

Learning Topology and Dynamics of Large Recurrent Neural Networks

Yiyuan She, Yuejia He, and Dapeng Wu, *Fellow, IEEE*

Abstract—Large-scale recurrent networks have drawn increasing attention recently because of their capabilities in modeling a large variety of real-world phenomena and physical mechanisms. This paper studies how to identify all authentic connections and estimate system parameters of a recurrent network, given a sequence of node observations. This task becomes extremely challenging in modern network applications, because the available observations are usually very noisy and limited, and the associated dynamical system is strongly nonlinear. By formulating the problem as multivariate sparse sigmoidal regression, we develop simple-to-implement network learning algorithms, with rigorous convergence guarantee in theory, for a variety of sparsity-promoting penalty forms. A quantile variant of progressive recurrent network screening is proposed for efficient computation and allows for direct cardinality control of network topology in estimation. Moreover, we investigate recurrent network stability conditions in Lyapunov’s sense, and integrate such stability constraints into sparse network learning. Experiments show excellent performance of the proposed algorithms in network topology identification and forecasting.

Index Terms—Recurrent networks, topology learning, shrinkage estimation, variable selection, dynamical systems, Lyapunov stability.

I. INTRODUCTION

There has been an increasing interest in identifying network dynamics and topologies in the emerging scientific discipline of network science. In a dynamical network, the evolution of a node is controlled not only by itself, but also by other nodes. For example, in gene regulatory networks [1], the expression levels of genes influence each other, following some dynamic rules, such that the genes are connected together to form a dynamical system. If the topology and evolution rules of the network are known, we can analyze the regulation between genes or detect unusual behaviors to help diagnose and cure genetic diseases. Similarly, the modeling and estimation of dynamical networks are of great importance in various domains including stock market, brain network and social network [2, 3, 4]. To accurately identify the topology and dynamics underlying those networks, scientists are devoted to developing appropriate mathematical models and corresponding estimation methods.

In the literature, linear dynamical models are commonly

used. For example, the human brain connectivity network [5] can be characterized by a set of linear differential equations, where the rate of change of activation/observation of any node is a weighted sum of the activations/observations of its neighbors: $dx_i/dt = \sum_{j \neq i} \alpha_{ij} x_j - d_i x_i, 1 \leq i \leq n$. Here α_{ij} provide the connection weights and d_i is the decay rate. Nevertheless, a lot of complex dynamical networks clearly demonstrate *nonlinear* relationships between the nodes. For instance, the strength of influence is unbounded in the previous simple linear combination, but the so-called “**saturation**” effect widely exists in physical/biological systems (neurons, genes, and stocks)—the external influence on a node, no matter how strong the total input activation is, cannot go beyond a certain threshold. To capture the mechanism, nonlinearity must be introduced into the network system: $dx_i/dt = l_i \pi(\sum_{j \neq i} \alpha_{ij} x_j + u_i) - d_i x_i + c_i$, where π denotes a nonlinear activation function typically taken to be the sigmoidal function $\pi(\theta) = 1/(1 + e^{-\theta})$. It has a proper shape to resemble many real-world mechanisms and behaviors.

The model description is associated with a *continuous-time recurrent neural network*. The existing feedback loops allow the network to exhibit interesting dynamic temporal behaviors to capture many kinds of relationships. It is also biologically realistic in say modeling the effect of an input spike train. Recurrent networks have been successfully applied to a wide range of problems in bioinformatics, financial market forecast, electric circuits, computer vision, and robotics; see, e.g., [6, 7, 8, 9, 10] among many others.

In practical applications, it is often necessary to include noise contamination: $dx_i = (l_i \pi(\sum_{j \neq i} \alpha_{ij} x_j + u_i) - d_i x_i + c_i) dt + \sigma d\mathcal{B}_t, 1 \leq i \leq n$, where \mathcal{B}_t stands for an n -dimensional Brownian motion. Among the very many unknown parameters, α_{ij} might be the most important: the zero-nonzero pattern of α_{ij} indicates if there exists a (direct) connection from node j to node i . Collecting all such connections results in a directed graph to describe the node interaction structure.

A fundamental question naturally arises: *Given a sequence of node observations (possibly at very few time points), can one identify all existing connections and estimate all system parameters of a recurrent network?*

This task becomes extremely challenging in modern big network applications, because the available observations are usually very noisy and only available at a relatively small number of time points (say T), due to budget or equipment limitations. One frequently faces applications with n^2 much larger than T . In addition, in this continuous time setting, no analytical formula of the likelihood exists for the stochastic

Copyright (c) 2014 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

Yiyuan She is with Department of Statistics, Florida State University, Tallahassee, FL 32306. Yuejia He and Dapeng Wu are with Department of Electrical and Computer Engineering, University of Florida, Gainesville, FL 32611. Correspondence author: Prof. Dapeng Wu, wu@ece.ufl.edu, <http://www.wu.ece.ufl.edu>. This work was supported in part by NSF grants CCF-1117012, CCF-1116447 and DMS-1352259.

model, which increases the estimation difficulty even in large samples [11]. Instead of considering multi-step *ad-hoc* procedures, this paper aims at learning the network system *as a whole*. Multivariate statistical techniques will be developed for identifying complete topology and recovering all dynamical parameters. To the best of our knowledge, automatic topology and dynamics learning in large-scale recurrent networks has not been studied before.

In this work, we are interested in networks that are **sparse** in topology. First, many real-world complex dynamical networks indeed have sparse or approximately sparse structures. For example, in regulatory networks, a gene is only regulated by a handful of others [12]. Second, when the number of nodes is large or very large compared with the number of observations, the sparsity assumption reduces the number of model parameters so that the system is estimable. Third, from a philosophical point of view, a sparse network modeling is consistent with the principle of Occam's razor.

Not surprisingly, there is a surge of interest of using compressive sensing techniques for parsimonious network topology learning and dynamics prediction. However, relying on sparsity alone seems to have only limited power in addressing the difficulties of large-scale network learning from big data. To add more prior knowledge and to further reduce the number of effective unknowns, we propose to study how to incorporate structural properties of the network system into learning and estimation, in addition to sparsity. In fact, real-life networks of interest usually demonstrate **asymptotic stability**. This is one of the main reasons why practitioners only perform limited number of measurements of the system, which again provides a type of parsimony or shrinkage in network learning.

In this paper we develop sparse sigmoidal network learning algorithms, with rigorous convergence guarantee in theory, for a variety of sparsity-promoting penalty forms. A quantile variant, the progressive recurrent network screening, is proposed for efficient computation and allows for direct cardinality control of network topology in estimation. Moreover, we investigate recurrent network stability conditions in *Lyapunov's* sense, and incorporate such stability constraints into sparse network learning. The remaining of this paper is organized as follows. Section II introduces the sigmoidal recurrent network model, and formulates a multivariate regularized problem based on the discrete-time approximate likelihood. Section III proposes a class of sparse network learning algorithms based on the study of sparse sigmoidal regressions. A novel and efficient recurrent network screening (RNS) with theoretical guarantee of convergence is advocated for topology identification in ultra-high dimensions. Section IV investigates asymptotic stability conditions in recurrent systems, resulting in a stable-sparse sigmoidal (\mathbf{S}^3) network learning. In Section V, synthetic data experiments and real applications are given. All proof details are left to the Appendices.

II. MODEL FORMULATION

To describe the evolving state of a continuous-time recurrent neural network, one usually defines an associated dynamical system. Ideally, without any randomness, the network behavior

can be specified by a set of ordinary differential equations:

$$\frac{dx_i}{dt} = l_i \pi \left(\sum_{j \neq i} \alpha_{ij} x_j + u_i \right) - d_i x_i + c_i, \quad i = 1, \dots, n, \quad (1)$$

where x_i , short for $x_i(t)$, denotes the dynamic process of node i . Throughout the paper, π is the *sigmoidal* activation function

$$\pi(\theta) = \frac{1}{1 + e^{-\theta}}, \quad (2)$$

which is most frequently used in recurrent networks. This function is smooth and strictly increasing. It has a proper shape to resemble many real-world mechanisms and behaviors. Due to noise contamination, a stochastic differential equation model is more realistic

$$dx_i = (l_i \pi \left(\sum_{j \neq i} \alpha_{ij} x_j + u_i \right) - d_i x_i + c_i) dt + \sigma dB_t, \quad (3)$$

where B_t is a standard Brownian motion and reflects the stochastic nature of the system. Typically $l_i > 0$, $d_i > 0$, and in some applications $\alpha_{ii} = 0$ (no self-regulation) is required.

In the sigmoidal recurrent network model, the coefficients α_{ij} characterize the between-node interactions. In particular, $\alpha_{ij} = 0$ indicates that node j does not *directly* influence node i ; otherwise, node j regulates node i , and is referred to as a regulator of node i in gene regulatory networks. Such a regulation relationship can be either excitatory (if α_{ij} is positive) or inhibitory (if α_{ij} is negative). In this way, $\mathbf{A} = [\alpha_{ij}]$ is associated with a directed graph that captures the Granger causal relationships between all nodes [13], and the topology of the recurrent network is revealed by the zero-nonzero pattern of the matrix. Therefore, to identify all significant regulation links, it is of great interest to estimate \mathbf{A} , given a sequence of node observations (snapshots of the network).

Denote the system state at time t by $\mathbf{x}(t) = [x_1(t) \cdots x_n(t)]^T$ or \mathbf{x}_t (or simply \mathbf{x} when there is no ambiguity.) Define $\mathbf{l} = [l_1, \dots, l_n]^T$, $\mathbf{u} = [u_1, \dots, u_n]^T$, $\mathbf{c} = [c_1, \dots, c_n]^T$, $\mathbf{D} = \text{diag}\{d_i\}$, $\mathbf{L} = \text{diag}\{l_i\}$, and $\mathbf{A} = [\alpha_{ij}] = [\boldsymbol{\alpha}_1 \cdots \boldsymbol{\alpha}_n]^T \in \mathbb{R}^{n \times n}$. Then (3) can be represented in a *multivariate* form

$$d\mathbf{x}_t = (\mathbf{L}\boldsymbol{\pi}(\mathbf{A}\mathbf{x} + \mathbf{u}) - \mathbf{D}\mathbf{x} + \mathbf{c}) dt + \sigma d\mathbf{B}_t, \quad (4)$$

where \mathbf{B}_t is an n -dimensional standard Brownian motion. While this model is specified in continuous time, in practice, the observations are always collected at discrete time points. Estimating the parameters of an SDE model from few discrete observations is very challenging. There rarely exists an analytical expression of the exact likelihood. A common treatment is to discretize (4) and use an approximate likelihood instead. We use the Euler discretization scheme (see, e.g., [14]):

$$\Delta \mathbf{x} = (\mathbf{L}\boldsymbol{\pi}(\mathbf{A}\mathbf{x} + \mathbf{u}) - \mathbf{D}\mathbf{x} + \mathbf{c})\Delta t + \sigma \Delta \mathbf{B}_t.$$

