} \, du \\
&= \frac{1}{M}\log\left(\frac{C_1 M}{2}\right) - \frac{C_1 M}{2}\left[\frac{1}{M} e^{-Mu}\right]\Big|_{\frac{1}{M}\log\left(\frac{C_1 M}{2}\right)}^{\infty} \\
&= \frac{1}{M}\log\left(\frac{C_1 M}{2}\right) + \frac{C_1}{2} e^{-\log\left(\frac{C_1 M}{2}\right)} = \frac{1}{M}\left(\log\left(\frac{C_1 M}{2}\right) + 1\right).
\end{aligned}
$$

This is the bound in expectation that we seek for in Equation (18).

In the second part of the proof that follows, C, C' denote universal constants. Essentially, we will want to apply Lemma 6.2 that is used in Appendix 6.1 to obtain our tail bound. In the lemma, the tail bound of a random variable X can be estimated once we know the moments of X. Therefore, we require the moments of the random variable $\max_l |\widetilde{z}_l|^2$. For this, for any $p > 0$, we use the following simple estimate:

$$
\mathcal{E}\left[\max_l |\widetilde{z}_l|^{2p}\right] \leq \frac{M}{2} \max_l \mathcal{E}\left[|\widetilde{z}_l|^{2p}\right] = \frac{M}{2}\mathcal{E}\left[|\widetilde{z}_1|^{2p}\right], \tag{19}
$$

where the first step comes from writing the expectation as an integral of the cumulative distribution (as seen in Equation (18)) and taking the union bound, and the second step comes from the fact that the $|\widetilde{z}_l|^2$ are identically distributed. Now, $|\widetilde{z}_1|^2$ is a sub-exponential random variable since it is a sum of squares of Gaussian random variables (Vershynin, 2012).[6] Therefore, for any $p > 0$, it's p-th moment can be bounded

---

[6] A sub-exponential random variable is a random variable whose tail probability is bounded by $\exp^{-Cu}$ for some constant C. Thus, a χ^2 random variable is a specific instance of a sub-exponential random variable.

by

$$\mathcal{E}\left[|\widetilde{z}_1|^{2p}\right]^{1/p} \leq \frac{C'}{M}C^{1/p}p,$$

where the division by M comes again from the variance of the Gaussian random variables that make up $|\widetilde{z}_1|^2$. Putting this bound with Equation (19), we have the following estimate for the p-th moments of $\max_l |\widetilde{z}_l|^2$:[7]

$$\mathcal{E}\left[\max_l |\widetilde{z}_l|^{2p}\right]^{1/p} \leq \frac{C'}{M}\left(\frac{CM}{2}\right)^{1/p} p.$$

Therefore, by Lemma 6.2 with $\alpha = \frac{C'}{M}$, $\beta = \frac{CM}{2}$, and $\gamma = 1$, we have

$$\mathcal{P}\left[\max_l |\widetilde{z}_l|^2 > \frac{eC'u}{M}\right] \leq \frac{CM}{2}e^{-u}.$$

By choosing $u = \log\left(\frac{CM}{2}\eta^{-1}\right)$, we have our desired tail bound of

$$\mathcal{P}\left[\max_l |\widetilde{z}_l|^2 > \frac{C_2 \log\left(C_2'M\eta^{-1}\right)}{M}\right] \leq \eta.$$

□

Armed with this lemma, we can now turn out attention to the main proof. As stated earlier, this follows essentially the same form as (Tropp et al., 2009) with the primary difference of including the results from Lemma 6.3. As before, because $\mathcal{P}\left[\|Z\|_K > \delta\right] \leq \mathcal{P}\left[\|Z_1\|_K > \delta/2\right] + \mathcal{P}\left[\|Z_2\|_K > \delta/2\right]$ with $Z_2 = Z_1^*$, we just have to consider bounding the tail bound $\mathcal{P}\left[\|Z_1\|_K > \delta/2\right]$. This proof differs from that in Appendix 6.1 in that here, we will first show that $\mathcal{E}\left[\|Z_1\|_K\right]$ is small when M is large enough and then show that Z_1 does not differ much from $\mathcal{E}\left[\|Z_1\|_K\right]$ with high probability.

---

[7] We remark that this bound gives a worse estimate for the expected value as that calculated before because of the crude bound given by Equation (19).

Expectation

In this section, we will show that $\mathcal{E}\left[\|Z_1\|_K\right]$ is small. This will basically follow from Lemma 6.1 in Appendix 6.1 and Equation (16) in Lemma 6.3. To be precise, the remainder of this section is to prove:

Theorem 6.1. *Choose any* $\delta' \leq \frac{1}{2}$. *If* $M \geq \frac{C_3 K \mu(\Psi)^2 \log^5 N}{\delta'^2}$, *then* $\mathcal{E}\left[\|Z\|_K\right] \leq \delta'$.