Suppose the system (4) is observed at $T + 1$ time points t_1, \dots, t_{T+1} . Let $\mathbf{x}_s = [x_1(t_s), \dots, x_n(t_s)]^T \in \mathbb{R}^n$ be the observed values of all n nodes at t_s . Define $\Delta \mathbf{x}_s = (\mathbf{x}_{s+1} - \mathbf{x}_s)$ and $\Delta t_s = (t_{s+1} - t_s)$, $1 \leq s \leq T$. Because $\Delta \mathbf{B}_t \sim \mathcal{N}(\mathbf{0}, \Delta t \mathbf{I})$, the negative conditional log-likelihood

for the discretized model is given by

$$\begin{aligned} \ell(\mathbf{x}_1, \dots, \mathbf{x}_{T+1} | \mathbf{x}_1) &= -\log P(\Delta \mathbf{x}_1, \dots, \Delta \mathbf{x}_T | \mathbf{x}_1) \\ &= -\log P(\Delta \mathbf{x}_1 | \mathbf{x}_1) \cdots P(\Delta \mathbf{x}_T | \mathbf{x}_T) \\ &= \sum_{s=1}^T \|\Delta \mathbf{x}_s / \Delta t_s - (\mathbf{L}\pi(\mathbf{A}\mathbf{x}_s + \mathbf{u}) - \mathbf{D}\mathbf{x}_s + \mathbf{c})\|_2^2 \Delta t_s / (2\sigma^2) \\ &\quad + C(\sigma^2) =: f(\mathbf{A}, \mathbf{u}, \mathbf{l}, \mathbf{d}, \mathbf{c}) / \sigma^2 + C(\sigma^2). \end{aligned}$$

$C(\sigma^2)$ is a function that depends on σ^2 only. The fitting criterion f is separable in $\alpha_1, \dots, \alpha_n$. To see this, let $x_{i,s} = x_i(t_s)$ and $\Delta x_{i,s} = x_i(t_{s+1}) - x_i(t_s)$. Then, it is easy to verify that

$$\begin{aligned} f(\mathbf{A}, \mathbf{u}, \mathbf{l}, \mathbf{d}, \mathbf{c}) &= \frac{1}{2} \sum_{i=1}^n \sum_{s=1}^T \left(\frac{\Delta x_{i,s}}{\Delta t_s} - (l_i \pi(\alpha_i^\top \mathbf{x}_s + u_i) - d_i x_{i,s} + c_i) \right)^2 \Delta t_s. \end{aligned} \quad (5)$$

Conventionally, the unknown parameters can then be estimated by minimizing f . In modern applications, however, the number of available observations (T) is often much smaller than the number of variables to be estimated ($n^2 + 4n$), due to, for example, equipment/budget limitations. Classical MLE methods do not apply well in this high-dimensional setting.

Fortunately, the networks of interest in reality often possess topology sparsity. For example, a stock price may not be directly influenced by all the other stocks in the stock market. A parsimonious network with only significant regulation links is much more interpretable. Statistically speaking, the sparsity in \mathbf{A} suggests the necessity of **shrinkage estimation** [15] which can be done by adding penalties and/or constraints to the loss function. The general penalized maximum likelihood problem is

$$\min_{\mathbf{A}, \mathbf{u}, \mathbf{l}, \mathbf{d}, \mathbf{c}} f(\mathbf{A}, \mathbf{u}, \mathbf{l}, \mathbf{d}, \mathbf{c}) + \sum_{i,j} P(\alpha_{ij}; \lambda_{ji}) \quad (6)$$

where P is a penalty promoting sparsity and λ_{ji} are regularization parameters. Among the very many possible choices of P , the ℓ_1 penalty is perhaps the popular nowadays to enforce sparsity:

$$P(t; \lambda) = \lambda |t|. \quad (7)$$

It provides a convex relaxation of the ℓ_0 penalty

$$P(t; \lambda) = \frac{\lambda^2}{2} 1_{t \neq 0}. \quad (8)$$

Taking both topology identification and dynamics prediction into consideration, we are particularly interested in the $\ell_0 + \ell_2$ penalty [16]

$$P(t; \lambda, \eta) = \frac{1}{2} \frac{\lambda^2}{1 + \eta} 1_{t \neq 0} + \frac{\eta}{2} t^2, \quad (9)$$

where the ℓ_2 penalty or Tikhonov regularization can effectively deal with large noise and collinearity [17, 18] to enhance estimation accuracy.

The shrinkage estimation problem (6) is however nontrivial. The loss f is nonconvex, π is nonlinear, and the penalty P may be nonconvex or even discrete, let alone the high-dimensionality challenge. Indeed, in many practical networks,

the available observations are usually quite limited and noisy. Effective and efficient learning algorithms are in great need to meet the modern big data challenge.

III. SPARSE SIGMOIDAL REGRESSION FOR RECURRENT NETWORK LEARNING

A. Univariate-response sigmoidal regression

As analyzed previously, to solve (6), it is sufficient to study a univariate-response learning problem

$$\begin{aligned} \min_{(\beta, l, \gamma)} \frac{1}{2} \sum_{s=1}^T w_s \left(y_s - l \pi(\tilde{\mathbf{x}}_s^\top \beta) - \mathbf{z}_s^\top \gamma \right)^2 + \sum_{k=1}^n P(\beta_k, \lambda_k) \\ \triangleq F(\beta, l, \gamma) \end{aligned} \quad (10)$$

where $\tilde{\mathbf{x}}_s, \mathbf{z}_s, y_s, w_s$ are given, and l, γ, β are unknown with β desired to be sparse. (10) is the recurrent network learning problem for node i when we set $\tilde{\mathbf{x}}_s = [1, \mathbf{x}_s]^\top$, $y_s = \Delta x_{i,s} / \Delta t_s$, $\mathbf{z}_s = [1, x_{i,s}]^\top$, and $w_s = \Delta t_s$ ($1 \leq s \leq T$) (note that the intercept β_1 is usually subject to no penalization corresponding to $\lambda_1 = 0$). For notational ease, define $\tilde{\mathbf{X}} = [\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_T]^\top$, $\mathbf{y} = [y_1, \dots, y_T]^\top$, $\mathbf{Z} = [\mathbf{z}_1, \dots, \mathbf{z}_T]^\top$, $\boldsymbol{\lambda} = [\lambda_k] \in \mathbb{R}^n$, and $\mathbf{W} = \text{diag}\{\mathbf{w}\} = \text{diag}\{w_1, \dots, w_s\}$ (to be used in this subsection only, unless otherwise specified).

We propose a simple-to-implement and efficient algorithm to solve the general optimization problem in the form of (10). First define two useful auxiliary functions

$$\xi(\theta, y) = \pi(\theta)(1 - \pi(\theta))(\pi(\theta) - y), \quad (11)$$

$$k_0(\mathbf{y}, \mathbf{w}) = \max_{1 \leq s \leq T} \frac{w_s}{16} \left(1 + \frac{(1 - 2y_s)^2}{2} \right). \quad (12)$$

The vector versions of π and ξ are defined componentwise: $\boldsymbol{\pi}(\boldsymbol{\theta}) = [\pi(\theta_1), \dots, \pi(\theta_T)]^\top$ and $\boldsymbol{\xi}(\boldsymbol{\theta}, \mathbf{y}) = [\xi(\theta_1, y_1), \dots, \xi(\theta_T, y_T)]^\top$, and the matrix versions are defined similarly. A prototype algorithm is described as follows, starting with an initial estimate $\beta^{(0)}$, a thresholding rule Θ (an odd, shrinking and nondecreasing function, cf. [19]), and $j = 0$.

repeat

0) $j \leftarrow j + 1$

1) $\boldsymbol{\mu}^{(j)} \leftarrow \boldsymbol{\pi}(\tilde{\mathbf{X}} \beta^{(j-1)})$

2) Fit a weighted least-squares model

$$\min_{l, \gamma} \|\mathbf{W}^{1/2}(\mathbf{y} - [\boldsymbol{\mu}^{(j)} \mathbf{Z}] [l \ \gamma^\top]^\top)\|_2^2, \quad (13)$$

with the corresponding solution denoted by $(l^{(j)}, \gamma^{(j)})$.
3) Construct $\tilde{\mathbf{y}}^{(j)} \leftarrow (\mathbf{y} - \mathbf{Z} \gamma^{(j)}) / l^{(j)}$, $\tilde{\mathbf{w}}^{(j)} \leftarrow (l^{(j)})^2 \mathbf{w}$, $\tilde{\mathbf{W}}^{(j)} = \text{diag}\{\tilde{\mathbf{w}}^{(j)}\}$. Let $K^{(j)}$ be any constant no less than $k_0(\tilde{\mathbf{y}}^{(j)}, \tilde{\mathbf{w}}^{(j)}) \|\tilde{\mathbf{X}}\|_2^2$, where $\|\cdot\|_2$ is the spectral norm.

4) Update β via thresholding:

$$\beta^{(j)} = \Theta(\beta^{(j-1)} - \frac{1}{K^{(j)}} \tilde{\mathbf{X}}^\top \tilde{\mathbf{W}}^{(j)} \boldsymbol{\xi}(\tilde{\mathbf{X}} \beta^{(j-1)}, \tilde{\mathbf{y}}^{(j)}); \boldsymbol{\lambda}^{(j)}), \quad (14)$$

where $\boldsymbol{\lambda}^{(j)} = [\lambda_k^{(j)}]_{n \times 1}$ is a scaled version of $\boldsymbol{\lambda}$ satisfying $P(t; \lambda_k) / K^{(j)} = P(t; \lambda_k^{(j)})$ for any $t \in \mathbb{R}$, $1 \leq k \leq n$.

until convergence

Before proceeding, we give some examples of Θ and $\lambda^{(j)}$. The specific form of the thresholding function Θ in (14) is related to the penalty P through the following formula [16]:

$$\begin{aligned} & P(t; \lambda) - P(0; \lambda) \\ &= \int_0^{|t|} (\sup\{s : \Theta(s; \lambda) \leq u\} - u) du + q(t; \lambda) \end{aligned} \quad (15)$$

with $q(\cdot; \lambda)$ nonnegative and $q(\Theta(s; \lambda)) = 0$ for all s . The regularization parameter(s) are rescaled at each iteration according to the form of P . Examples include: (i) the ℓ_1 penalty (7), and its associated *soft-thresholding* $\Theta_S(t; \lambda) = \text{sgn}(t)(|t| - \lambda)1_{|t| > \lambda}$, in which case $P(t; \lambda)/K^{(j)} = P(t; \lambda^{(j)})$, $\forall t$ implies $\lambda^{(j)} = \lambda/K^{(j)}$, (ii) the ℓ_0 penalty (8) and the *hard-thresholding* $\Theta_H(t) = t1_{|t| > \lambda}$, which determines $\lambda^{(j)} = \lambda/\sqrt{K^{(j)}}$, and (iii) the $\ell_0 + \ell_2$ penalty (9) and its associated *hard-ridge thresholding* $\Theta_{HR}(t; \lambda, \eta) = \frac{t}{1+\eta}1_{|t| > \lambda}$, where $\lambda^{(j)} = \frac{\lambda}{K^{(j)}}\sqrt{(\eta + K^{(j)})/(\eta + 1)}$, $\eta^{(j)} = \eta/K^{(j)}$. The ℓ_p ($0 < p < 1$) penalties, the elastic net penalty, and others [16] are also instances of this framework.

Theorem 1. *Given the objective function in (10), suppose Θ and P satisfy (15) for some nonnegative $q(t; \lambda)$ with $q(\Theta(t; \lambda)) = 0$ for all t and λ . Then given any initial point $\beta^{(0)}$, with probability 1 the sequence of iterates $(\beta^{(j)}, l^{(j)}, \gamma^{(j)})$ generated by the prototype algorithm satisfies*

$$F(\beta^{(j-1)}, l^{(j-1)}, \gamma^{(j-1)}) \geq F(\beta^{(j)}, l^{(j)}, \gamma^{(j)}). \quad (16)$$

See Appendix A for the proof.

Normalization is usually necessary to make all predictors equally span in the space before penalizing their coefficients using a single regularization parameter λ . We can center and scale all predictor columns but the intercept before calling the algorithm. Alternatively, it is sometimes more convenient to specify λ_k componentwise—e.g., in the ℓ_1 penalty $\sum \lambda_k |\beta_k|$ set $\lambda_k = \lambda \cdot \|\mathbf{X}[:, k]\|_2$ for non-intercept coefficients.

B. Cardinality constrained sigmoidal network learning

In the recurrent network setting, one can directly apply the prototype algorithm in Section III-A to solve (6), by updating the columns in \mathbf{B} one at a time. On the other hand, a multivariate update form is usually more efficient and convenient in implementation. Moreover, it facilitates integrating stability into network learning (cf. Section IV).

To formulate the loss in a multivariate form, we introduce

$$\begin{aligned} \mathbf{Y} &= [y_{i,s}] = [\Delta x_{i,s}/\Delta t_s] \in \mathbb{R}^{T \times n}, \\ \mathbf{X} &= [\mathbf{x}_1, \dots, \mathbf{x}_T]^\top \in \mathbb{R}^{T \times n}, \\ \mathbf{B} &= \mathbf{A}^\top = [\boldsymbol{\alpha}_1, \dots, \boldsymbol{\alpha}_n] \in \mathbb{R}^{n \times n}, \\ \mathbf{W} &= \text{diag}\{\mathbf{w}\} = \text{diag}\{\Delta t_1, \dots, \Delta t_T\}. \end{aligned}$$

Then f in (5) can be rewritten as

$$\begin{aligned} & f(\mathbf{B}, \mathbf{u}, \mathbf{l}, \mathbf{d}, \mathbf{c}) \\ &= \frac{1}{2} \|\mathbf{W}^{1/2} \{\mathbf{Y} - [\boldsymbol{\pi}(\mathbf{X}\mathbf{B} + \mathbf{1}\mathbf{u}^\top)\mathbf{L} - \mathbf{X}\mathbf{D} + \mathbf{1}\mathbf{c}^\top]\}\|_F^2, \end{aligned} \quad (17)$$

with $\|\cdot\|_F$ denoting the Frobenius norm, and the objective function to minimize becomes $\|\mathbf{W}^{1/2}\{\mathbf{Y} - [\boldsymbol{\pi}(\mathbf{X}\mathbf{B} +$

$$\mathbf{1}\mathbf{u}^\top)\mathbf{L} - \mathbf{X}\mathbf{D} + \mathbf{1}\mathbf{c}^\top]\}\|_F^2 + P(\mathbf{B}, \boldsymbol{\Lambda}).$$

One of the main issues is to choose a proper penalty form for sparse network learning. Popular sparsity-promoting penalties include ℓ_1 , SCAD, ℓ_0 , among others. The ℓ_0 penalty (8) is ideal in pursuing a parsimonious solution. However, the matter of parameter tuning cannot be ignored. Most tuning strategies (such as K -fold cross-validation) require computing a solution path for a grid of values of λ , which is quite time-consuming in large network estimation. Rather than applying the ℓ_0 penalty, we propose an ℓ_0 **constrained** sparse network learning

$$\|\mathbf{B}\|_0 \leq m, \quad (18)$$

where $\|\cdot\|_0$ denotes the number of nonzero components in a matrix. In contrast to the penalty parameter λ in (8), m is more meaningful and customizable. One can directly specify its value based on prior knowledge or availability of computational resources to have control of the network connection cardinality. To account for collinearity and large noise contamination, we add a further ℓ_2 penalty in \mathbf{B} , resulting in a new ' $\ell_0 + \ell_2$ ' regularized criterion

$$\begin{aligned} & \min_{\mathbf{B}, \mathbf{L}, \mathbf{D}, \mathbf{u}, \mathbf{c}} \frac{1}{2} \|\mathbf{W}^{1/2} \{\mathbf{Y} - [\boldsymbol{\pi}(\mathbf{X}\mathbf{B} + \mathbf{1}\mathbf{u}^\top)\mathbf{L} - \mathbf{X}\mathbf{D} + \mathbf{1}\mathbf{c}^\top]\}\|_F^2 \\ & + \frac{\eta}{2} \|\mathbf{B}\|_F^2 =: F_0, \text{ s.t. } \|\mathbf{B}\|_0 \leq m. \end{aligned} \quad (19)$$

Not only is the cardinality bound m convenient to set, because of the sparsity assumption, but the ℓ_2 shrinkage parameter η is easy to tune. Indeed, η is usually not a very sensitive parameter and does not require a large grid of candidate values. Many researchers simply fix it at a small value (say, $1e - 3$) which can effectively reduce the prediction error (e.g., [18]). Similarly, to handle the possible collinearity between $\boldsymbol{\pi}$ and $\mathbf{1}$, we recommend adding mild ridge penalties $\frac{\eta_l}{2} \|\mathbf{l}\|_2^2 + \frac{\eta_c}{2} \|\mathbf{c}\|_2^2$ in (19) (say, $\eta_l = 1e - 4$ and $\eta_c = 1e - 2$).

The constrained optimization of (19) does not apparently fall into the penalized framework proposed in Section III-A. However, we can adapt the technique to handle it through a *quantile* thresholding operator (as a variant of the hard-ridge thresholding (9)). The detailed recurrent network screening (**RNS**) algorithm is described as follows, where for notational simplicity $\tilde{\mathbf{X}} := [\mathbf{1} \ \mathbf{X}]$, $\tilde{\mathbf{B}} := [\mathbf{u} \ \mathbf{B}^\top]^\top$, and we denote by $\mathbf{A}[I, J]$ the submatrix of \mathbf{A} consisting of the rows and columns indexed by I and J , respectively.

Theorem 2. *Given any initial point $\tilde{\mathbf{B}}^{(0)}$, Algorithm 1 converges in the sense that with probability 1, the function value decreasing property holds:*

$$\begin{aligned} & F_0(\mathbf{B}^{(j-1)}, \mathbf{L}^{(j-1)}, \mathbf{D}^{(j-1)}, \mathbf{u}^{(j-1)}, \mathbf{c}^{(j-1)}) \\ & \geq F_0(\mathbf{B}^{(j)}, \mathbf{L}^{(j)}, \mathbf{D}^{(j)}, \mathbf{u}^{(j)}, \mathbf{c}^{(j)}), \end{aligned}$$

and all $\mathbf{B}^{(j)}$ satisfy $\|\mathbf{B}^{(j)}\|_0 \leq m, \forall j \geq 1$.

See Appendix B for the proof.

Step 2 can be achieved by solving n weighted least squares. Or, one can re-formulate it as a single-response problem to obtain the solution $(\mathbf{l}^{(j)}, \mathbf{d}^{(j)}, \mathbf{c}^{(j)})$ in one step. The latter way is usually more efficient. When there are ridge penalties imposed on \mathbf{l} and \mathbf{c} , the computation is similar.

Algorithm 1 Recurrent Network Screening (RNS)

given $\tilde{B}^{(0)}$ (initial estimate), m (cardinality bound), η (ℓ_2 shrinkage parameter).

$j \leftarrow 0$

repeat

0) $j \leftarrow j + 1$

1) $\boldsymbol{\mu}^{(j)} \leftarrow \boldsymbol{\pi}(\tilde{X}\tilde{B}^{(j-1)})$

2) Update $\mathbf{L} = \text{diag}\{\mathbf{l}\}$, $\mathbf{D} = \text{diag}\{\mathbf{d}\}$, and \mathbf{c} by fitting a weighted vector least-squares model:

$$\min_{\mathbf{L}, \mathbf{D}, \mathbf{c}} \|\mathbf{W}^{1/2}(\mathbf{Y} - [\boldsymbol{\mu}^{(j)}\mathbf{L} - \mathbf{X}\mathbf{D} + \mathbf{1}\mathbf{c}^T])\|_F^2, \quad (20)$$

with the solution denoted by $\mathbf{L}^{(j)} = \text{diag}\{\mathbf{l}^{(j)}\} = \text{diag}\{l_1^{(j)}, \dots, l_n^{(j)}\}$, $\mathbf{D}^{(j)} = \text{diag}\{d_1^{(j)}, \dots, d_n^{(j)}\}$, and $\mathbf{c}^{(j)}$. (This amounts to solving n separate weighted least squares problems.)

3) Construct $\tilde{\mathbf{Y}} \leftarrow (\mathbf{Y} + \mathbf{X}\mathbf{D}^{(j)} - \mathbf{1}(\mathbf{c}^{(j)})^T)(\mathbf{L}^{(j)})^-$, $\tilde{\mathbf{W}} \leftarrow \mathbf{w} \cdot (\mathbf{l}^{(j)} \circ \mathbf{l}^{(j)})^T$, where $^-$ denotes the Moore-Penrose pseudoinverse and \circ is the Hadamard product. Let $K_i^{(j)}$ be any constant no less than $k_0(\tilde{\mathbf{Y}}[:, i], \tilde{\mathbf{W}}[:, i])\|\tilde{\mathbf{X}}\|_2^2$, and $\mathbf{K}^{(j)} \leftarrow \text{diag}\{K_1^{(j)}, \dots, K_n^{(j)}\}$.

4) Update \mathbf{B} and \mathbf{u} :

4.1) $\tilde{\boldsymbol{\Xi}} \leftarrow \tilde{\mathbf{B}}^{(j-1)} - \tilde{\mathbf{X}}^T \{\tilde{\mathbf{W}} \circ \boldsymbol{\xi}(\tilde{\mathbf{X}}\tilde{\mathbf{B}}^{(j-1)}, \tilde{\mathbf{Y}})\}(\mathbf{K}^{(j)})^-$, $\mathbf{u}^{(j)} \leftarrow (\tilde{\boldsymbol{\Xi}}[1, :])^T$, $\boldsymbol{\Xi} \leftarrow \tilde{\boldsymbol{\Xi}}[2 : \text{end}, :]$ (the submatrix of $\tilde{\boldsymbol{\Xi}}$ without the first row), $\boldsymbol{\eta}^{(j)} = \boldsymbol{\eta} \cdot (\mathbf{1}\mathbf{1}^T)(\mathbf{K}^{(j)})^-$.

4.2) Perform the hard-ridge thresholding $\mathbf{B}^{(j)} \leftarrow \Theta_{HR}(\boldsymbol{\Xi}; \boldsymbol{\zeta}(\boldsymbol{\Xi}), \boldsymbol{\eta}^{(j)})$ (cf. Section III-A for the definition of Θ_{HR}) with an adaptive threshold matrix $\boldsymbol{\zeta}(\boldsymbol{\Xi})$. The entries of $\boldsymbol{\zeta}(\boldsymbol{\Xi})$ are all set to the medium of the m th and the $(m+1)$ th largest components of $\text{vec}(|\boldsymbol{\Xi}|)$. See (21) for other variants when certain links must be maintained or forbidden.

4.3) $\tilde{\mathbf{B}}^{(j)} \leftarrow [\mathbf{u}^{(j)} \quad (\mathbf{B}^{(j)})^T]^T$

until the decrease in function value is small

deliver $\hat{\mathbf{B}} = \mathbf{B}^{(j)}$, $\hat{\mathbf{u}} = \mathbf{u}^{(j)}$, $\hat{\mathbf{L}} = \mathbf{L}^{(j)}$, $\hat{\mathbf{D}} = \mathbf{D}^{(j)}$, $\hat{\mathbf{c}} = \mathbf{c}^{(j)}$.