Proof. Again, C is some universal constant that may not be the same from line to line. We follow the same symmetrization step found in the proof in Appendix 6.1 to arrive at:

$$E := \mathcal{E}\left[\|Z_1\|_K\right] \quad \leq \quad 2\mathcal{E}\left[\mathcal{E}\left[\left\|\sum_{l=1}^{M/2} \xi_l |\widetilde{z}_l|^2 X_l X_l^H\right\|_K \Big| \{w_l\}, \widetilde{z}\right]\right],$$

where the outer expectation is over the Rademacher sequence $\{\xi_l\}$ and the inner expectation is over the random "frequencies" $\{w_l\}$ and feed-forward vector $\widetilde{z}$. As before, for $l = 1, \cdots, \frac{M}{2}$, we set $Y_l = \widetilde{z}_l^* X_l$. Observe that by definition $\kappa := \max_{l=1,\cdots,M/2} \|Y_l\|_\infty = \max_l |\widetilde{z}_l| \|X_l\|_\infty$ and thus is a random variable. We then use Lemma 6.1 with $p = 1$ to get

$$
\begin{aligned}
E &\leq 2C\sqrt{K}\log(100K)\sqrt{\log(4N)\log(5M)}\mathcal{E}\left[\kappa\sqrt{\left\|\sum_{l=1}^{M/2}|\widetilde{z}_l|^2 X_l X_l^H\right\|_K}\right] \\
&\leq \sqrt{4CK\log^4(N)}\sqrt{\mathcal{E}\left[\kappa^2\right]}\sqrt{\mathcal{E}\left[\left\|\sum_{l=1}^{M/2}|\widetilde{z}_l|^2 X_l X_l^H\right\|_K\right]} \\
&\leq \sqrt{4CK\log^4(N)}\sqrt{\mathcal{E}\left[\kappa^2\right]}\sqrt{E + \frac{1}{2}}, \quad\quad (20)
\end{aligned}
$$

where the second line uses the Cauchy-Schwarz inequality for expectations and the third line uses triangle inequality. Again, to get to $\log^4 N$ in the second line, we used our assumption that $N \geq M$, $N \geq K$ and $N \geq O(1)$ in Theorem 4.1. It therefore remains to calculate $\mathcal{E}\left[\kappa^2\right]$. Now, $\kappa = \max_l |\widetilde{z}_l| \|X_l\|_\infty \leq \max_l |\widetilde{z}_l| \max_l \|X_l\|_\infty$.

First, we have $\max_l \|X_l\|_\infty = \max_{l,n} |\langle \boldsymbol{F}_l, \boldsymbol{\Psi}_n \rangle| \leq \mu(\boldsymbol{\Psi})$. Next, (16) in Lemma 6.3 tells us that $\mathcal{E}\left[\max_{l=1,\cdots,M/2} |\widetilde{z}_l|^2\right] \leq \frac{1}{M}\left(\log\left(\frac{C_1 M}{2}\right) + 1\right)$. Thus, we have $\mathcal{E}\left[\kappa^2\right] \leq \frac{\mu(\boldsymbol{\Psi})^2}{M}\left(\log\left(\frac{C_1 M}{2}\right) + 1\right)$. Putting everything together, we have

$$E = \mathcal{E}\left[\|Z_1\|_K\right] \leq \sqrt{\frac{CK \log^4(N)\left(\log\left(\frac{C_1 M}{2}\right) + 1\right)\mu(\boldsymbol{\Psi})^2}{M}}\sqrt{E + \frac{1}{2}}.$$

Now, the above can be written as $E \leq a\sqrt{E + \frac{1}{2}}$, where $a = \sqrt{\frac{CK \log^4(N)\left(\log\left(\frac{C_1 M}{2}\right)+1\right)\mu(\boldsymbol{\Psi})^2}{M}}$. By squaring it, rearranging the terms and completing the squares, we have $E \leq \frac{a^2}{2} + a\sqrt{\frac{1}{2} + \frac{a^2}{4}}$. By supposing $a \leq \frac{1}{2}$, this can be simplified as $E \leq a$. To conclude, let us choose M such that $a \leq \delta'$ where $\delta' \leq \frac{1}{2}$ is our pre-determined conditioning (which incidentally fulfills our previous assumption that $a \leq \frac{1}{2}$). By applying the formula for a, we have that if $M \geq \frac{C_3 K \mu(\boldsymbol{\Psi})^2 \log^5(N)}{\delta'^2}$, then $E \leq \delta'$. $\qquad\square$