In certain applications self-regulations are not allowed. Then $\boldsymbol{\zeta}(\boldsymbol{\Xi})$ should be the medium of the m th and the $(m+1)$ th largest elements of $|\boldsymbol{\Xi} - \boldsymbol{\Xi} \circ \mathbf{I}|$ ($\boldsymbol{\Xi}$ in absolute value after excluding its diagonal entries).¹ Other prohibited links can be similarly treated in determining the dynamic threshold. In general, given \mathcal{T} the index set of the links that must be maintained and \mathcal{F} the index set of the links that are forbidden, the threshold is constructed as follows

$$\begin{aligned} \boldsymbol{\Xi}[i] &\leftarrow 0, \forall i \in \mathcal{F} \\ \boldsymbol{\zeta}[i] &\leftarrow \begin{cases} 0, & \text{if } i \in \mathcal{T} \\ |\boldsymbol{\Xi}[\mathcal{T}^c]|_{(m+1)}, & \text{otherwise,} \end{cases} \end{aligned} \quad (21)$$

where $|\boldsymbol{\Xi}[\mathcal{T}^c]|_{(m+1)}$ is the $(m+1)$ th largest element (in absolute value) in $\boldsymbol{\Xi}[\mathcal{T}^c]$ (all entries of $\boldsymbol{\Xi}$ except those indexed by \mathcal{T}). It is easy to show that the convergence result still holds based on the argument in Appendix B.

¹Throughout the paper $|\mathbf{A}|$ is the absolute value of the elements of \mathbf{A} . That is, for $\mathbf{A} = [a_{ij}]$, $|\mathbf{A}| = [|a_{ij}|]$.

In implementation, we advocate reducing the network cardinality in a **progressive** manner to lessen greediness: at the j th step, m is replaced by $m(j)$, where $\{m(j)\}$ is a monotone sequence of integers decreasing from $n(n-1)$ (assuming no self-regulation) to m . Empirical study shows that the sigmoidal decay cooling schedule $m(j) = \lceil 2n(n-1)/(1 + e^{\alpha j}) \rceil$ with $\alpha = 0.01$, works well.

RNS involves no costly operations at each iteration, and is simple to implement. It runs efficiently for large networks. The RNS estimate can be directly used for analyzing the network topology. One can also use it for screening (in which case a relatively large value of m is specified), followed by a fine network learning algorithm restricted on the screened connections. In either case RNS substantially reduces the search space of candidate links.

IV. STABLE-SPARSE SIGMOIDAL NETWORK LEARNING

For a general dynamical system possibly nonlinear, stability is one of the most fundamental issues [20, 21]. If a system's equilibrium point is *asymptotically stable*, then the perturbed system must approach the equilibrium point as t increases. Moreover, one of the main reasons many real network applications only have limited number of observations measured after perturbation is that the associated dynamical systems stabilize fast (e.g., exponentially fast). This offers another important type of parsimony/shrinkage in network parameter estimation.

To design a new type of regularization, we first investigate stability conditions of sigmoidal recurrent networks in Lyapunov's sense [22, 23]. Then, we develop a stable sparse sigmoidal (\mathbf{S}^3) network learning approach.

A. Conditions for asymptotic stability

Recall the multivariate representation of (1)

$$\frac{d\mathbf{x}}{dt} = \mathbf{L}\boldsymbol{\pi}(\mathbf{A}\mathbf{x} + \mathbf{u}) - \mathbf{D}\mathbf{x} + \mathbf{c}. \quad (22)$$

Because \mathbf{A} is sparse and typically singular and the degradation rates d_i are not necessarily identical, in general (22) can not be treated as an instance of the Cohen-Grossberg neural networks [24]. We must first derive its own stability conditions to be considered in network estimation.

Hereinafter, $\mathbf{A} \succeq \mathbf{A}'$ and $\mathbf{A} \succ \mathbf{A}'$ stand for the positive semi-definiteness and positive definiteness of $\mathbf{A} - \mathbf{A}'$, respectively, and the set of eigenvalues of \mathbf{A} is $\text{spec}(\mathbf{A})$. Our conditions are stated below:

$$\mathbf{D} \succ \mathbf{0}, \quad (A1)$$

$$\mathbf{L} \succeq \mathbf{0}, \quad (A2)$$

$$\text{Re}(\lambda) < 0 \text{ for any } \lambda \in \text{Spec}(\mathbf{L}\mathbf{A} - \mathbf{4}\mathbf{D}), \quad (A3a)$$

$$\mathbf{L}\mathbf{A}/2 + \mathbf{A}^T\mathbf{L}/2 \prec \mathbf{4}\mathbf{D}, \quad (A3b)$$

Theorem 3. *Suppose (A1) & (A2) hold. Then (A3a) guarantees the network defined by (22) has a unique equilibrium point \mathbf{x}^* and is globally exponentially stable in the sense that $\|\mathbf{x}(t) - \mathbf{x}^*\|_2^2 \leq e^{-\varepsilon t} \|\mathbf{x}(0) - \mathbf{x}^*\|_2^2$ for any solution $\mathbf{x}(t)$. The same conclusion holds if (A3a) is replaced by (A3b).*

See Appendix C for the detailed proof.

Figure 1 shows an example of stochastic processes gener-

ated from a stable recurrent network and an unstable recurrent network respectively. In the upper panel, the recurrent network system parameters satisfy the stability condition (A3a), while those in the lower panel violate (A3a). (In both situations, the number of nodes is 10 and the diffusion parameter σ is fixed at 0.5.) The differences between the two models are obvious.

In reality, asymptotically stable systems are commonly observed. The stability conditions reflect structural characteristics. For example, when all l_i are equal and $d_i > 0$, then the *skew-symmetry* of \mathbf{A} , i.e., $\mathbf{A} = -\mathbf{A}^\top$, guarantees asymptotic stability. The information provided by the constrains can assist topology learning. This motivates the design of sparse recurrent network learning with stability.

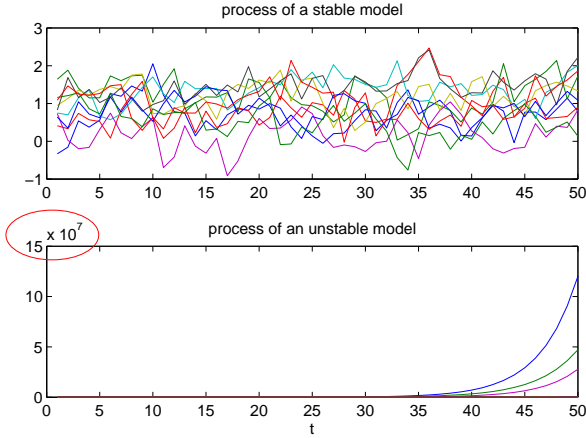


Figure 1: Stochastic processes generated from a stable recurrent network model and an unstable recurrent network model.

B. \mathbf{S}^3 network estimation

(A3a) is less restrictive than (A3b). In optimization, imposing (A3a) seems however difficult. We propose

$$\begin{aligned} \min_{\mathbf{B}, \mathbf{L}, \mathbf{D}, \mathbf{u}, \mathbf{c}} \quad & \frac{1}{2} \|\mathbf{W}^{1/2} \{\mathbf{Y} - [\pi(\mathbf{X}\mathbf{B} + \mathbf{1}\mathbf{u}^\top)\mathbf{L} - \mathbf{X}\mathbf{D} + \mathbf{1}\mathbf{c}^\top]\}\|_F^2 \\ & + P(\mathbf{B}, \mathbf{\Lambda}) =: F_1, \\ \text{s.t.} \quad & \mathbf{L} \succeq \mathbf{0}, \mathbf{D} \succeq \mathbf{0}, (\mathbf{L}\mathbf{B}^\top + \mathbf{B}\mathbf{L})/2 \preceq 4\mathbf{D}, \end{aligned} \quad (23)$$

referred to as the Stable-Sparse Sigmoidal (\mathbf{S}^3) network learning. The stability constraints are now imposed on \mathbf{B} as the transpose of the raw coefficient matrix \mathbf{A} . Similar to the discussion of (19), in implementation, we add mild ℓ_2 penalties on l_i and c_i to deal with possible collinearity, and it is common to replace $\mathbf{0}$ by $\epsilon\mathbf{I}$ with ϵ extremely small.

In this part, we focus on the ℓ_1 penalty $P(\mathbf{B}, \mathbf{\Lambda}) = \|\mathbf{\Lambda} \circ \mathbf{B}\|_1$ where matrix $\mathbf{\Lambda}$ usually has the form $\Lambda[k, k'] = \|\mathbf{X}[:, k]\|_2$ for any $k, k' \leq n$ (other options are possible.) Introducing the componentwise regularization matrix is helpful when one wants to maintain or forbid certain links based on prior knowledge or preliminary screening. For example, if no self-regulation is allowed, then all diagonal entries of $\mathbf{\Lambda}$ ought to be $+\infty$.

It turns out that one can modify Step 2 and Step 4 of the RNS algorithm to solve (23).

First, the non-negativity of l_i and d_i can be directly incorporated, if the weighted least-squares problem in Step 2 is replaced by the following quadratic programming with non-negativity constraints:

$$\begin{aligned} \text{Step 2) Solve } \quad & \min_{\mathbf{L}, \mathbf{D}, \mathbf{c}} \|\mathbf{W}^{1/2}(\mathbf{Y} - [\boldsymbol{\mu}^{(j)}\mathbf{L} - \mathbf{X}\mathbf{D} + \mathbf{1}\mathbf{c}^\top])\|_F^2 \\ \text{s.t.} \quad & \mathbf{L} = \text{diag}\{l_i\} \succeq \mathbf{0}, \mathbf{D} = \text{diag}\{d_i\} \succeq \mathbf{0}. \end{aligned} \quad (24)$$

A variety of algorithms and packages can be used [25].

Integrating the spectral constraint on \mathbf{B} into network learning is much trickier. We modify Step 4 of Algorithm 1 as follows.

Step 4) Update \mathbf{B} and \mathbf{u} :

$$\begin{aligned} \text{4.1) } \quad & \tilde{\Xi} \leftarrow \tilde{\mathbf{B}}^{(j-1)} - \tilde{\mathbf{X}}^\top \{\tilde{\mathbf{W}} \circ \xi(\tilde{\mathbf{X}}\tilde{\mathbf{B}}^{(j-1)}, \tilde{\mathbf{Y}})\}(\mathbf{K}^{(j)})^-, \\ & \mathbf{u}^{(j)} \leftarrow (\tilde{\Xi}[1, :])^\top, \Xi \leftarrow \tilde{\Xi}[2 : \text{end}, :], \text{ and } \mathbf{\Lambda}^{(j)} = \\ & \mathbf{\Lambda} \cdot (\mathbf{K}^{(j)})^-. \end{aligned}$$

4.2) Perform the inner loop iterations, starting with $\mathbf{B}_3 \leftarrow \Xi$, $\mathbf{C}_3 \leftarrow (\mathbf{L}^{(j)}\mathbf{B}_3^\top + \mathbf{B}_3\mathbf{L}^{(j)})/2$, $\mathbf{P} = \mathbf{0}$, $\mathbf{Q}_B = \mathbf{0}$, $\mathbf{Q}_C = \mathbf{0}$, $\mathbf{R} = \mathbf{0}$, and the operators $\mathcal{P}^1, \mathcal{P}^2, \mathcal{P}^3$ defined in Lemmas 4-6:

repeat

- i) $\mathbf{B}_1 \leftarrow \mathcal{P}^1(\mathbf{B}_3 + \mathbf{P}; \mathbf{\Lambda}^{(j)})$, $\mathbf{C}_1 \leftarrow \mathbf{C}_3$, $\mathbf{P} \leftarrow \mathbf{P} + \mathbf{B}_3 - \mathbf{B}_1$.
- ii) $[\mathbf{B}_2, \mathbf{C}_2] \leftarrow \mathcal{P}^2(\mathbf{B}_1 + \mathbf{Q}_B, \mathbf{C}_1 + \mathbf{Q}_C; \mathbf{L}^{(j)})$, $\mathbf{Q}_B \leftarrow \mathbf{Q}_B + \mathbf{B}_1 - \mathbf{B}_2$, $\mathbf{Q}_C \leftarrow \mathbf{Q}_C + \mathbf{C}_1 - \mathbf{C}_2$.
- iii) $\mathbf{B}_3 \leftarrow \mathbf{B}_2$, $\mathbf{C}_3 \leftarrow \mathcal{P}^3(\mathbf{C}_2 + \mathbf{R}; \mathbf{D}^{(j)})$, $\mathbf{R} \leftarrow \mathbf{R} + \mathbf{C}_2 - \mathbf{C}_3$.

until convergence

$$\mathbf{B}^{(j)} \leftarrow \mathbf{B}_3$$

$$\text{4.3) } \tilde{\mathbf{B}}^{(j)} \leftarrow [\mathbf{u}^{(j)} (\mathbf{B}^{(j)})^\top]^\top$$

Algorithm 1 with such modifications in Step 2 and Step 4 is referred to as the \mathbf{S}^3 estimation algorithm.