Tail Probability

To give a probability tail bound estimate to Z_1, we use the following lemma found in (Tropp et al., 2009; Rauhut, 2010):

Lemma 6.4. *Suppose Y_l for $l = 1, \cdots, M$ are independent, symmetric random variables such that $\|Y_l\|_K \leq \zeta < \infty$ almost surely. Let $Y = \sum_{l=1}^M Y_l$. Then for any $u, t > 1$, we have*

$$\mathcal{P}\left[\|Y\|_K > C(u\mathcal{E}\left[\|Y\|_K\right] + t\zeta)\right] \leq e^{-u^2} + e^{-t}.$$

The goal of this section is to prove:

Theorem 6.2. *Pick any $\delta \leq \frac{1}{2}$ and suppose $N^{-\log^4(N)} \leq \eta \leq \frac{1}{e}$. Suppose $M \geq \frac{C_4 K \mu(\boldsymbol{\Psi})^2 \log^5 N \log \eta^{-1}}{\delta^2}$, then $\mathcal{P}\left[\|Z_1\|_K > \delta\right] \leq 8\eta$.*

Proof. To use Lemma 6.4, we want Y_l to look like the summands of

$$Z_1 = \sum_{l=1}^{M/2} \left(|\widetilde{z}_l|^2 X_l X_l^H - \mathcal{E}\left[|\widetilde{z}_l|^2 X_l X_l^H\right] \right).$$

However, this poses several problems. First, they are *not* symmetric[8] and thus, we need to symmetrize it by defining

$$
\begin{aligned}
\widetilde{Y}_l &= |\widetilde{z}_l|^2 X_l X_l^H - |\widetilde{z}_l'|^2 X_l'(X_l')^H \\
&\sim \xi_l \left(|\widetilde{z}_l|^2 X_l X_l^H - |\widetilde{z}_l'|^2 X_l'(X_l')^H \right)
\end{aligned}
$$

where $\widetilde{z}'$, X_l' are independent copies of $\widetilde{z}$ and X_l respectively, and ξ_l is an independent Rademacher sequence. Here, the relation $X \sim Y$ for two random variables X, Y means that X has the same distribution as Y. To form $\widetilde{Y}_l$, what we have done is take each summand of Z_1 and take it's difference with an independent copy of itself. Because $\widetilde{Y}_l$ is symmetric, adding a Rademacher sequence does not change its distribution and this sequence is only introduced to resolve a technicality that will arise later on. If we let $\widetilde{Y} := \sum_{l=1}^{M/2} \widetilde{Y}_l$, then the random variables $\widetilde{Y}$ (symmetrized) and Z_1 (un-symmetrized) are related via the following estimates (Rauhut, 2010):

$$\mathcal{E}\left[\|\widetilde{Y}\|_K\right] \leq 2\mathcal{E}\left[\|Z_1\|_K\right], \tag{21}$$

$$\mathcal{P}\left[\|Z_1\|_K > 2\mathcal{E}\left[\|Z_1\|_K\right] + u\right] \leq 2\mathcal{P}\left[\|\widetilde{Y}\|_K > u\right]. \tag{22}$$

However, a second condition imposed on Y_l in Lemma 6.4 is that $\|Y_l\|_K \leq \zeta < \infty$ almost surely. Because of the unbounded nature of the Gaussian random variables $\widetilde{z}_l$ and $\widetilde{z}_l'$ in $\widetilde{Y}_l$, this condition is not met. Therefore, we need to define a Y_l that is conditioned on the event that these Gaussian random variables are bounded. To do so,

---

[8]A random variable X is symmetric if X and $-X$ has the same distribution.

define the following event:

$$F = \left\{ \max \left\{ \max_l |\widetilde{z}_l|^2, \ \max_l |\widetilde{z}_l'|^2 \right\} \leq \frac{C_2 \log \left(C_2' M \eta^{-1} \right)}{M} \right\}.$$

Using Equation (17) in Lemma 6.3, we can calculate $\mathcal{P}\left[F^c \right]$, where F^c is the complementary event of F:

$$
\begin{aligned}
\mathcal{P}\left[F^c \right] &= \mathcal{P}\left[\max \left\{ \max_l |\widetilde{z}_l|^2, \ \max_l |\widetilde{z}_l'|^2 \right\} > \frac{C_2 \log \left(C_2' M \eta^{-1} \right)}{M} \right] \\
&\leq \mathcal{P}\left[\max_l |\widetilde{z}_l|^2 > \frac{C_2 \log \left(C_2' M \eta^{-1} \right)}{M} \right] + \mathcal{P}\left[\max_l |\widetilde{z}_l'|^2 > \frac{C_2 \log \left(C_2' M \eta^{-1} \right)}{M} \right] \\
&\leq 2\eta.
\end{aligned}
$$

Conditioned on event F, the $\| \cdot \|_K$ norm of $\widetilde{Y}_l$ is well-bounded:

$$
\begin{aligned}
\left\| \widetilde{Y}_l \right\|_K &= \left\| |\widetilde{z}_l|^2 X_l X_l^H - |\widetilde{z}_l'|^2 X_l'(X_l')^H \right\|_K \leq 2 \max \left\{ \max_l |\widetilde{z}_l|^2, \ \max_l |\widetilde{z}_l'|^2 \right\} \left\| X_l X_l^H \right\|_K \\
&= \frac{2 C_2 \log \left(C_2' M \eta^{-1} \right)}{M} \sup_{y \text{ is } K\text{-sparse}} \left\{ \frac{y^H X_l X_l^H y}{y^H y} \right\} \\
&\leq \frac{2 C_2 \log \left(C_2' M \eta^{-1} \right)}{M} \sup_{y \text{ is } K\text{-sparse}} \left\{ \|X_l\|_\infty^2 \frac{\|y\|_1^2}{\|y\|_2^2} \right\} \\
&\leq \frac{2 K C_2 \log \left(C_2' M \eta^{-1} \right)}{M} \max_l \|X_l\|_\infty^2 \leq \frac{C K \mu(\mathbf{\Psi})^2 \log \left(C_2' M \eta^{-1} \right)}{M} := \zeta,
\end{aligned}
$$

where in the last line we used the fact that the ratio between the ℓ_1 and ℓ_2 norms of an K-sparse vector is K, and the estimate we derived for $\max_l \|X_l\|_\infty^2$ in Appendix 6.1.

We now define a new random variable that is a truncated version of $\widetilde{Y}_l$ which takes for value 0 whenever we fall under event F^c, i.e.,

$$Y_l := \widetilde{Y}_l \, \mathbb{I}_F = \xi_l \left(|\widetilde{z}_l|^2 X_l X_l^H - |\widetilde{z}_l'|^2 X_l'(X_l')^H \right) \mathbb{I}_F,$$

where $\mathbb{I}_F$ is the indicator function of event F_l. If we define $Y = \sum_{l=1}^{M/2} Y_l$, then the random variables Y (truncated) and $\widetilde{Y}$ (un-truncated) are related by (Tropp et al., 2009)

(see also Lemma 1.4.3 of (De La Peña & Giné, 1999))

$$\mathcal{P}\left[\|\widetilde{Y}\|_K > u\right] \leq \mathcal{P}\left[\|Y\|_K > u\right] + \mathcal{P}\left[F^c\right]. \tag{23}$$

When $\widetilde{z}, \widetilde{z}', X_l, X_l'$ are held constant so only the Rademacher sequence ξ_l is random, then the contraction principle (Tropp et al., 2009; Ledoux & Talagrand, 1991) tells us that $\mathcal{E}\left[\|Y\|_K\right] \leq \mathcal{E}\left[\|\widetilde{Y}\|_K\right]$. Note that the sole reason for introducing the Rademacher sequences is for this use of the contraction principle. As this holds point-wise for all $\widetilde{z}, \widetilde{z}', X_l, X_l'$, we have

$$\mathcal{E}\left[\|Y\|_K\right] \leq \mathcal{E}\left[\|\widetilde{Y}\|_K\right]. \tag{24}$$

We now have all the necessary ingredients to apply Lemma 6.4. First, by choosing $\delta' \leq \frac{1}{2}$, from Theorem 6.1, we have that $\mathcal{E}\left[\|Z\|_K\right] \leq \delta'$ whenever $M \geq \frac{C_3 K \mu(\mathbf{\Psi})^2 \log^5 N}{\delta'^2}$. Thus, by chaining (24) and (21), we have

$$\mathcal{E}\left[\|Y\|_K\right] \leq \mathcal{E}\left[\|\widetilde{Y}\|_K\right] \leq 2\mathcal{E}\left[\|Z_1\|_K\right] \leq 2\delta'.$$

Also, with this choice of M, we have

$$\zeta = \frac{CK\mu(\mathbf{\Psi})^2 \log\left(C_2' M \eta^{-1}\right)}{M} \leq \frac{C\delta'^2 \log\left(C_2' M \eta^{-1}\right)}{\log^5 N}.$$