Theorem 4. Given any initial point $\tilde{\mathbf{B}}^{(0)}$, the \mathbf{S}^3 algorithm converges in the sense that the function value decreasing property holds, and furthermore, $\mathbf{B}^{(j)}$, $\mathbf{L}^{(j)}$, and $\mathbf{D}^{(j)}$ satisfy $\mathbf{L}^{(j)} \succeq \mathbf{0}$, $\mathbf{D}^{(j)} \succeq \mathbf{0}$, $(\mathbf{L}^{(j)}(\mathbf{B}^{(j)})^\top + \mathbf{B}^{(j)}\mathbf{L}^{(j)})/2 \preceq 4\mathbf{D}^{(j)}$ for any $j \geq 1$.

The proof is given in Appendix D.

We observe that practically, the inner loop converges fast (usually within 100 steps). Moreover, to guarantee the functional value is decreasing, one does not have to run the inner loop till convergence. Although it is possible to apply the stable-sparse estimation directly, we recommend running the screening algorithm (RNS) first, followed by the fine \mathbf{S}^3 network learning.

V. EXPERIMENTS

A. Simulation Studies

In this subsection, we conduct synthetic data experiments to demonstrate the performance of the proposed learning framework in recurrent network analysis. An Erdős-Rényi-like scheme of generating system parameters, including a sparse regulation matrix \mathbf{A} , is described as follows. Given any node i , the number of its regulators is drawn from a binomial distribution $\mathcal{B}(n-1, \frac{1}{2n})$. The regulator set S_i is

chosen randomly from the rest $(n - 1)$ nodes (excluding node i itself). If $j \notin S_i$, $a_{ij} = 0$. Otherwise, a_{ij} follows a mixture of two Gaussians $\mathcal{N}(1.5, 0.1^2)$ and $\mathcal{N}(-1.5, 0.1^2)$ with probability $1/2$ for each. Then draw random $l, \mathbf{u}, \mathbf{c}$ from Gaussian distributions (independently) $l_i \sim \mathcal{N}(1.5, 0.1^2)$, $u_i \sim \mathcal{N}(0, 0.1^2)$, $c_i \sim \mathcal{N}(0, 0.1^2)$. Finally \mathbf{d} is generated so that the system satisfies the stability condition (A3a).

Topology identification. First, we test the performance of RNS in recurrent network topology identification. We compare it with TSNI [26] and QTIS [27]. TSNI is a popular network learning approach and applies principle component analysis for dimensionality reduction. QTIS is a network screening algorithm based on sparsity-inducing techniques. To avoid the ad-hoc issue of parameter tuning, ROC curves will be used to summarize link detection results in a comprehensive way, in terms of true positive rate (TPR) and false positive rate (FPR).

We simulate two networks according to the scheme introduced early. In the first example we set $n = 10, T = 100$, in the second example $n = 200, T = 500$, and in the third $n = 100, T = 1000$. Given all system parameters, we can call Matlab functions SDE and SDE.SIMULATE to generate continuous-time stochastic processes according to (4). The discrete-time observations are sampled from the stochastic processes with sampling period $\Delta T = 1$.

In this experiment, the number of unknowns in either case is larger than the sample size, especially in Ex.2 which has about **41K** variables but only 500 observations. Given any algorithm, we vary the target cardinality parameter m from 1 to $n(n-1)$, collect all estimates, and compute their associated TPRs and FPRs. The experiment is repeated for 50 times and the averaged rates are shown in the ROC curves in Figure 2. RNS beats TSNI and QTIS by a large margin. In fact, the ROC curve of RNS is *everywhere* above the TSNI curve and the QTIS curve.

System stability. Next, we show the necessity of stable learning in network dynamics forecasting. For simplicity, the ℓ_1 penalization is used. We compare the \mathbf{S}^3 network estimation with the approach based on sparse sigmoidal regression in Section III-A (with no stability guarantee), denoted by SigSpar.

We use two network examples (Ex.4 and Ex.5) with $n = 20, 40$ respectively. In each setting, we generate $T = 20$ samples for training, and 200 validation samples for parameter tuning. In this experiment, forecast error at a future time point is the major concern. Suppose $\mathbf{x}(T)$, the network snapshot at T , is given. With system parameter estimates obtained, one can simulate a stochastic process $\hat{\mathbf{x}}(t)$ ($t \geq T$) starting with $\mathbf{x}(T)$ based on model (4). The forecast error at time point $T + h$ is defined as $\text{FE} = \|\mathbf{x}(T + h) - \hat{\mathbf{x}}(T + h)\|_2^2/n$. Long-term forecasting corresponds to large values of h . We repeat the experiment for 50 times and show the average FE in Table I. The error of SigSpar becomes extremely large as h increases, because there is no stability guarantee of its network estimate, while \mathbf{S}^3 has excellent performance even in long-term forecasting.

B. Real data

Yeast gene regulatory network. We use RNS to study the transcriptional regulatory network in the yeast cell cycle. The

dataset is publicly available and a detailed description of the microarray experiment is in [28]. Following [29], we focus on the 20 genes whose expression patterns fluctuate in a periodic manner. The dataset contains their expression levels recorded at 18 time points during a cell cycle. In this regulatory (sub)network, five genes have been identified as transcription factors, namely SWI4, HCM1, NDD1, SWI5, ASH1, and 19 regulatory connections from them to target genes have been reported with direct evidence in the literature (cf. the YEASTRACT database at <http://yeastract.com/>). [29] found 55 connections from the transcription factors to the target genes, of which 14 have biological evidence (and so the true positive rate is $14/19 = 73.7\%$). For a fair comparison, we also let RNS detect 55 connections from the transcription factors to the target genes, and achieved a higher true positive rate 89.5%. A detailed comparison of the identified regulatory connections is shown in Table II.

fMRI data. The resting state fMRI data provided by the ADHD-200 consortium [30] have been preprocessed by the NIAK interface [31]. The dataset we are using has 954 ROIs (regions of interest) and 257 time points. In the experiment, the first 200 observations are for training (and so $T = 200$ and $n = 954$), and the following 57 observations are for testing. We applied RNS to get a network pattern, followed by the \mathbf{S}^3 network estimation for stability correction. Table III shows the results averaged over 10 randomly chosen subjects. Our learning algorithm can achieve much lower error than TSNI, and its performance is pretty robust to the choice of m .

APPENDIX A PROOF OF THEOREM 1

For notation simplicity, we introduce $\mu_s(\boldsymbol{\beta}) := \pi(\tilde{\mathbf{x}}_s^\top \boldsymbol{\beta})$. Then $\partial \mu_s(\boldsymbol{\beta}) / \partial \boldsymbol{\beta} = \mu_s(\boldsymbol{\beta})(1 - \mu_s(\boldsymbol{\beta}))\tilde{\mathbf{x}}_s^\top$, from which it follows that

$$\nabla \left(\sum_{s=1}^T w_s (y_s - \mu_s(\boldsymbol{\beta}))^2 / 2 \right) = \tilde{\mathbf{X}}^\top \mathbf{W} \boldsymbol{\xi}(\tilde{\mathbf{X}} \boldsymbol{\beta}, \mathbf{y}),$$

where $\mathbf{W} = \text{diag}\{w_s\}_{s=1}^T$ and $\boldsymbol{\xi}$ is the function defined in (11) applied componentwise. Moreover, we can compute its Hessian (details omitted)

$$\begin{aligned} \mathbf{H} &= D \left(\nabla \left(\sum_{s=1}^T w_s (y_s - \mu_s(\boldsymbol{\beta}))^2 / 2 \right) \right) \\ &= \tilde{\mathbf{X}}^\top \mathbf{W} \text{diag} \{ \mu_s(1 - \mu_s) \\ &\quad [\mu_s(1 - \mu_s) + (\mu_s - y_s)(1 - 2\mu_s)] \} \tilde{\mathbf{X}} \\ &\equiv \tilde{\mathbf{X}}^\top \mathbf{W} \boldsymbol{\Sigma}(\boldsymbol{\beta}) \tilde{\mathbf{X}}. \end{aligned}$$

Note that \mathbf{H} is not necessarily positive semi-definite.

Let $F_0(\boldsymbol{\beta}) = \frac{1}{2} \sum w_s (y_s - \mu_s(\boldsymbol{\beta}))^2 + \sum P(\beta_k; \lambda)$. Define a surrogate function as $G(\boldsymbol{\beta}, \boldsymbol{\beta}') = \frac{1}{2} \sum w_s (y_s - \mu_s(\boldsymbol{\beta}'))^2 + \sum P(\beta'_k; \lambda) + \frac{K}{2} \|\boldsymbol{\beta} - \boldsymbol{\beta}'\|_2^2 + \frac{1}{2} \sum w_s ((y_s - \mu_s(\boldsymbol{\beta}))^2 - (y_s - \mu_s(\boldsymbol{\beta}'))^2) + \sum w_s \xi(\mathbf{x}_s^\top \boldsymbol{\beta}, y_s) (\mathbf{x}_s^\top \boldsymbol{\beta}' - \mathbf{x}_s^\top \boldsymbol{\beta})$. Based on the




Figure 2: Comparison of ROC curves.

Table I: Forecasting error comparison.

	h	1	5	10	15	20
Ex.4	SigSpar	0.08	3.01	41.95	450.54	4.4×10^3
	\mathbf{S}^3	0.04	0.48	1.01	1.39	1.76
Ex.5	SigSpar	0.35	23.76	41.95	7.3×10^3	1.0×10^5
	\mathbf{S}^3	0.06	0.70	1.75	2.80	4.04

Table II: Regulatory connections of the Yeast cell cycle subnetwork. ‘ \times ’ stands for the regulatory connections identified by RNS, ‘_’ stands for those by [29], and ‘ \square ’ stands for the confirmed regulatory connections with published evidence in the literature.