Using these estimates for ζ and $\mathcal{E}\left[\|Y\|_K\right]$, and choosing $u = \sqrt{\log \eta^{-1}}$ and $t = \log \eta^{-1}$, Lemma 6.4 says that

$$\mathcal{P}\left[\|Y\|_K > C'\left(2\delta'\sqrt{\log \eta^{-1}} + \frac{C\delta'^2 \log\left(C_2' M \eta^{-1}\right) \log \eta^{-1}}{\log^5 N}\right)\right] \leq 2\eta.$$

Then, using the relation between the tail probabilities of Y and $\widetilde{Y}$ (23) together with

our estimate for $\mathcal{P}\left[F^c\right]$, we have

$$\mathcal{P}\left[\|\widetilde{Y}\|_K > C'\left(2\delta'\sqrt{\log\eta^{-1}} + \frac{C\delta'^2\log\left(C_2'M\eta^{-1}\right)\log\eta^{-1}}{\log^5 N}\right)\right] \leq 2\eta + \mathcal{P}\left[F^c\right] \leq 4\eta.$$

Finally, using the relation between the tail probabilities of $\widetilde{Y}$ and Z (22), we have

$$\mathcal{P}\left[\|Z_1\|_K > 2\delta' + 2C'\delta'\sqrt{\log\eta^{-1}} + \frac{CC'\delta'^2\log\left(C_2'M\eta^{-1}\right)\log\eta^{-1}}{\log^5 N}\right] \leq 8\eta,$$

where we used the fact that $\mathcal{E}\left[\|Z_1\|_K\right] \leq \delta'$. Then, for a pre-determined conditioning $\delta \leq \frac{1}{2}$, pick $\delta' = \frac{\delta}{3C''\sqrt{\log\eta^{-1}}}$ for a constant C'' which will be chosen appropriately later. With this choice of δ' and with our assumptions that $\delta \leq \frac{1}{2}$ and $\eta \leq \frac{1}{e}$, the three terms in the tail bound becomes

$$
\begin{aligned}
2\delta' &= \frac{\delta}{3C''\sqrt{\log\eta^{-1}}} \leq \frac{1}{C''}\frac{\delta}{3}, \\
2C'\delta'\sqrt{\log\eta^{-1}} &= \frac{2C'}{C''}\frac{\delta}{3}, \\
\frac{CC'\delta'^2\log\left(C_2'M\eta^{-1}\right)\log\eta^{-1}}{\log^5 N} &= \frac{CC'\delta^2(\log(C_2'M) + \log\eta^{-1})}{9\left(C''\right)^2\log^5 N} \\
&\leq \frac{CC'(\log(C_2'M) + \log\eta^{-1})}{3\left(C''\right)^2\log^5 N}\frac{\delta}{3}.
\end{aligned}
$$

As for the last term, if $\eta \geq \frac{1}{C_2'M}$, then $\frac{CC'(\log(C_2'M)+\log\eta^{-1})}{3(C'')^2\log^5 N} \leq \frac{2CC'\log(C_2'M)}{3(C'')^2\log^5 N} \leq \frac{2CC'}{3(C'')^2}$ (where we further supposed that $N \geq O(1)$). If $N^{-\log^4 N} \leq \eta \leq \frac{1}{C_2'M}$ (where the lower bound is from the theorem assumptions), then $\frac{CC'(\log(C_2'M)+\log\eta^{-1})}{3(C'')^2\log^5 N} \leq \frac{2CC'\log\eta^{-1}}{3(C'')^2\log^5 N} \leq \frac{2CC'}{3(C'')^2}$. By choosing C'' appropriately large, we then have

$$\mathcal{P}\left[\|Z_1\|_K > \frac{\delta}{3} + \frac{\delta}{3} + \frac{\delta}{3}\right] \leq 8\eta.$$

Putting the formula for δ' into $M \geq \frac{C_3 K\mu(\boldsymbol{\Psi})^2\log^5 N}{\delta'^2}$ completes the proof. $\qquad\square$

6.3   Derivation of recovery bound for infinite length inputs

In this appendix we derive the bound in Equation (11) of the main text. The approach we take is to bound the individual components of Equation (4) of the main text. As the noise term due to noise in the inputs is unaffected, we will bound the noise term due to the unrecovered signal (the first term in Equation (4) of the main text) by the component of the input history that is beyond the attempted recovery, and we will bound the signal approximation term (the second term in Equation (4) of the main text) by the quality of the signal recovery possible in the attempted recovery length. In this way we can observe how different properties of the system and input sequence affect signal recovery.