	SWI4	HCM1	NDD1	SWI5	ASH1
SWI4					
HCM1	\times				
NDD1	\times				
SWI5					
ASH1				\times	
CLN2	\times				
SVS1	\times				
SWE1	\times				
MNN1	\times				
CLB6	\times				
HTA1	\times				
HTB1	\times				
HHT1					
CLB4				\times	
CLB2	\times		\times		
CDC20					
SPO12			\square		
SIC1				\times	
CLN3	\times	\times		\times	\square
CDC46					

Table III: Forecast error for the fMRI data.

m/n^2	0.5	0.6	0.7	0.8	0.9
TSNI	0.981	1.021	1.044	1.056	1.060
\mathbf{S}^3	0.194	0.194	0.194	0.194	0.194

previous calculation, we have

$$\begin{aligned} & \frac{1}{2} \sum w_s ((y_s - \mu_s(\boldsymbol{\beta}))^2 - (y_s - \mu_s(\boldsymbol{\beta}'))^2) \\ &= (\tilde{\mathbf{X}}^\top \mathbf{W} \boldsymbol{\xi}(\tilde{\mathbf{X}} \boldsymbol{\beta}, \mathbf{y}))^\top (\boldsymbol{\beta} - \boldsymbol{\beta}') \\ & \quad - \frac{1}{2} (\boldsymbol{\beta} - \boldsymbol{\beta}')^\top (\tilde{\mathbf{X}}^\top \mathbf{W} \boldsymbol{\Sigma}(\theta \boldsymbol{\beta} + (1 - \theta) \boldsymbol{\beta}') \tilde{\mathbf{X}}) (\boldsymbol{\beta} - \boldsymbol{\beta}'), \end{aligned}$$

for some $\theta \in [0, 1]$. Let $\boldsymbol{\zeta} = \theta \boldsymbol{\beta} + (1 - \theta) \boldsymbol{\beta}'$. It follows that

$$\begin{aligned} & \frac{K}{2} \|\boldsymbol{\beta} - \boldsymbol{\beta}'\|_2^2 + \frac{1}{2} \sum w_s ((y_s - \mu_s(\boldsymbol{\beta}))^2 - (y_s - \mu_s(\boldsymbol{\beta}'))^2) \\ & \quad + \sum w_s \boldsymbol{\xi}(\tilde{\mathbf{x}}_s^\top \boldsymbol{\beta}, y_s) (\tilde{\mathbf{x}}_s^\top \boldsymbol{\beta}' - \tilde{\mathbf{x}}_s^\top \boldsymbol{\beta}) \\ &= \frac{1}{2} (\boldsymbol{\beta}' - \boldsymbol{\beta})^\top (K \mathbf{I} - \tilde{\mathbf{X}}^\top \boldsymbol{\Sigma}(\boldsymbol{\zeta}) \mathbf{W} \tilde{\mathbf{X}}) (\boldsymbol{\beta}' - \boldsymbol{\beta}) \\ & \geq \frac{K - \|\tilde{\mathbf{X}}\|_2^2 \|\boldsymbol{\Sigma}(\boldsymbol{\zeta}) \mathbf{W}\|_2}{2} \|\boldsymbol{\beta} - \boldsymbol{\beta}'\|_2^2. \end{aligned}$$

Because

$$\begin{aligned} & w_s \mu_s (1 - \mu_s) ((1 - 2\mu_s)(\mu_s - y_s) + \mu_s(1 - \mu_s)) \\ & \leq \frac{w_s}{4} \left(\frac{1}{2} \left(\frac{1 - 2\mu_s + 2\mu_s - 2y_s}{2} \right)^2 + \frac{1}{4} \right) \\ & = \frac{w_s}{4} \left(\frac{(1 - 2y_s)^2}{8} + \frac{1}{4} \right), \end{aligned}$$

the diagonal entries of $\mathbf{W} \boldsymbol{\Sigma}(\boldsymbol{\zeta})$ are uniformly bounded by

$$\max_s \frac{w_s}{4} \left(\frac{(1 - 2y_s)^2}{8} + \frac{1}{4} \right)$$

or $k_0(\mathbf{y}, \mathbf{w})$ (see (12)). Therefore, choosing $K \geq k_0(\mathbf{y}, \mathbf{w}) \|\tilde{\mathbf{X}}\|_2^2$, we have $G(\boldsymbol{\beta}, \boldsymbol{\beta}') = \frac{1}{2} \sum w_s (y_s - \mu_s(\boldsymbol{\beta}'))^2 + \sum P(\beta'_k; \lambda) + \frac{1}{2} (K - \|\tilde{\mathbf{X}}\|_2^2 \|\mathbf{W} \boldsymbol{\Sigma}\|_2^2) \geq F_0(\boldsymbol{\beta})$ for any $\boldsymbol{\beta}, \boldsymbol{\beta}'$.

On the other hand, based on the definition, it is easy to show (details omitted) that given $\boldsymbol{\beta}$, $\min_{\boldsymbol{\beta}'} G(\boldsymbol{\beta}, \boldsymbol{\beta}')$ is equivalent to:

$$\min_{\boldsymbol{\beta}'} \frac{K}{2} \|\boldsymbol{\beta}' - \boldsymbol{\beta} + \frac{1}{K} \tilde{\mathbf{X}}^\top \mathbf{W} \boldsymbol{\xi}(\tilde{\mathbf{X}} \boldsymbol{\beta}, \mathbf{y})\|_2^2 + \sum P(\beta'_k; \lambda). \quad (25)$$

Applying Lemma 1 in [16] without requiring uniqueness, there exists a globally optimal solution

$$\boldsymbol{\beta}'_o(\boldsymbol{\beta}) = \Theta(\boldsymbol{\beta} - \tilde{\mathbf{X}}^\top \mathbf{W} \boldsymbol{\xi}(\tilde{\mathbf{X}} \boldsymbol{\beta}, \mathbf{y})/K; \lambda'), \quad (26)$$

provided that $P(t; \lambda') = P(t; \lambda)/K$ for any t . In summary, we obtain $F_0(\boldsymbol{\beta}) = G(\boldsymbol{\beta}, \boldsymbol{\beta}) \geq G(\boldsymbol{\beta}, \boldsymbol{\beta}'_o(\boldsymbol{\beta})) \geq F_0(\boldsymbol{\beta}'_o(\boldsymbol{\beta}))$.

We are now in a position to prove (16). In fact, given γ and l , the optimization problem

$$\min_{\boldsymbol{\beta}} \frac{1}{2} \sum w_s (y_s - \mathbf{z}_s^\top \boldsymbol{\gamma} - l \pi(\tilde{\mathbf{x}}_s^\top \boldsymbol{\beta}))^2 + \sum P(\beta_k, \lambda),$$

is equivalent to $\min_{\boldsymbol{\beta}} \frac{1}{2} \sum w'_s (\tilde{y}_s - \pi(\tilde{\mathbf{x}}_s^\top \boldsymbol{\beta}))^2 + \sum P(\beta_k, \lambda)$ with $\tilde{y}_s = (y_s - \mathbf{z}_s^\top \boldsymbol{\gamma})/l$, $w'_s = l^2 w_s$. (Note that the l obtained from WLS is nonzero with probability 1.) Therefore, the function value decreasing property always holds during the iteration.

APPENDIX B PROOF OF THEOREM 2

Define a quantile thresholding rule $\Theta^\#(\cdot; m, \eta)$ as a variant of the hard-ridge thresholding rule (9). Given $1 \leq m \leq np$:

$\mathbf{A} \in \mathcal{R}^{n \times p} \rightarrow \mathbf{B} \in \mathcal{R}^{n \times p}$ is defined as follows: $b_{ij} = a_{ij}/(1 + \eta)$ if $|a_{ij}|$ is among the m largest in the set of $\{|a_{ij}| : 1 \leq i \leq n, 1 \leq j \leq p\}$, and $b_{ij} = 0$ otherwise.

Lemma 1. $\hat{\mathbf{B}} = \Theta^\#(\mathbf{A}; m, \eta)$ is a globally optimal solution to

$$\begin{aligned} & \min_{\mathbf{B}} f_0(\mathbf{A}) = \frac{1}{2} \|\mathbf{A} - \mathbf{B}\|_F^2 + \frac{\eta}{2} \|\mathbf{B}\|_F^2 \\ & \text{s.t. } \|\mathbf{B}\|_0 \leq m. \end{aligned}$$

Proof: Let $I \subset \{(i, j) | 1 \leq i \leq n, 1 \leq j \leq p\}$ with $|I| = m$. Assuming $\mathbf{B}_{I^c} = \mathbf{0}$, we get the optimal solution $\hat{\mathbf{B}}$ with $\hat{\mathbf{B}} = \frac{1}{1+\eta} \mathbf{A}_I$. It follows that $f_0(\hat{\mathbf{B}}) = \frac{1}{2} \|\mathbf{A}\|_F^2 - \frac{1}{2} \sum_{i,j \in I} a_{ij}^2$. Therefore, the quantile thresholding $\Theta^\#(\mathbf{A}; m, \eta)$ yields a global minimizer. ■

Using Lemma 1, we can prove the function value decreasing property; the remaining part follows similar lines of Theorem 1 because of the separability of F .

APPENDIX C PROOF OF THEOREM 3

Let $\mathbf{f}(\mathbf{x}) = \mathbf{L}\boldsymbol{\pi}(\mathbf{A}\mathbf{x} + \mathbf{u}) - \mathbf{D}\mathbf{x} + \mathbf{c}$, where \mathbf{x} is short for $\mathbf{x}(t)$. First, we prove the existence of an equilibrium. It suffices to show that there is a solution to $\mathbf{f}(\mathbf{x}) = \mathbf{0}$ or $\mathbf{x} = \mathbf{D}^{-1} \mathbf{L}\boldsymbol{\pi}(\mathbf{A}\mathbf{x} + \mathbf{u}) + \mathbf{D}^{-1} \mathbf{c} =: \varphi(\mathbf{x})$. Obviously, the mapping φ is continuous and bounded (say $\|\varphi\|_\infty \leq M$), Brouwer's fixed point theorem [32] indicates the existence of at least one equilibrium in $[-M, M]^n$.

Let \mathbf{x}^* be an equilibrium point, i.e., $\mathbf{f}(\mathbf{x}^*) = \mathbf{0}$. Construct a Lyapunov function candidate $V(\mathbf{x}) = \frac{1}{2} (\mathbf{x} - \mathbf{x}^*)^\top \mathbf{P} (\mathbf{x} - \mathbf{x}^*)$ with \mathbf{P} positive definite and to be determined. Then

$$\begin{aligned} & \frac{dV(\mathbf{x})}{dt} = V'(\mathbf{x}) \mathbf{f}(\mathbf{x}) \\ &= (\mathbf{x} - \mathbf{x}^*)^\top \mathbf{P} \mathbf{f}(\mathbf{x}) \\ &= (\mathbf{x} - \mathbf{x}^*)^\top \mathbf{P} (\mathbf{f}(\mathbf{x}) - \mathbf{f}(\mathbf{x}^*)) \\ &= -(\mathbf{x} - \mathbf{x}^*)^\top \mathbf{P} \mathbf{D} (\mathbf{x} - \mathbf{x}^*) \\ & \quad + (\mathbf{x} - \mathbf{x}^*)^\top \mathbf{P} \mathbf{L} (\boldsymbol{\pi}(\mathbf{A}\mathbf{x} + \mathbf{u}) - \boldsymbol{\pi}(\mathbf{A}\mathbf{x}^* + \mathbf{u})) \\ &= -(\mathbf{x} - \mathbf{x}^*)^\top (\mathbf{P} \mathbf{D} - \mathbf{P} \mathbf{L} \mathbf{G} \mathbf{A}) (\mathbf{x} - \mathbf{x}^*) \\ &= -(\mathbf{x} - \mathbf{x}^*)^\top \left(\frac{\mathbf{P} \mathbf{D} + \mathbf{D} \mathbf{P}^\top}{2} - \frac{\mathbf{P} \mathbf{L} \mathbf{G} \mathbf{A} + \mathbf{A}^\top \mathbf{G} \mathbf{L} \mathbf{P}^\top}{2} \right) \\ & \quad \cdot (\mathbf{x} - \mathbf{x}^*), \end{aligned}$$

where

$$\mathbf{G} = \text{diag} \left\{ \frac{\pi(\boldsymbol{\alpha}_i^\top \mathbf{x} + \mathbf{u}) - \pi(\boldsymbol{\alpha}_i^\top \mathbf{x}^* + \mathbf{u})}{\boldsymbol{\alpha}_i^\top \mathbf{x}_i - \boldsymbol{\alpha}_i^\top \mathbf{x}_i^*} \right\}.$$

It is easy to verify that $\mathbf{G} = \text{diag}\{\boldsymbol{\pi}'(\boldsymbol{\xi})\} \preceq \mathbf{I}/4$, and thus $\mathbf{P} \mathbf{L} \mathbf{G} \mathbf{A} + \mathbf{A}^\top \mathbf{G} \mathbf{L} \mathbf{P}^\top \preceq (\mathbf{P} \mathbf{L} \mathbf{A} + \mathbf{A}^\top \mathbf{L} \mathbf{P}^\top)/4$. It is well known [21] that under (A3a), the Lyapunov equation

$$\mathbf{P} \left(\mathbf{D} - \frac{\mathbf{L} \mathbf{A}}{4} \right) + \left(\mathbf{D} - \frac{\mathbf{L} \mathbf{A}}{4} \right)^\top \mathbf{P}^\top = -\mathbf{R} \quad (27)$$

is solvable and uniquely determines a positive definite \mathbf{P} for any positive definite \mathbf{R} . Therefore, V is indeed a Lyapunov function for the nonlinear dynamical system (22). Moreover,

(A3a) implies

$$\frac{dV(\mathbf{x})}{dt} \leq -\varepsilon_0 \|\mathbf{x} - \mathbf{x}^*\|_2^2 \leq -\varepsilon V(\mathbf{x})$$

for some $\varepsilon_0, \varepsilon > 0$. By the Lyapunov stability theory—see, e.g., [33] (Chapter 3), (22) must be globally exponentially stable. The uniqueness of the equilibrium is implied by the global exponential stability.

The second result can be shown by setting $\mathbf{P} = \mathbf{I}$ in (27); details are omitted.

APPENDIX D PROOF OF THEOREM 4

Based on the proof of Theorem 1, the modified Step 2 does not affect the convergence of function value because at each iteration a global optimum of (24) is obtained. It remains to show that $\tilde{\mathbf{B}}^{(j)}$ generated by the modified Step 4 improves $\tilde{\mathbf{B}}^{(j-1)}$ in terms of reducing the objective function value, and $\mathbf{B}^{(j)}$ obeys the stability condition.

Let $\tilde{\Xi} = \tilde{\mathbf{B}}^{(j-1)} - \tilde{\mathbf{X}}^\top \mathbf{W} \xi(\tilde{\mathbf{X}} \tilde{\mathbf{B}}^{(j-1)}, \tilde{\mathbf{Y}})(\mathbf{K}^{(j)})^-, \Lambda^{(j)} = \Lambda \cdot (\mathbf{K}^{(j)})^-$. (The modified Step 2 may result in zeros in $\tilde{\Xi}$ to make the associated activation terms in $\pi(\tilde{\mathbf{X}} \tilde{\mathbf{B}})$ vanish; using the pseudoinverse can handle the issue and maintain the decreasing property.) Based on the argument of Appendix A, the problem in Step 4 reduces to

$$\min_{\tilde{\mathbf{B}}=[\mathbf{u}, \mathbf{B}^\top]^\top} \frac{1}{2} \|\tilde{\mathbf{B}} - \tilde{\Xi}\|_F^2 + \|\Lambda^{(j)} \circ \mathbf{B}\|_1, \quad (28)$$

s.t. $(\mathbf{L}^{(j)} \mathbf{B}^\top + \mathbf{B} \mathbf{L}^{(j)})/2 \preceq 4\mathbf{D}^{(j)}$.

Therefore it suffices to prove Lemma 2 for any given $\Xi (= \tilde{\Xi}[2 : \text{end}, :])$, $\Lambda^{(j)}$, $\mathbf{L}^{(j)} \succeq \mathbf{0}$, $\mathbf{D}^{(j)} \succeq \mathbf{0}$, and $\mathbf{B}^{(j-1)}$. (In fact, the lemma holds given any initialization of \mathbf{C}_3 .)

Lemma 2. For any $j \geq 1$, $\frac{1}{2} \|\mathbf{B}^{(j)} - \Xi\|_F^2 + \|\Lambda^{(j)} \circ \mathbf{B}^{(j)}\|_1 \leq \frac{1}{2} \|\mathbf{B}^{(j-1)} - \Xi\|_F^2 + \|\Lambda^{(j)} \circ \mathbf{B}^{(j-1)}\|_1$, and $(\mathbf{L}^{(j)} \mathbf{B}^{(j)\top} + \mathbf{B}^{(j)} \mathbf{L}^{(j)})/2 \preceq 4\mathbf{D}^{(j)}$.

Proof: Define $f(\mathbf{B}) = \frac{1}{2} \|\mathbf{B} - \Xi\|_F^2 + \|\Lambda \circ \mathbf{B}\|_1$, and $g(\mathbf{B}', \mathbf{C}', \mathbf{B}, \mathbf{C}) = \frac{1}{2} \|\mathbf{B}' - \Xi\|_F^2 + \|\Lambda \circ \mathbf{B}\|_1 + \frac{1}{2} \|\mathbf{B} - \mathbf{B}'\|_F^2 + \frac{1}{2} \|\mathbf{C} - \mathbf{C}'\|_F^2 + \langle \mathbf{B}' - \Xi, \mathbf{B} - \mathbf{B}' \rangle$. Then $f(\mathbf{B}') = g(\mathbf{B}', \mathbf{C}', \mathbf{B}', \mathbf{C}')$ and $g(\mathbf{B}', \mathbf{C}', \mathbf{B}, \mathbf{C}) \geq f(\mathbf{B})$.

On the other hand, given $(\mathbf{B}', \mathbf{C}')$, we can write g as a function of (\mathbf{B}, \mathbf{C}) : $\frac{1}{2} \|\Xi, \mathbf{C}'\| - [\mathbf{B}, \mathbf{C}]\|_F^2 + \|\Lambda \circ \mathbf{B}\|_1$ (up to additive functions of \mathbf{B}' and \mathbf{C}'). Based on Lemma 3, with $\Xi_C = \mathbf{C}'$, $g(\mathbf{B}', \mathbf{C}', \mathbf{B}', \mathbf{C}') \geq g(\mathbf{B}', \mathbf{C}', \mathbf{B}^o, \mathbf{C}^o) \geq f(\mathbf{B}^o)$ and $(\mathbf{L} \mathbf{B}^o \mathbf{L}^\top + \mathbf{B}^o \mathbf{L})/2 = \mathbf{C}^o \preceq 4\mathbf{D}$.

Applying the result to the modified Step 4 in Section IV-B yields the desired result. \blacksquare

Lemma 3. Consider the sequence of $(\mathbf{B}_3, \mathbf{C}_3)$ generated by the following procedure, with the operators \mathcal{P}^1 , \mathcal{P}^2 , \mathcal{P}^3 defined in Lemma 4, Lemma 5 and Lemma 6, respectively:

- 0) $\mathbf{B}_3 \leftarrow \Xi_B, \mathbf{C}_3 \leftarrow \Xi_C, \mathbf{P} = \mathbf{0}, \mathbf{Q}_B = \mathbf{0}, \mathbf{Q}_C = \mathbf{0}, \mathbf{R} = \mathbf{0}$
- repeat
- 1) $\mathbf{B}_1 \leftarrow \mathcal{P}^1(\mathbf{B}_3 + \mathbf{P}; \Lambda), \mathbf{C}_1 \leftarrow \mathbf{C}_3, \mathbf{P} \leftarrow \mathbf{P} + \mathbf{B}_3 - \mathbf{B}_1$

- 2) $[\mathbf{B}_2, \mathbf{C}_2] \leftarrow \mathcal{P}^2(\mathbf{B}_1 + \mathbf{Q}_B, \mathbf{C}_1 + \mathbf{Q}_C; \mathbf{L}), [\mathbf{Q}_B, \mathbf{Q}_C] \leftarrow [\mathbf{Q}_B, \mathbf{Q}_C] + [\mathbf{B}_1, \mathbf{C}_1] - [\mathbf{B}_2, \mathbf{C}_2]$
- 3) $\mathbf{B}_3 \leftarrow \mathbf{B}_2, \mathbf{C}_3 \leftarrow \mathcal{P}^3(\mathbf{C}_2 + \mathbf{R}; \mathbf{D}), \mathbf{R} \leftarrow \mathbf{R} + \mathbf{C}_2 - \mathbf{C}_3$.

until convergence

Then, the sequence of iterates converges to a globally optimal solution $(\mathbf{B}^o, \mathbf{C}^o)$ to

$$\min_{\mathbf{B}, \mathbf{C}} \frac{1}{2} \|[\Xi_B, \Xi_C] - [\mathbf{B}, \mathbf{C}]\|_F^2 + \|\Lambda \circ \mathbf{B}\|_1 \quad (29)$$

s.t. $\mathbf{C} = (\mathbf{L} \mathbf{B}^\top + \mathbf{B} \mathbf{L})/2, \mathbf{C} \preceq 4\mathbf{D}$.

Proof: With the following three lemmas, applying Theorem 3.2 and Theorem 3.3 in [34] yields the strict convergence of the iterates and the global optimality of the limit point. \blacksquare

Lemma 4. Let $\mathcal{P}^1(\Phi)$ be the optimal solution to

$$\min_{\mathbf{B}} \frac{1}{2} \|\mathbf{B} - \Phi\|_F^2 + \|\Lambda \circ \mathbf{B}\|_1. \quad (30)$$

Then $\mathcal{P}^1(\Phi; \Lambda) = \Theta_S(\Phi; \Lambda)$ where Θ_S , applied component-wise on Φ , is the soft-thresholding rule given in Section III-A.

Proof: Apply Lemma 1 in [16]. \blacksquare

Lemma 5. The optimal solution to

$$\min_{\mathbf{B}, \mathbf{C}} \frac{1}{2} \|[\mathbf{B}, \mathbf{C}] - [\Phi_B, \Phi_C]\|_F^2 \text{ s.t. } \mathbf{C} = (\mathbf{L} \mathbf{B}^\top + \mathbf{B} \mathbf{L})/2. \quad (31)$$

is given by $\mathcal{P}^2(\Phi_B, \Phi_C; \mathbf{L}) = [\mathbf{B}^o, \mathbf{C}^o]$ with $\mathbf{B}^o = [b_{i,j}^o]$, $b_{i,j}^o = \psi_{i,j} \frac{2+l_i^2}{2+l_i^2+l_j^2} - \psi_{j,i} \frac{l_i l_j}{2+l_i^2+l_j^2}$, and $\mathbf{C}^o = (\mathbf{L} \mathbf{B}^o \mathbf{L}^\top + \mathbf{B}^o \mathbf{L})/2$, where $\Psi = [\psi_{i,j}] = \Phi_B + (\Phi_C + \Phi_C^\top) \mathbf{L}/2$.