To bound the first term in Equation (4) of the main text (i.e., the omission errors due to inputs beyond the recovery window), we first write the current state at any time N^* as

$$\boldsymbol{x}[N^*] = \sum_{n=0}^{N^*} \boldsymbol{W}^{N^*-n} \boldsymbol{z} s[n].$$

We only wish to recover the past $N \leq N^*$ time steps, so we break up the summation into components of the current state due to "signal" (i.e., signal we attempt to recover) and "noise" (i..e, older signal we omit from the recovery):

$$\boldsymbol{x}[N^*] = \sum_{n=N^*-N+1}^{N^*} \boldsymbol{W}^{N^*-n} \boldsymbol{z} s[n] + \sum_{n=0}^{N^*-N} \boldsymbol{W}^{N^*-n} \boldsymbol{z} s[n]$$

$$= \sum_{n=N^*-N+1}^{N^*} \boldsymbol{W}^{N^*-n} \boldsymbol{z} s[n] + \boldsymbol{\epsilon}$$

$$= \boldsymbol{A}\boldsymbol{s} + \boldsymbol{\epsilon}_2.$$

From here we can see that the first summation is the matrix multiply $\boldsymbol{A}\boldsymbol{s}$ as is discussed in the paper. The second summation here, $\boldsymbol{\epsilon}_2$, essentially acts as an additional noise term in the recovery. We can further analyze the effect of this noise term by un-

derstanding that ϵ_2 is bounded for well behaved input sequences $s[n]$ (in fact all that is needed is that the maximum value or the expected value and variance are reasonably bounded) when the eigenvalues of $\boldsymbol{W}$ are of magnitude $q \leq 1$. We can explicitly calculate the worst case scenario bounds on the norm of ϵ_2,

$$\left\| \sum_{n=0}^{N^*-N} \boldsymbol{W}^{N^*-n} \boldsymbol{z} s[n] \right\|_2 \leq \left\| \sum_{n=0}^{N^*-N} \boldsymbol{U} (q\boldsymbol{D})^{N^*-n} \boldsymbol{U}^{-1} \boldsymbol{z} s[n] \right\|_2$$

$$\leq \|\boldsymbol{U}\|_2 \left\| \sum_{n=0}^{N^*-N} (q\boldsymbol{D})^{N^*-n} \boldsymbol{U}^{-1} \boldsymbol{z} s[n] \right\|_2,$$

where $\boldsymbol{D} = \mathrm{diag}(d_1, \ldots, d_M)$ is the diagonal matrix containing the normalized eigenvalues of $\boldsymbol{W}$. If we assume that $\boldsymbol{z}$ is chosen as mentioned as in Section 3.2 so that $\boldsymbol{U}^{-1}\boldsymbol{z} = \left(1/\sqrt{M}\right)\boldsymbol{1}$, the eigenvalues of $\boldsymbol{W}$ are uniformly spread around a complex circle of radius q, and that $s[n] \leq s_{\max}$ for all n, then we can bound this quantity as

$$\left\| \sum_{n=0}^{N^*-N} \boldsymbol{W}^{N^*-n} \boldsymbol{z} s[n] \right\|_2 \leq \frac{s_{\max}}{\sqrt{M}} \|\boldsymbol{U}\|_2 \left\| \sum_{n=0}^{N^*-N} (q\boldsymbol{D})^{N^*-n} \boldsymbol{1} \right\|_2$$

$$= \frac{s_{\max}}{\sqrt{M}} \|\boldsymbol{U}\|_2 \left\| \begin{bmatrix} \sum_{n=0}^{N^*-N} q^{N^*-n} d_1^{N^*-n} \\ \vdots \\ \sum_{n=0}^{N^*-N} q^{N^*-n} d_M^{N^*-n} \end{bmatrix} \right\|_2$$

$$= \frac{s_{\max}}{\sqrt{M}} \|\boldsymbol{U}\|_2 \sqrt{\sum_{k=1}^{M} \left| \sum_{n=0}^{N^*-N} q^{N^*-n} d_k^{N^*-n} \right|^2}$$

$$\leq s_{\max} \|\boldsymbol{U}\|_2 \left| \sum_{n=0}^{N^*-N} q^{N^*-n} \right|$$

$$\leq s_{\max} \|\boldsymbol{U}\|_2 \left| \frac{q^N - q^{N^*}}{1 - q} \right|$$

where d_k is the k^{th} normalized eigenvalue of $\boldsymbol{W}$. In the limit of large input signal lengths ($N^* \to \infty$), we have $N^* \gg N$ and so $q^N \gg q^{N^*}$, which leaves the approximate

expression

$$||\boldsymbol{\epsilon}_2||_2 \leq s_{\max} \, ||\boldsymbol{U}||_2 \left| \frac{q^N}{1-q} \right|.$$

To bound the second term in Equation (4) of the main text (i.e., the signal approximation errors due to imperfect recovery), we must characterize the possible error between the signal (which is K-sparse) and the approximation to the signal with the K^* largest coefficients. In the worst case scenario, there are $K - K^* + 1$ coefficients that cannot be guaranteed to be recovered by the RIP conditions, and these coefficients all take the maximum value $s_{\max}$. In this case, we can bound the signal approximation error as stated in the main text:

$$\begin{aligned} \frac{\beta}{\sqrt{K^*}} \, ||\boldsymbol{s} - \boldsymbol{s}_{S^*}||_1 \; &\leq \; \frac{\beta}{\sqrt{K^*}} \sum_{n=K^*+1}^{K} |q^n s_{\max}| \\ &= \; \frac{\beta s_{\max}}{\sqrt{K^*}} \left(\frac{q^{K^*} - q^K}{1-q} \right). \end{aligned}$$

In the case where noise is present, we can also bound the total power of the noise term,

$$\alpha \left\| \sum_{k=0}^{N+N^*} \boldsymbol{W}^k \boldsymbol{z} \widetilde{\epsilon}[k] \right\|_2^2,$$

using similar steps. Taking $\epsilon_{\max}$ as the largest possible input noise into the system, we obtain the bound

$$\alpha \left\| \sum_{k=0}^{N+N^*} \boldsymbol{W}^k \boldsymbol{z} \widetilde{\epsilon}[k] \right\|_2^2 < \alpha \epsilon_{\max} \, ||\boldsymbol{U}||_2 \left| \frac{q}{1-q} \right|$$

References

Bajwa, W. U., Sayeed, A. M., & Nowak, R. (2009). A restricted isometry property for structurally-subsampled unitary matrices. In 47th Annual Allerton Conference on Communication, Control, and Computing, pp. 1005–1012. IEEE.

Balavoine, A., Romberg, J., & Rozell, C. J. (2012). Convergence and rate analysis of neural networks for sparse approximation. IEEE Transactions on Neural Networks and Learning Systems, 23, 1377–1389.

Balavoine, A., Rozell, C. J., & Romberg, J. (2013). Convergence speed of a dynamical system for sparse recovery. IEEE Transactions on Signal Processing, 61, 4259–4269.

Baum, E. B., Moody, J., & Wilczek, F. (1988). Internal representations for associative memory. Biological Cybernetics, 92, 217–228.

Becker, S., Candes, E. J., & Grant, M. (2011). Templates for convex cone problems with applications to sparse signal recovery. Mathematical Programming Computation, 3.

Buonomano, D. V. & Maass, W. (2009). State-dependent computations: spatiotemporal processing in cortical networks. Nature Reviews Neuroscience, 10, 113–125.

Büsing, L., Schrauwen, B., & Legenstein, R. (2010). Connectivity, dynamics, and memory in reservoir computing with binary and analog neurons. Neural computation, 22, 1272–1311.

Candes, E. J. (2006). Compressive sampling. Proc. Int. Congr. Mathematicians, 3, 1433–1452.

Candes, E. J., Romberg, J., & T, T. (2006). Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. IEEE Transactions on Information Theory, 52, 489–509.

Candes, E. J. & Tao, T. (2006). Near-optimal signal recovery from random projections: Universal encoding strategies? IEEE Transactions on Information Theory, 52, 5406–5425.

Davenport, M. A., Boufounos, P. T., Wakin, M. B., & Baraniuk, R. G. (2010). Signal processing with compressive measurements. IEEE J. Sel. Topics Signal Process., 4, 445–460.

De La Peña, V. H. & Giné, E. (1999). Decoupling: From Dependence to Independence. (Springer Verlag).

Diaconis, P. & Shahshahani, M. (1994). On the eigenvalues of random matrices. Journal of Applied Probability, 31, 49–62.

Donoho, D. & Tanner, J. (2005). Sparse nonnegative solution of underdetermined linear equations by linear programming. Proceedings of the National Academy of Sciences of the United States of America, 102, 9446.

Eftekhari, A., Yap, H. L., Rozell, C. J., & Wakin, M. B. (2012). The restricted isometry property for random block diagonal matrices. Submitted.

Elad, M., Figueiredo, M., & Ma, Y. (2008). On the role of sparse and redundant representations in image processing. IEEE Proceedings - Special Issue on Applications of Compressive Sensing & Sparse Representation.

Ganguli, S., Huh, D., & Sompolinsky, H. (2008). Memory traces in dynamical systems. Proceedings of the National Academy of Sciences, 105, 18970.

Ganguli, S. & Sompolinsky, H. (2010). Short-term memory in neuronal networks through dynamical compressed sensing. Conference on Neural Information Processing Systems.

Ganguli, S. & Sompolinsky, H. (2012). Compressed sensing, sparsity, and dimensionality in neuronal information processing and data analysis. Annual Review of Neuroscience, 35, 485–508.