Proof: Let $f(\mathbf{B}) = \|\mathbf{B} - \Phi_B\|_F^2/2 + \|(\mathbf{L} \mathbf{B}^\top + \mathbf{B} \mathbf{L})/2 - \Phi_C\|_F^2/2$. It is not difficult to obtain the gradient (details omitted): $\mathbf{B} - \Phi_B + (\mathbf{L} \mathbf{B}^\top \mathbf{L} + \mathbf{B} \mathbf{L}^2)/2 - (\Phi_C + \Phi_C^\top) \mathbf{L}/2$. The optimal \mathbf{B} and \mathbf{C} can be evaluated accordingly. \blacksquare

Lemma 6. Let $\mathcal{P}^3(\Phi; \mathbf{D})$ be the optimal solution to

$$\min_{\mathbf{C}} \frac{1}{2} \|\mathbf{C} - \Phi\|_F^2 \text{ s.t. } \mathbf{C} \text{ is symmetric and satisfies } \mathbf{C} \preceq 4\mathbf{D}. \quad (32)$$

Then it is given by $\mathbf{U} \text{diag}\{\min(s_i, 0)\} \mathbf{U}^\top + 4\mathbf{D}$, where $\mathbf{S} = \text{diag}\{s_1, \dots, s_n\}$ and \mathbf{U} are from the spectral decomposition of $(\Phi + \Phi^\top)/2 - 4\mathbf{D} = \mathbf{U} \mathbf{S} \mathbf{U}^\top$.

Proof: Because \mathbf{C} is symmetric (but Φ may not be), we have

$$\begin{aligned} & \|\mathbf{C} - \Phi\|_F^2 \\ &= \sum_{1 \leq i < j \leq p} [(c_{ij} - \phi_{ij})^2 + (c_{ij} - \phi_{ji})^2] + \sum_{i=1}^p (c_{ii} - \phi_{ii})^2 \\ &= \sum_{1 \leq i < j \leq p} 2(c_{ij} - \frac{\phi_{ij} + \phi_{ji}}{2})^2 + \sum_{i=1}^p (c_{ii} - \phi_{ii})^2 + \text{const}(\Phi) \\ &= \|\mathbf{C} - \frac{\Phi + \Phi^\top}{2}\|_F^2 + \text{const}(\Phi), \end{aligned} \quad (33)$$

where $\text{const}(\Phi)$ is a term that does not depend on \mathbf{C} .

Therefore, problem (32) is equivalent to

$$\min_C \frac{1}{2} \|C - \frac{\Phi + \Phi^T}{2}\|_F^2, \text{ s.t. } C - 4D \preceq 0. \quad (34)$$

The optimality of $\mathcal{P}^3(\Phi; D)$ can then be argued by von Neumann's trace inequality [35, 36]. ■

ACKNOWLEDGEMENT

The authors are grateful to the associate editor and the two anonymous referees for their careful comments and useful suggestions.


REFERENCES

- [1] J. J. Faith, B. Hayete, J. T. Thaden, I. Mogno, J. Wierzbowski, G. Cottarel, S. Kasif, J. J. Collins, and T. S. Gardner, "Large-scale mapping and validation of escherichia coli transcriptional regulation from a compendium of expression profiles," *PLoS Biol.*, vol. 5, no. 1, p. e8, Jan. 2007.
- [2] T. Mills and R. Markellos, *The econometric modelling of financial time series*. Cambridge University Press, 2008.
- [3] E. Bullmore and O. Sporns, "Complex brain networks: graph theoretical analysis of structural and functional systems," *Nature Reviews Neuroscience*, vol. 10, pp. 186–198, Mar. 2009.
- [4] S. Hanneke, W. Fu, and E. P. Xing, "Discrete temporal models of social networks," *Electronic Journal of Statistics*, vol. 4, pp. 585–605, 2010.
- [5] A. Roebroeck, E. Formisano, and R. Goebel, "Mapping directed influence over the brain using granger causality and fmri," *Neuroimage*, vol. 25, no. 1, pp. 230–242, 2005.
- [6] R. D. Beer, "The dynamics of adaptive behavior: A research program," *Robotics and Autonomous Systems*, vol. 20, no. 2, pp. 257–289, 1997.
- [7] J. C. Gallacher and J. M. Fiore, "Continuous time recurrent neural networks: a paradigm for evolvable analog controller circuits," in *National Aerospace and Electronics Conference 2000, Proceedings of the IEEE 2000*. IEEE, 2000, pp. 299–304.
- [8] H. Mayer, F. Gomez, D. Wierstra, I. Nagy, A. Knoll, and J. Schmidhuber, "A system for robotic heart surgery that learns to tie knots using recurrent neural networks," in *Intelligent Robots and Systems, 2006 IEEE/RSJ International Conference on*, 2006, pp. 543–548.
- [9] R. Xu, D. Wunsch II, and R. Frank, "Inference of genetic regulatory networks with recurrent neural network models using particle swarm optimization," *IEEE/ACM Transactions on Computational Biology and Bioinformatics (TCBB)*, vol. 4, no. 4, pp. 681–692, 2007.
- [10] T. T. Vu and J. Vohradsky, "Nonlinear differential equation model for quantification of transcriptional regulation applied to microarray data of saccharomyces cerevisiae," *Nucleic acids research*, vol. 35, no. 1, pp. 279–287, 2007.
- [11] S. C. Kou, B. P. Olding, M. Lysy, and J. S. Liu, "A multiresolution method for parameter estimation of diffusion processes," *Journal of the American Statistical Association*, vol. 107, no. 500, pp. 1558–1574, 2012.
- [12] A. Fujita, J. Sato, H. Garay-Malpartida, R. Yamaguchi, S. Miyano, M. Sogayar, and C. Ferreira, "Modeling gene expression regulatory networks with the sparse vector autoregressive model," *BMC Systems Biology*, vol. 1, no. 1, 2007.
- [13] C. W. J. Granger, "Investigating causal relations by econometric models and cross-spectral methods," *Econometrica*, vol. 37, no. 3, pp. 424–438, 1969.
- [14] V. Bally and D. Talay, "The law of the Euler scheme for stochastic differential equations," *Probability theory and related fields*, vol. 104, no. 1, pp. 43–60, 1996.
- [15] W. James and C. Stein, "Estimation with quadratic loss," in *Proceedings of the fourth Berkeley symposium on mathematical statistics and probability*, vol. 1, no. 1961, 1961, pp. 361–379.
- [16] Y. She, "An iterative algorithm for fitting nonconvex penalized generalized linear models with grouped predictors," *Computational Statistics and Data Analysis*, vol. 9, pp. 2976–2990, 2012.
- [17] H. Zou and T. Hastie, "Regularization and variable selection via the elastic net," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 67, no. 2, pp. 301–320, 2005.
- [18] M. Y. Park and T. Hastie, "L1-regularization path algorithm for generalized linear models," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 69, no. 4, pp. 659–677, 2007.
- [19] Y. She and A. Owen, "Outlier detection using nonconvex penalized regression," *Journal of the American Statistical Association*, vol. 106, no. 494, pp. 626–639, 2011.
- [20] S. Sastry, *Nonlinear systems: analysis, stability, and control*. Springer New York, 1999, vol. 10.
- [21] J. P. LaSalle and Z. Artstein, *The stability of dynamical systems*. SIAM, 1976, vol. 25.
- [22] A. Lyapunov, *The General Problem of the Stability Of Motion*. Kharkov Mathematical Society, 1892.
- [23] —, *The General Problem of the Stability Of Motion*, ser. Control Theory and Applications Series. Taylor & Francis, 1992.
- [24] M. A. Cohen and S. Grossberg, "Absolute stability of global pattern formation and parallel memory storage by competitive neural networks," *Systems, Man and Cybernetics, IEEE Transactions on*, vol. SMC-13, no. 5, pp. 815–826, 1983.
- [25] P. Gill, W. Murray, and M. Wright, *Practical optimization*. Academic Press, 1981.
- [26] M. Bansal, G. Della Gatta, and D. Di Bernardo, "Inference of gene regulatory networks and compound mode of action from time course gene expression profiles," *Bioinformatics*, vol. 22, no. 7, pp. 815–822, 2006.
- [27] Y. He, Y. She, and D. Wu, "Stationary-sparse causality network learning," *Journal of Machine Learning Research*, vol. 14, pp. 3073–3104, 2013.
- [28] P. T. Spellman, G. Sherlock, M. Q. Zhang, V. R. Iyer, K. Anders, M. B. Eisen, P. O. Brown, D. Botstein, and B. Futcher, "Comprehensive identification of cell cycle-regulated genes of the yeast saccharomyces cerevisiae by microarray hybridization," *Molecular biology of the cell*,

- vol. 9, no. 12, pp. 3273–3297, 1998.
- [29] H.-C. Chen, H.-C. Lee, T.-Y. Lin, W.-H. Li, and B.-S. Chen, “Quantitative characterization of the transcriptional regulatory network in the yeast cell cycle,” *Bioinformatics*, vol. 20, no. 12, pp. 1914–1927, 2004.
- [30] M. P. Milham, D. Fair, M. Mennes, and S. H. Mostofsky, “The adhd-200 consortium: a model to advance the translational potential of neuroimaging in clinical neuroscience,” *Frontiers in Systems Neuroscience*, vol. 6, p. 62, 2012.
- [31] S. Lavoie-Courchesne, P. Rioux, F. Chouinard-Decorte, T. Sherif, M.-E. Rousseau, S. Das, R. Adalat, J. Doyon, C. Craddock, D. Margulies *et al.*, “Integration of a neuroimaging processing pipeline into a pan-canadian computing grid,” in *Journal of Physics: Conference Series*, vol. 341, no. 1. IOP Publishing, 2012, p. 012032.
- [32] L. E. J. Brouwer, “Über abbildung von mannigfaltigkeiten,” *Mathematische Annalen*, vol. 71, no. 1, pp. 97–115, 1911.
- [33] W. M. Haddad and V. Chellaboina, *Nonlinear dynamical systems and control: a Lyapunov-based approach*. Princeton University Press, 2008.
- [34] H. H. Bauschke and P. L. Combettes, “A Dykstra-like algorithm for two monotone operators,” *Pacific Journal of Optimization*, vol. 4, no. 3, pp. 383–391, Sep. 2008.
- [35] J. von Neumann, “Some matrix inequalities and metrization of matrix space,” *Tomsk. Univ. Rev.*, vol. 1, pp. 153–167, 1937.
- [36] Y. She, “Reduced rank vector generalized linear models for feature extraction,” *Statistics and Its Interface*, vol. 6, pp. 197–209, 2013.



Dapeng Wu (S’98–M’04–SM’06–F’13) received Ph.D. in Electrical and Computer Engineering from Carnegie Mellon University, Pittsburgh, PA, in 2003. He is a professor at the Department of Electrical and Computer Engineering, University of Florida, Gainesville, FL. His research interests are in the areas of networking, communications, signal processing, computer vision, machine learning, smart grid, and information and network security.



Yiyuan She received B.S. in Mathematics and M.S. in Computer Science from Peking University in 2000 and 2003, respectively, and received Ph.D. in Statistics from Stanford University in 2008. He is currently an Associate Professor in the Department of Statistics at Florida State University. His research interests include high-dimensional statistics, machine learning, multivariate statistics, robust statistics, statistics computing, and network science.



Yuejia He received Ph.D. degree in Electrical and Computer Engineering from University of Florida, Gainesville, FL, in 2013. Her research interest is mainly in large scale network learning.