Haupt, J., Bajwa, W. U., Raz, G., & Nowak, R. (2010). Toeplitz compressed sensing matrices with applications to sparse channel estimation. IEEE Transactions on Information Theory, 56, 5862–5875.

Hermans, M. & Schrauwen, B. (2010). Memory in linear recurrent neural networks in continuous time. Neural Networks, 23, 341–355.

Hu, T., Genkin, A., & Chklovskii, D. B. (2012). A network of spiking neurons for computing sparse representations in an energy-efficient way. Neural Computation, 24, 2852–2872.

Isely, G., Hillar, C. J., & Sommer, F. T. (2011). Deciphering subsampled data: adaptive compressive sampling as a principle of brain communication. Proceedings of NIPS.

Jaeger, H. (2001). Short term memory in echo state networks. GMD Report 152 German National Research Center for Information Technology.

Jaeger, H. & Haas, H. (2004). Harnessing nonlinearity: predicting chaotic systems and saving energy in wireless communication. Science, 304, 78–80.

Krahmer, F., Mendelson, S., & Rauhut, H. (2012). Suprema of chaos processes and the restricted isometry property. arXiv preprint arXiv:1207.0235.

Ledoux, M. & Talagrand, M. (1991). Probability in Banach Spaces: isoperimetry and processes, vol. 23. (Springer).

Legenstein, R. & Maass, W. (2007). Edge of chaos and prediction of computational performance for neural circuit models. Neural Networks, 20, 323–334.

Maass, W., Natschläger, T., & Markram, H. (2002). Real-time computing without stable states: A new framework for neural computation based on perturbations. Neural computation, 14, 2531–2560.

Mayor, J. & Gerstner, W. (2005). Signal buffering in random networks of spiking neurons: Microscopic versus macroscopic phenomena. Physical Review E, 72, 051906.

Mongillo, G., Barak, O., & Tsodyks, M. (2008). Synaptic theory of working memory. Science, 319, 1543–1546.

Olshausen, B. A. & Field, D. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. Nature, 381, 607–609.

Park, J. Y., Yap, H. L., Rozell, C. J., & Wakin, M. B. (2011). Concentration of measure for block diagonal matrices with applications to compressive signal processing. IEEE Transactions on Signal Processing, 59, 5859–5875.

Rauhut, H. (2010). Compressive sensing and structured random matrices. Theoretical Found. and Numerical Methods for Sparse Recovery, pp. 1–92.

Rhen, M. & Sommer, F. T. (2007). A network that uses few active neurones to code visual input predicts the diverse shapes of cortical receptive fields. Journal of Computational Neuroscience, 22, 135–146.

Rozell, C. J., Johnson, D. H., Baraniuk, R. G., & Olshausen, B. A. (2010). Sparse coding via thresholding and local competition in neural circuits. Neural Computation, 20, 2526–2563.

Rudelson, M. & Vershynin, R. (2008). On sparse reconstruction from fourier and gaussian measurements. Comms. Pure and Applied Math., 61, 1025–1045.

Shapero, S., Charles, A. S., Rozell, C., & Hasler, P. (2011). Low power sparse approximation on reconfgurable analog hardware. IEEE Jour. on Emer. and Sel. Top. in Circ. and Sys., 2, 530–541.

Strauss, T., Wustlich, W., & Labahn, R. (2012). Design strategies for weight matrices of echo state networks. Neural Computation, 24, 3246–3276.

Tropp, J. A., Laska, J. N., Duarte, M. F., Romberg, J. K., & Baraniuk, R. G. (2009). Beyond Nyquist: efficient sampling of sparse bandlimited signals. IEEE Trans. Inform. Theory, 56.

Vapnik, V. N. & Chervonenkis, A. Y. (1971). On the uniform convergence of relative frequencies of events to their probabilities. Theory of Probability & Its Applications, 16, 264–280.

Vershynin, R. (2012). Introduction to the non-asymptotic analysis of random matrices. In Compressed Sensing, Theory and Applications, Y. Eldar & G. Kutyniok, eds. (Cambridge Univ. Pr.), chap. 5, pp. 210–268.

Wallace, E., Hamid, R. M., & Latham, P. E. (2013). Randomly connected networks have short temporal memory. Neural Computation, 25, 1408–1439.

White, O. L., Lee, D. D., & Sompolinsky, H. (2004). Short-term memory in orthogonal neural networks. Physical Review Lett., 92, 148102.

Zhu, M. & Rozell, C. (2013). Visual nonclassical receptive field effects emerge from sparse coding in a dynamical system. PLoS Computational Biology, 9, e1003191.